Small-Scale Mentoring, Large-Scale Innovation: Evidence from a Superstar Firm

Colleen Chien* Jillian Grennan† Jason Sandvik‡

Abstract

Exploiting the staggered rollout of an accelerator-style mentorship program at a superstar technology firm, we examine the implications of the program for innovative productivity and inclusion. Novel idea-level data for 20,000+ engineers show that mentees produce more ideas in the short- and long-term, with proportionally larger gains for ideas that are patentable. Underrepresented mentees exhibit larger increases. Non-mentored engineers who later collaborate with mentees also increase productivity. Mechanism tests suggest improved knowledge flows and expanded networks, rather than reallocated effort. Perceived and behavioral measures shift: stronger culture, satisfaction, and retention. Conservative valuations exceed costs, and scalability tests support external deployability.

JEL: J16, J24, M14, M54, O31, O32, O34

Keywords: Innovation, Mentoring, Organizational Economics, Human Capital, Intellectual Property, Corporate Finance, Corporate Culture, Labor & Finance

*Colleen Chien, University of California Berkeley, School of Law, cchien@berkeley.edu
†Emory University, Economics Department, jillian.grennan@emory.edu
‡University of Arizona, Eller College of Management, sandvik@arizona.edu

Innovation and research productivity have slowed markedly.[1] Even among superstar firms that historically powered U.S. productivity growth and now dominate R&D investment, knowledge diffusion is slowing and economy-wide gains have weakened (Autor et al., 2020; Kalyani et al., 2025), suggesting that there is unused inventive capacity within organizations. Where might that inventive capacity be trapped? External barriers such as unequal early-life opportunities (Cook, 2018; Bell et al., 2019) and frictions in patenting and commercialization (Aneja et al., 2024; Dossi, 2024; Mezzanotti and Simcoe, 2025; Cook et al., 2025; Hochberg et al., 2026) shape who and what enters the innovation pipeline. But inside firms, management and organizational practices remain central to translating initial ideas into inventive productivity and patents (Bloom et al., 2012a, 2016).

As firms grow, internal frictions can impede knowledge flows, erode innovative culture, and limit breakthrough output (Sandvik et al., 2020; Guiso et al., 2015).[2] Some firms respond with defensive strategies such as killer acquisitions and patent blocking (Cunningham et al., 2021; Abrams et al., 2023), rather than internal invention, further dampening aggregate innovation. Many organizational frictions stand between an initial insight and a granted patent. Engineers must know that a disclosure system exists and that their idea is patent-eligible; they must articulate the invention in the vocabulary that the firm's Patent Review Board approves; they need senior colleagues to vouch for the idea's novelty and commercial relevance; and they must be willing to expose an unpolished idea to bureaucratic scrutiny amid career concerns. Whether informational, linguistic, relational, or psychological, each friction disproportionately constrains engineers without prior patenting experience. Alleviating these frictions offers a scalable path to unlocking latent inventive capacity.

Whether mentorship can reduce these frictions in a superstar firm setting is, ex ante, unclear. Existing evidence that mentoring improves worker outcomes comes from settings

[1] This slowdown has been attributed to diminishing returns to inventive effort (Jones, 2009; Bloom et al., 2020), declining business dynamism, and weaker allocative efficiency (Goldin et al., 2024; Decker et al., 2020). Syverson (2017) rules out mismeasurement as an explanation for the slowdown.

[2] A salient example is Meta's "Year of Efficiency" initiative, where CEO Mark Zuckerberg explicitly acknowledged that they were less innovative because of bureaucratic impediments to information flows.

in which the final output is accessible and doable, such as call centers and sales jobs, so effort, confidence, and encouragement become the binding constraints (Rockoff, 2008; Lyle and Smith, 2014; Sandvik et al., 2025). Patentable inventions are different. A commonly held view is that inventive ability is a scarce, largely innate capacity and "either you have it, or you don't," which means that mentoring, however well designed, cannot generate new ideas that were not already there. Under this view, mentorship might raise satisfaction or retention but not change who invents or how much they produce. The contrasting view, motivated by the frictions we describe above, holds that a large share of the observed gap between innovators at firms and named inventors on patents reflects malleable organizational barriers, not immutable creative endowments, in which case a program that transfers tacit patenting know-how and expands professional networks could plausibly expand who invents and inventive output. Our setting allows us to distinguish between these views.

We study whether a structured mentorship program can reduce organizational frictions in the idea-to-patent pipeline within a leading technology firm valued at over $400 billion. A key advantage of our setting is access to the firm's full universe of pre-patent ideas, covering 66,397 idea disclosures submitted by 29,211 innovators, including many abandoned ideas.[3] These abandoned ideas are not economic dead ends: they represent productive search draws over the space of possible inventions, often provide the raw material for subsequent successful disclosures, and may be protected through other means such as trade secrets (Weitzman, 1998; Fleming, 2001; Heath and Mace, 2020). Observing this activity lets us study how mentorship influences both creative ideation and patenting outcomes. Finally, we link these data to demographic records, collaboration networks, internal culture measures, and retention outcomes, allowing us to provide a detailed description of the channels through which the mentorship intervention is linked to value creation.

[3]Ideas are disclosed through an online form and are highly heterogeneous in both substance and presentation. Some ideas read like polished technical memos, while others appear as rough sketches or bullet points. As a result, disclosures vary meaningfully in clarity, completeness, novelty, and technical maturity, including whether they include a problem statement, describe a proposed solution, or even have prototypes or drawings.

The mentorship program was rolled out in seven cohorts from 2020Q4 to 2023Q3, during which the firm operated under a flexible hybrid-work policy typical of large technology firms. Four or five mentees worked with one or two experienced inventors for three to four months. The program was widely publicized, and mentees and mentors had to indicate their interest to participate. Program administrators created mentor-mentee groups randomly after conditioning on technology area and meeting time preferences. Due to capacity constraints, some applicants were not selected for the first program cohort they applied to and instead were encouraged to participate in later cohorts. The program's design included features associated with best practices for human capital investment, including results-oriented meetings with feedback (Bloom and Reenen, 2007), structured accountability (Gonzalez-Uribe et al., 2026), peer learning and education (Hochberg and Fehder, 2015; Gonzalez-Uribe and Leatherbee, 2017), role models (Adams and Lowry, 2022), and a culminating event (Cohen et al., 2019).

For any field experiment that involves voluntary opt-in decisions from participants, one may worry about selection. We address selection concerns in several ways. First, we surveyed participants and non-participants and found that women were more likely to participate, but they were also more skeptical about the program's expected benefits, suggesting participation was not driven by enthusiasm bias.[4] Then, we evaluate and show balance in baseline idea submission patterns across cohorts. Then, we exploit the program's staggered adoption to compare mentored innovators to their not-yet-mentored peers using modern, stacked difference-in-differences (DiD) estimators. Parallel trends tests reveal no pre-trend divergence in idea submissions between the treated and control groups. Finally, we isolate like-for-like comparisons to assuage concerns about differential trends in unobserved innovation opportunities that less-represented groups may face (Ewens and Townsend, 2020).

To interpret the program's potential impact, we develop a theoretical framework that distinguishes between three hypotheses. Under the time-shift hypothesis, mentoring temporar-

[4] Sandvik et al. (2025) document a negative relation between selection into mentorship and realized gains, wherein those with a high desire to opt *out* of the program actually stand to benefit the most from mentorship. If a similar type of self-selection operated in our setting, it would bias against finding positive returns to the program.

ily reallocates effort toward patenting activities without changing underlying state variables, implying short-lived patenting gains during the program but no persistence after it ends. Under the innate-ability hypothesis, mentoring should produce little movement in inventive output, particularly among engineers with no prior patenting experience. In contrast, the process-learning hypothesis implies that mentoring transmits tacit knowledge and expands professional networks that expose innovators to diverse knowledge inputs, generating persistent gains in inventive output and a shift in the composition of ideas disclosed toward the more consistently high-quality mix produced by experienced inventors. Our model assumes that a person's innate creative ability is constant.

We also examine whether targeted mentoring narrows the "innovator-inventor gap," whereby women and underrepresented minorities are far less likely to become named inventors despite meaningful representation as STEM professionals (NCSES, 2023; Chien and Grennan, 2024). We test whether targeted mentorship can transform non-patenting STEM workers ("innovators") into named patent contributors ("inventors"), thereby improving both inventive productivity and inclusion.

Our analyses yield five main results. First, mentorship is associated with increases in innovative output in both the short and long run, consistent with the process-learning hypothesis and inconsistent with the time-shift hypothesis, wherein productivity gains would fade after the program ends. The point estimates suggest mentees produce roughly one additional disclosure per year in the short-term, and the difference persists after the program ends, with mentees continuing to generate about 0.33 additional disclosures annually. The economic magnitude represents a doubling of their baseline rates and exceeds comparable estimates of peer spillovers (Herbst and Mas, 2015). Decomposing the volume gains by idea quality, we find that mentees produce more ideas that both advance to patent applications and those that do not, but with proportionally larger gains for ideas approved for patenting. This finding is consistent with mentees, through process learning, being able to shift the mix of ideas they disclose to match the mix of ideas experienced inventors produce.

These results are similar across alternative estimators and methods for accounting for potential interference and collaboration-based exposure. Our robustness checks include explicitly removing future collaborators from the control group, restricting the control group to eventual mentees, accounting for cohort compositional changes through matching estimators and like-for-like comparisons, and clustering standard errors by mentor (or mentorship team) rather than by innovator. Importantly, normalizing by team size yields similar results, suggesting gains do not simply reflect mentees joining existing teams. The estimates also survive when we remove disclosures in which the mentor is listed as an inventor, indicating that mentees became more capable of producing ideas independently rather than merely joining their mentor's projects. We also provide additional evidence that helps mitigate "diversity-washing" concerns.[5] Finally, our main inferences are unchanged after making multiple-hypothesis-testing adjustments (Harvey et al., 2015).

Second, heterogeneity tests suggest that mentorship is associated with larger productivity gains for innovators with the greatest informational deficits and for those paired with the most capable mentors. Female and URM innovators, who have historically faced greater barriers to becoming named inventors, experience larger productivity gains than highly represented peers,[6] submitting 1.30 versus 1.00 additional disclosures annually. Affinity tests suggest the gains are due to reductions in information frictions, and not due to shared identity alone, as the gains are not confined to mentees with demographically similar mentors. Mentees without prior disclosure experience exhibit the largest and most persistent increases, as do mentees paired with high-ability mentors, as measured by pre-program inventive output. Notably, mentees matched with below-median mentors also show statistically significant increases. These patterns are consistent with scalable knowledge transfer.

[5] In principle, there are strict standards for inventorship, and patent attorneys would put the firm's patent in jeopardy for not accurately naming the inventors. While penalties have declined since the Leahy–Smith America Invents Act signed into law in 2011 for wrongful naming of inventors, there are still deterrents for misleading, and patent attorneys have been sued for adjusting claims to change inventorship.

[6] URM innovators are defined as those who are neither Asian nor White, the two most represented groups among inventors at the firm and its peers (Cook, 2018). The combination of female and URM innovators is defined as underrepresented groups (URGs).

Third, survey evidence helps to clarify the underlying mechanism, pointing toward mentorship improving knowledge about patenting, expanding professional networks, and increasing job satisfaction. To validate these stated mechanisms, we use our unique dataset to track subsequent collaborations and examine whether collaboration patterns are consistent with broader network expansion. We find that mentored innovators collaborate with more diverse inventor teams after participating in the mentorship program relative to pre-program baselines. Team diversity rises along many dimensions, including gender, race, and prior patenting success. These attributes are important for thinking about potential financial returns to the program, given recent evidence linking diversity to higher-value patents and reduced "homophily traps" (Bernstein et al., 2022b; Heath et al., 2023).

Fourth, we examine whether the program is accompanied by broader productivity changes among non-participants. We find that non-mentees who later collaborate with mentees increase their innovation efforts after those collaborations begin. Interestingly, the estimated increases are larger for well-represented groups, which is consistent with the view that informal information flows may operate through existing professional networks. These patterns should be interpreted cautiously, though, as such spillovers are difficult to isolate causally. While the patterns are suggestive of knowledge diffusion within the organization through project ties, they may also reflect changes in project assignment or other unobserved factors.

In addition, to assess whether these patterns merely reflect a reallocation of patenting activity across employees, we examine firm-level outcomes. We find increases in patent quantity and quality relative to counterfactual peer firms after the program began, consistent with broader gains in inventivity productivity from process learning at the firm rather than a simple reallocation of patenting work towards program participants (and away from non-participants). These firm-level comparisons, however, should be interpreted cautiously, as this design only allows for a single treated-firm contrast, albeit one with no observable pre-trend differences in outcomes between the focal firm and peer firms.

Fifth, we examine whether mentorship is accompanied by changes in workplace culture

and retention, two patterns that may help explain the persistence of the innovation results.[7] Following program implementation, engineers at the focal firm report a 0.12–0.14 higher star rating on Glassdoor out of a maximum of 5.0, and they are 3%–4% less likely to leave following program implementation compared to peers at similar firms. These retention patterns, which are also present when restricting the sample to named inventors on granted patents, point to potentially meaningful savings given the replacement costs of technical talent. Additional analyses using "Culture & Values" ratings and text-based culture measures show that improvements are concentrated among engineers and load most strongly on perceived gains in community-orientation, adaptability, collaboration, and integrity.

Our main findings suggest that mentorship is associated with increases in the volume and efficiency with which disclosed ideas are converted into patentable inventions. We use these estimates to revisit traditional ROI calculations and account for the many ways mentorship may influence value creation. Even conservative estimates based solely on patenting metrics, and excluding indirect gains from collaborator spillovers, retention, and culture, imply returns that substantially exceed implementation costs. The additional indirect channels are more illustrative, as they are harder to quantify precisely, although they plausibly raise returns even further. Because the program's costs were modest (two part-time administrators, volunteer mentors, and mentees' voluntary time over a four-month cohort), and because mentors realized productivity gains as well, the results suggest the program is not only scalable but readily deployable in other corporate contexts, with the binding constraint on adoption being managerial initiative, rather than capital constraints.

The ROI estimates also highlight why the setting is conceptually interesting. A common view in the innovation literature is that large, incumbent firms struggle to nurture novel ideas internally (Christensen, 1997), so the best practice is to move innovation efforts outside of the firm's boundaries. In contrast, our findings suggest that large firms may not be inherently poor environments for invention if they can overcome frictions in their innovation process.

[7] We use data from Glassdoor, which provides employee reviews of their jobs, and from Revelio, which tracks employee job transitions, to examine changes in workplace perceptions and retention, respectively.

The mentorship program's purposeful design features provide suggestive evidence of why this internal intervention may have succeeded.

A natural follow-up question, then, is whether the results from this large, R&D-intensive firm generalize to other settings. We view this context as a benchmark, as the frictions the program alleviates are widely documented at comparable technology firms. Specifically, the mechanisms empowering innovators to become inventors are likely portable (e.g., learning the best practices in taking an initial idea and developing it through their firm's patenting process). To support this point, we draw on best practices from experimental economics suggesting that shifts in near-universal constraints that firms face, rather than underlying preferences, are more likely to generalize (List, 2022). Viewed this way, our findings suggest that the focal firm's mentorship program may serve as a template for how firms in other settings might adapt and scale structured mentorship.

Finally, our findings raise a puzzle of adoption. Despite illustratively high returns, not all firms have implemented comparable programs. While contrary received wisdom plays a role, another reason may be that managerial frictions impede adoption, consistent with behavioral and cultural barriers preventing firms from implementing efficiency-enhancing practices (Benson et al., 2019; Mishra et al., 2022; Gertler et al., 2025). Further, Graham et al. (2022a) find that corporate leaders frequently perceive limited investor support for human-capital and culture investments, even when they believe such programs enhance value. In conclusion, our findings imply that, despite organizational barriers, structured, accelerator-style mentorship is a cost-effective and scalable intervention to unlock latent inventive capacity.

1 Related Literature and Contributions

Although workplace mentorship programs are ubiquitous, with roughly 70% of Fortune 500 firms operating one (Gutner, 2009), the empirical evidence supporting their efficacy is thin. Reviews of the mentoring literature emphasize that most existing studies rely on cross-sectional comparisons of participants to non-participants, retrospective self-reports, or se-

lected samples in which participation was confounded with unobserved motivation (Allen, 2007). Rigorous randomized or quasi-experimental evidence remains rare and is concentrated in a small number of settings: sales work (Sandvik et al., 2020, 2025), classroom teaching (Rockoff, 2008), military service (Lyle and Smith, 2014), and most closely related to our setting on the dimension of knowledge-intensive work, junior academic economists (Blau et al., 2010; Ginther et al., 2020). None of this evidence speaks to whether mentorship can raise inventive output inside superstar technology firms, where the strong presumption is that inventive ability is largely innate and where the relevant frictions are informational and relational, rather than motivational. In fact, received wisdom for inventive activity is typically not to try mentorship but to focus outside the firm via Corporate Venture Capital (CVC) or a separate venture-building unit (Vasquez-McCall et al., 2024).

We show that this outward-looking presumption is premature. Using data on over 20,000 engineers and 66,397 pre-patent idea disclosures—a scale that dwarfs prior randomized mentorship evidence—we find that structured mentorship meaningfully increases idea disclosures. These results also speak to the frictions that determine whose ideas become financially valuable. Prior work documents gender and racial barriers to entry into innovation and entrepreneurship (Bernstein et al., 2017; Cook, 2018; Bell et al., 2019; Ewens and Townsend, 2020; Koning et al., 2021; Calder-Wang and Gompers, 2021; Howell and Nanda, 2023; Dossi, 2024; Calaway, 2024; Cook et al., 2025; Gornall and Strebulaev, 2025; Hebert, 2025; Hochberg et al., 2026; Chien et al., 2026; Botelho et al., 2026). We complement it by showing that much of the innovator-to-inventor gap inside a superstar firm reflects removable organizational frictions—access to tacit patenting know-how and to input networks—rather than differences in inventive capacity. Incumbents can therefore unlock inventive capacity internally, rather than sourcing it through acquisitions, CVC, or entrepreneurship (Bena and Li, 2014; Chemmanur et al., 2014; Adelino et al., 2017; Ma, 2020; Bernstein et al., 2022a; Jeffers, 2023; Bernstein et al., 2023; Babina and Howell, 2024; Wang, 2025).

Second, our paper contributes to labor-and-finance research on how informal organiza-

tional practices shape corporate outcomes. For example, negative early workplace experiences, such as demeaning comments, can reduce effort and aspiration (Bennedsen et al., 2019; Azmat et al., 2025). We show that constructive feedback offered through group mentoring can have a counterveiling influence on productivity. Moreover, our findings on collaboration networks show how mentorship can substitute for informal ties that may otherwise be fragile or unequally distributed (Shue, 2013; Lerner and Malmendier, 2013; Grennan, 2019; Cho et al., 2026; Gertsberg, 2026). In this sense, we offer a potential solution for remedying workplace policies that lead to uneven labor-market outcomes (Bennedsen et al., 2022; Simintzi et al., 2024; Lagaras et al., 2025; Bennett et al., 2026; Benson et al., 2026).

Third, we contribute more broadly to the literature on organizational capital as a determinant of firm productivity (Eisfeldt and Papanikolaou, 2013). Over 80% of the assets of the dominant public firms are intangible (Grullon et al., 2019; Ayyagari et al., 2024; Kwon et al., 2024), yet our understanding of the unique labor challenges associated with work at these intangible-driven, superstar firms is still developing (Ewens et al., 2024; Nishesh et al., 2024; Hoberg and Phillips, 2025; Cai et al., 2026). Firms build organizational capital through investments in leadership (Bertrand and Schoar, 2003; Dessein and Prat, 2022), corporate culture (Guiso et al., 2015; Li et al., 2020; Graham et al., 2022b; Gorton et al., 2022; Grennan, 2023; Fitzgerald and Liu, 2025), and customer capital (Larkin, 2013; Campello and Gao, 2017; He et al., 2024). We extend this understanding by showing that formal programs that develop informal skills generate organizational capital returns through stronger culture, satisfaction, and retention.

Fourth, we contribute to the entrepreneurship and accelerator literature by examining an accelerator-style mentorship program launched within a large incumbent firm. Prior research studies how accelerators and pitch competitions influence early-stage venture ideas (Hochberg and Fehder, 2015; Gonzalez-Uribe and Leatherbee, 2017; Cohen et al., 2019; Howell, 2020, 2021; Hu and Ma, 2025; Gonzalez-Uribe et al., 2026). This literature emphasizes how structured accountability, peer learning, and enhanced visibility support nascent ven-

tures. We show that similar design features also strengthen corporate innovation, suggesting large firms can replicate some benefits of entrepreneurial ecosystems internally, consistent with earlier growth patterns (Babina, 2019; Babina et al., 2023).

A final contribution relates to measurement. Empirical work typically observes innovation outputs only after ideas have survived multiple filters. This creates a "dogs that did not bark" problem in that researchers rarely observe the abandoned counterfactual. Our data overcome this limitation by observing the universe of the firm's pre-patent ideas, including failed ones, which is an important contribution in its own right, given how much of the inventive process is ignored when researchers observe only granted patents. This distinction is important for interpreting patent-based measures of innovation (Lerner and Seru, 2022; Yang et al., 2025), especially in settings where organizational processes determine whether ideas become observable to firm outsiders. In this sense, we contribute to a growing body of research working to develop more nuanced measures of innovation (Kogan et al., 2017a; Mezzanotti, 2021; Abrams et al., 2023; Bena et al., 2022; Bena and Simintzi, 2025; Bena et al., 2026). Moreover, this distinction is especially consequential in our setting because ideation and conversion have different implications for the innate-ability null hypothesis. Our evidence is consistent with mentorship enabling fixed underlying creativity to be expressed more efficiently and provides suggestive evidence on whether ideas are harder to find and whether productivity growth suffers from the suppression of latent inventive capacity (Bloom et al., 2020; Fort et al., 2026).

2 Theoretical Framework: Mentoring, Creativity, and Patenting

This section develops a model of mentoring that accommodates both innate creativity and learned inventive skills, and that clarifies the specific margin on which mentoring is predicted to operate. The model adds a time-allocation constraint to on-the-job training models of human capital accumulation, and then we let mentoring enter through three channels centered on process-learning. First, mentoring accelerates the accumulation of tacit knowledge about

the firm's own patenting pipeline. Second, mentoring reduces procedural and informational frictions that block engineers from submitting the ideas they already have. Third, mentorship expands the variety and relevance of an engineer's knowledge inputs through exposure to more experienced colleagues. We contrast these channels of process-learning with two alternatives. The first is a pure time-shift mechanism that reallocates effort toward patenting without changing underlying state variables. The second alternative is an innate-ability null that inventive productivity is bounded primarily by a scarce and largely fixed creative capacity, such that even a purposefully designed mentoring program cannot generate inventive output that was not already latent in an engineer's endowment. Our model imposes this null as a modeling restriction. Specifically, creativity C is fixed by assumption and cannot be raised by the intervention. Any observed gains from mentoring must therefore operate through either the three process-learning channels or the time-shift channel, rather than through changes to an individual's underlying creative endowment.

In the model, time is discrete and indexed by t. We assume each engineer has a fixed workweek, normalized to one unit of time t, which they allocate between core production (t_C) and patent-related activities (t_P). Next, we assume each engineer has a creativity endowment $C \geq 0$. This creativity endowment is time-invariant and heterogeneous across engineers.[8] Each engineer also has a stock of process knowledge about the firm's patenting pipeline, $h_t \geq 0$, that is accumulated through work, training, and mentoring. This stock captures the tacit know-how that helps an engineer recognize which ideas are worth submitting, articulate them in the vocabulary the firm's review process rewards, and move them through the pipeline in the same way experienced inventors do. Finally, there is a friction index k_t which reflects procedural and informational barriers to disclosure, including both organizational hurdles (e.g., unfamiliarity with the disclosure system, uncertainty about the review process) and psychological hurdles (e.g., lack of confidence or reluctance to expose an unfinished idea

[8]While the psychology literature has long debated whether creativity is defined as a characteristic of the person, process, or eventual product, our conceptualization of creativity draws from a well-established strand of the literature that views creativity as a person-specific capacity to generate novel and useful ideas through a recombinant process (Amabile, 1983; Bol et al., 2023).

to institutional scrutiny). In our setup, idea production is as follows:

$$I_t \;=\; \Lambda(C,\, h_t)\; t_P(t), \qquad S_t \;=\; I_t \cdot q(h_t;\, k_t)\,, \tag{1}$$

where I_t is the expected number of invention disclosure attempts and S_t is the expected number of successful disclosures whose pipeline outcome (approval for application, defensive publication, or decline) reflects an informed judgment about where the idea belongs rather than a procedurally uncertain attempt. $\Lambda(C, h_t)$ is the mapping from creativity and process knowledge into disclosure attempts, and $q(h_t; k_t)$ is the probability that an attempt is well-formed in the sense just described. We assume that the idea-generation function, $\Lambda(C, h_t)$, is increasing in both creativity and process knowledge, with diminishing returns, and that they are complementary. This captures the commonsense notion that both novelty and practical know-how matter, and that one amplifies the other. Similarly, q increases with process knowledge and decreases with frictions. We deliberately do not interpret q as an unconditional per-disclosure approval rate. In our setting, the meaningful margin is not whether a disclosure is approved for a patent application in isolation, but whether the composition of an engineer's disclosure outcomes comes to resemble that of experienced inventors as process knowledge accumulates.

Next, let process knowledge evolve according to:

$$h_{t+1} \;=\; (1-\delta)\, h_t \;+\; \phi\, t_P(t) \;+\; \mu\, M_t, \qquad \delta \in (0,1),\; \phi > 0,\; \mu \geq 0, \tag{2}$$

where δ is depreciation, ϕ is the productivity of learning-by-doing, $M_t \in \{0, 1\}$ indicates mentoring, and μ is the direct increment from mentoring. In practice, μ reflects guided practice in which mentors and mentees jointly work through an actual inventive idea, with the mentor demonstrating how the firm's pipeline treats different classes of ideas which are refined into applications, defensive publications, or are declined. This contrasts with orientation videos on how to submit an idea, which are common in HR-based onboarding

sessions that may enter into the initial periods of h_t, but which do not convey the tacit, pipeline-specific knowledge that μ represents.

Next, we consider the two alternatives to the process-learning hypothesis. Under the time-shift hypothesis, mentoring does not convey knowledge but instead temporarily shifts attention to patenting. Formally, $M_t = 0$, neither frictions nor the learning technology change, and t_P increases for a bounded window with effects disappearing once time allocation reverts. Under the innate-ability null hypothesis, gains from any intervention are bounded above by the fixed distribution of C: because mentoring cannot raise C and the marginal engineer's C is (by hypothesis) low, the mentoring intervention should produce little movement in inventive output, particularly among engineers with no prior patenting record. Our model imposes fixed C throughout, so both nulls sit inside the model as restrictions on which state variables move, not as competing frameworks.

If mentoring is effective through the process-learning channels, then the three channels are jointly predicted to generate the following observable patterns. First, disclosure volume (I_t and S_t) rises persistently, because h_t accumulates and k_t falls; the persistence distinguishes this prediction from the time-shift alternative, which yields only transitory volume gains. Second, the composition of pipeline outcomes for mentees' disclosures shifts toward the composition observed for experienced inventors, because process knowledge is precisely the ability to distinguish among the pipeline's outcome categories—this is the empirically meaningful analog of q rising, and it does not require an increase in unconditional per-disclosure approval rates. Third, gains are largest among engineers with the lowest baseline h_t (i.e., those without prior patenting experience) because the marginal return to mentoring is highest among those with the least ex ante knowledge about the patenting process. Fourth, gains can appear even among engineers with modest C, because reducing k_t and raising h_t allows a larger share of latent creativity to be expressed through the pipeline. This last prediction is what distinguishes the process-learning hypothesis from the innate-ability null hypothesis: if the null were correct, engineers without prior patenting output should show

little response to mentoring; whereas if process-learning is the operative channel, these engineers should show the largest response. Fifth, if tacit know-how diffuses to collaborators, we may observe spillovers to non-mentees embedded in mentees' project networks.

In contrast, the time-shift hypothesis predicts only transitory volume gains due to mentorship, with no composition changes, no differential effects by baseline experience, and no spillovers. The innate-ability null generates little movement of any kind, particularly for engineers with no prior patenting output. The three predictions therefore partition the space of possible outcomes in a way our design can test.

This framework also connects to the external validity literature. Following best practices in experimental economics, we classify the elements of our model into the preferences, beliefs, and constraints taxonomy emphasized by List (2020). In our setting, preferences correspond to engineers' intrinsic creativity endowments (C) and career objectives, which mentoring does not alter. Beliefs capture expectations about whether ideas will succeed in the patent process, represented by the friction index (k_t), and mentoring reduces these frictions through process-learning. Constraints include both the allocation of time between core tasks and patenting, as well as the organizational processes developed by management and IP professionals. Because the shifts we identify operate on constraints and beliefs rather than on preferences, they are, by List's logic, the elements most likely to generalize to comparable settings.

Readers interested in the formal optimization problem and proofs that lead to these comparative statics can consult Internet Appendix C.

3 Firm Setting

We study a mentorship program at a publicly traded S&P 500 technology firm with more than 50,000 employees that develops AI, cloud, data storage, networking, and security products. The firm has over 20 offices in the U.S. and operates over 100 offices worldwide. The firm invests meaningfully in R&D, making human and organizational capital key aspects of its strategy. To illustrate the economic stakes of inventive output at this firm, it is one of

the largest corporate patent holders in the U.S., with an active patent stock in the tens of thousands and consistently ranking among the top corporate assignees by annual grants over the sample period. Over the sample period, the firm's annual R&D expenditure ranged from $6 to $8 billion, a scale comparable to the National Science Foundation's entire annual budget of approximately $8.9 billion, on average, over the same years. Using the market-based patent valuation approach of Kogan et al. (2017a), the firm's annual patent output alone contributes approximately $50 billion in 2025 U.S. dollars to the aggregate stock of innovation capital that subsequent research has linked directly to aggregate productivity growth (Akcigit and Kerr, 2018; Kogan et al., 2020; Kelly et al., 2021). Thus, the invention disclosures and mentoring intervention that we study are economically important, because a program that expands the flow of disclosures at this firm creates plausibly measurable implications for the firm's private returns and for the aggregate innovation flow that drives long-run growth (Romer, 1990; Jones, 1995).

To buoy their innovative output, the firm instituted a mentorship program to help its technical talent gain the skills and knowledge necessary to successfully develop and patent their inventions. Advertisements for the mentoring program stated two explicit goals: (1) to "develop talent with a focus on intellectual property and innovation" and (2) to "diversify invention submissions and patentable inventions from all parts of the [firm's] talent pool." For the first three cohorts of the program, advertisements for the program were sent to all employees involved in or potentially involved in the innovation process, including those in upper management, to encourage them to encourage their direct reports to apply. After the third cohort, program administrators relied on word-of-mouth publicity and reached back out to previous applicants who had not been able to participate in prior cohorts. While the mentorship program was open to all employees, those without invention disclosure experience and those from underrepresented groups (e.g., women and racial/ethnic minorities) were especially encouraged to participate, and advertisements were specifically sent to identity-based employee resource groups.

After the success of the first cohort, the mentee application process was continuously open with the website stating, "Express your initial interest in one of the best professional development programs at [the firm]! Our team will reach out to you as we lead up to the next cohort for the final enrollment/commitment. Your initial enrollment here does not commit you to anything at this time." The potential mentee then had to answer simple, routine questions such as name, business leader, location, technology area, and prior experience with the IP system, with no open-ended questions. The only non-routine question was to explicitly acknowledge that they had read about the commitment, which was described as "Mentees must be willing to commit four hours per week for the duration of the four-month cohort. Manager awareness of program participation is encouraged."

Mentors were identified and chosen in two ways. First, there was a call for people to register their interest in becoming mentors. The form paralleled features of the mentee form but also asked about prior inventions and the biggest technical challenges that the business must address in the next 12–24 months. Second, the mentoring program administrators targeted specific people. As the administrators explained, they felt the need to target when they needed to fill technological or geographic coverage gaps in the mentor community. On occasion, the administrators said they had more mentors than mentee teams, and in those cases, they favored mentors with more patent program experience and a strong track record as inventors. There were no instances in which those volunteering to be mentors were not chosen because of personality, preferred working style, or because one of their senior leaders felt they would not be a good fit.

The program operated in seven cohorts from 2020Q4 to 2023Q3, with the number of mentees ranging from 40 to 171 across cohorts. In the early cohorts, the program administrators increased the number of teams to accommodate more participants. But following cohort three, with 171 mentees and internal discussions about scaling too quickly, the administrators deliberately capped the number of teams. The administrators perceived that their intentional size reduction helped them to improve monitoring and accountability. Additional

constraints also led administrators to reduce the number of participants. These constraints included the number of available mentors and occasional managerial approval limits, such as when a new product launch was scheduled in the same quarter.

When there were more mentees than slots available, the administrators prioritized based on inventive experience, meaning they prioritized mentees with little experience with the patenting process, as assessed by the number of idea disclosure submissions they had previously completed. If the administrators had to turn down an applicant for a given cohort, they encouraged them to reapply to a later cohort, as they did not want non-participants to be discouraged. Or if the administrators turned them away because they already had too much prior inventive experience, they tried to identify other, more experienced people who had indicated similar technological interests and encouraged them to connect and potentially work together without a mentor.

A mentorship group generally consisted of four to five mentees, mentored by one experienced inventor. The mentor-mentee matches were made based on geographical overlap (e.g., both mentees and mentors indicated their preferred time zones on their forms) and common inventive interests, but the matches were otherwise random. The application forms listed 15 inventive interest topics, such as data analytics, AI in security, AI in customer experience, optics, automation, and multi-cloud. Participants were unable to request specific mentors, and administrators indicate that mentees likely did not know their mentors or fellow mentees before beginning the mentorship program. As the mentorship program was voluntary and not part of regular work assignments, participants were credited with the time required to complete the program to reduce the risk of falling behind on their regular tasks and responsibilities.

