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

Colleen Chien* Jillian Grennan† Jason Sandvik‡

Abstract

We use the staggered rollout of a mentorship program within a superstar technology firm to test whether a small-scale, targeted intervention can reduce frictions limiting latent innovation potential. Analyzing novel data capturing the complete innovation pipeline—from initial idea disclosures through patent applications—we study seven mentoring cohorts with 633 mentees and track outcomes for over 20,000 employees. First, we find a positive and significant relationship between participating in the mentorship program and innovative productivity in the short- and long-term, with underrepresented innovators realizing the largest gains. Second, we document significant indirect spillovers through team formation, with non-mentored peers experiencing productivity gains through collaboration. Survey evidence points to three mechanisms: knowledge transfer about patent processes, expanded professional networks, and increased satisfaction at work. Additional analyses show broader organizational improvements: engineers perceive stronger cultural norms, and attrition is lower than at peer firms. Accounting for these spillovers indicates a substantial return on investment, so we conclude by evaluating its scaling potential.

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

Keywords: Intellectual Property, Innovation, Patents, Diversity, Labor and Finance, Mentorship, Human Capital Investment, Intangible Assets, Corporate Culture, Employee Turnover

We thank Daniel Bias, Joan Farre-Mensa, Filippo Mezzanotti, Alex Oettl, Ben Pyle, Ting Xu, and seminar participants at Emory University, Georgia Institute of Technology, St. Andrews University, NBER Corporate Finance, SFS Cavalcade, the Conference on Empirical Legal Studies, the Innovator Diversity Pilots Initiative Conference, Organization Science Winter Conference, and the University of Illinois Chicago Finance Conference for helpful discussions and comments on an earlier draft of this study. We thank the Ciocca Center for Innovation and Entrepreneurship at Santa Clara University for its support. This research was approved by the Institutional Review Board (IRB) at the University of California Berkeley.

*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

1 Introduction

Superstar technology firms account for a disproportionate share of U.S. R&D investment, making their innovation efficiency a central determinant of aggregate productivity growth (Autor et al., 2020; Kalyani et al., 2025). As these firms mature, they typically transform from agile innovators that produce breakthrough inventions into bureaucratic organizations that produce more incremental innovations (Akcigit and Kerr, 2018). In addition, as firms grow organizational frictions can impede knowledge flows and reduce R&D returns (Sandvik et al., 2020; Akcigit and Ates, 2021).[1] The resulting decline in innovative capacity leads some firms to resort to anti-competitive strategies such as killer acquisitions (Cunningham et al., 2021) or strategic patent blocking (Abrams et al., 2023), further reducing aggregate innovation output.

What initiatives and interventions might help superstar firms maintain innovation efficiency at scale? Our analysis focuses on mentorship within a S&P 500 firm as a potentially scalable intervention (List, 2022). Survey evidence from engineers identifies patent mentoring as a high-priority need (Chien and Grennan, 2024), particularly for professionals from underrepresented groups (URGs) who face higher barriers to patenting (Cook, 2018; Bell et al., 2019).[2] This demand for mentorship aligns with established evidence on human capital formation (Hoffman and Stanton, 2024). A significant portion of skills development occurs within firms through internal training programs and on-the-job learning (Becker, 1965; Acemoglu and Pischke, 1998),[3] and mentorship represents a natural extension of firms' existing human capital investments (Sandvik et al., 2025).

We test whether a targeted patent mentoring program can increase innovation produc-

[1] 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 flow.

[2] As one survey respondent noted, "I sometimes hesitate to put my ideas forward. Writing a patent disclosure is intimidating, and I don't know what makes something patentable. Most inventors in our company are men, and I don't have close connections to them. A mentorship program would be great because it would connect me to someone who's been through the process and could validate my ideas are worth pursuing."

[3] U.S. firms invested over $100 billion on training in 2023 (Statista, 2024).

tivity while simultaneously addressing persistent disparities in inclusive patent production. These disparities, commonly referred to as the "innovator-inventor gap" (the systematic underrepresentation of certain demographic groups among patent inventors relative to their presence in STEM professions), are substantial[4], and evidence suggests they stem from impediments to participation in the innovation process (Budish et al., 2015; Koning et al., 2021; Heath et al., 2023; Koffi and Marx, 2023), indicating substantial untapped potential.