Mentors received training in mentoring and were instructed to teach research principles, ensure a fair division of labor within the team, connect mentees with resources, coach mentees on writing effective invention disclosures, and help teams navigate the inventive process more productively. As shown in Internet Appendix Figure A.1, the program's curriculum

was structured with five phases and required mentees to commit to each phase in the initial planning stage. The five phases include a kickoff event, a problem statement (with two check-ins), a solution draft (with two check-ins), an invention submission (with one check-in and a submission deadline), and a showcase event where the mentees presented their ideas to senior technical talent. In most cases, groups successfully generated and submitted an invention disclosure within four months of starting the mentorship program. While these invention submissions constitute successful R&D outcomes in and of themselves, the main purpose of the mentorship program was to cultivate within participants the skills and knowledge necessary to develop and patent their own inventions outside the mentorship program setting.

Over the last decade, and throughout the program period, the firm has offered stable monetary incentives for patenting (e.g., $3,000 at application filing and $2,000 at grant for the Silicon Valley inventors). The firm's internal Patent Review Board rejects most ideas, leading to a patent grant rate of about 90% conditional on patent application filing, which is higher than its peers.[9] We also confirmed that the Patent Review Board did not know whether submissions were associated with the mentorship program, consistent with impartiality. Finally, anecdotes suggest that leadership consistently described the mentorship initiative as valuable and credited it with expanding the pool of named inventors, a perspective that aligns with the program's stated objectives and which informs our later ROI illustration. The program continues today, albeit with some strategic changes.

Finally, given that the mentorship program spans several years, we reviewed the firm's annual reports and discussed with the program administrators whether its rollout coincided with other changes in innovation strategy, IP leadership, or R&D priorities that may have influenced patenting outcomes. The firm's core strategic direction remained stable. Throughout the period, the firm continued to integrate its technology capabilities while pursuing a pre-existing strategy toward more recurring revenue through subscriptions. Other topics mentioned were climate and sustainability issues becoming more prominent and AI receiving

[9]The average patent application acceptance rate (i.e., grant rate) is about 65% (Farre-Mensa et al., 2020).

greater emphasis. R&D expenditures increased by more than $1 billion, but R&D intensity remained approximately 13 percent of revenue, with no discrete break around program adoption. There was no CEO turnover or realignment of senior leadership. Our discussions with the IP professionals also indicate that there was no material change in the disclosure process, Patent Review Board, review standards, or inventor incentives. Consistent with this account, the firm's publicly disclosed IP policy remained substantively unchanged across annual reports, continuing to emphasize a broad portfolio of patents, rather than a new or more aggressive patenting policy. Thus, the program was introduced against a relatively stable period in terms of strategy, management, and IP policy.

4 Data Sources, Innovator Characteristics, and Empirical Specifications

4.1 Data Sources

Estimating the relation between mentorship and the quantity and quality of ideas submitted as invention disclosures requires a detailed panel of innovator-level data. The firm provided us with this data, which covers all invention disclosures submitted by its employees from its early days through 2023. The dataset contains 66,397 unique disclosures filed by 29,211 unique innovators, with most disclosures being developed and submitted by a team of innovators. This data includes demographic details for most individuals, including their gender, age, marital status, race/ethnicity, job role, and the country where they work. We use these data to form an innovator-quarter level dataset, in which we count the number of invention disclosures each innovator submits in a particular calendar quarter.

Our main innovation productivity tests consider all disclosures that an innovator submits in a given quarter, regardless of the disclosure's status in the review process. We also conduct tests with alternative outcome variables, such as focusing only on disclosures that become approved patent applications and scaling disclosures by team size.[10] Internet Appendix

[10]We refer to these disclosures as "advanced" disclosures, which have any of the following statuses: Approved Patent Application; Prior Art Search Requested; Returned for Comments. At times we also discuss "non-advanced" disclosures, which have any of the following statuses: Declined; Auto Declined; Defensive Publication Candidate; Approved Defensive Publication; Approved Trade Secret; Deleted; Merged. In In-

Figure A.2 highlights the patent process as taught to mentees by mentors as part of the mentorship program. The main terminal outcomes that an invention disclosure submission can realize are (i) being declined, (ii) being approved for defensive publication, and (iii) being approved for patent application filing. Submissions can also be returned for comments by the review committee.[11]

The firm provided us with data on which individuals were part of each cohort of the mentorship program. This dataset includes details on who each person's mentor was, allowing us to account for the observable characteristics of the innovator's mentor. This dataset included data on innovator tenure with the firm and within the industry more broadly. For innovators not in the mentorship program, we have to infer their tenure with the firm based on the timing of their first observed disclosure (we set their hire date to be one year before their first disclosure date). We do not impute industry tenure for innovators without this information.

4.2 Innovator Characteristics

Table 1 reports demographic characteristics of the mentorship program participants and the rest of the innovators in the idea disclosure database. The first cohort included 101 innovators, with an average age of 35, approximately 3.5 years of experience at the firm and 6.5 years in the industry. Most participants were engineers (85%), and the program drew heavily from historically underrepresented talent pipelines. Ninety percent were women and 16% were URMs, although all of the most common ethnicities within the firm were represented in the program. Subsequent cohorts exhibit similar profiles, with increased gender balance in cohorts four through seven. The broadening of participation was partially

ternet Appendix B, we provide evidence from a series of robustness tests based on other status classification approaches.

[11] For confidentiality reasons associated with our internal review board (IRB) approval, we did not get access to the title of the invention disclosure or patent application number. This was because we received other indirect Personally Identifiable Information (PII) data (i.e., gender, race, date of birth, citizenship status, etc.). While we can not look at individual patents, we can look at aggregate patent count and quality measures, which we do in tests in the Internet Appendix.

in response to requests from HR to be more inclusive.

To benchmark the representativeness of program participants, we compare mentees to the broader innovator population of over 28,500 non-participants. Mentees were slightly younger and had somewhat shorter firm tenure, which is consistent with newer employees seeking guidance on disclosure and patenting, but baseline inventive output was quite similar across groups.[12] In particular, mentees averaged about 0.08 quarterly disclosures pre-program, and over 93% had never received an approved patent, similar to their peers. These patterns suggest participants were neither top performers nor a remedial group of underachievers. To further assess whether our estimates are sensitive to observable selection and cohort-by-cohort compositional changes, Internet Appendix B reports additional analyses using propensity-score and nearest-neighbor matching on pre-treatment observables, reweighting, like-for-like comparisons, and alternative DiD estimators. These exercises do not rule out selection on unobservables, but they indicate that the estimates are not sensitive to observable differences in cohort composition or to particular comparison-group and weighting choices.

4.3 Empirical Specifications

To estimate the relation between mentorship program participation and productivity, we compare participants' quarterly disclosure outcomes before and after they begin the program. We use data from Q1 of 2020 onwards, as the first cohort of the mentorship program was in Q4 of 2020 and as this ensures all of our data is in the during-COVID or post-COVID period, mitigating concerns that using pre-COVID data may bias estimates (though all our inferences hold if we include pre-COVID data). We employ a stacked DiD estimator that isolates within-cohort treatment effects while avoiding contamination from differential

[12] A logistic regression of program participation on observable characteristics yields similar patterns. Younger innovators are more likely to participate and women and Black/Latinx innovators are more likely to participate than are men and White innovators.

timing, which can bias results from two-way fixed effects estimations (Baker et al., 2022).[13]

For each mentorship cohort, we compare treated innovators with two control groups: (1) those never participating in the program (never-mentored innovators) and (2) those participating in later cohorts (not-yet-mentored innovators). The latter serves as potentially the most apples-to-apples counterfactual because they ultimately self-select into treatment, differing only in timing. This design mitigates concerns about selection on unobservables, relying instead on the assumption that, conditional on observables, the timing of participation is as at least quasi-random.[14] The consistent curriculum and mentor training across cohorts further minimize treatment heterogeneity.

We aggregate innovator-quarter observations within each cohort and stack them across all seven mentorship cohorts. Our exact empirical specification is as follows:

$$\mathrm{y}_{i,t,g} = \beta_1 \mathrm{Mentored}_{i,t,g} + \alpha_{i,g} + \lambda_{t,g} + \varepsilon_{i,t,g}, \tag{3}$$

where $y_{i,t,g}$ is the disclosure outcome variable for innovator i in year-quarter t for mentorship cohort g. $\alpha_{i,g}$ is an innovator-by-cohort fixed effect, $\lambda_{t,g}$ is a year-quarter-by-cohort fixed effect, and $\varepsilon_{i,t,g}$ is an idiosyncratic error term. Standard errors are double-clustered by innovator-by-cohort and year-quarter-by-cohort to account for repeated observations.[15] Conceptually, this approach estimates cohort-specific DiD effects using both control groups and averages the coefficients across cohorts.

4.3.1 Parallel Trends Tests

Before turning to the main results, we assess the parallel trends assumption underlying our DiD design. We compare pre-program disclosure trajectories for mentored innovators

[13] Specifically, a critique of two-way fixed effects estimations with staggered treatment timing is that bias can arise from comparing already-treated units with newly-treated ones. This bias can be especially problematic when treatment effects are heterogeneous over time

[14] Some innovators were deferred to later cohorts due to capacity limits, creating quasi-random variation in treatment timing.

[15] Our inferences are robust if we double-cluster the standard errors by mentor-by-cohort or mentor-team-by-cohort (instead of innovator-by-cohort) and year-quarter-by-cohort.

and not-yet-mentored (and never-mentored) peers. While parallel trends cannot be proven, these tests provide suggestive evidence on credibility (Rambachan and Roth, 2023). Figure 1, subfigure (a), plots quarterly differences in disclosure submissions around mentorship start. The pre-treatment estimates are small, statistically indistinguishable from zero, and display tight confidence intervals. Formal tests also fail to reject equality of pre-treatment coefficients (p-values 0.613–0.958; see figure notes). Similarly, in subfigure (b), where we restrict the data to only include disclosure submissions where a mentor is not a collaborator, we find that the pre-treatment estimates are statistically indistinguishable from zero. We find similar pre-trend patterns when scaling disclosures by team size (Internet Appendix Figure A.3), when separating advanced and non-advanced disclosures (Internet Appendix Figure A.4a and Figure A.4b, respectively), when including never-mentored innovators as controls while using unscaled and team size-scaled outcomes (Internet Appendix Figure A.5a and Figure A.5b, respectively), and when separately considering early cohorts (Cohorts 1–3) and late cohorts (Cohorts 4–7) (Internet Appendix Figure A.6a and Figure A.6b, respectively). Across these checks, pre-program trajectories are comparable, supporting the identifying assumption.

4.3.2 Survey Evidence on Non-Participation and Self-Selection

To assess potential self-selection concerns, we surveyed a representative subsample of innovators defined by the firm's HR team as engineers and technical staff eligible to contribute to IP development. The survey, which was part of a broader multi-firm study by Chien and Grennan (2024), achieved a 15% response rate, which is typical for corporate surveys (Graham et al., 2005, 2022a). The limited total number of responses reflects the firm's deliberate choice to sample representatively to avoid survey fatigue. Chien and Grennan (2024) evaluate the data validity and show internally consistent responses that align closely to HR records.

As shown in Internet Appendix Table A.1, 43% of respondents reported participating in the program, and women were significantly more likely to participate. Among non-

participants, lack of awareness was the leading reason for non-participation and there was no gender difference. Notably, women were more likely than men to believe mentorship would not benefit them (11% versus 3%). This pattern suggests that ex-ante skepticism among some women works against us finding results, as this type of self-selection attenuates estimates toward zero (Sandvik et al., 2025). While the survey evidence provides some confidence that self-selection is not the sole driver of our ultimate findings, our econometric design addresses remaining concerns.

To interpret the observed gender differences in the expected program benefits, we rely on additional questions. Specifically, women did not disproportionately cite avoiding activities highlighting weaknesses, suggesting their skepticism does not stem from differences in confidence (Coffman, 2014; Bordalo et al., 2016; Angeli et al., 2025). Moreover, they reported similar rates of ideation, awareness of the invention process, and attendance at trainings.[16] Importantly, few respondents indicated time constraints (only 6%, with no gender gap), and few respondents reported already receiving adequate IP support. This suggests organizational frictions associated with management practices and culture induce skepticism.

5 Results

This section presents our empirical findings in six parts. First, we document the relationship between mentorship participation and subsequent innovative activity, showing substantial short- and long-term increases in disclosure submissions and a compositional shift toward the types of higher-quality ideas that experienced inventors submit (Section 5.1). Second, we evaluate heterogeneity across mentees and mentors (Section 5.2), identifying groups with the largest gains and testing predictions about informational deficits. Third, we summarize survey evidence on potential mechanisms and use our administrative data to validate the importance of expanded professional networks (Section 5.3). Fourth, we analyze the potential

[16]Interestingly, fewer than half recalled formal invention training despite firmwide IP training for engineers during onboarding. This divergence underscores that formal training alone does not sustain inventive engagement, whereas mentoring appears to activate innovative behavior that standard training does not.

for externalities by studying productivity changes among collaborators (Section 5.4). Fifth, we assess engineers' perceptions of their workplace culture and consider their turnover rates (Section 5.5). Sixth, we consider alternative explanations for our findings and use firm-wide patenting metrics to further support our main findings (Section 5.6)

5.1 Mentorship Program Participation and Innovator Productivity

5.1.1 Overall Productivity and Idea Advancement

We report the relation between participation in the mentorship program and innovator productivity in Table 2 and Table 3. Together, these results let us distinguish between the three posited hypotheses: the process-learning hypothesis predicts persistent gains in disclosure volume and a shift in the composition of inventive outcomes of mentees toward that of experienced inventors; the time-shift hypothesis predicts only transitory volume gains and no composition shift; and the innate-ability null hypothesis predicts negligible changes in productivity, particularly among engineers with no prior patenting record.

Column (1) of Table 2 reports the overall relation between mentorship and inventive productivity, showing a statistically significant increase in disclosure output. The coefficient estimate on *Mentored* implies that mentorship increases productivity by roughly 0.25 disclosure submissions per quarter, or about one additional invention per year, and the estimate is significant at the 1% level. The pre-program average level of productivity for mentored workers, 0.077 disclosures per quarter, is displayed in the penultimate row, implying innovation gains of approximately 330% (reported in the bottom row). Because the baseline level of inventive activity was low, as in a sample that is disproportionately composed of engineers without prior patenting experience, modest absolute increases translate into large proportional gains.

We next examine whether the mentorship-productivity relation persists beyond the program end. Although the curriculum guides mentees through the invention development and patenting process, disclosure submissions are an expected short-run outcome. The key ques-

tion is whether gains continue once formal mentoring ends, which would help rule out the time-shift hypothesis. To test persistence, we re-estimate the model excluding data from the program quarter and the following quarter. Column (2) shows that the positive association remains statistically significant at the 1% level. The point estimate suggests that mentorship is associated with a 0.082 increase in ideas submitted in the long-term per quarter, implying sustained gains of about 0.33 additional disclosures per innovator per year, which represents more than a 100% increase relative to baseline productivity. Our findings are robust to the use of a Poisson Pseudo-Maximum Likelihood estimator suggested by (Chen and Roth, 2024), which accounts for the count-like nature of quarterly disclosure submissions (see Internet Appendix B). The results are also robust if we exclude from the outcome variable the first disclosure an innovator submits after starting the mentorship program (which is most likely the disclosure associated with the program curriculum). In particular, when making this restriction the estimates in Columns (1) and (2) of Table 2 become 0.132 (t-statistic = 4.085) and 0.077 (t-statistic = 3.346), respectively.

In Columns (3) and (4), we repeat our estimation process for the overall and long-term relation, respectively, but we use an alternative dependent variable that scales each disclosure by the collaborator team size before aggregating to the innovator-quarter level. In doing so, we can assess whether the observed productivity gains reflect actual increases in innovative output or if they simply reflect team composition decisions (e.g., joining pre-existing teams). The estimates on *Mentored* remain positive and statistically significant, reinforcing that mentorship is linked to increases in individual innovative capacity rather than simply expanding team sizes or redistributing credit.

Finally, Columns (5) and (6) address another concern, namely, that the mentorship relation is being driven by innovators who are simply co-authoring with their mentors. If mentored innovators appear more productive only because they are collaborating on their mentors' disclosures, then the relation should disappear once mentor-collaborated disclosures are removed. Columns (5) and (6) restrict the outcome to disclosures on which the mentors

are not listed as a collaborator. The coefficient on *Mentored* remains positive and statistically significant, indicating that mentees became more capable of developing their own disclosures independently of their mentor, not merely more visible on their mentor's work.

Next, in Table 3, we decompose the higher inventive output result into two economically distinct components: (i) disclosures that advance through the pipeline toward a patent application in Columns (1)–(3) (i.e., advanced disclosures), and (ii) disclosures that do not advance in Columns (4)–(6) (i.e., non-advanced disclosures). This decomposition is the empirical analog of the composition prediction in our theoretical framework. The process-learning hypothesis does not predict that mentees learn to submit only the disclosures that will be approved, as an unconditional per-disclosure conversion rate is not the meaningful margin since inexperienced innovators have no baseline rate to raise. What our framework predicts instead is that as process knowledge accumulates, the composition of a mentee's disclosures across potential outcomes (advancing, not advancing) will shift toward the composition observed for experienced inventors. If mentees learn what the firm's IP professionals believe are patentworthy versus what they decline, we should see gains along both margins, with the advanced margin gaining proportionally more.

This composition shift toward higher-quality ideas that align with the Patent Review Board is what the estimates show. Columns (1)–(3) of Table 3 reveal that mentored innovators produce significantly more advanced disclosures overall, in the long term, and when we remove mentor-collaborated disclosures. The proportional gain relative to the pre-program mean is 520% overall and 211% in the long term. Columns (4)–(6) show corresponding increases in non-advanced disclosures, at proportional gains of 245% overall and 86% in the long term. Two patterns are worth emphasizing. First, the advanced-disclosure gains are proportionally larger than the non-advanced-disclosure gains, consistent with the composition shift predicted by the model. Specifically, we find that the mix of pipeline outcomes for mentees' disclosures moves toward the mix observed for experienced inventors. Second, once mentor-collaborated disclosures are removed, in Columns (3) and (6), the advanced-

disclosure result strengthens in significance while the non-advanced-disclosure result becomes less precise. This asymmetry is informative. Mentees' independent output is disproportionately concentrated on ideas that ultimately advance, which is precisely what would be predicted by a process-learning mechanism operating through improved recognition of what the firm values.[17] We caution that we do not interpret these patterns as evidence of a rising per-disclosure approval rate; disclosure-level tests of that margin yield small, positive but imprecise estimates. Therefore, we interpret the data as supporting one of learning what to submit and a shift in composition driven by the extensive margin.

5.1.2 Robustness of the Mentorship-Productivity Relation

We conduct a series of robustness tests to assess the stability of our main findings. Results appear in Internet Appendix Table A.2 (overall productivity) and Table A.3 (long-term productivity). Finally, Internet Appendix Table A.4 and Internet Appendix Table A.5 provide additional insights about the extent to which endogenous selection into specific cohorts changes the point estimates.

First, in Table A.2 and Table A.3, alongside our preferred stacked DiD specification (even columns), we report two-way fixed-effects (TWFE) estimates (odd columns). Replicating the main results, the TWFE coefficients in Column (1) closely match their stacked DiD counterparts in Column (2). Second, a natural concern in collaborative environments is potential interference across units (i.e., violations of the Stable Unit Treatment Value Assumption (SUTVA)), whereby an innovator's potential outcomes may depend on peers' treatment status rather than solely on their own.[18] In our setting, such spillovers are an intended mechanism. The mentorship program was designed to diffuse tacit knowledge and strengthen the inventive network. Thus, in the language of List (Forthcoming), the es-

[17] The focal firm reports a patent grant rate of 90%, which is significantly higher than industry averages (Farre-Mensa et al., 2020). Strictly speaking, patent applications are never irrevocably rejected by the USPTO. Instead, they are abandoned by applicants following what technically are appealable rejections issued by examiners. In our case, ideas approved for patent application are rarely abandoned.

[18] See List (Forthcoming) for a discussion of interference and total causal effects in networked settings.

timand is a total causal effect rather than a narrow direct effect. Nonetheless, to probe whether cross-unit spillovers bias inference, we re-estimate the models after excluding data from the future collaborators of mentored innovators. Columns (3) and (4) show that estimates remain unchanged, indicating that knowledge diffusion does not artificially inflate measured treatment effects. Similarly, in subfigure (b) of Figure 1, we show that the increase in disclosure submissions after the program was not just driven by mentees increasing their ex post collaborations with their mentors, as we find that disclosure counts increase when we remove all mentor-collaborative submissions. This finding mitigates the concern one might have that the mentorship program simply taught innovators how to join pre-existing projects (e.g., via their new connection with their mentor). Rather, the results indicate that mentored innovators became more capable of developing their own invention disclosures after the program (i.e., independent of their mentor).

Third, to address selection concerns, Columns (5) and (6) restrict controls to only not-yet-mentored innovators. The economic magnitude of the coefficients remain similar and, if anything, are slightly larger. Fourth, to verify that mentorship improves learning about what makes an idea high-quality rather than simply increasing low-quality activity, Columns (7) and (8) of Table A.2 and Table A.3 focus on advanced disclosures that become approved patent applications. Consistent with the process-learning hypothesis, the results are positive and stable across specifications.

Fourth, to address multiple-hypothesis testing concerns in experimental settings (List et al., 2019), we follow Anderson (2008) and report sharpened q-values in Table A.2 and Table A.3, which are analogous to p-values after adjusting for the False Discovery Rate (FDR). We report the q-values in brackets, below the t-statistics (which are reported in parentheses). In all of the tests in Table A.2, the sharpened q-values are less than 0.01, suggesting that our inferences are robust to holding fixed the proportion of false positives as the number of tests increases.[19] Similarly, the sharpened q-values in Table A.3 highlight the

[19]The q-values are adjusted for all sixteen tests reported in Table A.2 (i.e., the eight tests in Panel A and the eight tests in Panel B). The estimated sharpened q-values are conservative because they do not account

robustness of the long-term findings, respectively.[20]

Fifth, we present additional evidence on whether changes in cohort composition or cohort-specific innovation opportunities drive our main findings. To examine whether the estimates reflect differential trends in innovation opportunities faced by less-represented groups, we re-estimate the stacked DiD using only female innovators. Across all specifications in Internet Appendix Table A.4, the results remain significant and of comparable magnitude to our estimates when using the full sample, suggesting changes in gender composition across cohorts do not drive our results. In addition, Internet Appendix Table A.5 presents summary statistics on ideas submitted by cohort along with the associated cohort-specific point estimates. We observe no statistically significant difference in idea submission rates across cohorts during the program. After the program, we observe a drop-off in idea submissions beginning in Cohort 5, which is to be expected mechanically, given that less time has elapsed since the mentees' last idea submission. Finally, estimates of cohort-specific ATTs are positive and statistically significant across all cohorts, and their magnitudes are broadly comparable.

Lastly, we explore whether our findings are sensitive to organizational-unit differences that could proxy for differential managerial support. The estimates are essentially unchanged when we include job role $\times$ work country fixed effects, which absorb common differences across employees in similar roles and locations. Moreover, participation did not require managerial nomination or recommendation, only awareness, as described in Section 3.

5.2 Heterogeneity Tests Based on Mentee and Mentor Characteristics

We next test whether mentorship differentially benefits innovators based on their demographic traits and experience, as predicted if the program alleviates informational frictions and network barriers that disproportionately constrain certain groups of innovators. While

for the positive correlations across tests.

[20] We also conduct a comprehensive series of additional robustness checks summarized in Internet Appendix B, including (i) employing different estimators (e.g., the Callaway and Sant'Anna (2021) estimator or a Poisson estimator), (ii) using alternative disclosure classifications (e.g., defensive publications and other IP categories), and (iii) applying alternative patent-quality metrics (forward citations, claims, and scope).

the mentorship program was available to all employees, a goal from the firm leadership's perspective was to assist inexperienced employees and those from demographic groups that have historically been underrepresented to be more successful in the patenting process.

In Table 4, we report the results from these heterogeneity tests. In Column (1), we interact *Mentored* with the binary indicator $Mentee_{URG}$, which equals one if the innovator is from an underrepresented demographic group (i.e., either female or a URM (neither Asian nor White), or both), and zero otherwise. The sample size is smaller than that of our main tests due to incomplete data on gender and race/ethnicity for all innovators. The estimate on the interaction term shows that mentorship yields larger gains for innovators from URGs. URGs gain an extra 0.08 ideas per quarter, producing a total of 1.30 additional annual submissions. This pattern implies mentorship helps innovators from URGs overcome organizational frictions in the invention process, consistent with evidence on targeted interventions closing gender gaps (Goldin and Rouse, 2000).

Importantly, though, the coefficient estimate on *Mentored* remains positive and statistically significant, suggesting that mentorship program participation corresponds to changes in productivity for all innovators, including White and Asian males. In Column (2) of Table 4, we test whether matching mentees and mentors by demographic background alters the productivity relation. The estimates suggest positive correlations across all pairs, but URG mentees paired with URG mentors experience the largest increases in overall innovative output. These results suggest that both a shared background and cross-group network expansion help contribute to knowledge exchange and reduce the organizational frictions in idea development.

To distinguish between our three competing hypotheses, in Column (3) we interact *Mentored* with an indicator for whether the mentee has no prior idea submission experience, *No Experience*. The estimates provide suggestive evidence that productivity increases are concentrated among inexperienced innovators, consistent with the process-learning hypothesis rather than the time-shift or innate-ability hypotheses. Thus, this finding helps assuage

concerns that the observed productivity gains are simply due to mentored workers shifting their time allocation away from other responsibilities and towards innovative activities. If this were the case, we would expect to see substantial productivity gains from experienced mentees, too, but the estimate on *Mentored* in Column (3) is much smaller than the baseline estimate and is statistically insignificant. That the greatest estimated economic magnitude is associated with mentorship among inexperienced innovators suggests that knowledge transfer from mentors to mentees likely played an important role in the observed increase in innovative output.

We plot a visual representation of the increased productivity of previously inexperienced innovators in Figure 2. We plot the post-mentorship point estimates using all disclosure outcomes in subfigure (a) and when removing collaborations with mentors in subfigure (b), both of which show that inexperienced innovators exhibited persistent and statistically significant increases in productivity out to nearly the two-year horizon.

In Column (4) of Table 4, we consider whether mentored innovators benefited more or less from the program depending on their age. We find no support for heterogeneous treatment by age when we interact an indicator *Old*, equal to one for workers who are above the median age of 42 and zero otherwise, with our main specification. This test also speaks to the potential concern that mentorship's relationship to productivity during the sample period simply fills the gap created by reduced in-person interactions under work-from-home (WFH). If that mechanism were operative, gains should concentrate differentially by age given widespread generational differences in adapting to hybrid work (Barrero et al., 2023).

The findings in Columns (5)–(8) demonstrate that heterogeneity based on mentee URG status and pre-program experience is even more pronounced in the long run, when we remove from the data disclosures that occur during the mentorship program period. In particular, the findings in Columns (5) and (6) suggest that only URG mentees, who had less pre-program experience developing and submitting invention disclosures, are associated with long-term performance gains in relation to mentorship, and the results in Column (7) similarly indicate

that only ex ante inexperienced mentees exhibited long-term productivity increases.

Next, we explore whether heterogeneity in mentor quality changes the relation between mentorship and productivity. To delineate between high and low levels of mentor quality, we aggregate mentors' productivity levels in the years before the mentorship program began. We then split mentors into above versus below the median value of pre-program mentor productivity, creating a *Q4/Q3-Mentor* indicator (for innovators whose mentors are in the highest, Q4, or second-highest, Q3, productivity quartiles), and we interact this indicator with the *Mentored* indicator and re-estimate our main regression model.

This allows us to examine whether mentor quality influences mentees' inventive productivity, a key test for distinguishing the time-shift hypothesis from the process-learning hypothesis. Column (1) of Table 5 displays the coefficient estimates, which show that mentees paired with above-median-productivity mentors experience significantly greater gains. The results in Column (2) show that top-quartile mentors are associated with mentees with the greatest increases in innovative output. Importantly, though, even mentees paired with lower-productivity mentors see positive productivity responses, as evidenced by the significant estimates for *Mentored*. These estimates imply that effective mentoring is not confined to a narrow set of individuals. The findings in Columns (3)–(4) show that the heterogeneous relation between mentor quality and inventive productivity persists in the long term, which also strengthens support for mentorship working through a mechanism that transfers tacit knowledge.

In conclusion, examining heterogeneity in mentor quality supports the interpretation that mentorship works by reducing frictions that block the expression of latent creative talent, not merely by reallocating effort toward patenting activity. The findings also underscore the program's potential scalability, as the benefits persist even when mentor quality varies.

5.3 Exploring the Mechanisms Driving the Productivity Gains

To examine the mechanisms underlying the observed productivity gains, we combine survey evidence with administrative data. Table 6 summarizes our survey evidence.[21] The program was positively received, and most mentees continued meeting with mentors after the formal program ended, reporting close to two substantive interactions per month. Nearly all respondents strongly agree that they plan to pursue more invention disclosures at work (mean response = 1.2 on a –2 to 2 scale).

The surveyed mentees' responses highlight three mechanisms: (i) knowledge transfer, (ii) expanded networks, and (iii) increased satisfaction at work. Two other mechanisms received weaker support: confidence in applying mentors' guidance and confidence in navigating the internal IP review process. Because survey responses were randomized to mitigate order bias, these results are most consistent with mentoring reducing organizational constraints to understanding the invention process, rather than individuals changing their beliefs about their own innovative ability. Furthermore, gender differences in perceived impact mirror the program's intent to reduce such constraints for URGs. Female mentees express significantly greater increases in perceived learning about patenting (mean of 1.5 vs. 0.9 for males) and job satisfaction (1.4 vs. 1.0). These results reinforce the idea that mentorship is particularly useful for employees facing informational or social access constraints.

Overall, the survey evidence highlights three mechanisms through which mentorship transforms non-patenting STEM workers into active inventors: (i) enhanced understanding of the invention process, (ii) expanded professional networks, and (iii) higher satisfaction at work. The first two mechanisms map directly to our process-learning framework. Mentorship increases the conversion rate of ideas into patentable inventions by transferring tacit procedural knowledge and broadening access to collaborative opportunities. Thus, creativ-

[21] We surveyed innovators involved in the mentorship program using the same instrument described in Section 4.3.2. To avoid survey fatigue, the firm routinely surveys representative subsamples. In this case, the participant and non-participant surveys received 125 responses (a 15% response rate), which is consistent with organizational survey norms (Graham et al., 2005).

ity is not being produced, but rather latent creativity has the opportunity to be more fully expressed. The third mechanism of workplace satisfaction hints that this mentoring intervention may also change engineers' perception of the firm's cultural values and norms, which we explore in Section 5.5.

5.3.1 Expanded Social Networks as a Mechanism.

To further explore the expanded social network mechanism, we test whether mentored innovators subsequently form more diverse teams. The program could plausibly influence team composition by widening mentees' collaborator sets or by making the benefits of heterogeneous teams salient. To test for this influence, we estimate stacked-DiD models at the innovator-quarter level, and we consider three outcomes: (i) the mean fraction of female teammates, (ii) a gender-by-race/ethnicity HHI (with lower values capturing greater diversity), and (iii) a patent-experience HHI based on teammates' prior approved applications.[22]

Table 7 summarizes our estimates, revealing that participation in the mentorship program is positively and significantly associated with all definitions of team diversity (gender, ethnicity, and experience), although some results for gender diversity are insignificant in the long term. These changes to team composition are economically meaningful, as diverse teams are linked to stronger performance and to avoiding certain pitfalls associated with homophily in team innovation (Heath et al., 2023). In our setting, the shift in team composition supports the survey interpretation that expanded social networks are an important mechanism through which mentorship functions. In this sense, these tests also align with the process-learning hypothesis, because mentoring appears to provide more than awareness or on-the-job training; rather, it broadens access to collaborators and lowers friction in assembling teams with complementary skills. From our discussion with program administrators, neither mentees nor mentors knew each other ahead of time, suggesting a potentially important design choice for scaling any such intervention.

[22]Variable construction details are included in Internet Appendix B.

5.4 Responses Among Non-Mentees

Next, we attempt to document whether non-participants experience any changes in productivity through their collaboration networks with mentees. This addresses scalability and speaks to whether the mentorship program generates complementary knowledge flows among engineers or merely redistributes innovative activity. In the model, mentoring can produce externalities of either sign: positive responses if tacit procedural knowledge diffuses through collaboration networks, or negative responses if attention and resources shift toward mentored workers and away from others. We test these alternative channels.

We restrict the sample to innovators who never participated in the program as a mentor or mentee. For each non-participant, we identify the first quarter in which they appear on a disclosure team with a previously mentored colleague and define *Post First Exposure* as one from that quarter onward (and zero otherwise). The window spans four quarters before and after the first exposure. We then estimate stacked DiD regressions of disclosure outcomes on this exposure indicator. We make sure to omit from the dependent variable the disclosure that determines the value of *Post First Exposure*, which ensures our findings are not mechanically driven by the variable construction.

Column (1) of Table 8 shows the point estimates for non-participants who later collaborate with mentored workers. Disclosure counts increase by 0.375, or approximately 38% relative to these workers' pre-exposure baseline. Because these workers were already relatively active innovators, the pattern does not appear to be driven only by low-output employees moving toward the mean. We interpret this positive point estimate cautiously. It is consistent with the mentorship program's reach extending beyond direct mentees, potentially through collaboration-based transmission of procedural know-how or other inputs into the idea-to-patent process. Given the endogenous formation of these collaborations, however, we view the evidence as suggestive of broader responses among non-mentees rather than definitive evidence of spillovers.

Next, we examine whether these broader program-related responses vary with non-

participants' demographic characteristics. Interaction terms with *Female* and/or *URM* indicators are uniformly negative and often significant, suggesting that the positive association between collaboration with mentored workers and subsequent disclosure activity is stronger for historically well-represented groups. This pattern is consistent with prior evidence that women and URM engineers face weaker access to informal invention networks and that managers are less likely to share important information with them (Chien and Grennan, 2024). In Column (5), we examine heterogeneity by non-participants' pre-exposure disclosure experience, where *High-Pre* equals one for those with above-median baseline productivity levels. The insignificant coefficient on the interaction term, together with the positive and significant coefficient on *Mentored*, suggests that the positive association is not confined to workers with low or high baseline disclosure activity.

Conceptually, these findings point to an important boundary condition. While the mentorship program is associated with reduced innovation frictions for direct participants, the broader reach of the program appears uneven. One interpretation is that program-related know-how may travel primarily through existing collaboration networks, which can limit access for workers who are less integrated into those networks. At the same time, the positive estimates for non-participants speak against a pure crowd-out story in which mentored workers simply absorb scarce organizational attention. Therefore, we view the preponderance of evidence as consistent with broader program-related responses within the firm, while remaining cautious about interpreting these estimates as causal spillovers. We note this broader positive interpretation is also reinforced by the firm-level tests in Section 5.6, which demonstrate that the start of the mentorship program is associated with increases in both patent quantity and quality relative to peer firms.

5.5 Perceptions of Corporate Culture and Engineer Turnover

The survey results revealed that engineers perceive the program positively and as a contributor to their increased workplace satisfaction. This heightened satisfaction may serve as a

mechanism enabling innovators to become productive inventors. To assess whether these self-reported patterns are reflected in broader workplace outcomes, we turn to external data on employee perceptions and turnover decisions. This mirrors our earlier approach of triangulating mechanisms using complementary data sources.

Using data from crowd-sourced employee reviews on Glassdoor, a career intelligence website that attempts to provide transparency about jobs, salaries, and companies, we examine the relation between the mentorship program and employees' perceptions of their workplace environment.[23] We estimate a triple-differences (DDD) model relating mentorship program timing to changes in employee perceptions of the firm. The DDD compares (i) pre- versus post-program periods, (ii) engineers versus non-engineers, and (iii) the focal firm versus peers. Identification requires a single parallel-trends assumption (Olden and Møen, 2022), and Internet Appendix Figure B.1 shows no meaningful pre-trend divergence.

Column (1) in Panel A of Table 9 shows that engineers' overall ratings of their workplace increase significantly at the focal firm after program initiation, relative to engineers at peer firms.[24] This pattern is consistent with the internal survey results and suggests that mentorship strengthens perceived organizational support. Similarly, Column (2) shows that perceptions of the focal firm's corporate culture improve as well (using Glassdoor's "Culture & Values" rating). The culture patterns are intriguing, given prior work showing how persistent devaluation often drives members of URGs to exit STEM roles (Spoon et al., 2023).

In Columns (3)–(6), we examine which cultural values and norms underlie these overall cultural changes by looking at text-based measures derived from the open-ended responses in Glassdoor reviews.[25] The largest change occurs in the community-orientation value, with

[23]Previous studies have used the Glassdoor database to evaluate non-pecuniary benefits and in the process have shown that the crowd-sourced data is reliable and consistent with internal firm surveys (Graham et al., 2022a; Grennan, 2023; Liu et al., 2023; Martellini et al., 2024; Gornall et al., 2025).

[24]For brevity, we only report the coefficients on the main interaction terms of interest, which is the triple interaction term in most models with the exception of Columns (3) and (4) of Panel B, but the regression specifications include all the relevant individual regressors and pairwise interaction terms. We report the results of alternative specifications, including firm specific DiD estimations, in Internet Appendix A.

[25]These text-based measures are described in more detail in Internet Appendix E. Glassdoor has open-ended text for pros, cons, advice to management, and title. We parse the text into individual sentences and train separate hierarchical BERT-based multiclassifier models for each value that includes sentence type

more modest gains in adaptability, collaboration, and integrity.[26] These patterns align with the mechanisms identified in administrative data regarding expanded networks, stronger interpersonal ties, and reduced informational frictions.

Next, we repeat this exercise when examining patterns of engineer turnover between the focal firm and its peers. Using data from Revelio, a labor market intelligence platform that tracks employee movements across firms, we examine the relation between the mentorship program and the firm's ability to retain its human capital talent.[27] The measure we construct captures the number of employee departures that the firm experiences in a particular month within a specific occupation-region-gender-ethnicity-seniority group. We also develop these turnover metrics for the focal firm's peers and, as an alternative measure, we take the inverse hyperbolic sine transformation of these values.

Column (1) in Panel B of Table 9, which uses the non-transformed number of departures value, shows a decline in departures among the focal firm's engineers after the program begins, relative to engineers in peer firms. The results in Column (2), which use the transformed measure, yield a similar pattern and imply a roughly 3 p.p. decline in attrition. This estimate is economically meaningful given the high replacement costs associated with specialized engineering talent. In Columns (3) and (4), we restrict the sample to only named inventors on granted patents at either the focal firm or peer firms. All of these inventors are engineers, so we only conduct DiD tests for this sample. The estimates again point to lower attrition at the focal firm after the mentorship program, relative to its peer firms. Consistent with these patterns, among innovators tracked forward by the focal firm, they departed the firm at a rate 2.5 p.p. below comparable non-mentored innovators. While our empirical design does not allow us to attribute the gains in retention of the full population of engineers

information (i.e., pros vs. cons) and oversampling (to account for imbalance in categories). The individual cultural values align with the principal components of culture defined by O'Reilly et al. (1991) and used in Graham et al. (2022a).

[26]These estimates are robust to alternative specifications such as clustering by firm.

[27]The Revelio data is compiled using LinkedIn profiles and data from other employment-related social media platforms like Indeed. Using this data, we construct a measure of the number of employees who exit the firm each month within specific categories of individuals based on their occupation, the region in which they work, their gender, their ethnicity, and their level of seniority.

at the focal firm to the mentorship program, this evidence is consistent with the internal survey evidence and the notion that the program contributed to improved perceptions of the workplace and a lower desire to look elsewhere for work.

5.6 Additional Robustness and Alternative Explanations

As with our main results, we examine the sensitivity of our subsequent tests to alternative specifications. These analyses assess whether the patterns are similar across estimator choice, sample composition, and outcome definition. We report the mentor and mentee heterogeneity tests in Table A.6 and Table A.7. Table A.8 shows the results of similar tests when gender and ethnicity are evaluated individually or intersectionally. We also consider other variants of mentee heterogeneity, including age, geographic location, and timing of participation. These results are included in Table A.9 and Table A.10. Overall, the estimates are broadly similar across age, geography, and cohort timing. This evidence helps mitigate concerns that the results primarily reflect natural learning-by-doing among younger employees or unobserved regional factors.

In Table A.11, we report the robustness of the mentor-quality findings. In addition, we consider an alternative measure of mentor quality (i.e., using mentor fixed effects) in Table A.12, which shows that the results hold when using this alternative approach. We report mentor demographic characteristics in Table A.13, and we consider alternative heterogeneity analyses in Table A.14, which establish that mentor gender, job status (e.g., "top engineer" designation), and prior program involvement do not meaningfully alter the productivity gains due to mentorship, further suggesting that effective mentoring is not confined to a narrow set of individuals. We also verify the robustness of the social network mechanism in Table A.15, and we verify the robustness of the heterogeneity in spillovers to non-participants by demographic groups in Table A.16.

Additional robustness tests for the workplace perception and turnover analyses appear in Table A.17–Table A.20. Specifically, we show that workplace perceptions improved and

engineer turnover declined at the focal firm after the mentorship program began. These results are robust to clustering standard errors at the firm level and to alternative non-count based definitions of departure. Revelio produces these scaled external outflow measures to account for the fact that not all employees have LinkedIn profiles, so measures cannot aggregate into a firm-wide turnover rate. Finally, in Internet Appendix B, we conduct falsification tests to further support the robustness of our workplace perceptions results and our findings on engineer turnover.

We also assess alternative explanations, including whether the results simply reflect mentored workers being added to existing teams rather than generating new innovation. However, the patterns we observe in disclosure submissions suggest actual increases in firm-wide innovation productivity rather than inflation in team sizes on existing projects. While we show that our main results are robust to scaling by team size (see Columns (3) and (4) of Table 2), we conduct additional tests and find no support for the alternative team-size explanation.

First, team sizes for accepted patent applications remained statistically unchanged, averaging 4.62 before, 4.59 during, and 4.31 after the program (p-values from pairwise comparisons range from 0.317 to 0.920). Second, analysis of U.S. PTO data confirms that there was no differential increase in team sizes at the focal firm relative to its peer firms (see Table A.21). Third, we explore measures of aggregate innovative productivity and find a significant increase at the focal firm relative to its peers, coinciding with the implementation of the mentoring program (see Table A.22).

We interpret this aggregate analysis narrowly. Table A.22 is intended as a corroborative test of whether the participant-level gains reflect a reallocation of patenting activity from other employees, rather than as a causal estimate of changes to firm-wide activity. Because the analysis contains a single treated firm and public patent records cannot be linked to individual program participants, it cannot rule out every contemporaneous firm-specific shock. Importantly, however, common firm-wide changes in strategy, R&D investment, or

management cannot account for the main employee-level estimates, which compare mentored employees with control employees in the same firm and calendar quarter and absorb common shocks through year-quarter-by-cohort fixed effects. Moreover, as discussed in Section 3, we find no discrete change around program adoption in the firm's core strategy, R&D intensity, senior leadership, inventor incentives, patent-review process, or IP policy. We therefore view the aggregate results as evidence that the gains among mentored employees were not accompanied by an offsetting decline in firm-wide patent output. We do not attribute the full increase in aggregate patenting reported in Table A.22 to the mentorship program.

Finally, we examine two additional predictions of the process-learning mechanism using disclosure timing and collaborator composition. If mentorship accumulates process knowledge, mentees should generate subsequent disclosures more quickly because process-related uncertainty about how to draft and submit a disclosure has been reduced. In addition, mentees should rely less on their original collaborators, because they no longer need those collaborators to navigate the pipeline on their behalf. Consistent with both predictions, mentored innovators submit subsequent ideas about one quarter earlier, on average, and are more likely to file independently of their initial collaborators. These patterns are consistent with the process-learning hypothesis and are difficult to reconcile with a pure time-shift alternative, which predicts no change in either timing or collaborator dependence once the program ends.

6 An Illustrative Return-on-Investment Calculation

This section presents an illustrative ROI calculation using a framework that builds on established methodologies for assessing human-capital investments. We emphasize the *illustrative* nature of this analysis because, while our empirical design plausibly identifies the mentorship program's relation to invention disclosures and patents, valuing the broader organizational changes we document, including collaborator spillovers, culture, and retention, requires assumptions that lie outside what the experiment identifies. We therefore proceed in two steps:

we first report a direct calculation grounded in the identified patenting estimates, and we then present additional channels as bounded illustrations of the magnitudes those channels would need to reach to matter, and we are transparent about the assumptions each requires.

First, we estimate direct innovation-related value. Our most conservative estimates focus on patent productivity gains. Our estimates imply approximately one additional invention disclosure annually among mentored innovators, with approximately 31% converting to patent applications and 90% of those to granted patents. Valuing patents conservatively at $38,460 per patent based on application and attorney fees, this translates to approximately $6.8 million (633 mentees $\times$ 1 submissions $\times$ 0.28 conversion $\times$ $38,460 per patent). Accounting for average team size (3.08 named inventors), a baseline lower-bound estimate of the direct benefit is $2.2 million. Alternative valuation approaches for patents such as market-based metrics would yield much larger estimates.

Second, we estimate the value associated with broader program-related responses. We document two measurable channels. First, non-mentored innovators who later collaborate with mentored colleagues exhibit a 38% increase relative to their mean pre-exposure productivity. This pattern is consistent with knowledge diffusion through project ties, though we cannot rule out changes in project assignment or other unobserved factors. This pattern corresponds to approximately $8.4 million in additional value under the same conservative per-patent valuation used above. We treat this figure as an illustrative upper contribution from this channel rather than an identified estimate.

Third, we consider whether the program's positive relation with workplace satisfaction and retention might be economically meaningful. Engineer turnover at the firm declined by 3.3% following program implementation. Of course, we cannot causally attribute this aggregate change to the mentorship program, because the design does not include a comparison group of engineers exposed to variation in program availability.[28] As an illustrative

[28]We note that for mentored individuals tracked forward by the program administrators, the retention rate among them exceeded the average retention rate among similar inventors at the firm over the sample period, which is qualitatively consistent with retention being related to the program.

benchmark, however, if even 1% of the aggregate decline were attributable to the program, the implied replacement-cost savings would total approximately $45.9 million; if the estimated relation were confined to the mentored engineers alone, which data from the program administrators supports, the corresponding figure would be approximately $2.5 million. Additionally, our survey and Glassdoor evidence are consistent with enhanced organizational capital, an important component of firm value (Bloom et al., 2012b; Eisfeldt and Papanikolaou, 2014).

The program had low costs, just two part-time administrators, and mentor participation. One could worry that mentors' productivity would decrease, but contrary to that concern, we find that mentors' productivity temporarily increased from 2.6 to 3.7 submissions per quarter before returning to baseline levels after the program. This pattern is consistent with limited opportunity costs from mentor time allocation. Nevertheless, to present the most conservative estimates possible, we instead assume that participation replaced work hours and count that portion of the mentor's salary as a cost, which totals $0.75 million.

Thus, even using conservative assumptions and focusing primarily on direct benefits grounded in the estimates identified through our experimental design, the program's implied returns substantially exceed the measured implementation costs even before any of the additional channels are considered. The additional channels illustrate the magnitudes at which spillovers, retention, and culture would need to operate to bolster the total. Whether the indirect gains to organizational capital in fact operate at those magnitudes is a question the present design cannot fully resolve and one we flag as a productive direction for future work.

7 Scaling and Public Policy Considerations

Having documented patterns consistent with improvements in inventive output, process knowledge, workplace perceptions, retention, and ROI, we now turn to whether similar programs could scale. Replicating elements of this intervention at other firms appears feasible based on our analysis of deployability, external validity, and common frictions across firms.

7.1 Deployability

Deployability is itself a dimension of external validity. A mentorship intervention generalizes only if other organizations can feasibly implement it, and on this dimension the program's requirements were minimal. Implementation relied on two part-time administrators and volunteer mentors drawn from the firm's existing inventor community. Mentees committed four hours per week over a four-month cohort, but that time too was voluntary, and mentors' time, even under the conservative assumption that participation fully displaced work hours, totaled $0.75 million, an order of magnitude below our most conservative benefit estimates. The program required no dedicated technology, external consultants, or capital expenditure; its core components were a brief application form, group-based mentoring at a ratio of one mentor to four or five mentees, structured check-ins, and a closing showcase, which are administrative practices available to any organization with experienced inventors on staff. Nor does the design depend on scarce inputs. As we demonstrated, even mentees paired with below-median mentors in terms of quality also exhibit significant productivity gains in the short and long term, indicating that success does not hinge on access to exceptional mentoring talent. The binding constraint on adoption therefore appears to be managerial initiative rather than resources or specialized talent, consistent with behavioral rather than resource-based explanations for the limited diffusion of value-enhancing management practices (Benson et al., 2019; Mishra et al., 2022; Gertler et al., 2025). Indeed, replicating elements of this intervention at other firms appears feasible. Since 2021, more than forty major patent-filing firms have signed the "Diversity Pledge" committing to measurable steps to broaden participation in innovation (Harrison, 2021).

7.2 External Validity

But before recommending greater adoption of structured mentorship programs, it is essential to evaluate external validity and the "voltage effect" of proposed solutions (Chien, 2019, 2024). That is, whether the patterns observed in this setting would persist under broader

implementation (Muralidharan and Niehaus, 2017; Al-Ubaydli et al., 2017; List, 2022, 2024).

Our theory helps discipline which mechanisms are more likely to travel across settings. In the List (2020) framework, preferences such as creativity and career goals may be less malleable in the short run, while beliefs, such as confidence in one's idea, and constraints, such as institutional support, are more likely to vary across firms. Our evidence is most consistent with mentoring relaxing constraints rather than altering underlying preferences, suggesting potential portability to other settings where similar constraints bind.

We evaluate six dimensions of external validity. First, the estimates are broadly similar across estimation approaches and sample restrictions, which helps mitigate concerns about false positives. The persistence of the productivity estimates over longer horizons is also consistent with the program changing more than short-run disclosure timing. Although the full economic value of innovation typically unfolds over many years, our evidence suggests that knowledge transfer leads to significantly greater disclosure output in the medium term, higher conversion ratios from ideas to high-quality patents, and broader responses to the program among non-mentored employees who become later collaborators. These outcomes are useful leading indicators of longer-run gains to innovative productivity, even if the full realized value of the underlying inventions or program will take longer to observe.

Second, because the program was open to all innovators, we can examine whether the estimates differ across demographic groups, baseline productivity, age, and geographic location. This breadth allows us to assess whether the main patterns are concentrated among a narrow subset of workers or appear across a wider range of innovators.

Third, our analysis of non-participants who later collaborate with mentored workers examines whether program-related changes extend beyond direct participants. Non-participants who later collaborate with mentees exhibit higher disclosure activity, which is more consistent with broader program reach than with a simple crowd-out interpretation. One important boundary condition is that the positive association among non-participants is larger for well-represented workers than for underrepresented workers. This contrasts with the direct

participant results, where underrepresented mentees exhibit somewhat larger increases. The contrast suggests that passive diffusion through collaboration networks may not be sufficient to broaden access, and that targeted outreach may remain important as firms scale similar programs.

Fourth, the estimates across cohorts provide little evidence of attenuation as the program expands, a central concern in the literature on scaling interventions (DellaVigna and Linos, 2022). Fifth, the estimates are not concentrated only among mentees paired with the highest-output mentors: mentees paired with below-median mentors also exhibit positive and statistically significant increases in productivity. This pattern suggests that successful program implementation is not necessarily dependent on having a large number of especially high-performing mentors, although mentor quality remains an important design margin.

Sixth, we provide firm-level evidence that helps distinguish broader capacity expansion from within-firm reallocation. After program implementation, the focal firm exhibits increases in both the quantity and quality of granted patents relative to peer firms, including a higher share of top-cited inventions (see Internet Appendix Table A.22). Approximately 10% of patents produced during the program period fall into the highest-quality category based on forward-citation concentration. These firm-level patterns are consistent with the individual-level results and help mitigate the concern that the program merely reallocates inventive activity across workers. While the results suggest that the gains associated with a targeted, small-scale mentoring intervention can diffuse through collaboration networks and contribute to broader increases in the firm's inventive capacity, because the aggregate analysis involves a single treated firm, we view this evidence more narrowly and mostly as supportive, rather than definitive, evidence of firm-level scalability.

Finally, a related external-validity concern is that the program ran during a period in which remote work was unusually common, and a program that pairs employees with a senior mentor may be at its most valuable when informal in-person interactions are scarcest. If so, the estimates may reflect a period-specific complementarity between mentoring and

remote work rather than a portable feature of the intervention. We view this concern as legitimate but mitigated by three pieces of evidence. First, the firm's work-from-home (WFH) policy during the sample period was one of flexible hybrid work rather than fully remote arrangements: the official policy emphasized outcomes over location, but managers retained discretion to set team expectations, and the feasible arrangement varied with role, laboratory or equipment requirements, and security or regulatory constraints. This policy is broadly representative of large technology firms during the sample period (Barrero et al., 2023). Second, most mentoring sessions during the program were hybrid, indicating that the intervention was not designed for and did not depend on a fully remote environment. Third, and most substantively, we examine whether our estimates differ for employees more or less impacted by the shift to remote work. Age has been shown to be a factor in how well workers adapted to and could succeed in hybrid arrangements (Makridis and Schloetzer, 2026). We find no difference in program effectiveness by age (see Column (4) of Table 4 and Appendix Table A.9). Taken together, this evidence suggests the productivity gains associated with the mentorship program are unlikely to be an artifact of the WFH environment in which it was implemented.

7.3 Are the Alleviated Frictions Unusual? Benchmarking Against Comparables

A potential scaling concern is that our estimates reflect conditions specific to the focal firm, such as unusually severe organizational frictions or an unusually large stock of latent ideas, rather than constraints common to R&D-intensive firms. Therefore, we benchmark these constraints and find little support for such arguments. Specifically, the frictions our mechanism tests identify, including limited procedural knowledge of the disclosure process, restricted access to experienced inventor networks, and weak feedback on early-stage ideas, are pervasive across large technology firms, though they manifest through different channels across firm types. Among mature, diversified incumbents with long-established patenting infrastructures (e.g., IBM), documented frictions include the rationing of in-house patent

counsel across tens of thousands of engineers, invention review committees whose expertise skews toward legacy product lines, the integration of serially acquired engineering teams into unfamiliar disclosure processes, and races between filing and public disclosure through standards-body participation (Cohen et al., 2000; Hall and Ziedonis, 2001). Among younger, software-centric platforms (e.g., Meta), a complementary set of frictions exist such as promotion systems that reward shipped products rather than invention disclosures, product cycles shorter than patent prosecution timelines, open-source norms that discourage engagement with the patent system, and a preference for trade secrecy where infringement is difficult to detect (Arora et al., 2018; Contigiani et al., 2018). That substantial idea-to-patent frictions arise through distinct mechanisms in both types of superstar firms supports the notion that such frictions are a near-universal feature of large-firm innovation, precisely the class of constraint that List (Forthcoming) argues is most likely to generalize, rather than a peculiarity of our setting.

In addition, observable characteristics of the focal firm suggest that its frictions were below, not above, those of the typical superstar firm. Throughout the sample period, the firm maintained stable monetary incentives for patenting, an experienced and disciplined internal screening process, and mature prosecution capabilities. A firm suffering atypically severe organizational dysfunction in its inventive pipeline would be unlikely to exhibit this combination of sustained incentives, institutionalized review, and above-average prosecution outcomes. This is the sense in which we view the setting as a demanding benchmark. The program generated measurable gains in an environment where the formal process for converting ideas into patents was already comparatively strong, implying that the binding constraints it relaxed were informational and relational rather than procedural, and thus likely to bind at least as tightly elsewhere.

The dynamics of our estimates are similarly difficult to reconcile with the alternative hypothesis that the firm held an unusually abundant reservoir of latent ideas awaiting release. If the program had merely drawn down an atypically large pre-existing backlog, we would

expect our results to be concentrated in early cohorts and to mean revert toward zero as the backlog depleted. Instead, we observe the opposite pattern. Our estimates show little attenuation across the seven cohorts and productivity gains that persist over long horizons. Together with the firm-level evidence that patent quantity and quality rose relative to peer firms rather than reflecting reallocation, these patterns indicate that the program expanded inventive capacity by relaxing common frictions in the idea-to-patent pipeline, not by harvesting an idiosyncratic surplus of dormant ideas.

In conclusion, our findings suggest that the mentorship program helped participants overcome procedural and informational constraints, and it helped them update their beliefs regarding which ideas are likely to succeed through the patenting process. In the List (2020) framework, this distinction between preferences, beliefs, and constraints provides a way to think about where the program is most likely to travel. It reveals that settings in which workers have latent ideas but face organizational frictions to moving those ideas through the patenting process are the ones where our findings are the most likely to generalize.

8 Conclusion

This study examines whether targeted mentoring can improve innovative productivity and inclusion within a superstar technology firm. Using proprietary data on the innovation pipeline from initial idea disclosure through patent application and the staggered implementation of a mentoring program, we document several findings.

First, mentorship is associated with economically meaningful increases in innovative productivity in both the short and long run. These findings are consistent with a process-learning interpretation and are difficult to reconcile with a pure temporary effort-reallocation explanation. Second, mentees exhibit higher idea-to-patent conversion rates and stronger measures of innovation quality. Third, the estimates are largest among innovators without prior patenting experience and those paired with high-output mentors, with especially large increases for innovators from underrepresented demographic groups. That the largest gains

accrue to innovators without prior patenting experience is difficult to reconcile with a view in which inventive productivity is bounded primarily by an immutable distribution of individual creative ability, and points instead to organizational and informational constraints as first-order. Fourth, non-participants who later collaborate with mentees also exhibit increases in disclosure activity, a pattern more consistent with genuine knowledge transfer than with a simple redistribution of attention or credit. Fifth, the program is accompanied by broader organizational changes, including improved perceptions of corporate culture and lower engineer turnover.

These findings point to structured mentorship as a potential way to relax the specific organizational frictions that prevent many "innovators" from becoming "named inventors": limited procedural knowledge of the disclosure process, unfamiliarity with the writing and framing that the firm's patent review process rewards, weak access to senior colleagues who can vouch for an idea's novelty and relevance, and reluctance to expose an unfinished idea to institutional scrutiny. Evidence from mentor-mentee matching, idea-quality tests, and collaboration patterns is consistent with two related mechanisms: targeted knowledge transfer and expanded access to invention networks, which map directly onto the internal frictions. The program's design may be important for understanding these patterns. In addition to group mentoring, the program incorporates accelerator-like features such as structured accountability, peer learning, credible role models, constructive feedback, and a grand-pitch-style event. These features help innovators navigate the patenting process for the first time and gain the skills to repeatedly advance their ideas within the organization. In this sense, targeted mentoring may offer a scalable way to unlock latent inventive capacity in R&D-intensive firms, particularly where informational and procedural frictions limit participation in innovation.

From an applied perspective, even under conservative valuation assumptions applied only to the identified patenting outcomes, the program realized a sizable ROI, and the accompanying improvements in workplace perceptions and retention strengthen the economic

case. Because the invention flow at the focal firm contributes measurably to the aggregate stock of innovation capital linked to long-run productivity growth (Kogan et al., 2017b; Bloom et al., 2020), the frictions we document and the modest, scalable intervention that alleviates them have implications that extend beyond the private returns to a single firm. More broadly, the results speak to a central tension that corporate leaders face when deciding where and how to invest in innovation. Incumbent firms are often viewed as poorly positioned to nurture new ideas internally because existing processes can favor established projects over emerging opportunities. Our research points to a complementary possibility. Internal innovation may be limited not only by the dilemmas caused by strategic inertia but also by fixable organizational frictions.

Targeted mentoring offers one way to mitigate these frictions and increase the quantity and quality of inventions within large incumbent firms. Moreover, because the frictions our intervention alleviates are widely documented at comparable technology firms while the focal firm's own patenting process and infrastructure were, if anything, comparatively strong, this setting likely represents a demanding rather than favorable benchmark for the intervention. For firms and policymakers seeking to increase the returns to R&D, the evidence suggests that targeted mentoring is a promising, low-cost approach to reducing frictions in the inventive pipeline and broadening participation in innovation. An important question for future work is whether the short- and medium-run patterns documented here translate into breakthrough inventions and other long-run outcomes.

References

Abrams, David, Ufuk Akcigit, Jillian Grennan. 2023. Patent value and citations: Creative destruction or strategic disruption. *University of Chicago Working Paper* .

Adams, Renee, Michelle Lowry. 2022. What's good for women is good for science: Evidence from the american finance association. *Review of Corporate Finance Studies* (11) 554–604.

Adelino, Manuel, Song Ma, David Robinson. 2017. Firm age, investment opportunities and job creation,. *Journal of Finance* **72** 999–1038.

Akcigit, Ufuk, William R Kerr. 2018. Growth through heterogeneous innovations. *Journal of Political Economy* **126**(4) 1374–1443.

Al-Ubaydli, Omar, John A. List, Dana L. Suskind. 2017. What can we learn from experiments? understanding the threats to the scalability of experimental results. *American Economic Review* **107**(5) 282—-286.

Alesina, Alberto, Eliana La Ferrara. 2005. Ethnic diversity and economic performance. *Journal of Economic Literature* **43**(3) 762—-800.

Allen, Tammy D. 2007. Mentoring relationships from the perspective of the mentor. *The handbook of mentoring at work: Theory, research, and practice* 123–147.

Amabile, Teresa M. 1983. The social psychology of creativity: A componential conceptualization. *Journal of Personality and Social Psychology* **45**(2) 357–376.

Anderson, Michael L. 2008. Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American Statistical Association* **103**(484) 1481–1495.

Aneja, Abhay, Oren Reshef, Gauri Subramani. 2024. Attrition and the gender patenting gap. *Review of Economics and Statistics* .

Angeli, Deivis, Ieda Matavelli, Fernando Secco. 2025. Expected discrimination and job search.

Arora, Ashish, Sharon Belenzon, Andrea Patacconi. 2018. The decline of science in corporate R&D. *Strategic Management Journal* .

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, John Van Reenen. 2020. The Fall of the Labor Share and the Rise of Superstar Firms. *Quarterly Journal of Economics* **135**(2) 645–709.

Ayyagari, Meghana, Asli Demirguc-Kunt, Vojislav Maksimovic. 2024. The rise of star firms: Intangible capital and competition. *Review of Financial Studies* **37**(3) 882–949.

Azmat, Ghazala, Vicente Cuñat, Emeric Henry. 2025. Gender promotion gaps and career aspirations. *Management Science* **71**(3) 2127–2141.

Babina, Tania. 2019. Destructive creation at work: How financial distress spurs entrepreneurship. *Review of Financial Studies* **33**(9) 4061–4101.

Babina, Tania, Asaf Bernstein, Filippo Mezzanotti. 2023. Financial disruptions and the organization of innovation: Evidence from the great depression. *Review of Financial Studies* **36**(11) 4271–4317.

Babina, Tania, Sabrina T. Howell. 2024. Entrepreneurial spillovers from corporate ramp;d. *Journal of Labor Economics* **42**(2) 469–509.

Baker, Andrew C, David F Larcker, Charles CY Wang. 2022. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* **144**(2) 370–395.

Barrero, José María, Nicholas Bloom, Steven J. Davis. 2023. The evolution of work from home. *Journal of Economic Perspectives* **37**(4) 23–49.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, John Van Reenen. 2019. Who becomes an inventor in america? The importance of exposure to innovation. *Quarterly Journal of Economics* 647–713.

Bena, Jan, Isil Erel, Daisy Wang, Michael Weisbach. 2026. Relationship-specific investments and firms' boundaries: Evidence from textual analysis of patents. *Working paper* .

Bena, Jan, Kai Li. 2014. Corporate innovations and mergers and acquisitions. *Journal of Finance* **69** 1923–1960.

Bena, Jan, Hernán Ortiz-Molina, Elena Simintzi. 2022. Shielding firm value: Employment protection and process innovation. *Journal of Financial Economics* **146**(2) 637–664.

Bena, Jan, Elena Simintzi. 2025. Machines could not compete with chinese labor. evidence from u.s. firms' innovation. *Review of Finance* **29** 1619–1661.

Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, Daniel Wolfenzon. 2022. Do firms respond to gender pay gap transparency? *Journal of Finance* **77**(4) 2051–2091.

Bennedsen, Morten, Margarita Tsoutsoura, Daniel Wolfenzon. 2019. Drivers of effort: Evidence from employee absenteeism. *Journal of Financial Economics* **133**(3) 658–684.

Bennett, Benjamin, Isil Erel, Lea Stern, Zexi Wang. 2026. Paid leave pays off: The effects of paid family leave on firm performance. *Review of Finance* .

Benson, Alan, Danielle Li, Kelly Shue. 2019. Promotions and the peter principle. *Quarterly Journal of Economics* **134**(4) 2085–2134.

Benson, Alan, Danielle Li, Kelly Shue. 2026. "potential" and the gender promotion gap. *American Economic Review* **116**(2) 375—417.

Berk, Ales, Robert Kase. 2010. Establishing the value of flexibility created by training: Applying real options methodology to a single hr practice. *Organization Science* **21**(3) 765–780.

Bernstein, Shai, Emanuele Colonnelli, Davide Malacrino, Tim McQuade. 2022a. Who creates new firms when local opportunities arise? *Journal of Financial Economics* **143**(1) 107–130.

Bernstein, Shai, Rebecca Diamond, Abhisit Jiranaphawiboon, Timothy McQuade, Beatriz Pousada. 2022b. The contribution of high-skilled immigrants to innovation in the united states. Working Paper 30797, National Bureau of Economic Research.

Bernstein, Shai, Arthur Korteweg, Kevin Laws. 2017. Attracting early-stage investors: Evidence from a randomized field experiment. *Journal of Finance* **72**(2) 509–538.

Bernstein, Shai, Richard R Townsend, Ting Xu. 2023. Flight to safety: How economic downturns affect talent flows to startups. *Review of Financial Studies* **37**(3) 837–881.

Bertrand, Marianne, Antoinette Schoar. 2003. Managing with style: The effect of managers on firm policies. *Quarterly Journal of Economics* **118**(4) 1169–1208.

Blau, Francine D., Janet M. Currie, Rachel T. A. Croson, Donna K. Ginther. 2010. Can mentoring help female assistant professors? interim results from a randomized trial. *American Economic Review* **100**(2) 348—-352.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, Michael Webb. 2020. Are ideas getting harder to find? *American Economic Review* **110**(4) 1104—-1144.

Bloom, Nicholas, John Van Reenen. 2007. Measuring and explaining management practices across firms and nations. *Quarterly Journal of Economics* **122** 1351–1408.

Bloom, Nicholas, Raffaella Sadun, John Van Reenen. 2012a. Americans do IT better: US multinationals and the productivity miracle. *American Economic Review* **102**(1) 167—201.

Bloom, Nicholas, Raffaella Sadun, John Van Reenen. 2012b. The organization of firms across countries. *Quarterly Journal of Economics* **127**(4) 1663–1705.

Bloom, Nicholas, Raffaella Sadun, John Van Reenen. 2016. Management as a technology? Tech. rep., National Bureau of Economic Research.

Bol, Jasmijn, Lisa LaViers, Jason Sandvik. 2023. Creativity contests: An experimental investigation of eliciting employee creativity. *Journal of Accounting Research* **61**(1) 47–94.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, Andrei Shleifer. 2016. Stereotypes. *The Quarterly Journal of Economics* **131**(4) 1753–1794.

Botelho, Tristan L., Daniel C. Fehder, Yael V. Hochberg. 2026. Innovation-driven entrepreneurship. *Journal of Economic Literature* **64**(1) 89—-140.

Cai, Wei, Andrea Prat, Jiehang Yu. 2026. Measuring organizational capital. *NBER Working Paper No. 35039* .

Calaway, Ian. 2024. Early mentors for exceptional students. *Working paper* .

Calder-Wang, Sophie, Paul A. Gompers. 2021. And the children shall lead: Gender diversity and performance in venture capital. *Journal of Financial Economics* **142**(1) 1–22.

Callaway, Brantly, Pedro HC Sant'Anna. 2021. Difference-in-differences with multiple time periods. *Journal of econometrics* **225**(2) 200–230.

Campello, Murillo, Janet Gao. 2017. Customer concentration and loan contract terms. *Journal of Financial Economics* **123**(1) 108–136.

Chemmanur, Thomas J., Elena Loutskina, Xuan Tian. 2014. Corporate venture capital, value creation, and innovation. *Review of Financial Studies* **27**(8) 2434–2473.

Chen, Jiafeng, Jonathan Roth. 2024. Logs with zeros? some problems and solutions **139**(2) 891–936.

Chien, Colleen. 2019. Rigorous policy pilots: Experimentation in the administration of patent law. *Iowa Law Review* **104** 2313.

Chien, Colleen. 2024. Redefining progress and the case for diversity in innovation and inventing. *UCLA Law Review* **70** 540–624.

Chien, Colleen, Joan Farre-Mensa, Jillian Grennan. 2026. The returns to recognition: Patent grants and inventors' life trajectories. *Working Paper* .

Chien, Colleen, Jillian Grennan. 2024. Unpacking the innovator-inventor gap: Evidence from engineers. *Working Paper* .

Cho, DuckKi, Lyungmae Choi, Michael Hertzel, Jessie Jiaxu Wang. 2026. It's not who you know—it's who knows you: Employee social capital and firm performance. *Management Science* **72**(5) 3948–3973.

Christensen, Clayton M. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail.*. Harvard Business School Press, Bosotn, MA.

Coffman, Katherine Baldiga. 2014. Evidence on self-stereotyping and the contribution of ideas. *Quarterly Journal of Economics* **129**(4) 1625–1660.

Cohen, Susan, Daniel C. Fehder, Yael V. Hochberg, Fiona Murray. 2019. The design of startup accelerators. *Research Policy* **48**(7) 1781–1797.

Cohen, Wesley M, Richard R Nelson, John P Walsh. 2000. Protecting their intellectual assets: Appropriability conditions and why u.s. manufacturing firms patent (or not). Working Paper 7552, National Bureau of Economic Research.

Contigiani, Andrea, David H. Hsu, Iwan Barankay. 2018. Trade secrets and innovation: Evidence from the "inevitable disclosure" doctrine. *Strategic Management Journal* **39**(11) 2921–2942.

Cook, Lisa, Matthew Marx, Emmanuel Yimfor. 2025. Funding black high-growth startups. *Journal of Finance* .

Cook, Lisa D. 2018. *Unequal Opportunity: The Innovation Gap in Pink and Black*. Does America Need More Innovators?, MIT Press.

Cunningham, Colleen, Florian Ederer, Song Ma. 2021. Killer acquisitions. *Journal of Political Economy* **129**(3) 649–702.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, Javier Miranda. 2020. Changing business dynamism and productivity: Shocks versus responsiveness. *American Economic Review* **110**(12) 3952—-3990.

DellaVigna, Stefano, Elizabeth Linos. 2022. Rcts to scale: Comprehensive evidence from two nudge units. *Econometrica* **90**(1) 81–116.

Dessein, Wouter, Andrea Prat. 2022. Organizational capital, corporate leadership, and firm dynamics. *Journal of Political Economy* **130**(6) 1477–1536.

Dossi, Gaia. 2024. Race and science. *EIEF Working Paper* .

Eisfeldt, Andrea, Dimitris Papanikolaou. 2013. Organization capital and the cross-section of expected returns. *Journal of Finance* **68**(4) 1365–1406.

Eisfeldt, Andrea L, Dimitris Papanikolaou. 2014. The value and ownership of intangible capital. *American Economic Review* **104**(5) 189–194.

Ewens, Michael, Ryan H. Peters, Sean Wang. 2024. Measuring intangible capital with market prices. *Management Science* .

Ewens, Michael, Richard R. Townsend. 2020. Are early stage investors biased against women? *Journal of Financial Economics* **135**(3) 653–677.

Farre-Mensa, Joan, Deepak Hegde, Alexander Ljungqvist. 2020. What is a patent worth? Evidence from the U.S. patent "lottery". *Journal of Finance* **75** 639–682.

Fitzgerald, Tristan, Xiaoding Liu. 2025. Shared culture and technological innovation: Evidence from corporate rd teams. *Working paper* .

Fleming, Lee. 2001. Recombinant uncertainty in technological search. *Management Science* **47**(1) 117–132.

Fort, Teresa C, Nathan Goldschlag, Jack Liang, Peter K Schott, Nikolas Zolas. 2026. Growth is getting harder to find, not ideas. *NBER Working Paper No. 35182* .

Gertler, Paul, Sean Higgins, Ulrike Malmendier, Waldo Ojeda. 2025. Do behavioral frictions prevent firms from adopting profitable opportunities? Working Paper 33387, National Bureau of Economic Research.

Gertsberg, Marina. 2026. The unintended consequences of metoo: Evidence from research collaborations in economics and finance. *Journal of Finance* .

Ginther, Donna K, Janet M Currie, Francine D Blau, Rachel TA Croson. 2020. Can mentoring help female assistant professors in economics? an evaluation by randomized trial. *AEA Papers and Proceedings*, vol. 110. 205–09.

Goldin, Claudia, Cecilia Rouse. 2000. Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review* **90**(4) 715—741.

Goldin, Ian, Pantelis Koutroumpis, François Lafond, Julian Winkler. 2024. Why is productivity slowing down? *Journal of Economic Literature* **62**(1) 196—268.

Gonzalez-Uribe, Juanita, Edgar Kausel, Michael Letherbee. 2026. Better to flee from freedom? testing the effects of structured accountability on new venture performance. *LSE Working Paper* .

Gonzalez-Uribe, Juanita, Michael Leatherbee. 2017. The effects of business accelerators on venture performance: Evidence from start-up chile. *Review of Financial Studies* **31**(4) 1566–1603.

Gornall, Will, Oleg R. Gredil, Sabrina T. Howell, Xing Liu, Jason Sockin. 2025. Do employees cheer for private equity? the heterogeneous effects of buyouts on job quality. *Management Science* **0**(0) null.

Gornall, Will, Ilya A. Strebulaev. 2025. Gender, race, and entrepreneurship: A randomized field experiment on venture capitalists and angels. *Management Science* **71**(6) 5308–5327.

Gorton, Gary B., Jillian Grennan, Alexander K. Zentefis. 2022. Corporate culture. *Annual Review of Financial Economics* **14**(1) 535–561.

Graham, John R., Jillian Grennan, Campbell R. Harvey, Shivaram Rajgopal. 2022a. Corporate culture: Evidence from the field. *Journal of Financial Economics* **146**(2) 552–593.

Graham, John R., Jillian A. Grennan, Campbell R. Harvey, Shivaram Rajgopal. 2022b. Corporate culture: The interview evidence. *Journal of Applied Corporate Finance* **34**(4) 22–41.

Graham, John R., Campbell R. Harvey, Shiva Rajgopal. 2005. The economic implications of corporate financial reporting. *Journal of Accounting and Economics* **40**(1) 3–73.

Grennan, Jillian. 2019. Dividend payments as a response to peer influence. *Journal of Financial Economics* **131**(3) 549–570.

Grennan, Jillian. 2023. A corporate culture channel: How increased shareholder governance reduces firm value. Working Paper.

Grullon, Gustavo, Yelena Larkin, Roni Michaely. 2019. Are us industries becoming more concentrated? *Review of Finance* **23**(4) 697–743.

Gruner, Richard. 2021. Filtering innovation wheat from chaff: Quick citations as indicators of patented invention value. *UIC Review of Intellectual property Law* **20**(3) 202–231.

Guiso, Luigi, Paola Sapienza, Luigi Zingales. 2015. The value of corporate culture. *Journal of Financial Economics* **117**(1) 60–76.

Gutner, Toddi. 2009. Finding anchors in the storm: Mentors. *The Wall Street Journal* .

Hall, Bronwyn H., Rosemarie Ham Ziedonis. 2001. The patent paradox revisited: An empirical study of patenting in the us semiconductor industry, 1979–1995. *RAND Journal of Economics* **32**(1) 101–128.

Harrison, Suzanne. 2021. Increasing diversity in innovation. URL increasingdii.org/pledge/.

Harvey, Campbell R., Yan Liu, Heqing Zhu. 2015. ... and the cross-section of expected returns. *Review of Financial Studies* **29**(1) 5–68. doi:10.1093/rfs/hhv059.

He, Bianca, Lauren I Mostrom, Amir Sufi. 2024. Investing in customer capital. Working Paper 33171, National Bureau of Economic Research.

Heath, Davidson, Christopher Mace. 2020. The strategic effects of trademark protection. *Review of Financial Studies* **33**(4) 1848–1877.

Heath, Davidson, Nathan Seegert, Jeffrey Yang. 2023. Team production and the homophily trap: Evidence from open source software. *Available at SSRN 4655458* .

Hebert, Camille. 2025. Gender stereotypes and entrepreneur financing. *Review of Financial Studies* .

Herbst, Daniel, Alexandre Mas. 2015. Peer effects on worker output in the laboratory generalize to the field. *Science* **350**(6260) 545–549.

Hoberg, Gerard, Gordon Phillips. 2025. Scope, scale, and concentration: The 21st-century firm. *Journal of Finance* **80**(1) 415–466.

Hochberg, Yael, Ali Kakhbod, Peiyao Li, Kunal Sachdeva. 2026. Are patents with female inventors undercited? evidence from text estimation. *Journal of Financial Economics* .

Hochberg, Yael V., Daniel C. Fehder. 2015. Accelerators and ecosystems. *Science* **348**(6240) 1202–1203. doi:10.1126/science.aab3351.

Hong, Lu, Scott E. Page. 2001. Problem solving by heterogeneous agents. *Journal of Economic Theory* **97**(1) 123–163.

Howell, Sabrina T. 2020. Reducing information frictions in venture capital: The role of new venture competitions. *Journal of Financial Economics* **136**(3) 676–694.

Howell, Sabrina T. 2021. Learning from feedback: Evidence from new ventures. *Review of Finance* **25**(3) 595–627.

Howell, Sabrina T., Ramana Nanda. 2023. Networking frictions in venture capital, and the gender gap in entrepreneurship. *Journal of Financial and Quantitative Analysis* 1–29.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, Peter J. Klenow. 2019. The allocation of talent and U.S. economic growth. *Econometrica* **87**(5) 1439–1474.

Hu, Allen, Song Ma. 2025. Persuading investors: A video-based study. *Journal of Finance* **80**(5) 2639–2688.

Jeffers, Jessica S. 2023. The impact of restricting labor mobility on corporate investment and entrepreneurship. *Review of Financial Studies* **37**(1) 1–44. doi:10.1093/rfs/hhad054.

Jones, Benjamin F. 2009. The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? *Review of Economic Studies* **76**(1) 283–317.

Jones, Charles I. 1995. R&d-based models of economic growth. *Journal of Political Economy* **103**(4) 759–784.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, Ahmed Tahoun. 2025. The diffusion of new technologies. *The Quarterly Journal of Economics* .

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, Matt Taddy. 2021. Measuring technological innovation over the long run. *American Economic Review: Insights* **3**(3) 303–20.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, Noah Stoffman. 2017a. Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics* **132**(2) 665–712.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, Noah Stoffman. 2017b. Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics* **132**(2) 665–712.

Kogan, Leonid, Dimitris Papanikolaou, Noah Stoffman. 2020. Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy* **128** 855–906.

Koning, Rembrand, Sampsa Samila, John-Paul Ferguson. 2021. Who do we invent for? patents by women focus more on women's health, but few women get to invent. *Science* **372**(6548) 1345–1348.

Kuhn, Jeffrey, Neil Thompson. 2019. How to measure and draw causal inferences with patent scope. *International Journal of the Economics of business* **26** 5–38.

Kwon, Spencer Y., Yueran Ma, Kaspar Zimmermann. 2024. 100 years of rising corporate concentration. *American Economic Review* **114**(7) 2111–40.

Lagaras, Spyridon, Maria-Teresa Marchica, Elena Simintzi, Margarita Tsoutsoura. 2025. Women in the financial sector .

Larkin, Yelena. 2013. Brand perception, cash flow stability, and financial policy. *Journal of Financial Economics* **110**(1) 232–253.

Lazear, Edward P, Kathryn L Shaw, Christopher T Stanton. 2015. The value of bosses. *Journal of Labor Economics* **33**(4) 823–861.

Lerner, Josh, Ulrike Malmendier. 2013. With a little help from my (random) friends: Success and failure in post-business school entrepreneurship. *Review of Financial Studies* **26**(10) 2411–2452.

Lerner, Josh, Amit Seru. 2022. The use and misuse of patent data: Issues for finance and beyond. *Review of Financial Studies* **35**(6) 2667–2704.

Lerner, Joshua. 1994. The importance of patent scope: An empirical analysis. *RAND Journal of Economics* **25**(2) 319–333.

Li, Kai, Feng Mai, Rui Shen, Xinyan Yan. 2020. Measuring corporate culture using machine learning. *Review of Financial Studies* **34**(7) 3265–3315.

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

List, John. Forthcoming. *Experimental Economics: Theory and Practice*. University of Chicago Press.

List, John A. 2020. Non est disputandum de generalizability? A glimpse into the external validity trial. *NBER Working Paper No. 27535* .

List, John A. 2022. *The voltage effect: How to make good ideas great and great ideas scale*. Crown Currency.

List, John A. 2024. Optimally generate policy-based evidence before scaling. *Nature* **626**(7999) 491–499.

List, John A., Azeem M. Shaikh, Yang Xu. 2019. Multiple hypothesis testing in experimental economics. *Experimental Economics* **22**(4) 773–793.

Liu, Tim, Christos A Makridis, Paige Ouimet, Elena Simintzi. 2023. The distribution of nonwage benefits: maternity benefits and gender diversity. *Review of Financial Studies* **36**(1) 194–234.

Lyle, David S, John Z Smith. 2014. The effect of high-performing mentors on junior officer promotion in the us army. *Journal of Labor Economics* **32**(2) 229–258.

Ma, Song. 2020. The life cycle of corporate venture capital. *Review of Financial Studies* **33**(1) 358–394.

Makridis, Christos. 2018. Does culture pay? evidence from crowdsourced employee engagement data. *Stanford Working Paper* .

Makridis, Christos A., Jason Schloetzer. 2026. How do different remote work arrangements impact employee job satisfaction and retention? *Management Science* **0**(0) null.

Martellini, Paolo, Todd Schoellman, Jason Sockin. 2024. The global distribution of college graduate quality. *Journal of Political Economy* **132**(2) 434–483.

Mezzanotti, Filippo. 2021. Roadblock to innovation: The role of patent litigation in corporate R&D. *Management Science* 7291–7950.

Mezzanotti, Filippo, Timothy Simcoe. 2025. Research and/or development? financial frictions and innovation investment. *Journal of Finance* .

Mishra, Prachi, Nagpurnanand Prabhala, Raghuram G Rajan. 2022. The relationship dilemma: Why do banks differ in the pace at which they adopt new technology? *Review of Financial Studies* **35**(7) 3418–3466.

Muralidharan, Karthik, Paul Niehaus. 2017. Experimentation at scale. *Journal of Economic Perspectives* **31**(4) 103—124.

NCSES. 2023. Diversity and stem: Women, minorities, and persons with disabilities 2023. *Special Report National Science Foundation 23-315* .

Nishesh, Naman, Paige Ouimet, Elena Simintzi. 2024. Labor and corporate finance. *Handbook of Corporate Finance*, chap. 17. Edward Elgar Publishing, 647–673.

Olden, Andreas, Jarle Møen. 2022. The triple difference estimator. *The Econometrics Journal* **25**(3) 531–553.

O'Reilly, Charles A., Jennifer Chatman, David F. Caldwell. 1991. People and organizational culture: A profile comparison approach to assessing person-organization fit. *The Academy of Management Journal* **34**(3) 487–516.

Pakes, Ariel. 1986. Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica* **54**(4) 755–784.

Rambachan, Ashesh, Jonathan Roth. 2023. A more credible approach to parallel trends. *Review of Economic Studies* **90**(5) 2555–2591.

Rockoff, Jonah E. 2008. Does mentoring reduce turnover and improve skills of new employees? evidence from teachers in new york city. Tech. rep., National Bureau of Economic Research.

Romer, Paul M. 1990. Endogenous technological change. *Journal of Political Economy* **98**(5) S71–S102.

Sandvik, Jason J, Richard E Saouma, Nathan T Seegert, Christopher T Stanton. 2020. Workplace knowledge flows. *Quarterly Journal of Economics* **135**(3) 1635–1680.

Sandvik, Jason J, Richard E Saouma, Nathan T Seegert, Christopher T Stanton. 2025. Should human capital development programs be voluntary or mandatory? Evidence from a field experiment on mentorship. *Management Science* .

Serrano, Carlos J. 2010. The dynamics of the transfer and renewal of patents. *RAND Journal of Economics* **41**(4) 686–708.

Shue, Kelly. 2013. Executive networks and firm policies: Evidence from the random assignment of mba peers. *Review of Financial Studies* **26**(6) 1401–1442.

Simintzi, Elena, Sheng-Jun Xu, Ting Xu. 2024. The effect of childcare access on women's careers and firm performance. *CEPR Working Paper* .

Spoon, Katie, Nicholas LaBerge, K. Hunter Wapman, Sam Zhang, Allison C. Morgan, Mirta Galesic, Bailey K. Fosdick, Daniel B. Larremore, Aaron Clauset. 2023. Gender and retention patterns among u.s. faculty. *Science Advances* **9**(42) eadi2205. doi:10.1126/sciadv.adi2205.

Syverson, Chad. 2017. Challenges to mismeasurement explanations for the us productivity slowdown. *Journal of Economic Perspectives* **31**(2) 165—186.

Vasquez-McCall, Belkis, Daniel Aminetzah, Markus Berger-de León, Paul Jenkins, Shaun Collins. 2024. How CEOs are turning corporate venture building into outsize growth. McKinsey & Company Global Survey.

Wang, Daisy. 2025. Dancing with the giants: Cvc investments and startup outcomes. *USC Marshall Working Paper* .

Weitzman, M. L. 1998. Recombinant growth. *Quarterly Journal of Economics* **113**(2) 331–360.

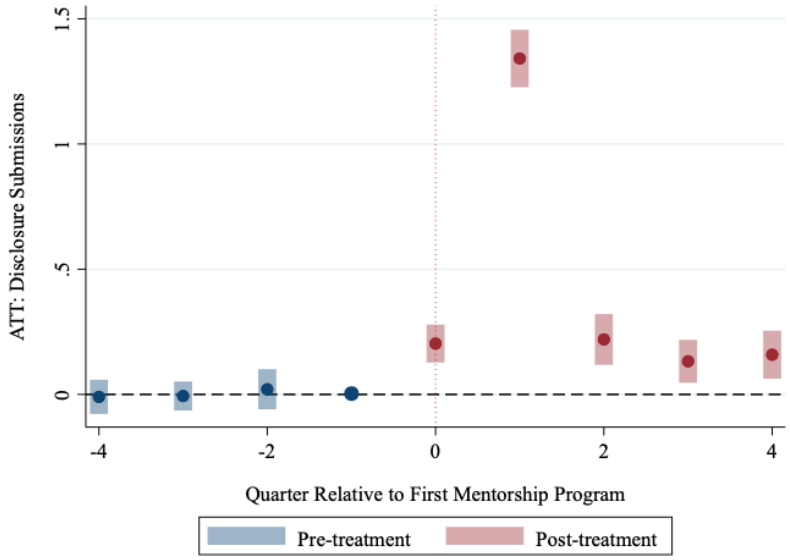
Wuchty, Stefan, Benjamin Jones, Brian Uzzi. 2007. The increasing dominance of teams in the production of knowledge. *Science* **316** 1036—-1039.

Yang, Jiho, Paola Criscuolo, Dmitry Sharapov. 2025. Autonomous inventions, problem formulation, and inventive outcomes. *Research Policy* **54**(7) 105271.

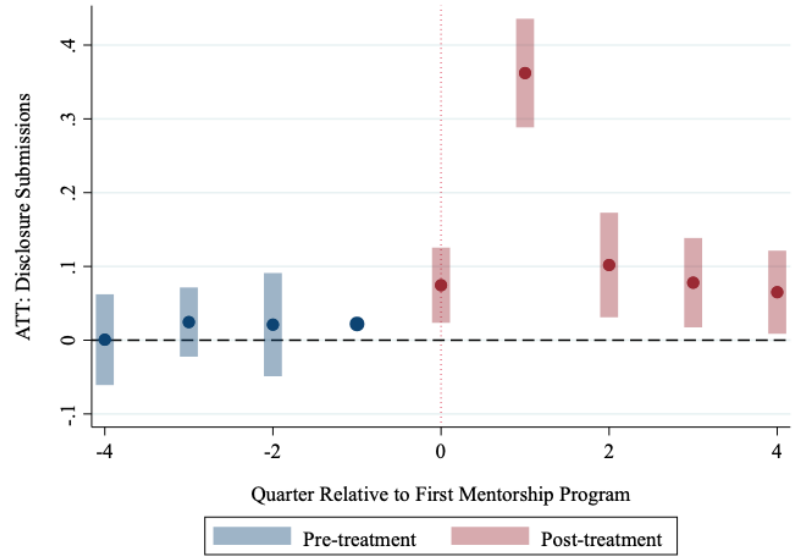
Tables and Figures

Figure 1: **Test of Parallel Pre-Trends in Disclosure Submission Activity**

(a) All Submissions



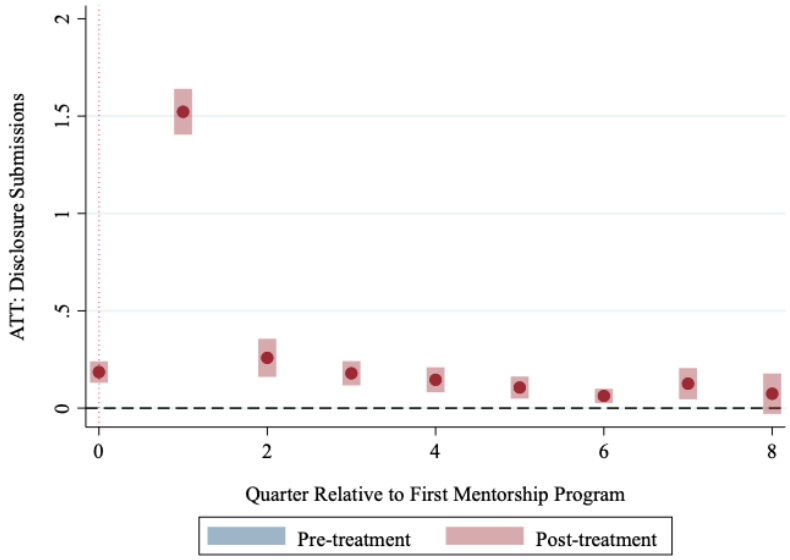
(b) Removing Mentor-Collaborative Submissions



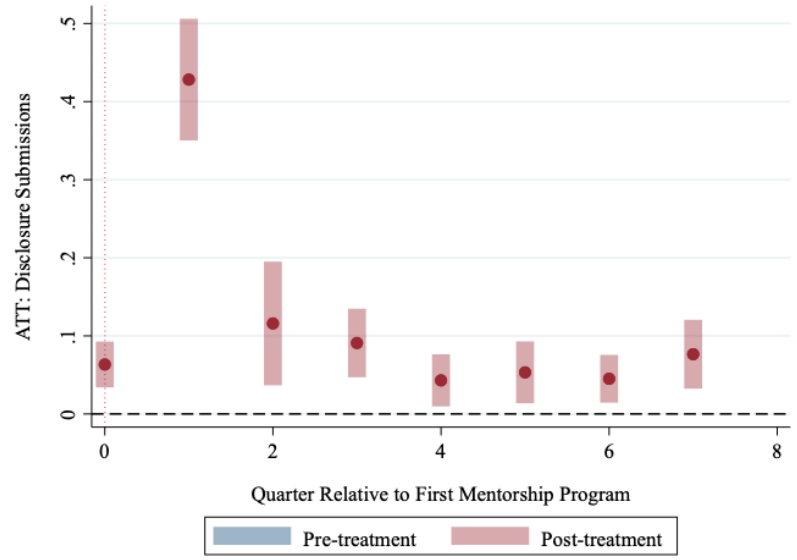
Notes. These figures plot the estimates from our parallel trends test. The point estimates capture the quarterly differences in disclosure submission outcomes between mentored and not-yet-mentored innovators. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the innovator begins the mentorship program as the baseline reference period. In subfigure (a), we aggregate together all disclosure submissions across all status categories. The average pre-period coefficient is 0.002 (p-value = 0.958), and the p-values on the individual pre-period coefficients range from 0.613–0.829. In subfigure (b), we remove all disclosure submissions for which a mentor is a collaborator. The average pre-period coefficient is 0.015 (p-value = 0.532), and the p-values on the individual pre-period coefficients range from 0.305–0.982.

Figure 2: **Disclosure Submission Outcomes of Non-Experienced Participants**

(a) All Submissions



(b) Removing Mentor-Collaborative Submissions



Notes. These figures plot post-program submission outcomes of innovators who had no disclosure submission experience before the program. The point estimates capture the quarterly differences in disclosure submission outcomes between mentored and not-yet-mentored innovators, restricting to those with no pre-program disclosure experience. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the innovator begins the mentorship program as the baseline reference period. In subfigure (a), we aggregate together all disclosure submissions across all status categories. In subfigure (b), we remove all disclosure submissions for which a mentor is a collaborator. The restriction in subfigure (b), removing disclosures with mentor collaborators, limits the amount of data available to estimate effects out to a full eight quarters.

Table 1: Demographics Characteristics of Innovators

	Cohort 1	Cohort 2	Cohort 3	Cohort 4	Cohort 5	Cohort 6	Cohort 7	No Mentor	(1)–(3)	(1)–(7)	t-stat
Age	34.62	36.28	33.88	34.00	31.60	30.16	34.22	42.69	0.164	0.003	0.001
	(8.60)	(7.94)	(7.22)	(9.05)	(8.79)	(6.83)	(7.54)	(9.19)			
Tenure (Firm)	3.54	4.61	3.40	3.97	2.58	1.08	2.51	9.53	0.228	0.002	0.001
	(4.83)	(6.33)	(4.74)	(5.04)	(4.31)	(1.91)	(3.92)	(6.67)			
Tenure (Ind.)	6.52	7.99	6.85	6.44	4.49	2.92	6.01	-	0.294	0.001	
	(6.37)	(7.06)	(5.96)	(6.63)	(6.19)	(4.39)	(5.83)	-			
Female	0.90	0.78	0.79	0.52	0.55	0.58	0.63	0.12	0.081	0.001	0.001
Asian	0.71	0.64	0.84	0.85	0.62	0.81	0.78	0.64	0.004	0.001	0.001
Black	0.02	0.02	0.00	0.00	0.01	0.00	0.04	0.00	0.181	0.194	0.001
Latinx	0.13	0.08	0.05	0.02	0.11	0.00	0.03	0.02	0.059	0.009	0.001
Multi	0.00	0.00	0.01	0.02	0.01	0.03	0.00	0.00	0.629	0.517	0.133
White	0.13	0.25	0.11	0.12	0.25	0.16	0.16	0.34	0.022	0.036	0.001
URM	0.16	0.10	0.05	0.04	0.13	0.03	0.07	0.03	0.028	0.033	0.001
Executive	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	-	-	0.006
Top Engineer	0.04	0.00	0.03	0.11	0.06	0.00	0.02	0.12	0.352	0.048	0.001
Engineer	0.81	0.72	0.73	0.74	0.58	0.47	0.65	0.88	0.429	0.002	0.001
Data Scientist	0.03	0.02	0.02	0.00	0.01	0.00	0.02	0.01	0.831	0.894	0.492
Management	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.01	0.620	0.814	0.254
Any Status	0.09	0.05	0.11	0.08	0.08	0.07	0.09	0.12	0.336	0.818	0.098
	(0.29)	(0.11)	(0.32)	(0.22)	(0.21)	(0.10)	(0.29)	(0.58)			
Advanced Status	0.03	0.02	0.01	0.01	0.00	0.01	0.01	0.04	0.360	0.469	0.044
	(0.16)	(0.06)	(0.06)	(0.04)	(0.01)	(0.04)	(0.10)	(0.30)			
Num. Innovators	101	73	171	55	88	40	105	28,579			

Notes. This table presents summary statistics of the demographic characteristics of the innovators in the firm. We report characteristics separately for individuals based on the cohort in which they first participated in the mentoring program. The column labeled (1)–(3) reports p-values from F-tests that the average value of the characteristics are equal across cohorts 1, 2, and 3. The column labeled (1)–(7) reports p-values from F-tests that the average value of the characteristics are equal across all seven cohorts. The column labeled t-stat reports p-values from T-tests that the average value of the characteristics are equal between mentored innovators and non-mentored innovators. p-values less than 0.001 are set to 0.001.

Table 2: **Mentorship and Innovator Productivity**

	Quarterly Disclosures (Raw Counts)		Quarterly Disclosures (Scaled by Team Size)		Quarterly Disclosures (Excluding Mentors)	
	Overall	Long-Term	Overall	Long-Term	Overall	Long-Term
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.254***	0.082***	0.059***	0.018**	0.078***	0.024**
	(3.753)	(3.459)	(3.521)	(2.324)	(3.094)	(2.002)
Stacked Diff-in-Diff	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.402	0.395	0.331	0.326	0.208	0.203
Observations	3,272,940	2,884,767	3,272,940	2,884,767	3,272,940	2,884,767
Pre-Program Mean	0.077	0.077	0.032	0.032	0.057	0.057
% Gain to Mean	330%	106%	184%	56%	137%	42%

Notes. This table presents estimates of the relation between mentorship and innovator productivity, as measured by submitting an inventive idea through an invention disclosure form (IDF). The data are innovator-quarter submissions. In Columns (1) and (2), the dependent variable is the total number of IDFs submitted in a given quarter. In Columns (3) and (4), each submission is first scaled by team size before aggregating them to the innovator-quarter level. In Columns (5) and (6), the dependent variable is the total number of IDFs that do not have a mentor as part of the collaborator team in a given quarter. To estimate the long-term relation in Columns (2), (4), and (6), observations are removed if they occur in the quarter of mentorship and the first quarter after mentorship for the focal cohort. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively. The penultimate row reports the pre-mentorship program mean of the dependent variable for the main treatment group of focus, and the bottom row reports the percentage gain relative to that mean for the focal group of workers.

Table 3: **Mentorship, Disclosure Advancement, and Process Learning**

	Advanced Disclosures			Non-Advanced Disclosures		
	Overall	Long-Term	Sans-Mentor	Overall	Long-Term	Sans-Mentor
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.052***	0.019***	0.011***	0.157***	0.054***	0.013
	(3.932)	(3.309)	(4.064)	(3.755)	(2.730)	(1.171)
Stacked Diff-in-Diff	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.309	0.301	0.107	0.340	0.335	0.191
Observations	3,272,940	2,884,767	2,884,767	3,272,940	2,884,767	2,884,767
Pre-Program Mean	0.010	0.009	0.004	0.064	0.063	0.052
% Gain to Mean	520%	211%	275%	245%	86%	25%

Notes. This table decomposes the innovator productivity estimates by pipeline outcome, distinguishing between higher- and lower-quality idea disclosures that advance toward a patent application, in Columns (1)–(3), and those that do not, in Columns (4)–(6). The data are innovator-quarter submissions. In Columns (1) and (4), the dependent variable is the total number of advanced and non-advanced disclosures, respectively, submitted in a given quarter. Columns (2) and (5) report long-term estimates that exclude the quarter of mentorship and the first quarter after mentorship for the focal cohort. Columns (3) and (6) restrict the outcome to disclosures on which the mentors are not listed as an inventor, addressing the concern that observed gains might reflect co-authorship with the mentor rather than the mentee's independent inventive capacity. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively. The penultimate row reports the pre-mentorship program mean of the dependent variable for the main treatment group of focus, and the bottom row reports the percentage gain relative to that mean.

Table 4: **Heterogeneous Effects to Test the Process-Learning Hypothesis**

	Quarterly Disclosures							
	Overall				Long-Term			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.240***	0.227***	0.097	0.291***	0.049	0.035	-0.065	0.121***
	(3.125)	(2.884)	(1.222)	(4.126)	(1.354)	(0.772)	(-1.080)	(4.807)
Mentored × Mentee_{URG}	0.083**				0.115***			
	(2.272)				(2.846)			
Mentored × Mentee_{URG} × M_{URG}		0.139**				0.155**		
		(2.314)				(2.593)		
Mentored × Mentee_{URG} × M_{NonURG}		0.101**				0.144***		
		(2.277)				(2.919)		
Mentored × Mentee_{NonURG} × M_{URG}		0.028				0.029		
		(0.326)				(0.306)		
Mentored × No Experience			0.213***				0.197***	
			(4.273)				(3.467)	
Mentored × Old				0.067				0.080
				(1.097)				(0.080)
Stacked Diff-in-Diff	✓	✓	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.422	0.422	0.428	0.426	0.422	0.422	0.429	0.426
Observations	1,139,562	1,138,466	3,129,536	876,430	1,138,372	1,137,362	3,128,284	875,378

Notes. This table presents estimates of potential heterogeneity in the relation between mentorship and innovator productivity. The data are innovator-quarter submissions. In all specifications, the dependent variable is the total number of IDFs submitted in a given quarter. In Column (1), the key variable of interest is the binary indicator $Mentee_{URG}$, which equals one if the innovator is either female or a URM (i.e., from an under-represented demographic group, URG), and zero otherwise. In Column (2), the key variables of interest are the binary indicators: $Mentee_{URG/NonURG} \times M_{URG/NonURG}$, which equal one or zero depending on the pairwise combination of the mentee's and mentor's URG classification. The key variable of interest in Column (3) is the binary indicator *No Experience*, which equals one if the innovator had no observed disclosure submission experience before the mentorship program, and zero otherwise. The key variable of interest in Column (4) is the binary indicator *Old*, which equals one if the innovator's age was above the median at the time the mentorship program first began, and zero otherwise. The indicators are interacted with the mentorship program participation variable. To estimate the long-term relation in Columns (5)–(8), observations are removed if they occur in the quarter of mentorship and the first quarter after mentorship for the focal cohort. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table 5: **Heterogeneity Based on Mentor Quality**

	Quarterly Disclosures			
	Overall		Long-Term	
	(1)	(2)	(3)	(4)
Mentored	0.208***	0.208***	0.051**	0.051**
	(2.966)	(2.966)	(2.209)	(2.209)
Mentored × Q4/Q3-Mentor	0.090**		0.064	
	(2.146)		(1.644)	
Mentored × Q4-Mentor		0.131**		0.118***
		(2.365)		(2.714)
Mentored × Q3-Mentor		0.027		-0.022
		(0.576)		(-0.439)
Stacked Diff-in-Diff	✓	✓	✓	✓
Adj. R-Square	0.402	0.402	0.403	0.403
Observations	3,272,452	3,272,452	3,271,214	3,271,214

Notes. This table presents our analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity, focusing on mentor quality. The data are innovator-quarter submissions. In Column (1), the key variable of interest is the binary indicator *Q4/Q3-Mentor*, which equals one for innovators whose mentor had above median pre-program disclosure productivity, relative to all other mentors, and zero otherwise. In Column (2), indicators are used to denote innovators whose mentor had pre-program disclosure productivity in the fourth (highest) or third (second highest) quartile of the mentor productivity distribution, *Q4-Mentor* and *Q3-Mentor*, respectively. These indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across mentor types. To estimate the long-term relation in Columns (3)–(4), observations are removed if they occur in the quarter of mentorship and the first quarter after mentorship for the focal cohort. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table 6: **Mentored Innovators' Perspectives on Mentorship Program**

	Full sample		Specific Subgroups					
			RG		URG	Male	Female	URM
Mentored Innovators' Perspectives	Obs (1)	Mean (2)	Obs (3)	Mean (4)	Mean (5)	Mean (6)	Mean (7)	Mean (8)
Since your formal mentorship relationship ended, on average, how many times per month have you met with your mentor outside of your regularly scheduled work meetings?	91	1.7	34	1.8	1.7	1.7	1.7	1.8
To what extent do you disagree or agree with the following statements about mentorship? (-2 = Strongly disagree, 2 = Strongly agree)								
I am planning to work on more invention disclosures at work.	89	1.2	33	1.1	1.2	1.1	1.3	0.8
I learned a lot about patenting and becoming a better inventor.	85	1.2	30	1.1	1.4	0.9	1.5**	0.8
Being mentored increased my satisfaction at work.	85	1.2	31	1.1	1.2	1.0	1.4*	0.5**
Mentorship has helped me develop professionally, gain leadership skills, and think more broadly.	89	1.2	33	1.2	1.3	1.1	1.3	1.2
I have benefited from my mentoring relationship.	86	1.2	31	1.3	1.2	1.1	1.2	0.5*
My mentor introduced me to new inventors that I anticipate connecting with in the future.	86	1.1	30	1.0	1.1	1.0	1.2	1.2
I feel confident that I can incorporate my mentor's tips into my work process and develop new inventions.	87	0.9	32	1.0	0.9	0.9	0.9	0.5
I feel confident that I have the social connections necessary to get my invention ideas accepted by IP professionals.	83	0.9	31	1.0	0.8	0.9	0.9	1.2

Notes. This table provides descriptive statistics of the engineers' perceptions of mentorship, separately by demographic groups. Columns (1)–(2) summarize the frequency and mean for the full sample, while Columns (3)–(8) report statistics for specific subgroups based on gender and ethnicity. Asian and White male engineers define the represented group (RG), while females and underrepresented minorities (URMs) define the underrepresented group (URG). The stars in Columns (7) and (8) denote the significance level from a t-test, indicating the likelihood that the observed differences in subgroup means are not due to chance. ***, ** and * indicate p-values under the assumption of a single test of 1%, 5%, and 10%, respectively.

Table 7: **Team Composition and a Social Network Mechanism**

	Overall Sample			Long-Term Sample		
	Fraction Female	Gender-Race HHI	Experience HHI	Fraction Female	Gender-Race HHI	Experience HHI
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.151***	-0.160***	-0.184***	0.056	-0.146**	-0.150***
	(4.924)	(-4.008)	(-5.141)	(1.347)	(-2.604)	(-2.636)
Stacked Diff-in-Diff	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.602	0.479	0.443	0.587	0.482	0.440
Observations	73,031	74,036	74,794	62,012	62,877	63,559

Notes. This table presents tests of potential changes in innovator team composition associated with mentorship. The data used is a panel of innovator-quarter team diversity metrics. Columns (1)–(3), utilize the full sample, while Columns (4)–(6) exclude observations of mentored workers in the quarter of first mentorship and the subsequent quarter. In Columns (1) and (4), the dependent variable is the average fraction female of the innovator teams the innovator is a part of in a given quarter. The dependent variable in Columns (2) and (5) is the average gender-race HHI values of the innovator teams the innovator is a part of in a given quarter. The dependent variable in Columns (3) and (6) is the average patenting experience HHI values of the innovator teams the innovator is a part of in a given quarter. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table 8: **Productivity of Innovators After Exposure to Mentored Workers**

	Quarterly Disclosures				
	(1)	(2)	(3)	(4)	(5)
Post First Exposure	0.375**	0.423**	0.411**	0.422**	0.137*
	(2.064)	(2.303)	(2.159)	(2.522)	(1.714)
Post First Exposure × Female/URM		-0.326**			
		(-2.393)			
Post First Exposure × Female			-0.283		
			(-1.575)		
Post First Exposure × URM				-0.380*	
				(-1.918)	
Post First Exposure × High-Pre					0.355
					(1.514)
Stacked Diff-in-Diff	✓	✓	✓	✓	✓
Adj. R-Square	0.329	0.346	0.344	0.341	0.329
Observations	3,588,318	1,208,808	918,971	1,167,235	3,588,318
Pre-Program Mean	0.975	0.605	0.264	1.033	1.465
% Gain to Mean	38%	16%	48%	4%	34%

Notes. This table reports our analyses of potential productivity changes among individuals who did not directly participate in the mentorship program but who were potentially exposed to it via their mentored colleagues. For these tests, we restrict the sample to innovators who were never formally associated with the mentorship program, either as mentors or mentees. For this group of non-participating innovators, we analyze their disclosure submission records to determine when they were first part of an invention disclosure team with an individual who had previously been mentored as part of the program. We introduce a new variable, *Post First Exposure*, which equals one in the quarter a non-participating innovator first collaborates with a mentored innovator and in all subsequent quarters, and zero otherwise. Thus the treatment group includes 520 non-participating innovators who we observe collaborating with mentored innovators. In all specifications, the dependent variable is the total number of disclosure submissions of any status that an innovator submitted in a given quarter. We make sure to omit from the dependent variable the disclosure that determines the value of *Post First Exposure*, which ensures our findings are not mechanically driven by the variable construction. In Columns (2)–(4), we examine heterogeneity in these indirect effects by interacting *Post First Exposure* with *Female* and/or *URM* indicators. In Column (5), we consider effect heterogeneity based on the innovator's pre-exposure level of disclosure experience, with *High-Pre* equaling one for those with above median baseline productivity levels. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively. The penultimate row reports the pre-mentorship program mean of the dependent variable for the main treatment group of focus, and the bottom row reports the percentage gain relative to that mean for the focal group of workers.

Table 9: **Engineers' Workplace Perceptions and Departures**

Panel A: Engineers' Perceptions of the Workplace and Corporate Culture

	Overall Rating	Culture Rating	Adaptability	Collaboration	Comm.-Orientation	Integrity
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Focal × Eng.	0.123***	0.082*	0.071***	0.084**	0.209***	0.114***
	(3.131)	(1.830)	(3.088)	(2.525)	(4.109)	(4.271)
Country FE	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓
Adj. R-squared	0.057	0.073	0.021	0.037	0.033	0.037
Observations	269,293	206,580	241,754	241,754	241,754	241,754

Panel B: Engineers' Departures

	Number of Departures	IHS(Number of Departures)	Number of Departures	Number of Departures
	(1)	(2)	(3)	(4)
Post × Focal			-0.025**	-0.024**
			(-2.211)	(-2.533)
Post × Focal × Eng.	-0.220***	-0.033***		
	(-5.745)	(-7.856)		
Region/Gender FEs	✓	✓	✓	✓
Ethnic./Senior. FEs	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓
Additional Controls				✓
Adj. R-Square	0.036	0.223	0.011	0.311
Observations	8,706,796	8,706,796	38,428	38,428

Notes. This table presents estimates of changes in engineers' perceptions of and departure from their workplace in relation to the patent mentorship program. Panel A focuses on workplace perceptions and Panel B focuses on departures. The data for workplace perceptions is individual reviews posted from 2018 to 2024 on Glassdoor. The analysis is restricted to reviews with non-missing job title data for the focal firm and its peers. In Columns (1) and (2) of Panel A, the dependent variable is the overall star rating and corporate culture rating, respectively, of the firm provided by the reviewer. In Columns (3)–(6), the dependent variables capture various cultural values as extracted from the open-ended text in the Glassdoor reviews. In Panel B, the dependent variable in Column (1) is the number of employees leaving the company within a specific occupation-region-gender-ethnicity-seniority group in a given month. The dependent variable in Column (2) is the inverse hyperbolic sine of this value. In Columns (3) and (4), we restrict the sample to only named inventors on granted patents at either the focal firm or peer firms. All of these inventors are engineers, so we only conduct DiD tests for this sample. *Post* equals one from when the mentorship program began and onward, and zero otherwise. *Engineer (Eng.)* is one for engineers, and zero otherwise. *Focal* is one if the data is about the focal firm, and zero if its about its peer firms. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

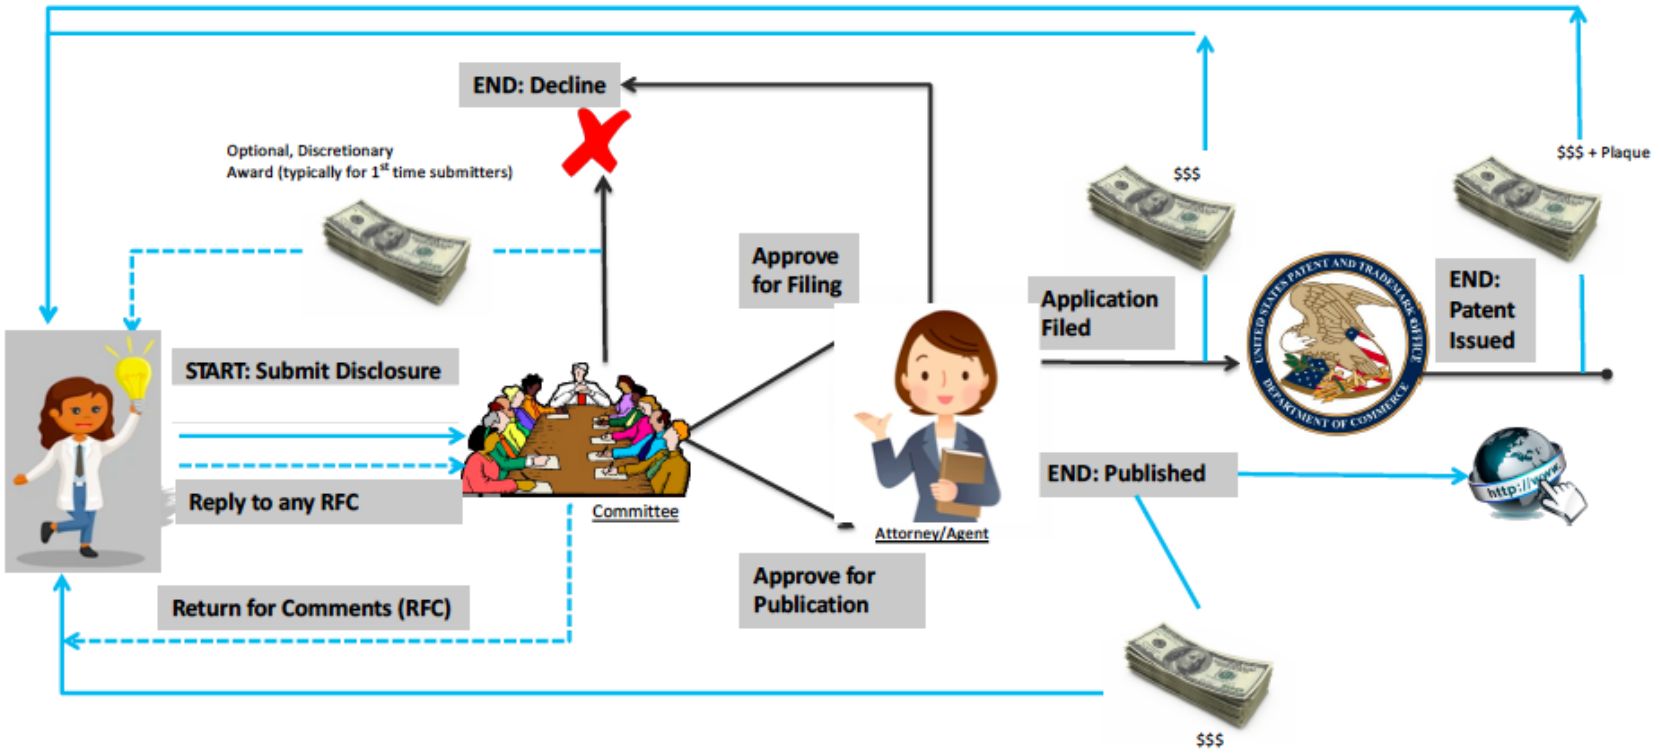
A Internet Appendix: Figures and Tables

Figure A.1: Representative Timeline of Events for a Mentorship Cohort



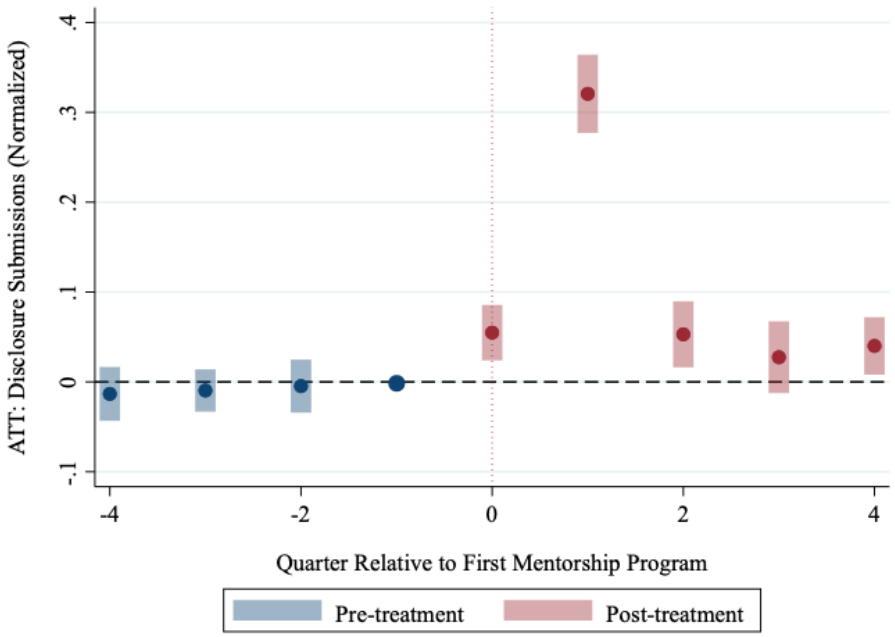
Notes. This figure highlights the typical timeline that mentorship cohorts followed throughout the program. Mentor-mentee teams generally committed to engage in the program curriculum for 3–4 hours a week for about two (four) months in the early (later) cohorts.

Figure A.2: Patent Process



Notes. This figure highlights the patent process as taught to mentees by mentors as part of the mentorship program. The main outcomes that a disclosure submission can realize are (i) being declined, (ii) being approved for defensive publication, and (iii) being a filed or approved patent.

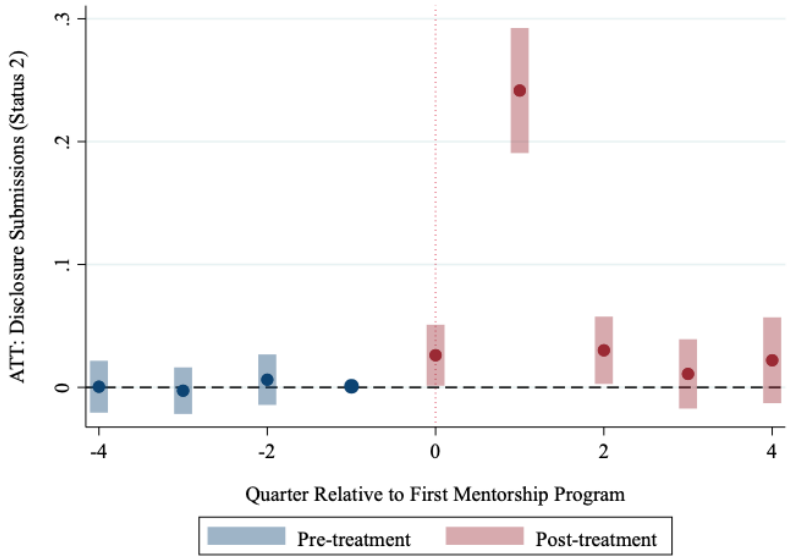
Figure A.3: Test of Common Pre-Trends in Disclosure Submission Activity



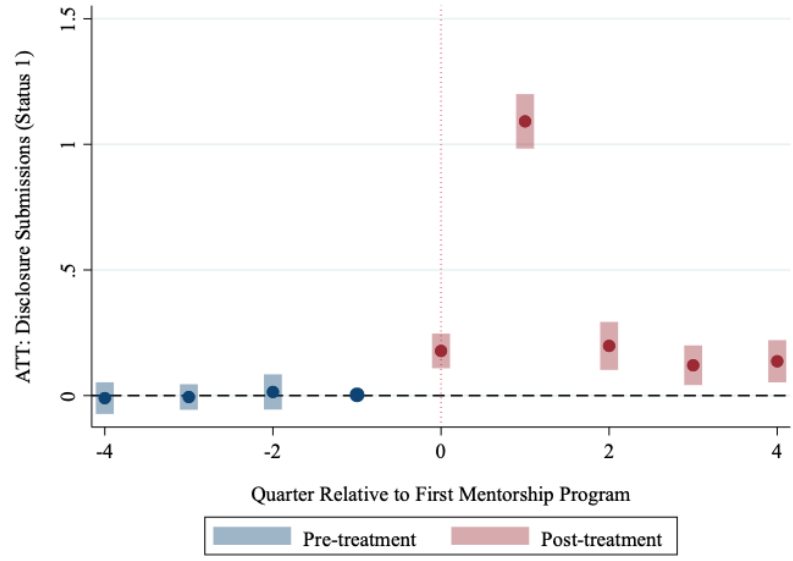
Notes. This figure plots the estimates from our parallel trends test. The point estimates capture the quarterly differences in disclosure submission outcomes (scaled by team size) between mentored and not-yet-mentored innovators. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the innovator begins the mentorship program as the baseline reference period. We aggregate together all disclosure submissions across all status categories. The average pre-period coefficient is -0.009 (p-value = 0.429), and the p-values on the individual pre-period coefficients range from 0.381–0.754.

Figure A.4: Test of Common Pre-Trends by Submission Type

(a) Number of Submissions: Advanced Status



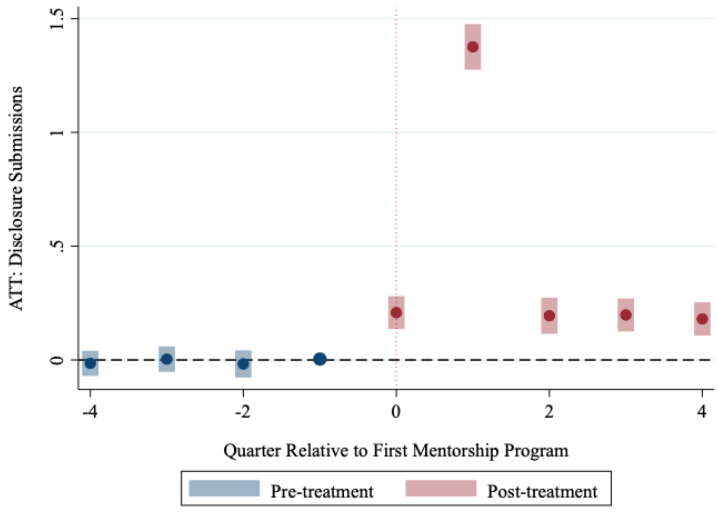
(b) Number of Submissions: Non-Advanced Status



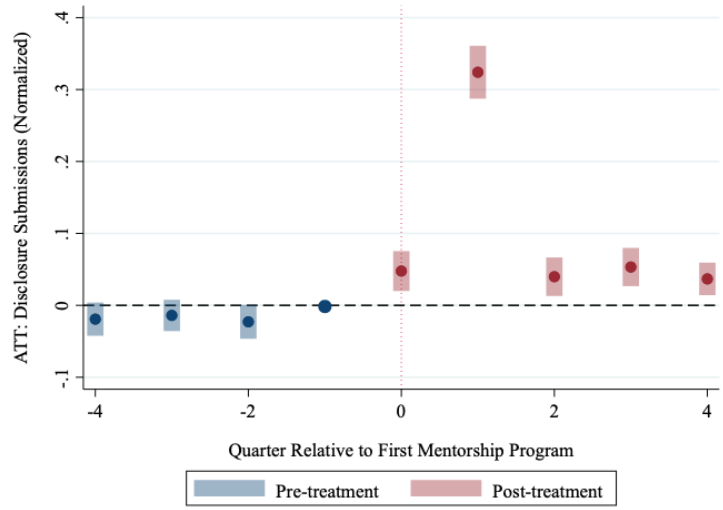
Notes. These figures plot the estimates from our parallel trends tests. The point estimates capture the quarterly differences in disclosure submission outcomes between mentored and not-yet-mentored innovators. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the innovator begins the mentorship program as the baseline reference period. Advanced status submissions are those with any of the following statuses: Approved - Patent Application; Prior Art Search Requested; RFC. Non-advanced status submissions are those with any of the following statuses: Declined; Auto Declined; Defensive Publication Candidate; Approved - Defensive Publication; Approved - Trade Secret; Deleted; Merged. The average pre-period coefficient in figure (a) is 0.001 (p-value = 0.849), and the p-values on the individual pre-period coefficients range from 0.551–0.960. The average pre-period coefficient in figure (b) is -0.001 (p-value = 0.984), and the p-values on the individual pre-period coefficients range from 0.687–0.823.

Figure A.5: Test of Common Pre-Trends (Including Never Mentored Control Innovators)

(a) Number of Submissions



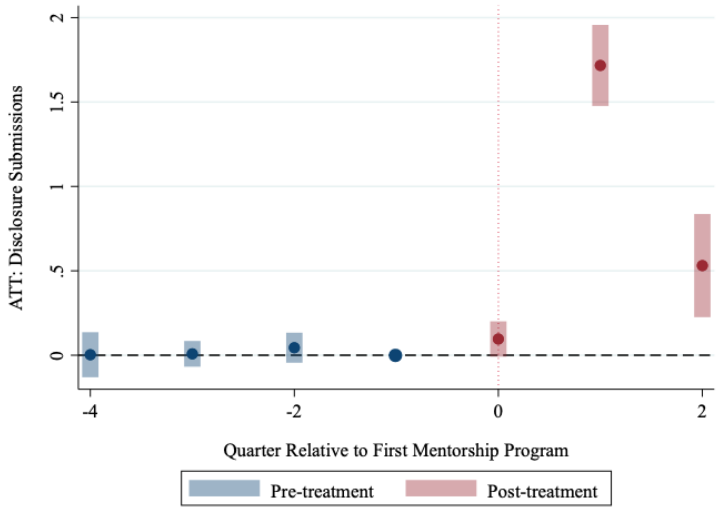
(b) Number of Scaled Submissions (Normalized by Team Size)



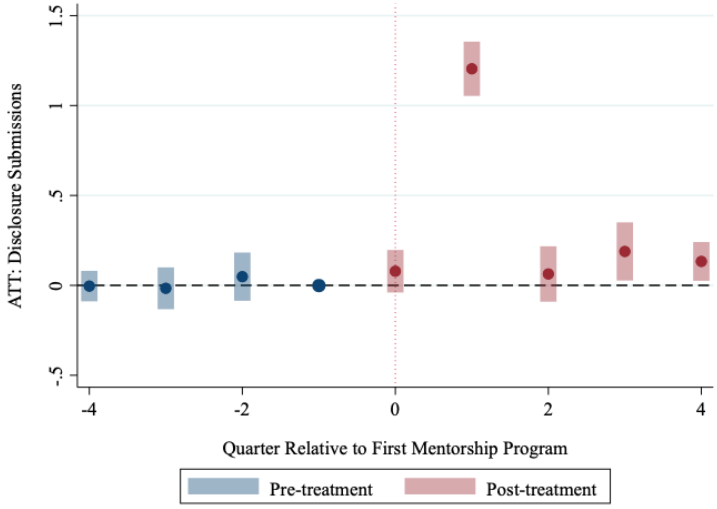
Notes. These figures plot the estimates from our parallel trends test. The point estimates capture the quarterly differences in disclosure submission outcomes between mentored and both not-yet-mentored and never mentored innovators. Specifically, we include into the control group the subset of never mentored innovators who are the most comparable to mentored innovators based their observable characteristics (based on propensity score matching discussed in Appendix B). We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the innovator begins the mentorship program as the baseline reference period. We aggregate together all disclosure submissions across all status categories. In figure (a), we report quarterly disclosure estimates from unscaled submission values. The average pre-period coefficient in figure (a) is -0.010 (p-value = 0.698), and the p-values on the individual pre-period coefficients range from 0.563–0.899. In figure (b), we report quarterly disclosure estimates based on scaled submission values, where we normalize each disclosure submission by the size of the innovator team. The average pre-period coefficient in figure (b) is -0.019 (p-value = 0.064), and the p-values on the individual pre-period coefficients range from 0.056–0.207.

Figure A.6: Test of Common Pre-Trends (Separately for Early vs. Late Cohorts)

(a) Early Cohorts Only



(b) Late Cohorts Only



Notes. These figures plot the estimates from our parallel trends test. The point estimates capture the quarterly differences in disclosure submission outcomes between mentored and both not-yet-mentored innovators. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the innovator begins the mentorship program as the baseline reference period. We aggregate together all disclosure submissions across all status categories. In figure (a), we report quarterly disclosure estimates for only early cohorts (i.e., Cohorts 1, 2, and 3). In figure (b), we report quarterly disclosure estimates for only late cohorts (i.e., Cohorts 4, 5, 6, and 7). Since Cohorts 1 and 3 are only four quarters apart, we have limited post-program data for these cohort comparisons, leading to only two quarters of post-program estimates. The p-value of the average pre-period coefficient in figure (a) is 0.642, and the p-values on the individual pre-period coefficients range from 0.328–0.967. The p-value of the average pre-period coefficient in figure (b) is 0.835, and the p-values on the individual pre-period coefficients range from 0.481–0.924.

Table A.1: Survey Evidence on Mentorship Program Participation

	Full sample		Specific Subgroups					
			RG		URG	Male	Female	URM
	Obs	Mean	Obs	Mean	Mean	Mean	Mean	Mean
Panel A. Mentorship	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Have you ever participated in our patent mentorship program?	230	43%	101	37%	48%	36%	52%**	38%
Panel B. Perspective of those without a mentor								
What are the reasons you have not sought patent mentorship?								
I wanted to participate but was unable to	125	3%	63	5%	2%	4%	2%	0%
I did not believe patent mentorship would benefit me	125	6%	63	3%	9%	3%	11%**	0%
I did not have time for patent mentorship	125	6%	63	6%	9%	5%	7%	10%
I did not know about it	125	86%	63	87%	84%	90%	82%	90%
I want to avoid activities that highlight my weaknesses to leaders	125	1%	63	2%	0%	1%	0%	0%
My co-workers already help me with IP development, so I don't need a mentor	125	1%	63	0%	0%	0%	0%	0%
I felt that the program would benefit my employer more than it would benefit me	125	1%	63	2%	0%	1%	0%	0%
Panel C. Overall Awareness of invention process								
Are you aware that we have an invention submission process and tools where you can submit inventions to be reviewed for potential patenting? (1 = Yes, 0 = No)	340	84%	148	85%	84%	84%	84%	78%
Have you attended any of our trainings on inventing? (1 = Yes, 0 = No)	339	45%	147	44%	44%	43%	47%	41%
Have you ever had an idea that you thought might be patentable? (1 = Yes, 0 = No)	338	51%	147	51%	47%	53%	47%	48%

Notes. This table provides descriptive statistics of the innovators' decisions to participate in patent mentorship and the perspectives of those not participating. Columns (1)–(2) summarize the frequency and mean for the full sample, while Columns (3)–(8) report statistics for specific subgroups based on gender and ethnicity. Asian and White male engineers define the represented group (RG), while females and underrepresented minorities (URMs) define the underrepresented group (URG). The stars in Column (7) denote the significance level from a t-test, indicating the likelihood that the observed differences in subgroup means are not due to chance. ***, ** and * indicate p-values under the assumption of a single test of 1%, 5%, and 10%, respectively. To avoid survey fatigue among the full innovator population, the firm's established practice is to field surveys to a representative subsample. This survey yielded a 15 percent response rate, which is consistent with standard organizational surveys.

Table A.2: Mentorship and Innovator Productivity

Panel A: Submissions

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.256***	0.254***	0.256***	0.253***	0.555***	0.348***	0.052***	0.052***
(t-stats)	(3.372)	(3.753)	(3.324)	(3.732)	(4.269)	(3.983)	(4.023)	(3.932)
[q-values]	[0.002]	[0.001]	[0.002]	[0.001]	[0.001]	[0.001]	[0.001]	[0.001]
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.206	0.225	0.308	0.309
Observations	473,939	3,272,940	465,778	3,215,813	9,549	22,210	473,939	3,272,940
Control Mean	0.072	0.072	0.062	0.061	0.084	0.077	0.072	0.072

Panel B: Scaled Submissions (Normalized by Team Size)

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.060***	0.059***	0.060***	0.059***	0.132***	0.082***	0.011***	0.011***
(t-stats)	(3.061)	(3.521)	(3.024)	(3.505)	(4.304)	(3.696)	(3.896)	(3.853)
[q-values]	[0.002]	[0.001]	[0.003]	[0.001]	[0.001]	[0.001]	[0.002]	[0.001]
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.328	0.331	0.312	0.316	0.148	0.155	0.252	0.253
Observations	473,939	3,272,940	465,778	3,215,813	9,549	22,210	473,939	3,272,940
Control Mean	0.023	0.023	0.020	0.020	0.033	0.032	0.023	0.023

Notes. This table presents estimates of the relation between mentorship and innovator productivity, as measured by submitting an inventive idea through an invention disclosure form (IDF). The data are innovator-quarter submissions. In Columns (1)–(6), the dependent variable is the total number of IDFs submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of IDFs that become (or are likely to become) approved patent applications (i.e., high-quality ideas). In Panel B, the submissions are first scaled by team size before aggregating them to the innovator-year-quarter level. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We use a TWFE estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked DiD estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(4) and Columns (7)–(8). In Columns (5)–(6), the control group is only not-yet mentored innovators. In Columns (3)–(4), to reduce concerns about SUTVA violations, we remove innovators who mentored innovators collaborate with after the end of the mentorship program. For the TWFE model, standard errors are clustered by innovator and quarter. In the stacked DiD specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Sharpened q-values, following Anderson (2008), are reported in brackets. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.3: Mentorship and Long-Term Innovator Productivity

Panel A: Submissions

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.084**	0.082***	0.084**	0.082***	0.165**	0.104***	0.019***	0.019***
(t-stats)	(2.576)	(3.459)	(2.453)	(3.375)	(2.731)	(2.802)	(3.623)	(3.309)
[q-values]	[0.020]	[0.007]	[0.020]	[0.007]	[0.017]	[0.010]	[0.007]	[0.007]
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.401	0.395	0.395	0.389	0.213	0.193	0.309	0.301
Observations	472,671	2,884,767	464,510	2,834,485	8,281	18,282	472,671	2,884,767
Control Mean	0.072	0.072	0.062	0.062	0.072	0.076	0.072	0.072

Panel B: Scaled Submissions (Normalized by Team Size)

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.018*	0.018**	0.018	0.018**	0.034**	0.021*	0.005***	0.005***
(t-stats)	(1.784)	(2.324)	(1.719)	(2.281)	(2.401)	(1.847)	(3.418)	(3.121)
[q-values]	[0.036]	[0.020]	[0.036]	[0.020]	[0.021]	[0.036]	[0.008]	[0.007]
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.330	0.326	0.314	0.311	0.162	0.147	0.253	0.251
Observations	472,671	2,884,767	464,510	2,834,485	8,281	18,282	472,671	2,884,767
Control Mean	0.023	0.023	0.020	0.020	0.023	0.032	0.023	0.023

Notes. This table presents estimates of the relation between mentorship and long-term innovator productivity, as measured by submitting an inventive idea through an invention disclosure form (IDF). The data are innovator-quarter submissions. In Columns (1)–(6), the dependent variable is the total number of IDFs submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of IDFs that become (or are likely to become) approved patent applications (i.e., high-quality ideas). To estimate the long-term relation, observations are removed if they occur in the quarter of mentorship and the first quarter after mentorship for the focal cohort. In the stacked DiD specifications, we also remove the corresponding quarters of data from each cohort's control group observations. In Panel B, the submissions are first scaled by team size before aggregating them to the innovator-year-quarter level. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We use a TWFE estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked DiD estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(4) and Columns (7)–(8). In Columns (5)–(6), the control group is only not-yet mentored innovators. In Columns (3)–(4), to reduce concerns about SUTVA violations, we remove innovators who mentored innovators collaborate with after the end of the mentorship program. For the TWFE model, standard errors are clustered by innovator and quarter. In the stacked DiD specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Sharpened q-values, following Anderson (2008), are reported in brackets. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.4: Mentorship and Innovator Productivity Among a Female-Only Sample

	Quarterly Disclosures (Raw Counts)		Quarterly Disclosures (Scaled by Team Size)		Quarterly Disclosures (Excluding Mentors)	
	Overall	Long-Term	Overall	Long-Term	Overall	Long-Term
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.321***	0.156***	0.081***	0.041***	0.120***	0.064***
	(4.618)	(5.775)	(4.821)	(6.179)	(4.269)	(6.158)
Stacked Diff-in-Diff	✓	✓	✓	✓	✓	✓
Adj. R-Square	0.225	0.227	0.191	0.194	0.128	0.127
Observations	106,652	93,185	106,652	93,185	106,652	93,185

Notes. This table presents estimates of the relation between mentorship and innovator productivity, as measured by submitting an inventive idea through an invention disclosure form (IDF). The data are innovator-quarter submissions. To mitigates concerns about gender-related selection bias, we only include data from female innovators in these tests. In Columns (1) and (2), the dependent variable is the total number of IDFs submitted in a given quarter. In Columns (3) and (4), each submission is first scaled by team size before aggregating them to the innovator-quarter level. To estimate the long-term relation in Columns (2), (4), and (6), observations are removed if they occur in the quarter of mentorship and the first quarter after mentorship for the focal cohort. In Columns (5) and (6), the dependent variable is the total number of IDFs that do not have a mentor as part of the collaborator team in a given quarter. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ a stacked DiD estimator in all columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. All specifications include innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Standard errors are clustered at the innovator-by-cohort and quarter-by-cohort levels, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.5: Cohort-by-cohort Idea Disclosures

	Quarterly Disclosures		
	During Program	After Program	Overall
	(1)	(2)	(3)
Cohort 1	89.1%	51.5%	0.294***
			(4.44)
Cohort 2	94.5%	45.2%	0.262***
			(4.67)
Cohort 3	88.9%	30.4%$^{+}$	0.203***
			(3.78)
Cohort 4	85.5%	40.0%	0.307***
			(5.24)
Cohort 5	90.9%	35.2%$^{+}$	0.336***
			(4.03)
Cohort 6	95.0%	27.5%$^{+}$	0.707***
			(7.01)
Cohort 7	92.3%	N.A.	N.A.

Notes. This table presents summary statistics on ideas submitted by cohort along with the associated cohort-specific point estimates for the average treatment effect on the treated (ATT). Column (1) reports the percentage of mentees from each cohort who submit at least one idea disclosure during the program (i.e., in quarter zero or quarter one). Column (2) shows the percentage of mentees from each cohort who submitted at least one idea disclosure after the program (i.e., after quarter one). In Columns (1) and (2), + indicates a statistically significant difference relative to the first cohort. Column (3) presents the Callaway Sant'Anna (2021) cohort-by-cohort estimator for the ATT. We cannot estimate one for Cohort 7 because there is no later cohort to serve as a control. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.6: Mentorship and Innovator Productivity from Underrepresented Groups

Panel A: Innovator Female/URM Categorization

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.246**	0.240***	0.249**	0.242***	0.512***	0.306***	0.042***	0.041***
	(2.882)	(3.125)	(2.855)	(3.144)	(3.838)	(3.261)	(3.007)	(2.657)
Mentored × Female/URM	0.082*	0.083**	0.081*	0.082**	0.068	0.077**	0.027*	0.027**
	(1.867)	(2.272)	(1.852)	(2.262)	(1.613)	(1.987)	(1.885)	(2.467)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.418	0.422	0.418	0.423	0.212	0.233	0.352	0.353
Observations	168,642	1,139,562	162,380	1,095,728	8,898	21,354	168,642	1,139,562
Control Mean	0.193	0.193	0.166	0.166	0.132	0.113	0.071	0.071

Panel B: Innovator × Mentor Female/URM (i.e., URG) Categorization

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.233**	0.227***	0.236**	0.230***	0.488***	0.287***	0.047***	0.046***
	(2.677)	(2.884)	(2.663)	(2.906)	(3.513)	(2.998)	(3.170)	(2.696)
Mentored × URG × URG	0.138**	0.139**	0.137**	0.138**	0.120**	0.120**	0.010	0.010
	(2.417)	(2.314)	(2.410)	(2.304)	(2.199)	(2.010)	(0.572)	(0.593)
Mentored × URG × Non-URG	0.100**	0.101**	0.100**	0.100**	0.087*	0.101**	0.034**	0.034***
	(2.143)	(2.277)	(2.129)	(2.272)	(1.842)	(2.197)	(2.780)	(3.177)
Mentored × Non-URG × URG	0.028	0.028	0.029	0.029	0.014	0.032	-0.011	-0.011
	(0.301)	(0.326)	(0.312)	(0.340)	(0.147)	(0.363)	(-0.305)	(-0.342)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.419	0.422	0.419	0.423	0.217	0.233	0.352	0.353
Observations	167,952	1,138,466	161,690	1,094,632	8,208	20,258	167,952	1,138,466
Control Mean	0.193	0.193	0.166	0.166	0.130	0.109	0.071	0.071

Notes. This table presents estimates of potential heterogeneity in the relation between mentorship and innovator productivity, focusing on gender and racial/ethnic diversity. The data are innovator-quarter submissions. In Columns (1)–(6), the dependent variable is the total number of IDFs submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of IDFs that become (or are likely to become) approved patent applications (i.e., high-quality ideas). Due to limitations in demographic data availability, the sample size is reduced compared to previous analyses. In Panel A, the key variable of interest is the binary indicator *Female/URM*, which equals one if the innovator is either female or a URM (i.e., from an under-represented demographic group, URG), and zero otherwise. The indicator is interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across demographic groups. In Panel B, the key variables of interest are the binary indicators: *URG* × *URG*, which equals one if the innovator and mentor are both from a URG, and zero otherwise; *URG* × *Non-URG*, which equals one if the innovator is from a URG and the mentor is not from a URG, and zero otherwise; and *Non-URG* × *URG*, which equals one if the innovator is not from a URG and the mentor is from a URG, and zero otherwise. The indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across mentor-mentee demographic group pairs. The empirical specification structure matches that used in Table A.2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.7: Mentorship and Innovator Productivity based on Pre-Program Experience

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.100	0.097	0.099	0.097	0.413***	0.189**	0.035*	0.035*
	(1.063)	(1.222)	(1.048)	(1.214)	(3.343)	(2.079)	(1.885)	(1.718)
Mentored × No Experience	0.213***	0.213***	0.213***	0.213***	0.198***	0.220***	0.021	0.021
	(3.385)	(4.273)	(3.377)	(4.271)	(3.159)	(4.205)	(1.224)	(1.324)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.424	0.428	0.424	0.430	0.209	0.227	0.348	0.350
Observations	453,421	3,129,536	446,043	3,077,890	9,512	22,173	453,421	3,129,536
Control Mean	0.251	0.251	0.219	0.219	0.263	0.226	0.095	0.095

Notes. This table presents estimates of potential heterogeneity in the relation between mentorship and innovator productivity, focusing on the disclosure submission experience level of mentored workers before the mentorship program. The data are innovator-quarter submissions. In Columns (1)–(6), the dependent variable is the total number of IDFs submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of IDFs that become (or are likely to become) approved patent applications (i.e., high-quality ideas). The key variable of interest is the binary indicator *No Experience*, which equals one if the innovator had no observed disclosure submission experience before the mentorship program, and zero otherwise. The indicator is interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association based on pre-program disclosure experience. The empirical specification structure matches that used in Table A.2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.8: Innovator Productivity and Heterogeneity by Race and Gender Individually

Panel A: Female Mentees

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.267**	0.260***	0.273**	0.265***	0.521***	0.314***	0.052***	0.051***
	(2.905)	(3.119)	(2.874)	(3.152)	(3.907)	(3.232)	(2.948)	(2.623)
Mentored × Female	0.060	0.061	0.059	0.060	0.044	0.069	0.016	0.016
	(1.031)	(1.239)	(1.017)	(1.227)	(0.790)	(1.382)	(0.888)	(1.042)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.423	0.426	0.424	0.429	0.221	0.237	0.352	0.354
Observations	130,280	876,517	124,825	838,332	7,863	19,598	130,280	876,517

Panel B: URM Mentees

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.290***	0.285***	0.293***	0.287***	0.544***	0.341***	0.055***	0.054***
	(3.737)	(3.937)	(3.686)	(3.954)	(4.204)	(3.922)	(4.597)	(4.008)
Mentored × URM	0.079	0.079	0.079	0.080	0.062	0.073	0.022	0.022
	(0.778)	(0.853)	(0.783)	(0.859)	(0.654)	(0.818)	(0.912)	(0.957)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.397	0.401	0.393	0.398	0.215	0.235	0.327	0.329
Observations	162,602	1,099,259	156,565	1,057,000	8,489	20,468	162,602	1,099,259

Panel C: Intersectionality (i.e., Female-URM Mentees)

	Full Sample		Primary Robustness Considerations					
			SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.299***	0.292***	0.304***	0.297***	0.538***	0.342***	0.059***	0.059***
	(3.833)	(3.990)	(3.786)	(4.033)	(4.242)	(4.001)	(4.507)	(3.999)
Mentored × Female-URM	0.172	0.175	0.172	0.176	0.133	0.189	0.020	0.020
	(1.112)	(1.183)	(1.113)	(1.184)	(0.965)	(1.349)	(0.592)	(0.581)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.401	0.405	0.399	0.404	0.225	0.240	0.328	0.329
Observations	124,240	836,214	119,010	799,604	7,454	18,712	124,240	836,214

Notes. This table presents estimates of potential heterogeneity in the relation between mentorship and innovator productivity, focusing on gender and racial/ethnic diversity individually. The data comprises a panel of innovator-quarter IDF (Invention Disclosure Form) submissions. Due to limitations in demographic data availability, the sample size is reduced compared to previous analyses. The key variables of interest are binary indicators: *Female* equals one for female innovators, *URM* equals one for innovators from underrepresented racial/ethnic minority groups, and *Female-URM* equals one if the innovator belongs to both categories. These indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across demographic groups. The empirical specification structure matches that used in Table A.2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.9: Innovator Productivity and Heterogeneity by Age

	Age 20–27 Years		Age 28–33 Years		Age 34–39 Years		Age 40+ Years	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.466***	0.429***	0.346***	0.368***	0.420***	0.416***	0.336***	0.365***
	(5.089)	(3.913)	(3.716)	(4.014)	(5.248)	(4.560)	(3.613)	(4.168)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.101	0.278	0.371	0.449	0.101	0.286	0.378	0.450
Observations	4,440	12,046	23,988	89,122	24,302	76,936	158,609	612,302

Notes. This table presents additional analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity, specifically based on innovator age. The data comprises a panel of innovator-quarter IDF (Invention Disclosure Form) submissions. The dependent variable is the total number of disclosure submissions of any status that an innovator submitted in a given quarter. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ two estimation strategies: a two-way fixed effects estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked difference-in-differences estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.10: Innovator Productivity and Heterogeneity by Country and Cohort

Panel A: Work Country

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.303***	0.297***	0.303***	0.298***	0.576***	0.366***	0.058***	0.057***
	(3.882)	(4.383)	(3.815)	(4.372)	(4.443)	(4.094)	(4.634)	(4.768)
Mentored × USA	-0.035	-0.033	-0.036	-0.034	-0.060	-0.041	0.001	0.001
	(-0.690)	(-0.735)	(-0.704)	(-0.754)	(-1.235)	(-0.898)	(0.046)	(0.084)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.408	0.412	0.405	0.410	0.206	0.225	0.343	0.345
Observations	226,933	1,543,898	219,160	1,489,487	9,549	22,210	226,933	1,543,898

Panel B: Cohort Timing

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.309**	0.306**	0.310**	0.306**	0.605***	0.375**	0.080***	0.080***
	(2.240)	(2.435)	(2.215)	(2.420)	(3.975)	(2.435)	(3.041)	(3.166)
Mentored × Early	-0.094	-0.092	-0.095	-0.093	-0.121	-0.040	-0.049*	-0.049*
	(-0.613)	(-0.640)	(-0.618)	(-0.643)	(-0.852)	(-0.212)	(-1.771)	(-1.775)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.207	0.225	0.308	0.309
Observations	473,939	3,272,940	465,778	3,215,813	9,549	22,210	473,939	3,272,940

Notes. This table presents additional analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity. The data comprises a panel of innovator-quarter IDF (Invention Disclosure Form) submissions. The tests in Panel A require data on the work country of the innovator. We do not have data on this characteristic for all individuals, leading to the reductions in observation counts. The variable *USA* is a binary indicator that equals one if the innovator works in the United States. The variable *Early* is a binary indicator that equals one if the innovator was part of the first, second, or third mentorship cohorts. These indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across groups. The empirical specification structure matches that used in Table A.2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.11: Heterogeneity Based on Mentor Quality

Panel A: Mentor Quality By Median

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.211**	0.208***	0.210**	0.208***	0.514***	0.301***	0.029***	0.029***
	(2.790)	(2.966)	(2.752)	(2.949)	(3.888)	(3.601)	(2.989)	(3.037)
Mentored × Q4/Q3-Mentor	0.090*	0.090**	0.090*	0.090**	0.082*	0.091**	0.043**	0.043***
	(2.014)	(2.146)	(2.016)	(2.149)	(1.831)	(2.127)	(2.828)	(2.779)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.207	0.226	0.308	0.309
Observations	473,723	3,272,452	465,562	3,215,325	9,333	21,722	473,723	3,272,452
Control Mean	0.072	0.072	0.061	0.061	0.082	0.078	0.027	0.027

Panel B: Mentor Quality By Top-Two Quartiles

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.211**	0.208***	0.210**	0.208***	0.514***	0.299***	0.029***	0.029***
	(2.790)	(2.966)	(2.751)	(2.949)	(3.889)	(3.573)	(2.989)	(3.037)
Mentored × Q4-Mentor	0.131*	0.131**	0.131*	0.131**	0.118*	0.137***	0.057**	0.057**
	(2.023)	(2.365)	(2.026)	(2.368)	(1.877)	(2.708)	(2.295)	(2.499)
Mentored × Q3-Mentor	0.027	0.027	0.027	0.027	0.027	0.022	0.021*	0.021
	(0.620)	(0.576)	(0.619)	(0.575)	(0.579)	(0.442)	(1.770)	(1.513)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.207	0.227	0.308	0.309
Observations	473,723	3,272,452	465,562	3,215,325	9,333	21,722	473,723	3,272,452
Control Mean	0.072	0.072	0.061	0.061	0.082	0.078	0.027	0.027

Notes. This table presents our analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity, focusing on mentor quality. The data comprises a panel of innovator-quarter IDF (Invention Disclosure Form) submissions. In Panel A, the key variable of interest is the binary indicator *Q4/Q3-Mentor*, which equals one for innovators whose mentor had above median pre-program disclosure productivity, relative to all other mentors, and zero otherwise. In Panel B, indicators are used to denote innovators whose mentor had pre-program disclosure productivity in the fourth (highest) or third (second highest) quartile of the mentor productivity distribution, *Q4-Mentor* and *Q3-Mentor*, respectively. These indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across mentor types. The empirical specification structure matches that used in Table A.2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.12: Heterogeneity Based on Mentor Quality – Alternative Quality Measure

Panel A: Mentor Fixed Effects By Median

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.133**	0.131**	0.133**	0.130**	0.424***	0.229***	0.024***	0.024***
	(2.482)	(2.425)	(2.447)	(2.409)	(3.781)	(2.868)	(3.720)	(3.223)
Mentored × Q4/Q3-Mentor	0.202***	0.202***	0.202***	0.202***	0.213***	0.207***	0.049**	0.049***
	(3.599)	(4.741)	(3.597)	(4.744)	(3.809)	(4.739)	(2.416)	(2.780)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.211	0.226	0.308	0.309
Observations	473,723	3,272,464	465,562	3,215,337	9,333	21,734	473,723	3,272,464

Panel B: Mentor Fixed Effects By Quartile

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.133**	0.131**	0.133**	0.130**	0.425***	0.229***	0.024***	0.024***
	(2.482)	(2.425)	(2.447)	(2.409)	(3.792)	(2.870)	(3.721)	(3.223)
Mentored × Q4-Mentor	0.272***	0.271***	0.272***	0.271***	0.271***	0.271***	0.061*	0.061**
	(3.490)	(4.516)	(3.491)	(4.518)	(3.483)	(4.456)	(2.080)	(2.360)
Mentored × Q3-Mentor	0.143***	0.142***	0.143***	0.142***	0.162***	0.148***	0.038*	0.038**
	(2.962)	(3.249)	(2.964)	(3.253)	(3.558)	(3.612)	(1.840)	(2.033)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.211	0.227	0.308	0.309
Observations	473,723	3,272,464	465,562	3,215,337	9,333	21,734	473,723	3,272,464

Notes. This table presents our analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity, focusing on mentor quality. The data comprises a panel of innovator-quarter IDF (Invention Disclosure Form) submissions. In Panel A, the key variable of interest is the binary indicator *Q4/Q3-Mentor*, which equals one for innovators whose mentor has a mentor fixed effect above the median, relative to all mentors in the program, and zero otherwise. In Panel B, indicators are used to denote innovators whose mentor has a mentor fixed effect in the fourth (highest) or third (second highest) quartile of the mentor fixed effect distribution, *Q4-Mentor* or *Q3-Mentor*, respectively. These indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across mentor types. The empirical specification structure matches that used in Table A.2. *t*-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.13: Demographics Characteristics of Mentors

	Mentor Fixed Effect Quartile			
	1st (Lowest)	2nd	3rd	4th (Highest)
Age	47.22	46.04	45.83	46.43
Tenure (Firm)	8.63	10.48	9.79	10.34
Female	0.26	0.30	0.04	0.22
Asian	0.58	0.56	0.43	0.48
Black	0.00	0.00	0.00	0.00
Latinx	0.04	0.04	0.04	0.08
Multi	0.00	0.00	0.00	0.00
White	0.38	0.40	0.52	0.44
URM	0.04	0.04	0.04	0.08
Executive	0.00	0.00	0.00	0.00
Top Engineer	0.48	0.70	0.79	0.61
Engineer	0.83	0.96	1.00	0.91
Data Scientist	0.09	0.00	0.00	0.04
Management	0.04	0.00	0.00	0.00
Any Status Discl.	2.16	1.01	2.76	3.44
Number Inventors	28	28	28	29

Notes. This table presents summary statistics of the demographic characteristics of the mentors who were involved in the program, split by their estimated mentor fixed effects.

Table A.14: Innovator Productivity and Mentor Heterogeneity

Panel A: Mentor Gender

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.278***	0.276***	0.277***	0.275***	0.541***	0.352***	0.061***	0.061***
	(3.583)	(4.006)	(3.530)	(3.983)	(4.194)	(3.999)	(4.133)	(4.024)
Mentored × Female	-0.022	-0.022	-0.022	-0.022	-0.023	-0.033	-0.005	-0.005
	(-0.445)	(-0.399)	(-0.441)	(-0.396)	(-0.493)	(-0.652)	(-0.326)	(-0.324)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.400	0.403	0.393	0.397	0.205	0.209	0.308	0.310
Observations	472,315	3,270,183	464,154	3,213,056	7,925	19,453	472,315	3,270,183

Panel B: Mentor Engineer Status

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.259**	0.257***	0.259**	0.256***	0.540***	0.319***	0.064***	0.064***
	(2.708)	(2.831)	(2.667)	(2.811)	(4.129)	(3.227)	(3.571)	(3.827)
Mentored × Top Eng.	0.020	0.020	0.020	0.021	-0.005	0.036	-0.006	-0.006
	(0.373)	(0.380)	(0.378)	(0.386)	(-0.119)	(0.655)	(-0.299)	(-0.342)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.400	0.403	0.393	0.397	0.205	0.209	0.308	0.310
Observations	472,315	3,270,183	464,154	3,213,056	7,925	19,453	472,315	3,270,183

Panel C: Mentor Formerly in Program

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.252***	0.249***	0.251***	0.249***	0.550***	0.348***	0.053***	0.053***
	(3.483)	(3.939)	(3.436)	(3.920)	(4.308)	(3.989)	(4.054)	(3.913)
Mentored × In Program	0.061	0.060	0.061	0.060	0.056	0.008	-0.016	-0.016
	(0.453)	(0.461)	(0.452)	(0.459)	(0.492)	(0.100)	(-0.945)	(-0.946)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.206	0.225	0.308	0.309
Observations	473,939	3,272,940	465,778	3,215,813	9,549	22,210	473,939	3,272,940

Notes. This table presents additional analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity. The data comprises a panel of innovator-quarter IDF (Invention Disclosure Form) submissions. The tests in Panel A require data on the gender of the innovator's mentor. The tests in Panel B require data on the job title of the innovator's mentor. We do not have data on these characteristics for all individuals, leading to the reductions in observation counts. The variable *Female* is a binary indicator that equals one if the innovator's mentor is female. The variable *Top Eng.* is a binary indicator that equals one if the innovator's mentor's job title is "Top Engineer." The variable *In Program* is a binary indicator that equals one if the innovator's mentor was a mentee in a prior cohort of the mentorship program. These indicators are interacted with the mentorship program participation variable to explore potential differences in the mentorship-productivity association across groups. The empirical specification structure matches that used in Table A.2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.15: Validating the Social Network Mechanism: Changes in Innovator Team Composition After Mentorship

Panel A: Full Sample of Data

			Primary Robustness Considerations							
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mentored	0.156***	0.151***	0.156***	0.150***	0.239***	0.158***	-0.161***	-0.160***	-0.182***	-0.184***
	(5.372)	(4.924)	(5.416)	(4.895)	(7.137)	(2.957)	(-4.026)	(-4.008)	(-5.538)	(-5.141)
Two-Way Fixed Effects	✓		✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓		✓
Adj. R-Square	0.674	0.602	0.684	0.609	0.538	0.626	0.489	0.479	0.438	0.443
Observations	11,129	73,031	9,781	63,595	853	1,080	11,280	74,036	11,390	74,794

Panel B: Long-Term Team Composition

			Primary Robustness Considerations							
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Mentored	0.056	0.056	0.056	0.055	0.161**	0.072	-0.143**	-0.146**	-0.147**	-0.150***
	(1.255)	(1.347)	(1.251)	(1.330)	(2.245)	(0.930)	(-2.449)	(-2.604)	(-2.690)	(-2.636)
Two-Way Fixed Effects	✓		✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓		✓
Adj. R-Square	0.622	0.587	0.631	0.592	0.593	0.638	0.484	0.482	0.439	0.440
Observations	10,688	62,012	9,340	53,825	412	562	10,834	62,877	10,944	63,559

Notes. This table presents an analysis of potential changes in innovator team composition associated with participation in the mentorship program. The data used is a panel of innovator-quarter team diversity metrics. Panel A utilizes the full sample, while Panel B excludes observations of mentored workers in the quarter of first mentorship and the subsequent quarter. In Columns (1)–(6), the dependent variable is the average fraction female of the innovator teams the innovator is a part of in a given quarter. The dependent variable in Columns (7)–(8) is the average gender-race HHI values of the innovator teams the innovator is a part of in a given quarter. The dependent variable in Columns (9)–(10) is the average patenting experience HHI values of the innovator teams the innovator is a part of in a given quarter. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ two estimation strategies: a two-way fixed effects estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked difference-in-differences estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(4) and Columns (7)–(10). In Columns (5)–(6), the control group is composed of not-yet mentored innovators only. In Columns (3)–(4), to reduce concerns about SUTVA violations, we remove from the sample the 520 innovators who mentored innovators collaborate with after the end of the mentorship program. For the two-way fixed effects model, standard errors are clustered by innovator and quarter. In the stacked difference-in-differences specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.16: Heterogeneity Among Non-Participants After Exposure to Mentored Colleagues

Panel A: Gender

			Primary Robustness Considerations			
	Full Sample		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)
Post First Exposure	0.610*	0.656***	0.596**	0.329*	0.246	0.209*
	(1.935)	(2.820)	(2.899)	(1.849)	(1.596)	(1.878)
Post First Exposure × Female	-0.384	-0.301	-0.410	-0.259*	-0.180	-0.131
	(-1.374)	(-1.613)	(-1.552)	(-1.749)	(-1.265)	(-1.312)
Two-Way Fixed Effects	✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓
Adj. R-Square	0.324	0.344	0.391	0.487	0.259	0.272
Observations	120,811	918,971	6,137	23,328	120,811	918,971

Panel A: URM

			Primary Robustness Considerations			
	Full Sample		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)
Post First Exposure	0.600*	0.664***	0.578***	0.372**	0.248*	0.220**
	(2.052)	(3.195)	(2.988)	(2.020)	(1.748)	(2.301)
Post First Exposure × URM	-0.467**	-0.389*	-0.483**	-0.405	-0.254**	-0.220
	(-2.527)	(-1.963)	(-2.325)	(-1.537)	(-2.596)	(-1.223)
Two-Way Fixed Effects	✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓
Adj. R-Square	0.315	0.341	0.382	0.484	0.256	0.269
Observations	152,439	1,167,235	5,962	22,388	152,439	1,167,235

Panel A: Gender or URM

			Primary Robustness Considerations			
	Full Sample		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)
Post First Exposure	0.624*	0.668***	0.627***	0.349*	0.253	0.218**
	(1.983)	(2.959)	(3.095)	(1.965)	(1.680)	(2.131)
Post First Exposure × Female/URM	-0.432*	-0.341**	-0.455**	-0.322**	-0.215**	-0.169**
	(-1.988)	(-2.410)	(-2.142)	(-2.337)	(-2.121)	(-2.186)
Two-Way Fixed Effects	✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓
Adj. R-Square	0.321	0.347	0.391	0.487	0.259	0.273
Observations	157,968	1,208,808	6,350	23,816	157,968	1,208,808

Notes. This table reports the results of tests of heterogeneous spillover effects of the mentorship program onto the productivity of non-participants. We restrict the sample to only include innovators who were not directly associated with the mentorship program at any time, either as a mentor or mentee. We create an indicator variable, *Post First Exposure*, which equals one in the first quarter a non-participating innovator is observed to collaborate with a mentored innovator and each subsequent quarter, and zero otherwise. These tests require data on the gender and/or race/ethnicity of the individual. We do not have data on these characteristics for all individuals, leading to the reductions in observation counts. The variables *Female*, *URM*, and *Female/URM* are binary indicators that equal one if the innovator is female, an underrepresented racial/ethnic minority, or either, respectively. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.17: **Engineers' Workplace Perceptions and Departures After the Mentorship Program**

Panel A: Engineers' Perceptions of the Firm

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	0.048***	-0.009	0.029***	0.027***	0.027***
	(2.658)	(-0.378)	(4.976)	(4.095)	(4.095)
Engineer		-0.193***		0.001	0.001
		(-5.846)		(0.051)	(0.069)
Post × Engineer		0.138***		0.009	0.009
		(3.751)		(0.660)	(0.673)
Post × Engineer × Focal Firm					0.123***
					(3.131)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.006	0.009	0.056	0.056	0.057
Observations	14,091	14,091	255,200	255,200	269,293

Panel B: Engineers' Departures

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	-0.054***	-0.021***	0.060***	0.060***	0.060***
	(-9.679)	(-8.335)	(15.927)	(19.865)	(19.872)
Engineer		0.945***		0.832***	0.829***
		(33.575)		(67.220)	(67.070)
Post × Engineer		-0.222***		-0.001	-0.001
		(-6.398)		(-0.068)	(-0.068)
Post × Engineer × Focal Firm					-0.220***
					(-5.745)
Region and Gender FEs	✓	✓	✓	✓	✓
Ethnicity and Seniority FEs	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.061	0.076	0.032	0.035	0.036
Observations	706,860	706,860	7,999,936	7,999,936	8,706,796

Notes. This table presents estimates of changes in engineers' perceptions of and departure from their workplace in relation to the patent mentorship program. Panel A focuses on workplace perceptions and Panel B focuses on departures. The data for workplace perception is individual reviews posted from 2018 to 2024 on Glassdoor. The analysis is restricted to reviews with non-missing job title data for the focal firm and its peers. In Panel A, the dependent variable is the overall star rating of the firm provided by the reviewer. In Panel B, the dependent variable of interest is the number of employees leaving the company within a specific occupation-region-gender-ethnicity-seniority group in a given month. *Post* equals one from when the mentorship program began and onward. Columns (1)–(2) focus on the single DiD based on the job title of the reviewer at the focal firm. Columns (3)–(4) focus on a single placebo-DiD using the industry peers. Column (5) uses data from all firms to estimate the DDD model. All models include country fixed effects, and the models in Columns (3)–(5) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.18: **Engineers' Perceptions of Corporate Culture After the Mentorship Program**

Panel A: Engineers' Perceptions of Firm Culture

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	0.078***	0.042	0.039***	0.040***	0.040***
	(3.782)	(1.553)	(5.510)	(5.013)	(5.012)
Engineer (Eng.)		-0.111***		-0.022	-0.022
		(-3.001)		(-1.432)	(-1.411)
Post × Eng.		0.089**		-0.002	-0.001
		(2.132)		(-0.100)	(-0.087)
Post × Eng. × Focal Firm					0.082*
					(1.830)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.014	0.015	0.070	0.071	0.073
Observations	11,293	11,293	195,288	195,288	206,580

Panel B: Text-Based Perceptions of Firm Culture

	Adapt-ability	Collab-oration	Community-orientation	Customer-orientation	Detail-orientation	Integ-rity	Results-orientation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post	-0.007**	0.109***	0.076***	0.131***	-0.035***	0.029***	0.007***
	(-2.010)	(22.389)	(10.499)	(38.861)	(-12.023)	(7.103)	(5.343)
Engineer (Eng.)	-0.012*	0.069***	0.092***	0.025***	0.042***	0.064***	-0.005**
	(-1.732)	(6.973)	(6.155)	(3.766)	(7.498)	(7.787)	(-2.001)
Post × Eng.	0.013*	0.026**	0.01	0.009	-0.006	0.001	-0.004*
	(1.865)	(2.438)	(0.607)	(1.267)	(-1.067)	(0.168)	(-1.663)
Post × Eng. × Focal Firm	0.071***	0.084**	0.209***	-0.035*	0.067***	0.114***	0.01
	(3.088)	(2.525)	(4.109)	(-1.760)	(3.606)	(4.271)	(1.379)
Country FE	✓	✓	✓	✓	✓	✓	✓
Firm FE	✓	✓	✓	✓	✓	✓	✓
Adj. R-squared	0.021	0.037	0.033	0.249	0.071	0.037	0.047
Observations	241,754	241,754	241,754	241,754	241,754	241,754	241,754

Notes. This table presents estimates of changes in engineers' perceptions of their workplace's culture in relation to the patent mentorship program. The data for workplace culture comes from individual reviews posted from 2018 to 2024 on Glassdoor. The analysis is restricted to reviews with non-missing job title data for the focal firm and its peers. In Panel A, the dependent variable is the star rating of the firm's "Culture & Values." In Panel B, the dependent variable is the cultural values extracted from the open-ended text in the Glassdoor reviews. *Post* equals one from when the mentorship program began and onward. Columns (1)–(2) focus on the single DiD based on the job title of the reviewer at the focal firm. Columns (3)–(4) focus on a single placebo-DiD using the industry peers. Column (5) uses data from all firms to estimate the DDD model. All models include country fixed effects, and the models in Columns (3)–(5) include firm fixed effects. In Panel A, the sum of observation counts between Columns (1) and (3) is smaller than the value in Column (5) due to the omission of singleton observations. The number of observations in Panel A is smaller than in Panel B because the star ratings for subcategories are often left blank, whereas the text and overall rating must be filled in. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.19: Engineers' Departures After the Mentorship Program (Alt. Specification)

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	-0.015***	-0.010***	0.007***	0.008***	0.008***
	(-17.264)	(-12.747)	(23.805)	(27.814)	(27.813)
Engineer		0.252***		0.169***	0.169***
		(78.506)		(203.505)	(203.378)
Post × Engineer		-0.040***		-0.006***	-0.006***
		(-9.740)		(-5.696)	(-5.696)
Post × Engineer × Focal Firm					-0.033***
					(-7.856)
Region and Gender FEs	✓	✓	✓	✓	✓
Ethnicity and Seniority FEs	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.206	0.244	0.207	0.223	0.223
Observations	706,860	706,860	7,999,936	7,999,936	8,706,796

Notes. This table presents a robustness analysis of potential changes in employee departure patterns coinciding with the implementation of the patent mentorship program. The data used is employee departure measures from 2018 to 2024 from Revelio, a labor market intelligence firm that tracks employee movements across firms. This table examines the inverse hyperbolic sine transformation of this scaled external outflow measure instead of just analyzing employee outflows as in Table 9. We employ ordinary least squares regressions to analyze these departure measures in relation to the binary indicator *Post*, which equals one from when the mentorship program began and onward. Some models include interactions between *Post* and the binary indicator *Engineer*, while others further incorporate interactions with the binary indicator *Focal Firm*. Columns (1)–(2) use only departure measures from the focal firm. Columns (3)–(4) use only departure measures from the firm's industry peers. Column (5) uses data from all firms. All models include region, gender, ethnicity, and seniority fixed effects, and the models in Columns (3)–(5) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.20: Named Inventors' Departures After the Mentorship Program

	Focal Firm Only		Peer Firms Only		Both	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	-0.035***	-0.031***	-0.033***	-0.015***	-0.033***	-0.015***
	(-3.792)	(-3.687)	(-9.360)	(-5.054)	(-9.619)	(-5.034)
Post × Focal Firm					-0.025**	-0.024**
					(-2.211)	(-2.533)
Region and Gender FEs	✓	✓	✓	✓	✓	✓
Ethnicity FEs	✓	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓	✓
Additional controls		✓		✓		✓
Adjusted R^2	0.002	0.220	0.012	0.319	0.011	0.311
Observations	4,381	4,381	36,023	36,023	38,428	38,428

Notes. This table presents a robustness test of potential changes in engineers' departure patterns coinciding with the implementation of the patent mentorship program. The data used is LinkedIn position-level data for 2018 to 2024 from Revelio, a labor market intelligence firm that tracks employee movements across firms. This table limits the sample to only named inventors on granted patents at either the focal firm or peer firms. We use a multi-step algorithm to match the disambiguated inventor data from the U.S. PTO to LinkedIn users, allowing us to match to more inventors than just those who put their patents on their profiles. We then use the position start and end dates to determine if the inventor leaves their job in a given year. We employ ordinary least squares regressions to analyze these job departures in relation to the binary indicator *Post*, which equals one from when the mentorship program began and onward. Some models include interactions between *Post* and the binary indicator *Focal Firm*. Columns (1)–(2) use only departure measures from the focal firm. Columns (3)–(4) use only departure measures from the firm's industry peers. Column (5)–(6) use data from all firms. All models include region, gender, and ethnicity fixed effects, and the models in even Columns include additional controls for compensation, bonus, and any additional positions held (i.e., board seats). t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.21: Innovator Team Size Dynamics

Panel A: After First Mentorship Cohort

	Total Innovators			IHS(Total Innovators)		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.249***	0.220***	0.220***	0.077***	0.056***	0.056***
	(5.027)	(11.118)	(11.216)	(5.846)	(10.854)	(10.926)
Post × Focal Firm			0.029			0.021
			(0.469)			(1.320)
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.004	0.025	0.024	0.005	0.033	0.032
Observations	5,735	88,497	94,232	5,735	88,497	94,232

Panel B: After Second Mentorship Cohort

	Total Innovators			IHS(Total Innovators)		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.252***	0.216***	0.216***	0.078***	0.055***	0.055***
	(4.904)	(10.013)	(10.101)	(5.724)	(9.825)	(9.891)
Post × Focal Firm			0.036			0.023
			(0.559)			(1.388)
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.004	0.024	0.023	0.005	0.033	0.032
Observations	5,735	88,497	94,232	5,735	88,497	94,232

Panel C: After Third Mentorship Cohort

	Total Innovators			IHS(Total Innovators)		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.232***	0.206***	0.206***	0.069***	0.051***	0.051***
	(3.886)	(7.968)	(8.038)	(4.484)	(7.573)	(7.623)
Post × Focal Firm			0.025			0.019
			(0.345)			(0.975)
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.003	0.024	0.023	0.003	0.032	0.031
Observations	5,735	88,497	94,232	5,735	88,497	94,232

Notes. The data used is patent grant-level data from the U.S. PTO from 2018 to the present. The dependent variable in Columns (1)–(3) is the total number of innovators listed on a particular patent, and the dependent variable in Columns (4)–(6) is the inverse hyperbolic sine of the total number of innovators listed on a particular patent. We employ ordinary least squares regressions to analyze how these patent team size measures change in relation to the binary indicator *Post*, which equals one from when the first mentorship program cohort began and onward (in Panel A). In Panel B (Panel C), *Post* equals one from when the second (third) mentorship program cohort and onward. Some models include interactions between *Post* and the binary indicator *Focal Firm*. Columns (1) and (4) use only team size measures from the focal firm. Columns (2) and (5) use only team size measures from the firm's industry peers. Columns (3) and (6) use data from all firms. The models in Columns (2), (3), (5), and (6) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.22: Aggregate Patent Characteristics

	Patent Count		Patent Value	
	All	Hot	Nominal	Real
	(1)	(2)	(3)	(4)
Post × Focal Firm	401.626***	38.336**	10987.316*	3742.552*
	(3.943)	(2.065)	(1.704)	(1.682)
Firm and Time FEs	✓	✓	✓	✓
Adj. R-Square	0.672	0.564	0.796	0.790
Observations	262	262	262	262

Notes. This table presents estimates of aggregate changes in granted patent characteristics based on patent grant-level data from the U.S. PTO from 2018 to the present. The dependent variables in Columns (1)–(2) are patent count measures. The first is the total number of patents applied for in a given quarter that get granted. The second is the total number of patents applied for in a given quarter that get granted and receive forward citations in the upper quartile of citations in the first three years, which are commonly referred to as "hot" patents and are associated with long-term value creation. The dependent variables in Columns (3)–(4) are patent value measures, nominal and real value, respectively. The measures are estimated using the stock-market approach pioneered by Kogan et al. (2017). All specifications are at the firm-quarter level and focus on the focal firm and its peers. *Post* equals one from when the mentorship program began and onward. Controls for application quarter average claims, backward citations, examiner citations, the total number of inventors, and having a female inventor as well as average firm characteristics, including size, market-to-book, R&D to revenue, cash holdings, leverage, and profitability, are included. Firm and quarter fixed effects are included. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

B Internet Appendix: Additional Robustness and Falsification Tests

This appendix demonstrates the robustness of our findings on the relation between patent mentorship and innovation productivity. Specifically, we validate our main results when using alternative control groups, employing different estimators, using different disclosure status classifications, using different idea quality metrics, and excluding collaborations with mentors. In addition, we conduct placebo tests for employee satisfaction and turnover outcomes. Finally, we provide additional details and sensitivity checks on our ROI calculations.

B.1 Alternative Control Groups

In Table 2 (and Table A.2–Table A.3), we presented our main estimates using a combined control group of both not-yet-mentored innovators and never-mentored innovators as the basis for comparison. We also showed the robustness of our main findings when using only not-yet-mentored innovators in the control group. To further assess robustness, we examine alternative control group specifications. In Columns (1) and (2) of Table B.1, we restrict the control group to never-mentored innovators, and we employ the stacked difference-in-differences estimator. In Column (1), we use the full sample of never-mentored innovators, while in Column (2), we restrict this sample to those most comparable to mentored innovators based on observable characteristics. We identify comparable innovators using the Mahalanobis matching procedure and our rich set of demographic data as observable characteristics (innovator tenure, gender, race/ethnicity, marital status, job role, and work country). We keep only the never-mentored innovators with above-median propensity scores as our final comparison group in Column (2). The estimates in Panel A, which uses the full sample of data, are all positive and statistically significant at the 1% level, with magnitudes comparable to the corresponding estimates in our main tables. Similarly, the results in Panel B, wherein we remove observations in the quarter of mentorship and the quarter after mentorship, demonstrate robust long-term productivity gains across alternative control groups and matching procedures. Collectively, these robustness tests support a positive and significant relation between patent mentorship program participation and innovation productivity as measured by invention disclosure submissions.

B.2 Alternative Estimators

In Columns (3) and (4) of Table B.1, we show that the estimated relation between mentorship and productivity is robust if we use two alternative estimators: the Poisson pseudo-likelihood estimator, which accounts for the count-like nature of disclosure submissions, and the Callaway and Sant'Anna (2021) estimator. We use the Poisson Pseudo-Maximum Likelihood estimator in Column (3), and we employ the Callaway and Sant'Anna (2021) estimator

in Column (4).[29] As in our main specifications, we use both not-yet-mentored and never-mentored innovators in the control group. We find consistent evidence of a positive and statistically significant relation between mentorship and overall submission outcomes, regardless of the choice in estimator. The estimate in Column (3) of Panel A, derived from the Poisson estimator, can be interpreted as mentorship leading to a level of disclosure output that is about six times the baseline rate, $e^{1.808} = 6.098$. Given a baseline of 0.09 submissions per quarter, this would translate into mentored workers submitting about 0.50 additional disclosures per quarter. The magnitude of this effect is comparable to the larger estimates reported in Table A.2. The estimates in Column (4), where the Callaway and Sant'Anna (2021) estimator is used, are quite similar in magnitude to our baseline estimates.

B.3 Alternative Disclosure Status Classifications

In Table B.2, we show the robustness of our results when using more granular disclosure status classifications. Specifically, we focus on the four most common types of disclosures: "Approved - Patent Application" (31% of the disclosures); "Declined" or "Auto-Declined" (65% of the disclosures); "Defensive Publication Candidate" or "Approved - Defensive Publication" (2% of the disclosures).[30] We use the full sample of data in Panel A. In Panel B, we consider the long-term productivity effects by removing the observations of mentored workers in the quarter of mentorship and the first quarter after mentorship. The dependent variable in Columns (1)–(2) is the number of disclosure submissions with the status "Declined" or "Auto Declined." The dependent variable in Columns (3)–(4) is the number of disclosure submissions with the status "Approved - Patent Application." The dependent variable in Columns (5)–(6) is the number of disclosure submissions with the status "Defensive Publication Candidate" or "Approved - Defensive Publication." We employ the two-way fixed effects estimator in the odd-numbered columns, and we employ the stacked difference-in-differences estimator in even-numbered columns. We find that the relation between mentorship and innovation outcomes is positive and statistically significant across all disclosure statuses, both overall and in the long-term.

[29] Using this approach, we separately consider the treatment effects across treatment cohorts, where in our case a treatment cohort designation is based on the year-quarter in which an employee first participates in the mentorship program, g. We estimate cohort- and time-specific average treatment effects on the treated employees, $\widehat{ATT}(g,t)$, using two-period/two-group difference-in-differences estimators. Then we compute summary treatment effect estimates by aggregating together the cohort- and time-specific estimates, weighting them by the size of each treatment cohort.

[30] The remaining 2% have statuses such as "Submitted," "RFC," and "Deleted."

B.4 Alternative Idea Quality Measures

In Table B.3, we evaluate the robustness of our results when using traditional measures of patent quality rather than the full innovation pipeline. Although we cannot directly link U.S. PTO data to the mentorship data due to security concerns related to PII and GDPR (which led us to pursue an IRB enabling consent and personnel data collection but not actual IP data linkage to U.S. PTO datasets), we can examine changes in patent quality over time from the U.S. PTO data. The table presents three distinct regression samples: *Focal*, Columns (1) and (4), include only granted patents from the focal firm, *Peer*, Columns (2) and (5), include only patents from peer firms, and *Both*, Columns (3) and (6), combine these samples to estimate the DiD specification, which is our specification of primary interest. The data spans 2018 to the present, with the latest iteration of the U.S. patent grant data allowing us to include patents applied for in 2023. *Post* equals one from when the mentorship program began and onward, zero otherwise.

When we examine the DiD specification, we find evidence consistent with a statistically significant increase in granted patent quality after the mentorship program implementation. This specification follows best practices outlined by Lerner and Seru (2022). Fixed effects include those for patent classification, application year, application state, and firm. Patent controls include backward citations, examiner citations, total number of inventors, and an indicator for female inventors. Firm controls include firm size, market-to-book, R&D to revenue, cash holdings, leverage, and profitability.

In Panel A, which focuses on forward citations, which are citations to a patent or a published patent application that are later or forward in time from the cited patent, we see that post-mentorship patents at the focal firm have increased forward citations (0.211 more per year, significant at 5%). This is also true when we look at hot citations, which is a proxy for long-term value creation. The hot citations measure is the number of forward citations received in the first three years (Gruner, 2021).

In Panel B, we see negligible changes in scope. We consider both the measure of scope proposed by Lerner (1994) based on the number of patent classifications as well as the first claim word count proposed by Kuhn and Thompson (2019), which measures patent scope inversely by counting words in the patent's first claim. Longer claims contain more specific details and limitations, each of which must be present for infringement to occur. This increased specificity effectively narrows the patent's legal protection, creating more opportunities for competitors to design non-infringing alternatives. Finally, in Panel C, we see a small increase in total claims at the focal firm, but the finding loses statistical significance in the DiD setting. Moreover, when we look at independent claims, of which you need one to survive in prosecution for a patent to be valid, we see no changes.

These findings suggest the mentorship program is associated with improvements in patent quality as proxied by forward citations, but it is not significantly associated with changes in scope or claims. This suggests that rather than diluting the pool of patent applications with a bunch of weak ideas, the focal firm's patent quality improves, thereby providing robust support for our main conclusion about an increase in innovator productivity.

B.5 Additional Restrictions to Mitigate SUTVA Violation Concerns

In Table B.4, we show the robustness of our main findings when removing all non-mentored innovators based in participating countries from the control group. Specifically, in Columns (1)–(4) and (7)–(8), where we include never-treated innovators within the control group, we restrict this group of innovators to be based in countries in which no individuals participated in the program. This exclusion helps to further reduce concerns about SUTVA violations, as never-mentored innovators in countries where no one participated in the program would be very unlikely to have experienced spillover effects related to the program.

B.6 Additional Restrictions on Collaborations with Mentors

In Table B.5, we show the robustness of our main findings when removing all disclosures for which a mentor (past, present, or future) is involved as a collaborator. By removing these disclosures, we mitigate the concern that pre-existing relationships between mentors and mentees could influence the relation between program participation and disclosure submission outcomes. We then account for the fact that some mentees participate in multiple cohorts of the mentorship program. One may be concerned that the program is only beneficial after an innovator's second (or even third) time participating. We do not find that this is the case. Specifically, in Table B.6, for innovators who participate in a second cohort of the mentorship program, we remove all data for them from the year-quarter of the second program onward. By doing so, we estimate (exclusively) the relation between first-time mentorship program participation and subsequent productivity. The results in Table B.6 are very similar to our main results, mitigating the concern that additional mentorship is necessary for benefits to be realized.

B.7 Alternative Employee Perception and Retention Tests

The results in Panel A of Table 9 (and those in Table A.17 and Table A.18) suggest that engineers' perceptions of the focal firm improved after the establishment of the mentorship program, highlighting an additional benefit of the program. In Table B.7, we show that the estimated relation between the mentorship program and engineers' perceptions of the focal firm is robust to alternative definitions of *Post*. In particular, just as we would expect, the

positive link remains statistically significant, but attenuated, when *Post* is set to switch from zero to one at a later time (after either the second or third mentorship cohorts begin). This is expected, as the positive perceptions of early cohort participants may be captured in the pre-event period when *Post* switches to one at later times. Again, if the estimated relation is due to forces independent of the mentorship program, we would not be likely to observe this expected pattern in the data.

We also conduct falsification tests using the employee sentiment data to show that the baseline results are not likely driven by some other factors. Specifically, we replace *Post* with the binary indicator *Pseudo-Post*, which equals one in and after October of 2018 (two years before the first cohort of the mentorship program), and we omit data after December of 2019 (to perform the falsification tests solely in the pre-COVID period). We report the results in Table B.8. The null interaction effects on *Pseudo-Post* × *Engineer* in Column (2) and on *Pseudo-Post* × *Engineer* × *Focal Firm* in Column (5) suggest that the corresponding interaction effects in our main tests, which are significant, are not driven by some alternative temporal variation that impacts engineers in the focal firm differently than engineers in peer firms.

Similarly, the results in Table B.9 show that the estimated relation between the mentorship program and engineers' departures from the focal firm is robust to alternative definitions of *Post*. In these tests, we note that the magnitudes of the coefficients in Column (5) are very similar to those in Columns (1) and (2) in Panel B of Table 9. So, the relation between the onset of the mentorship program and engineer turnover does not attenuate as immediately as does the relation between the mentorship program and employee sentiment. In addition, the results in Table B.10, which report null interaction effects on *Pseudo-Post* × *Engineer* in Column (2) and on *Pseudo-Post* × *Engineer* × *Focal Firm* in Column (5), suggest that the corresponding interaction effects in our main tests are not driven by some alternative factors.

Finally, in Figure B.1a and Figure B.1b we plot the estimates from parallel trends tests. The point estimates capture the quarterly differences in employee perceptions and departure outcomes, respectively, between engineers and non-engineers in the focal firm, before and after the onset of the mentorship program. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the first cohort of the mentorship program as the baseline reference period. We formally test whether the pre-mentorship coefficients differ from zero. The statistical tests fail to reject the null hypothesis that these coefficients equal zero, with p-values ranging from 0.185 to 0.751 (see figure notes).

B.8 Alternative ROI Calculation Assumptions

This appendix section provides detailed calculations and an extended analysis supporting our ROI analysis of the mentorship program in Section 6 of the main paper.

Our approach begins with directly quantifiable outcomes to establish a baseline estimate of the mentorship program's value and then extends to broader organizational benefits. Because returns to human capital investments are difficult to measure precisely, we report a range of estimates that vary in their degree of conservativeness.

For the direct productivity channel, we first assign value to patents using a conservative cost-based method akin to classic renewal-fee approaches (Pakes, 1986). Patent attorney fees (approximately $25,000) and required renewal fees ($13,460) total $38,460 per patent. This method almost certainly understates economic value but provides a transparent lower bound. Using market-based valuation following Kogan et al. (2017a), the median value of a granted patent between 2018 and 2023 is $19.7 million, and the mean is $68.1 million. For comparison, the focal firm's median patent is roughly 2.5 times larger, implying that applying firm-specific values would yield patent valuations over 1,000 times higher than the cost-based estimate. Even assigning a blended value across peer firms yields estimates roughly 400 times larger than those derived from attorney and renewal fees.

A third valuation benchmark comes from patent reassignment markets. Roughly 15% of patents change owners and are priced at transfer (Serrano, 2010). For example, Facebook acquired 650 patents and licensed 275 more from Microsoft for $550 million in 2012, or roughly $600,000 per patent. Using this mid-range figure implies returns roughly 16 times larger than the cost-based estimate, corresponding to $106.3 million in direct gains (or $34.5 million when scaled by team size).

Although these valuation metrics differ substantially in magnitude, all indicate that the mentorship program generates a meaningful direct innovation gains.

To value the benefits of diversity, we implicitly assume that engineering talent is not optimally allocated (Hsieh et al., 2019) and that more diverse innovator teams move the allocation of talent closer to optimal. While it is possible that diversity can increase coordination costs and reduce innovative efficiency (Alesina and Ferrara, 2005) rather than improve problem-solving (Hong and Page, 2001), the empirical evidence is consistent with collaboration among diverse teams being important for innovative efficiency and the production of high-value patents (Jones, 2009; Wuchty et al., 2007; Li and Wang, 2023). In the main specification, we assumed a 5% increase in patent value from the $38,460 baseline. The 5% increase is conservative relative to estimates in the literature and could arguably be increased.

The mentorship program likely generates additional value through improved retention

and cultural enhancements, though these channels are harder to quantify precisely. Following program implementation, engineer turnover falls by 3.3%. Even attributing only a small share of this decline to mentoring can yield meaningful savings, given the high cost of replacing technical workers. Recruiting engineers typically requires substantial expenditures on search, interviews, onboarding, and role-specific training, and new hires often take up to six months to reach full productivity. Industry estimates place total replacement costs at roughly one to one and a half times an engineer's annual salary (Berk and Kase, 2010).

To provide a conservative estimate, we follow an approach similar to Lazear et al. (2015) and assume that only one percent of the 3.3% turnover decline is attributable to mentoring. Multiplying this by the firm's 29,212 inventors and by one year of average engineer salary ($157,297, drawn from Revelio and weighted by region and seniority) yields implied savings of $45.9 million. Using an even more conservative benchmark that restricts attention to mentored engineers alone, who depart at a rate of 6.3% compared to 8.8% among non-mentored innovators, implies roughly 16 retained workers (2.5% $\times$ 633 mentees) and cost savings of approximately $2.5 million (16 $\times$ 157,297).

Employees also value non-pecuniary workplace attributes. Using Makridis (2018)'s estimate that workers are willing to forego $1,159 annually for a stronger corporate culture, and applying this to the IT sector's 19 percent turnover rate (BLS), we estimate additional savings of $19.2 million across a three-year period (29,212 inventors $\times$ 0.19 $\times$ $1,159 $\times$ 3).

These retention and satisfaction effects relate to a growing literature emphasizing organizational capital. Graham et al. (2022a) show that improvements in corporate culture predict increases in Tobin's Q. The focal firm's Tobin's Q rose after program adoption, and attributing even one percentage point of this change to culture improvements implies benefits of more than $650 million given the firm's market capitalization. Similarly, Eisfeldt and Papanikolaou (2013) argue that organizational capital is embodied in key talent and forms a significant component of intangible assets. Applying their 0.3 scaling parameter to the salaries of participating mentees yields an estimated $89.6 million in organizational-capital accumulation over the three-year program (633 mentees $\times$ $157,297 $\times$ 0.3 $\times$ 3).

A potential cost of the program is the opportunity cost of mentor time. However, mentors do not appear to reduce their own inventive output. Before mentoring, mentors submitted 2.6 disclosures per quarter; during the program their output rose to 3.7, and afterward it returned to baseline levels. To be conservative, we treat all hours devoted to mentoring as fully crowding out existing work. Mentees devoted 70 hours (3.6% of annual work time), generating an implied cost of $3.58 million (633 mentees $\times$ $157,297 $\times$ 0.036). Mentors devoted 40 hours (2.1 percent of calendar year hours), and assuming mentors earn twice the average engineer salary, this adds $0.75 million (113 mentors $\times$ $314,594 $\times$ 0.021).

Administrative expenses were small: roughly $14,000 for seven cohorts.

In total, the program produced direct gains from mentees' increased innovation output and indirect gains from collaborator spillovers, increased patent value from more diverse teams, and retention and culture benefits (ranging from $2.5 million to tens of millions depending on scenario). Even under the most conservative assumptions of only team-scaled direct and indirect benefits of innovation, and no cultural or retention gains, the total benefits substantially exceed labor costs. Under more realistic valuation assumptions, the gains are considerably larger. Across scenarios, the implied ROI remains positive, indicating that mentorship is a cost-effective human capital investment.

These calculations rest on transparent assumptions and observable data, but several considerations merit note. Patent valuation approaches capture only a fraction of inventions' long-run economic contribution. Measured spillovers include only direct collaborators and exclude wider network diffusion. Retention effects are estimated in a setting with potential confounders, though conversations with program administrators indicate no concurrent initiatives. Finally, converting cultural gains to monetary values is inherently difficult, though prior research suggests culture is a first-order determinant of firm performance (Graham et al., 2022a).

B.9 Measures of Collaborator Team Diversity

To study disclosure team diversity, as we do in Table 7, we consider three separate outcomes, each at the innovator-quarter level. First, we consider the average level of female representation among an innovator's collaborator teams, based on the disclosure submissions they had in a particular quarter. If, for example, innovator i submitted two disclosures in quarter t, with two different teams, where one team was composed of 50% women and the other team was composed of 25% women, then the value of *Fraction Female* for innovator i in quarter t would be 37.5% = (50% + 25%) / 2. This average fraction female teammates outcome becomes the dependent variable in Columns (1) and (4) of Table 7 (and Columns (1)–(6) of Table A.15). We also consider the average gender-by-racial/ethnic composition of an innovator's collaborator teams each quarter. Specifically, we compute gender-by-racial/ethnic HHI metrics for each collaborator team, and then we take the average value across all collaborator teams for a given innovator-quarter. Note, smaller HHI values capture larger amounts of diversity (i.e., less concentration in representation among a single group). This HHI metric becomes the dependent variable in Columns (2) and (5) of Table 7 (and Columns (7) and (8) of Table A.15). Finally, we also consider the average patenting experience composition of an innovator's collaborator teams each quarter. Specifically, we compute patenting experience HHI metrics for each collaborator team, based on the number of approved patent

filings each team member had contributed to up to that point in time (e.g., 0, 1–2, 3–5, 6–10, or 11+), and then we take the average value across all collaborator teams for a given innovator-quarter. This HHI metric becomes the dependent variable in Columns (3) and (6) of Table 7 (and Columns (9) and (10) of Table A.15).

Table B.1: Additional Robustness Tests of the Mentor-Productivity Relation

Panel A: Full Sample of Data

	Alt. Control Group Never Mentored		Alt. Estimators Poisson	Alt. Estimators CSDID
	(1)	(2)	(3)	(4)
Mentored	0.254***	0.282***	1.808***	0.220***
	(3.754)	(4.231)	(5.839)	(9.690)
Two-Way Fixed Effects			✓	
Stacked Diff-in-Diff	✓	✓		
PS-Matched Sample		✓		
Adj./Pse. R-Square	0.403	0.415	0.408	-
Observations	3,260,279	372,891	95,213	473,468

Panel B: Long-Term Productivity

	Alt. Control Group Never Mentored		Alt. Estimators Poisson	Alt. Estimators CSDID
	(1)	(2)	(3)	(4)
Mentored	0.083***	0.112***	0.940***	0.082***
	(3.462)	(4.537)	(4.117)	(3.400)
Two-Way Fixed Effects			✓	
Stacked Diff-in-Diff	✓	✓		
PS-Matched Sample		✓		
Adj./Pse. R-Square	0.396	0.412	0.411	–
Observations	2,874,766	328,248	89,772	471,066

Notes. The data used is a panel of innovator-quarter IDF submissions. The dependent variable in all specifications is the number of disclosure submissions of any status that an innovator submitted in a given quarter. *Mentored* equals one in the quarter in which an innovator is part of the mentorship program, and each subsequent quarter, and zero otherwise. Panel A utilizes the full sample, while Panel B excludes observations of mentored workers in the quarter of first mentorship and the subsequent quarter. In Columns (1) and (2), we use the stacked difference-in-differences estimator, and the control group includes only never-mentored innovators. We use all never-mentored innovators in Column (1), and we use only those with above median mentorship program participation propensity scores in Column (2). In Columns (3) and (4), we use both eventually-mentored and never-mentored innovators in the control group. In Column (3), we employ the Poisson pseudo-likelihood estimator with innovator and year-quarter fixed effects. In Column (4), we employ the Callaway and Sant'Anna (2021) estimator. t-statistics (z-statistics) are reported in parentheses. The Callaway and Sant'Anna (2021) estimator does not produce R-squared values. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.2: Mentor-Productivity Relation: Alternative Disclosure Status Classifications

Panel A: Full Sample of Data

	Declined		Approved Patent		Publication	
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.128**	0.126***	0.045***	0.045***	0.025**	0.025***
	(2.823)	(3.770)	(3.824)	(3.781)	(2.828)	(3.276)
Two-Way Fixed Effects	✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓
Adj. R-Square	0.299	0.306	0.308	0.309	0.123	0.130
Observations	473,939	3,272,940	473,939	3,272,940	473,939	3,272,940

Panel B: Long-Term Productivity

	Declined		Approved Patent		Publication	
	(1)	(2)	(3)	(4)	(5)	(6)
Mentored	0.048*	0.047***	0.016***	0.016***	0.008*	0.008**
	(1.844)	(2.627)	(3.091)	(2.870)	(1.901)	(2.344)
Two-Way Fixed Effects	✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓
Adj. R-Square	0.303	0.302	0.308	0.301	0.128	0.131
Observations	472,671	2,884,767	472,671	2,884,767	472,671	2,884,767

Notes. The data used is a panel of innovator-quarter IDF submissions. We use the full sample of data in Panel A. In Panel B, we remove the observations of mentored workers in the quarter of mentorship and the first quarter after mentorship. The dependent variable in Columns (1)–(2) is the number of disclosure submissions with the status "Declined" or "Auto Declined." The dependent variable in Columns (3)–(4) is the number of disclosure submissions with the status "Approved - Patent Application." The dependent variable in Columns (5)–(6) is the number of disclosure submissions with the status "Defensive Publication Candidate" or "Approved - Defensive Publication." *Mentored* equals one in the quarter in which an innovator is part of the mentorship program, and each subsequent quarter, and zero otherwise. We employ two estimation strategies: a two-way fixed effects estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked difference-in-differences estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators. For the two-way fixed effects model, standard errors are clustered by innovator and quarter. In the stacked difference-in-differences specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.3: Patent Characteristics and the Mentorship Program

Panel A: Patent Citations

	Forward Citations			Hot Citations		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.231*	0.006	0.012	0.206	-0.066	-0.064
	(1.803)	(0.142)	(0.311)	(1.281)	(-0.876)	(-0.883)
Post × Focal Firm			0.211**			0.375**
			(2.197)			(2.131)
CPC, App. Year, and State FEs	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.009	0.090	0.078	0.012	0.097	0.088
Observations	1,424	13,482	14,909	1,424	13,482	14,909

Panel B: Patent Scope

	Scope (CPC)			Scope (Text)		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	-0.036	-0.033	-0.044*	0.036	-0.077***	-0.069***
	(-0.493)	(-1.183)	(-1.658)	(0.704)	(-3.492)	(-3.293)
Post × Focal Firm			0.069			0.034
			(1.061)			(0.668)
CPC, App. Year, and State FEs	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.116	0.111	0.110	0.041	0.076	0.077
Observations	1,424	13,482	14,909	1,424	13,482	14,909

Panel C: Patent Claims

	Total Claims			Independent Claims		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.474*	-0.054	-0.032	0.004	0.007	0.007
	(1.918)	(-0.563)	(-0.345)	(0.107)	(0.375)	(0.403)
Post × Focal Firm			0.234			-0.009
			(1.040)			(-0.196)
CPC, App. Year, and State FEs	✓	✓	✓	✓	✓	✓
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.018	0.108	0.104	0.018	0.081	0.079
Observations	1,424	13,482	14,909	1,424	13,482	14,909

Notes. This table presents our analysis of changes in granted patent characteristics based on patent applications that coincide with the implementation of the patent mentorship program. Consistent with our main analysis, we use data from Q1 of 2020 onwards to ensure all of our data is the during-COVID or post-COVID era. Panel A focuses on forward citations, Panel B on patent scope, and Panel C on patent claims. *Post* equals one from when the mentorship program began and onward. Some models include interactions between *Post* and the binary indicator *Focal Firm*. All specifications include controls for backward citations, examiner citations, the total number of inventors, and the presence of female inventors, as well as firm characteristics, including firm size, market-to-book, R&D to revenue, cash holdings, leverage, and profitability. All models include fixed effects for patent class, application year, and application state. The models with more than one firm include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.4: Innovator Productivity After Mentorship (Country-Based SUTVA Tests)

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.234***	0.229***	0.234***	0.229***	0.555***	0.348***	0.048***	0.047***
	(2.928)	(3.335)	(2.929)	(3.336)	(4.269)	(3.983)	(3.216)	(3.469)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.099	0.071	0.100	0.072	0.206	0.225	0.041	0.039
Observations	266,952	1,824,031	266,499	1,820,860	9,549	22,210	266,952	1,824,031

Notes. This table presents an analysis of the relationship between participation in the mentorship program and innovator productivity, as measured by invention disclosure form (IDF) submissions. The data comprises a panel of innovator-quarter submissions. In Columns (1)–(6), the dependent variable is the total number of disclosure submissions of any status that an innovator submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of advanced status submissions. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ two estimation strategies: a two-way fixed effects estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked difference-in-differences estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(4) and Columns (7)–(8). In Columns (5)–(6), the control group is composed of not-yet mentored innovators only. In Columns (3)–(4), to reduce concerns about SUTVA violations, we remove from the sample the 520 innovators who mentored innovators collaborate with after the end of the mentorship program. Furthermore, in Columns (1)–(4) and (7)–(8), where we include never-treated innovators within the control group, we further restrict this group of innovators to be based in countries in which no individuals participated in the program. This exclusion helps to further reduce concerns about SUTVA violations, as never-mentored innovators in countries where no one participated in the program would be very unlikely to have experienced spillover effects related to the program. For the two-way fixed effects model, standard errors are clustered by innovator and quarter. In the stacked difference-in-differences specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.5: Innovator Productivity After Mentorship (Removing IDFs with Mentors)

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.079***	0.078***	0.078***	0.078***	0.141***	0.083***	0.019***	0.019***
	(3.056)	(3.094)	(2.992)	(3.072)	(3.750)	(3.092)	(4.134)	(4.185)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.207	0.208	0.162	0.163	0.075	0.111	0.111	0.111
Observations	473,939	3,272,940	465,778	3,215,813	9,549	22,210	473,939	3,272,940

Notes. This table presents an analysis of the relationship between participation in the mentorship program and innovator productivity, as measured by invention disclosure form (IDF) submissions. The data comprises a panel of innovator-quarter submissions. We exclude all submissions for which a mentor (past, present, or future) was part of the innovator team. In Columns (1)–(6), the dependent variable is the total number of disclosure submissions of any status that an innovator submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of advanced status submissions. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ two estimation strategies: a two-way fixed effects estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked difference-in-differences estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(4) and Columns (7)–(8). In Columns (5)–(6), the control group is composed of not-yet mentored innovators only. In Columns (3)–(4), to reduce concerns about SUTVA violations, we remove from the sample the 520 innovators who mentored innovators collaborate with after the end of the mentorship program. For the two-way fixed effects model, standard errors are clustered by innovator and quarter. In the stacked difference-in-differences specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.6: Innovator Productivity After Mentorship (Removing Subsequent Enrollments)

			Primary Robustness Considerations					
	Full Sample		SUTVA		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.251***	0.249***	0.251***	0.249***	0.510***	0.311***	0.053***	0.020***
	(3.450)	(3.159)	(3.403)	(3.144)	(4.238)	(3.268)	(3.630)	(3.956)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.400	0.403	0.394	0.397	0.208	0.241	0.309	0.111
Observations	472,798	3,271,799	464,637	3,214,672	8,408	21,069	472,798	3,271,799

Notes. This table presents an analysis of the relationship between participation in the mentorship program and innovator productivity, as measured by invention disclosure form (IDF) submissions. The data comprises a panel of innovator-quarter submissions. For mentees who enrolled in more than one mentorship program, we remove all data for those individuals from their second program enrollment onward. In Columns (1)–(6), the dependent variable is the total number of disclosure submissions of any status that an innovator submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of advanced status submissions. The variable of interest, *Mentored*, is an indicator equal to one in the quarter in which an innovator is part of the mentorship program and all subsequent quarters, and zero otherwise. We employ two estimation strategies: a two-way fixed effects estimator with innovator and time (calendar quarter) fixed effects in the odd-numbered columns, and we employ a stacked difference-in-differences estimator in even-numbered columns. The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(4) and Columns (7)–(8). In Columns (5)–(6), the control group is composed of not-yet mentored innovators only. In Columns (3)–(4), to reduce concerns about SUTVA violations, we remove from the sample the 520 innovators who mentored innovators collaborate with after the end of the mentorship program. For the two-way fixed effects model, standard errors are clustered by innovator and quarter. In the stacked difference-in-differences specification, clustering is at the innovator-by-cohort and quarter-by-cohort levels. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.7: Engineers' Workplace Perceptions: Alternative Timing

Panel A: After the Second Mentorship Cohort

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	0.069***	0.020	-0.003	-0.007	-0.007
	(4.054)	(0.904)	(-0.592)	(-1.177)	(-1.136)
Engineer		-0.171***		-0.006	-0.005
		(-5.792)		(-0.500)	(-0.470)
Post × Engineer		0.118***		0.018	0.018
		(3.449)		(1.467)	(1.470)
Post × Engineer × Focal Firm					0.091**
					(2.498)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.007	0.010	0.056	0.056	0.056
Observations	14,091	14,091	255,200	255,200	269,293

Panel B: After the Third Mentorship Cohort

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	0.056***	0.020	0.003	0.004	0.004
	(3.608)	(0.984)	(0.748)	(0.784)	(0.868)
Engineer		-0.133***		0.012	0.012
		(-5.667)		(1.363)	(1.418)
Post × Engineer		0.084***		-0.004	-0.004
		(2.723)		(-0.387)	(-0.402)
Post × Engineer × Focal Firm					0.084***
					(2.580)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.007	0.009	0.056	0.056	0.056
Observations	14,091	14,091	255,200	255,200	269,293

Notes. The data used are individual reviews from 2018 to 2024 posted on Glassdoor by current employees of the focal firm and the firm's industry peers. We only use reviews for individuals with non-missing job title data (necessary to identify engineers). The dependent variable is the overall rating of the firm left by the individual. We estimate ordinary least squares regressions of these ratings on the binary indicator *Post*, which equals one in and after January of 2021 (July of 2021) in Panel A (Panel B). In some models, we interact the binary indicator *Engineer* with *Post*, and in others we also interact the binary indicator *Focal Firm* into the model. Columns (1)–(2) use only reviews about the focal firm. Columns (3)–(4) use only reviews about the firm's industry peers. Column (5) uses data from all firms. All models include country fixed effects, and the models in Columns (3)–(5) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.8: Engineers' Workplace Perceptions: Falsification Tests

Two-Year Lagged Timing

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Pseudo-Post	0.061	0.046	0.047***	0.042**	0.041**
	(1.415)	(0.841)	(3.072)	(2.484)	(2.423)
Engineer		-0.205***		0.001	0.004
		(-2.830)		(0.038)	(0.124)
Pseudo-Post × Engineer		0.030		0.024	0.023
		(0.350)		(0.658)	(0.635)
Pseudo-Post × Engineer × Focal Firm					0.017
					(0.186)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.013	0.020	0.069	0.069	0.072
Observations	2,385	2,385	24,423	24,423	26,823

Notes. The data used are individual reviews posted on Glassdoor by current employees of the focal firm and the firm's industry peers. We only use reviews for individuals with non-missing job title data (necessary to identify engineers). The dependent variable is the overall rating of the firm left by the individual. We estimate ordinary least squares regressions of these ratings on the binary indicator *Pseudo-Post*, which equals one in and after October of 2018, and we omit data after December of 2019 (to perform the falsification tests solely in the pre-COVID period). In some models, we interact the binary indicator *Engineer* with *Post*, and in others we also interact the binary indicator *Focal Firm* into the model. Columns (1)–(2) use only reviews about the focal firm. Columns (3)–(4) use only reviews about the firm's industry peers. Column (5) uses data from all firms. All models include country fixed effects, and the models in Columns (3)–(5) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.9: Engineers' Departure Decisions: Alternative Timing

Panel A: Outflows (After the Second Mentorship Cohort)

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	-0.064***	-0.027***	0.061***	0.059***	0.059***
	(-11.673)	(-10.522)	(16.087)	(19.131)	(19.137)
Engineer		0.954***		0.824***	0.821***
		(34.997)		(70.232)	(70.082)
Post × Engineer		-0.253***		0.014	0.014
		(-7.462)		(0.810)	(0.810)
Post × Engineer × Focal Firm					-0.267***
					(-7.074)
Region and Gender FEs	✓	✓	✓	✓	✓
Ethnicity and Seniority FEs	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.061	0.077	0.032	0.035	0.036
Observations	706,860	706,860	7,999,936	7,999,936	8,706,796

Panel B: Inverse Hyperbolic Sine of Outflows (After the Second Mentorship Cohort)

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	-0.018***	-0.012***	0.007***	0.009***	0.009***
	(-20.608)	(-15.649)	(24.464)	(28.628)	(28.625)
Engineer		0.253***		0.169***	0.169***
		(82.073)		(212.368)	(212.247)
Post × Engineer		-0.044***		-0.006***	-0.006***
		(-11.060)		(-6.008)	(-6.008)
Post × Engineer × Focal Firm					-0.038***
					(-9.021)
Region and Gender FEs	✓	✓	✓	✓	✓
Ethnicity and Seniority FEs	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.206	0.245	0.207	0.223	0.223
Observations	706,860	706,860	7,999,936	7,999,936	8,706,796

Notes. This table presents an analysis of potential changes in employee departure patterns coinciding with the implementation of the patent mentorship program. The data used is employee departure measures from 2018 to 2024 from Revelio, a labor market intelligence firm that tracks employee movements across firms. Panel A analyzes *Outflows*, defined as the number of employees leaving the company within a specific occupation-region-gender-ethnicity-seniority group in a given month. Panel B examines the inverse hyperbolic sine transformation of this external outflows measure. We employ ordinary least squares regressions to analyze these departure measures in relation to the binary indicator *Post*, which equals one in and after January of 2021 (when the second mentorship cohort began). Some models include interactions between *Post* and the binary indicator *Engineer*, while others further incorporate interactions with the binary indicator *Focal Firm*. Columns (1)–(2) use only departure measures from the focal firm. Columns (3)–(4) use only departure measures from the firm's industry peers. Column (5) uses data from all firms. All models include region, gender, ethnicity, and seniority fixed effects, and the models in Columns (3)–(5) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table B.10: Engineers' Departure Decisions: Falsification Tests

Panel A: Outflows (Two-Year Lagged Timing)

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Pseudo-Post	0.010	0.005	0.016***	0.015***	0.015***
	(0.847)	(0.956)	(2.807)	(3.841)	(3.842)
Engineer		0.980***		0.813***	0.812***
		(17.706)		(36.236)	(36.222)
Pseudo-Post × Engineer		0.032		0.003	0.003
		(0.455)		(0.120)	(0.120)
Pseudo-Post × Engineer × Focal Firm					0.029
					(0.378)
Region and Gender FEs	✓	✓	✓	✓	✓
Ethnicity and Seniority FEs	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.064	0.082	0.035	0.040	0.041
Observations	220,320	220,320	2,493,480	2,493,480	2,713,800

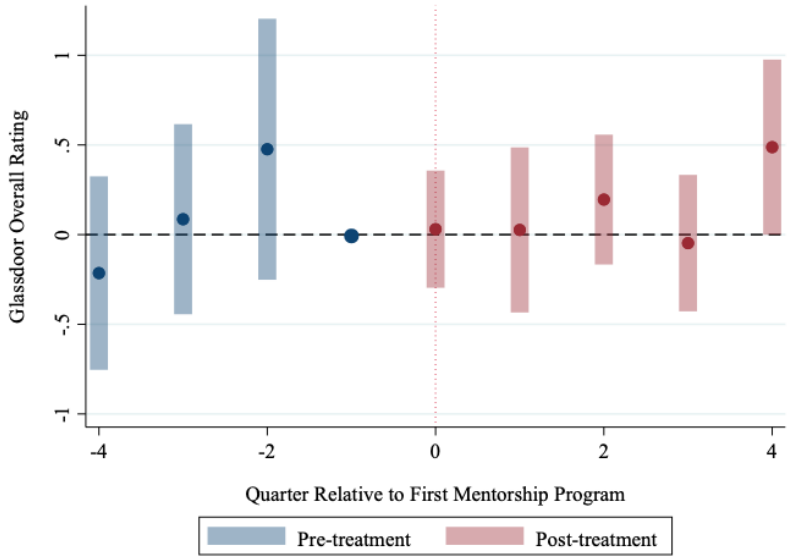
Panel B: Inverse Hyperbolic Sine of Outflows (Two-Year Lagged Timing)

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Pseudo-Post	0.003**	0.003*	0.004***	0.004***	0.004***
	(2.052)	(1.878)	(7.379)	(7.483)	(7.480)
Engineer		0.256***		0.167***	0.167***
		(41.909)		(106.357)	(106.456)
Pseudo-Post × Engineer		0.006		0.001	0.001
		(0.714)		(0.282)	(0.282)
Pseudo-Post × Engineer × Focal Firm					0.005
					(0.609)
Region and Gender FEs	✓	✓	✓	✓	✓
Ethnicity and Seniority FEs	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.215	0.259	0.203	0.220	0.221
Observations	220,320	220,320	2,493,480	2,493,480	2,713,800

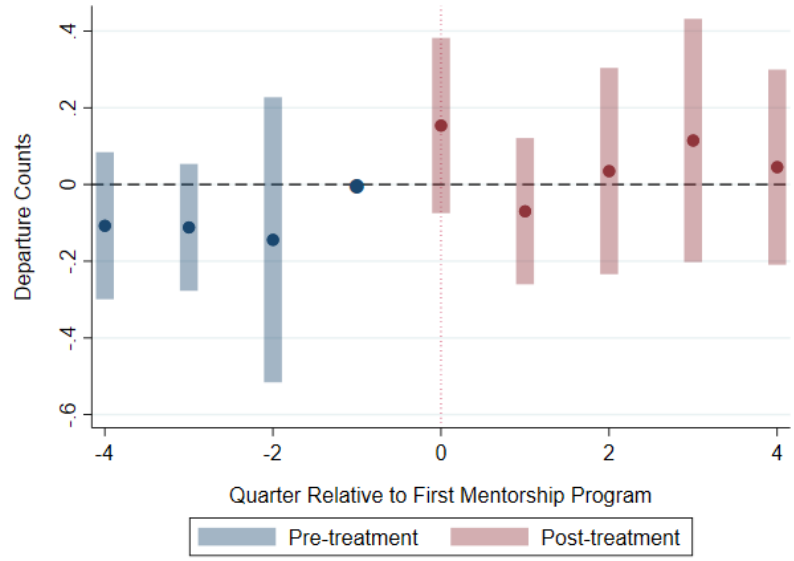
Notes. This table presents an analysis of potential changes in employee departure patterns coinciding with the implementation of the patent mentorship program. The data used is employee departure measures from Revelio, a labor market intelligence firm that tracks employee movements across firms. Panel A analyzes *Outflows*, defined as the number of employees leaving the company within a specific occupation-region-gender-ethnicity-seniority group in a given month. Panel B examines the inverse hyperbolic sine transformation of this external outflows measure. We employ ordinary least squares regressions to analyze these departure measures in relation to the binary indicator *Pseudo-Post*, which equals one in and after October of 2018, and we omit data after December of 2019 (to perform the falsification tests solely in the pre-COVID period). Some models include interactions between *Post* and the binary indicator *Engineer*, while others further incorporate interactions with the binary indicator *Focal Firm*. Columns (1)–(2) use only departure measures from the focal firm. Columns (3)–(4) use only departure measures from the firm's industry peers. Column (5) uses data from all firms. All models include region, gender, ethnicity, and seniority fixed effects, and the models in Columns (3)–(5) include firm fixed effects. t-statistics from robust standard errors are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Figure B.1: Test of Common Pre-Trends in Employee Perceptions and Departures

(a) Employee Perceptions



(b) Employee Departures



Notes. These figures plot the estimates from parallel trends tests. The point estimates capture the quarterly differences in employee perceptions and departure outcomes between engineers and non-engineers in the focal firm. We plot 95% confidence intervals around each estimate, which are generated using the estimator proposed by Callaway and Sant'Anna (2021), with the quarter before the first cohort of the mentorship program as the baseline reference period. The average pre-period coefficient in figure (a) is 0.116 (p-value = 0.615), and the p-values on the individual pre-period coefficients range from 0.200–0.751. The average pre-period coefficient in figure (b) is -0.121 (p-value = 0.319), and the p-values on the individual pre-period coefficients range from 0.185–0.447.

C Internet Appendix: Theory Details

This appendix provides additional details on our theory of mentoring, including the full optimization problem, proofs, and functional form assumptions omitted from the main text. Throughout, we follow the main text in treating the creativity endowment C as fixed by assumption, which operationalizes the innate-ability null as a modeling restriction: any predicted gains from mentoring must operate through the process-knowledge stock h_t, the friction index k_t, or the idea mapping Λ, and not through changes to C.

First, let the engineer (or manager) choose $t_P(t)$ each period to maximize the present value of inventive and core output net of effort costs. Per-period payoff is:

$$U_t = \pi S_t + A_C H t_C(t) - \chi t_P(t), \tag{C.1}$$

where π is the marginal value of a successful disclosure, $A_C H$ is the productivity of core work, and $\chi \geq 0$ is the mental/organizational cost of patenting effort. Using $t_C(t) = 1 - t_P(t)$ and the production equations, we can write

$$U_t = \pi \Lambda(C, h_t)\, q(h_t; k_t)\, t_P(t) + A_C H(1 - t_P(t)) - \chi t_P(t). \tag{C.2}$$

Given (h_t, k_t), the static best response for t_P solves

$$t_P^*(h_t, k_t) = \begin{cases} 1 & \text{if } \pi\Lambda(C, h_t) q(h_t; k_t) \geq A_C H + \chi, \\ 0 & \text{if } \pi\Lambda(C, h_t) q(h_t; k_t) \leq A_C H + \chi, \\ \text{interior} & \text{otherwise.} \end{cases}$$

Thus, the marginal return to patenting is $\pi\Lambda(C, h_t)q(h_t; k_t) - (A_C H + \chi)$. Mentoring shifts this marginal return upward by raising h_t and lowering k_t. Time-shift raises t_P temporarily but leaves (h_t, k_t) unchanged. The innate-ability null holds C fixed, as does our specification throughout; the null is distinguished empirically from the process-learning story not by any different treatment of C in the model, but by the prediction that its restriction on which state variables move should imply little response to mentoring, particularly among engineers with low baseline h_t.

As we stated in the main body of the text, process knowledge evolves as follows:

$$h_{t+1} = (1 - \delta)h_t + \phi t_P(t) + \mu M_t,$$

where $\mu > 0$ reflects direct learning-by-mentoring about the firm's own patenting pipeline. Similarly, frictions, which include both psychological (e.g., lack of confidence, reluctance to

expose an unfinished idea to institutional scrutiny) and organizational (e.g., unfamiliarity with the disclosure system, uncertainty about the review process) barriers, evolve as:

$$k_{t+1} = k_t - \rho\kappa M_t + \eta_t,$$

where $\kappa > 0$ captures the extent to which mentoring reduces frictions, $\rho \in (0, 1]$ governs persistence, and η_t represents slow-moving institutional drift.

We now state the key results formally. Proofs are straightforward and follow directly from the primitives.

Proposition 1 (Quantity). Under mentoring, the number of ideas submitted, I_t, and the number of well-formed disclosures, S_t, increase persistently:

$$\frac{\partial I_{t+s}}{\partial M_t} > 0, \quad \frac{\partial S_{t+s}}{\partial M_t} > 0 \quad \forall s \geq 0.$$

By contrast, under pure time-shift, effects are positive only for $s \in \mathcal{T}$ (the reallocation window), and zero thereafter. Under the innate-ability null, the effect on I_t and S_t is bounded above by the fixed distribution of C and is smallest for engineers with the lowest baseline h_t, because such engineers' output is (under the null) constrained by their creative endowment rather than by process knowledge. *Proof.* Mentoring raises h_{t+1} via $\mu > 0$ and lowers k_{t+1} via $\kappa > 0$, both of which persist into all future periods. Time-shift changes t_P only contemporaneously and does not affect the state variables. The innate-ability null holds C fixed while implying that C is the operative constraint; because inexperienced engineers by construction have low baseline h_t and unknown C, the null predicts these engineers are the least likely to respond to mentoring.

Proposition 2 (Composition of pipeline outcomes). Mentoring shifts the composition of an engineer's disclosure outcomes — across approvals for patent application, defensive publications, and declines — toward the composition observed for experienced inventors, by raising h_t and lowering k_t. Time-shift has no effect on this composition. The innate-ability null predicts no such shift, since composition is determined by h_t under our specification and the null implies h_t is not the operative constraint. *Proof.* Let $F(h_t; k_t)$ denote the outcome distribution over the three pipeline categories for a disclosure produced by an engineer with process knowledge h_t and frictions k_t, and let F^* denote the outcome distribution for experienced inventors. As h_t rises and k_t falls, $F(h_t; k_t)$ converges toward F^* in the sense that the engineer's ability to distinguish among the pipeline's outcome categories improves. Mentoring raises h_t and lowers k_t; time-shift changes neither. We deliberately state Proposition 2 in terms of the composition of outcomes rather than as a claim about an unconditional

per-disclosure approval rate, because the unconditional rate confounds movement along the extensive margin (whether the engineer submits at all) with movement along the conversion margin (whether a submitted disclosure is approved). The composition statement isolates the conversion margin and is well-defined even for engineers whose baseline submission rate is zero.

Proposition 3 (Persistence). Mentoring produces lasting effects on inventive output, since $\mu > 0$ raises the state h_t and $\kappa > 0$ reduces the state k_t, both of which persist as long as $\delta < 1$ and $\rho > 0$. By contrast, time-shift generates no persistence. The innate-ability null makes no direct prediction about persistence conditional on an initial effect, but its restriction that initial effects are small implies that persistent effects are also small. *Proof.* Follows from the law of motion for h_t and k_t.

Proposition 4 (Quality, variety, and potential spillovers). Mentoring expands the effective idea mapping $\Lambda(C, h_t)$ via $\Delta_\Lambda > 0$, increasing the quality and variety of ideas an engineer can articulate for a given creativity endowment. These effects can spill over to collaborators if templates or other forms of tacit knowledge that help reduce organizational and psychological frictions diffuse among collaborators. Time-shift does not affect $\Lambda(C, h_t)$, so quality and spillover effects are zero. The innate-ability null holds C fixed as our model does, but attributes observed inventive output primarily to C; under the null, expanding Λ through mentoring should have only a limited effect on output because C is presumed to be the binding factor. *Proof.* Δ_Λ shifts Λ upward for $M_t = 1$, raising average quality. Spillovers follow if reductions in k and increases in h are shared. Time-shift leaves Λ unchanged.

Proposition 5 (Heterogeneity by baseline experience). The effect of mentoring on inventive output is decreasing in baseline process knowledge h_0:

$$\frac{\partial^2 I_{t+s}}{\partial M_t \, \partial h_0} < 0.$$

That is, engineers with the lowest baseline process knowledge — those without prior patenting experience — exhibit the largest response to mentoring. This prediction sharply distinguishes the process-learning story from the innate-ability null, which predicts the opposite pattern: under the null, engineers with no prior patenting record are precisely those whose creative endowment C is expected to be low (as revealed by their prior non-participation), so they should show the smallest response to any intervention that leaves C unchanged. *Proof.* Λ and q exhibit diminishing returns in h_t, so the marginal return to a given increment μ is largest where h_t is smallest. The prediction under the null follows from the null's implicit selection interpretation: engineers who have not previously submitted disclosures are selected

on low C under the null, and C is fixed by assumption.

To summarize the results shown here, we include Table C.1 below, which describes the comparative statics between the time-shift mechanism, the process-learning mechanism through mentoring, and the innate-ability null. As the table shows, mentoring alters the state variables h_t and k_t, yielding persistence, a compositional shift in pipeline outcomes, and potential spillovers, with the largest effects among engineers with low baseline h_0. The time-shift mechanism only reallocates effort but leaves the underlying states unchanged, and as a result produces only temporary gains in disclosure quantity with no compositional shift and possible crowd-out of core output $Y_{C,t}$. The innate-ability null, which fixes C as our model does but treats C as the operative constraint, predicts small effects overall and the smallest effects among engineers with no prior patenting record.

Table C.1 summarizes these predictions.

Table C.1: Comparative Predictions: Mentoring (Process Learning), Time-Shift, and Innate-Ability Null

Outcome	**Mentoring (Process Learning)**	**Pure Time-Shift**	**Innate-Ability Null**
Quantity of disclosure attempts (I_t)	Persistent increase ($h_t \uparrow$, $\Lambda \uparrow$).	Temporary increase only while t_P is elevated.	Small increase at most; bounded above by fixed distribution of C.
Well-formed disclosures (S_t)	Increase due to higher I_t and pipeline composition shifting toward that of experienced inventors.	Temporary increase only while t_P is elevated.	Small; constrained by fixed C.
Composition of pipeline outcomes	Shifts toward the composition observed for experienced inventors as h_t rises and k_t falls.	No change: composition unaffected.	No change: composition is a function of h_t, which the null implies is not the operative constraint.
Persistence over time	Yes: higher h_t and lower k_t endure after program ends.	No: effects vanish once t_P reverts.	Small if any: bounded by small initial effect.
Quality and variety of ideas	Higher: exposure and networks expand $\Lambda(C, h_t)$.	No change: same creativity and know-how applied to more hours.	Little movement: C is fixed and treated as binding.
Heterogeneity by baseline experience	Largest effects for engineers with low baseline h_0 (no prior patenting record).	No heterogeneity by baseline experience.	Smallest effects for engineers with no prior patenting record, since such engineers are selected on low C under the null.
Spillovers to collaborators	Possible: reduced frictions and tacit know-how diffuse through teams.	No: no transferable knowledge created.	No: intervention leaves C unchanged and C is the operative constraint.

Notes. This table summarizes the comparative statics for three mechanisms: the process-learning mechanism through mentoring, the pure time-shift mechanism, and the innate-ability null. Column (1) lists the empirical

D Internet Appendix: Mentorship Survey Questions

1. Have you in your time at [Company] ever acted as a patent mentor to a less experienced employee?

   - Yes
   - No

2. Have you in your time at [Company] ever received patent mentorship from a more experienced employee?

   - Yes
   - No

   If the answer to mentorship is no, ask the following two questions.

3. If the answer to mentorship is no, what are the reasons you have not sought patent mentorship? (check all that apply)

   - I could not identify a potential mentor with overlapping interest
   - I did not believe patent mentorship would benefit me
   - I did not have time for patent mentorship or it was offered at inconvenient times
   - I want to avoid activities that highlight my weakness to leaders
   - I felt that the program would benefit my employer more than it would benefit me
   - I did not think about it
   - My co-workers already help me with IP so I don't need a mentor
   - I want to participate in patent mentorship but was unable to

4. If you were to be paired with a mentor, what type of affinity would you like to share? (Check all that apply)

   - Same gender
   - Same age
   - Same ethnicity
   - Same university
   - Overlapping affinity does not matter

   If the answer to mentorship is yes, ask the following questions.

5. For the next set of questions, think of your most recent patent mentorship relationship where you were the protégé (i.e., the person receiving mentorship).

   Slider with range-2 = Strongly disagree, 0 = Neutral, 2 = Strongly agree

   - I have learned about patenting and become a better inventor as a result of being part of the program
   - Patent mentorship increased my satisfaction at work
   - I am planning to work on more invention disclosures at work
   - I feel confident that I can incorporate my mentor's tips into my work process and develop new inventions
   - I feel confident that I have the social connections necessary to get my invention ideas accepted by IP professionals

   How often would your patent mentor provide you with advice, suggestions, or support?

   - Daily
   - Several times a week
   - Weekly
   - Several times a month
   - Monthly

6. Who first initiated mentor-protégé like contact?

   - Your mentor by offering unsolicited advice, suggestions, or support.
   - You by asking the person for advice, suggestions, or support.
   - Other (please specify)

7. Did you and your mentor have an affinity? (Check all that apply)

   - Yes, we are the same gender
   - Yes, we are the same age
   - Yes, we are the same ethnicity
   - Yes, we went to the same university
   - No, we do not have an overlapping affinity

If [Company] has a formal mentorship program, ask the following question.

8. Since your formal mentorship relationship ended, on average how many times per month have you met with your mentor/mentee outside of your regularly scheduled work meetings, even if just briefly?

    - 0 times
    - 1-2 times
    - 3-4 times
    - 5 or more times

E Internet Appendix: Text-based Measures of Culture

This appendix provides definitions of cultural values and norms and discusses the procedure for generating the text-based measures of corporate culture. The seven cultural values represent the principal components of culture derived by O'Reilly et al. (1991), and they are supplemented to include community-oriented, which was added to reflect changes in cultural values in the last thirty years. The seven cultural values are adaptability, collaboration, community-oriented, customer-oriented, detail-oriented, integrity, and results-oriented. The same cultural values used in Graham et al. (2022a) were benchmarked to executives' descriptions of their culture and website-based values. The terms pillars and cultural values are used interchangeably. Each cultural value or pillar has norms or sub-pillars described below. These norms are elements of the culture that are subcomponents of the main cultural values. Note this is a shortened word list meant to help familiarize the reader. Our approach did not use a word list or dictionary-based method. Instead, we used a Bidirectional Encoder Representations from Transformers (BERT)-based model.

The dataset consists of Glassdoor reviews for a given organization. The reviews have been parsed into multiple, non-overlapping sentences. Each sentence is associated with one of the four open text areas where they could have been entered. The four open text areas include title, pros, cons, and advice. We use a semi-supervised machine learning (ML) approach, which means for a subsample of the sentences, they have been annotated by three reviewers with their actual cultural values and norms. Specifically, each sentence in this training data is annotated for three classes +1, 0, and -1, for 7 different pillars and 34 different sub-pillars. Glassdoor review sentences tend to be highly imbalanced, with most of the data belonging to the 0 class and more data belonging to the +1 class than to the -1 class. Therefore, we artificially increased the number of samples in the minority classes through oversampling to balance the dataset. The value for each pillar is directly related to the subpillars, as such we only run our classifier models on the subpillars and aggregate to the pillar level using the predicted subpillars. The model used for sentence classification is a hierarchical BERT-based model that includes sentence type information with oversampling and a trained separate multiclassifier (-1, 0, +1) for each subpillar. A weighted cross-entropy loss function was applied, further addressing class imbalance by giving higher weights to minority classes.

We experimented with other variations, and selected the model that performed the best across the 34 subpillars. Our experiments with sentence type showed that incorporating this information improved the model performance. We also experimented with different transformer-based language models, including RoBERTa and Sentence Transformers, which are optimized for generating sentence embeddings. The other transformer models did not yield meaningful performance improvements in precision, recall, or F-1 scores, so we stuck

with the BERT model. Finally, we also considered a k-shot classifier based on exemplary sentences and two-stage predictions where presence (+1/-1) vs. absence (0) and positive (+1) vs. negative (-1) sentiment were classified. While an individual subpillar may outperform using these alternatives, we found that the average F-1 score across all 34 subpillars was better with the multiclassifier model.

To aggregate from the predicted subpillars to the pillars, we sum the total predicted scores across all subpillars. We then aggregate the sentence-level scores to review-level scores by weighting sentences according to their word count. This makes the implicit assumption that more detailed explanations are more valuable. The main inference reported in Panel A of Table 9 about community-orientation improving in a statistically and economically significant way remains robust across alternative aggregation techniques, including taking a simple mean rather than a sentence-weighted one or taking the mean across subpillars rather than the sum.

Cultural Value 1 (i.e., Pillar 1) Adaptability

- Subpillar 1 Adaptability
  - Words: adaptable, turnaround, change, dynamic, upcoming, flexible, spin-off
  - Phrases: changing environment, constantly trying to improve, weathered downturn, plan for shifting, willing to change, constant evaluation of advancement, follow latest trends
- Subpillar 2 Fast Moving
  - Words: fast, rapid, quick, new
  - Phrases: fast-paced, constantly changing, always something new, growth mindset
- Subpillar 3 Takes initiative
  - Words: initiative, entrepreneurial, proactive, independent, self-motivated, self-starter,
  - Phrases: new projects, forward-thinking, seizes opportunities, takes the lead, makes hard decisions, takes charge, takes a stance, drives action, goes the extra mile
- Subpillar 4 Autonomy vs. Bureaucracy
  - Words: autonomy, independence, empowerment, access

  - Phrases: individual contributions, no micromanagement, opportunity to contribute, value ideas, removing bureaucracies, ability to interact with leaders,

- Subpillar 5 Innovative
  - Words: technology, innovative, inventive, creative
  - Phrases: new solutions, technology leader, deep technology, exposed to lots of technology

Cultural Value 2 (i.e., Pillar 2) Collaboration

- Subpillar 6 Collaborative
  - Words: collaboration, cooperation, coordination, partnership, complementary, aligned objectives, similar goals,
  - Phrases: joint effort, collective effort, mutually beneficial, synergistic, cross-functional, strategic alliance, in partnership with
- Subpillar 7 Team-oriented
  - Words: teamwork, teammates, united or unified approach, good team chemistry, well-integrated,
  - Phrases: team-oriented, work together, rely on each other, team player, come together, working hand in hand
- Subpillar 8 Talented Individuals
  - Words (people): talented, impressive, smart, capable, professional, hardworking, great, intelligent
  - Phrases: terrific people, good people to work with
- Subpillar 9 Supportive (family-like)
  - Words: friendly, supportive, family, generous, helpful, respectful, sharing
  - Phrases: everyone is nice, I've made friends, the best people, friendly people, everyone shares expertise and resources
- Subpillar 10 Not aggressive
  - Words: happy, fun, open-minded

  - Phrases: take their time with you, help you understand, willing to listen

- Subpillar 11 Low conflict
  - Words: nice, kind, welcoming, caring, non-confrontational
  - Phrases: work well together, easy to ex press ideas, optimistic, nothing to fear, no harsh criticism

Cultural Value 3 (i.e., Pillar 3) Community-oriented

- Subpillar 12 Inclusive / embraces diversity
  - Words: diverse, respectful, equality, inclusive, empowering, DEI
  - Phrases: people from all over, cross-cultural environment, many different backgrounds

- Subpillar 13 Management
  - Words: management, leadership,
  - Phrases: good / great managers, supportive managers

- Subpillar 14 Personal growth
  - Words: challenging, learn, grow, skills, teach, develop, fulfilling, opportunities
  - Phrases: push you, personal growth, learning experience, challenging problems

- Subpillar 15 Career advancement
  - Words: opportunity, advance, promoted, climb, career
  - Phrases: good place to start a career, work your way up, opportunities available if you want them, hire from within, potential for growth, great resume addition

- Subpillar 16 Longevity vs. high turnover
  - Words: stable, stay, hiring, layoff, turnover, fired, unpredictable, steady
  - Phrases: People have been here forever, low turnover, don't know what to expect, cost-cutting, don't worry about your job

- Subpillar 17 Good work-life balance
  - Words: work-life, flexible, hours, vacation, no overtime

  - Phrases: good work-life balance, flexible environment, laid-back, good hours, no time

- Subpillar 18 Work environment
  - Words: playful, fun, casual, relaxed, games, prizes, free food, gym, perks, dog-friendly,
  - Phrases (environment/office/atmosphere): good, great, lots of amenities, beautiful campus, lots of comforts, plenty of chances to socialize, designed to inspire

- Subpillar 19 Sustainability
  - Words: philanthropy, environmentally conscious, sustainable, impact, donating
  - Phrases: committed to social responsibilities, community service, work with charities

Cultural value 4 (i.e., Pillar 4) Customer-oriented

- Subpillar 20 Customer-oriented
  - Words: customers, clients, service, satisfaction, exceed expectations,
  - Phrases: customer is always right, great sales job, customer service, good customer experience, help people live better lives, tailored to customer needs, great products, competitive pricing, affordably priced, easy to sell

- Subpillar 21 Listens to and satisfies customers
  - Words: listen, satisfy, fulfilling, customer-first, customer-centric
  - Phrases: attuned to consumer desires, adapts to customer preferences, prioritizes the customer, consistently high ratings, cold calling transparent with clients

- Subpillar 22 Brand & market driven
  - Words: brand, competitive, market, pride, leader, trendsetter, recognized,
  - Phrases: great vision, great brand, building the future, broad portfolio, top company

Cultural value 5 (i.e., Pillar 5) Detail-oriented

- Subpillar 23 Detail-oriented

  - Words: Focused, detail, careful, meticulous, exact, rigorous, standards, flawless, perfection, benchmark
  - Phrases: detail-oriented, put care into efforts, masterful execution

- Subpillar 24 Analytical
  - Words: analyze, consider, interesting, difficult, math, data-driven, predictive, forecasting
  - Phrases: analytical work, precision-engineered, improving processes, sophisticated tools

- Subpillar 25 Safety
  - Words: safety, careful, risk-averse, protocols, no accidents, no injuries
  - Phrases: safety is number 1, safety commitment, safer together

- Subpillar 26 High quality
  - Words: quality, cutting edge, care, superior, craftsmanship, standards
  - Phrases: high quality, no compromises, nothing overlooked, elevated standards

Cultural value 6 (i.e., Pillar 6) Integrity

- Subpillar 27 Integrity
  - Words: integrity, fair
  - Phrases: high integrity, not questionable

- Subpillar 28 High ethical standards
  - Words: ethical, respect, reasonable,
  - Phrases: high ethical standards, respect opinions, do the right thing, fair hiring / promoting

- Subpillar 29 Honest (trustworthy)
  - Words: honest, trust, open, predictable
  - Phrases: open management, open communication, free to share opinion

- Subpillar 30 Transparent / non-political

  - Words: transparent, not political, no disconnect, predictable, no favoritism
  - Phrases: playing politics, you need to drink to get promoted, only hire their favorites, self-centered, everyone is on the same page

Cultural value 7 (i.e., Pillar 7) Results-oriented

- Subpillar 31 Results-oriented
  - Words: results, revenue, top-line, growth, sales, bottom line, earnings, goals, targets, achievement, driven, performance-focused, metrics-driven, goal-oriented
  - Phrases: always focused on generating profits, get work done, increasing top line, work hard / play hard, number 1, always the best, continuously striving,
- Subpillar 32 Competitive
  - Words: work, competitive, challenging, outperform, high-stakes, merit-based, tenacious, ambitious
  - Phrases: few spots for many employees, highly competitive environment, value achievement, outperforming the competition, opportunities to excel
- Subpillar 33 Long hours
  - Words: long hours, hours
  - Phrases: no time for rest, over X hours per week, always need to be available, no time for family, always working, terrible hours, work up to Xper week
- Subpillar 34 Accountable
  - Words: accountable, consistency, diligent, responsible, reliable,
  - Phrases: rewarded for hard work, every action counts, acknowledge accomplishments, follow through, meets commitments, pull your weight, everyone contributes, clear expectations, constructive accountability, feedback loops,