By studying a mentorship program designed to improve innovation efficiency and inclusion, we contribute to debates on U.S. technological competitiveness and how to address the innovator-inventor gap (Bloom et al., 2019b; Chien, 2024). The program's goal was to transform non-patenting STEM workers ("innovators") into active patent contributors ("inventors") by enhancing their exposure to and understanding of the invention process. By analyzing multiple channels through which the program operates, we provide detailed evidence on the returns to human capital investments in an R&D-intensive setting, overcoming measurement challenges that have historically complicated the evaluation of training programs (Black and Lynch, 1996; Ichniowski and Shaw, 2013).

The mentorship program had a staggered rollout and involved 633 innovators participating in seven cohorts, with the participants being selected from a set of applicant innovators. The program was open to all of the firm's employees, but those from underrepresented groups (e.g., women and racial minorities) were especially encouraged to apply. Applicants not selected initially, due to caps on the size of each cohort, were encouraged to re-apply to subsequent cohorts. The modal participant was a 33-year-old Asian female engineer who had worked in the tech industry for 5 years and at this firm for 2 years. The program's curriculum paired experienced inventors with proven track records of patenting success (i.e., mentors) with small groups of mentees. The typical group had 5 mentees and 1–2 mentors who met regularly over a calendar quarter. The mentors guided their mentees through the

[4]For instance, while Black workers comprise 9% of STEM professionals, they represent only 1.2% of patent inventors; similarly, women constitute 35% of STEM workers but only 13% of inventors (NCSES, 2023; Akcigit and Goldschlag, 2022).

inventive idea development and patenting process.

We surveyed participants and non-participants to understand the potential for bias due to self-selection into the program. Despite women being more likely than men to participate in the mentorship program, they were also more skeptical about its benefits. This reduces concerns about selection effects, especially among women (who made up most of the mentees), as many participants opted into the program despite being uncertain about its likely efficacy. Similarly, Sandvik et al. (2025) document a negative relation between selection into mentorship and gains from mentorship in a call center setting, so self-selection likely biases us against observing returns to mentorship.

To estimate the relation between mentorship and innovation outcomes, we build on recent methodological advances in the evaluation of staggered treatment designs (Goodman-Bacon, 2021; Baker et al., 2025), while addressing specific challenges in measuring innovation outcomes (Lerner and Seru, 2022). Our primary specifications include a two-way fixed effects (TWFE) estimator and a stacked difference-in-differences (DID) estimator. The latter addresses potential biases arising from staggered treatment timing and evolving treatment effects (Baker et al., 2022; Roth et al., 2023).[5] We construct our stacked DID estimator using two control groups: never-mentored innovators and not-yet-mentored innovators. This approach allows us to exploit variation in the timing of program participation and to control for selection on unobservables. We examine the parallel trends assumption underlying our DID approach (Sun and Shapiro, 2022; Rambachan and Roth, 2023). These tests suggest no diverging pre-trends in inventive idea submissions between mentored and non-mentored innovators, supporting the validity of our empirical strategy.

Our analysis leverages unique access to proprietary data from the focal firm, including the complete universe of 66,397 inventive ideas submitted by 29,211 innovators, regardless of whether the ideas even become patent applications. This comprehensive innovation pipeline data linked to detailed demographic characteristics and mentorship participation across mul-

[5]Our inferences are also robust to the use of the estimator proposed by Callaway and Sant'Anna (2021).

tiple program cohorts enables us to estimate the relation between the mentoring program and innovation productivity. By observing pre-patent idea submissions, including those rejected during internal review, we overcome a fundamental limitation in existing studies that capture only successful patent applications or grants (Kogan et al., 2017; Lerner and Seru, 2022; Mkrtchyan et al., 2024), and this allows us to speak to our broader understanding of what drives ideation and innovation within firms (Farre-Mensa et al., 2020; Bhaskarabhatla et al., 2021).

Our analyses yield four main results. First, mentorship significantly increases innovative output in the short- and long-run. Using either estimator, we find that mentored innovators produce approximately one additional invention annually, which is statistically significant at the 1% level. This finding is robust to accounting for potential Stable Unit Treatment Value Assumption (SUTVA) violations by explicitly removing participants' future non-mentored collaborators from the analysis, restricting the control group to only those who eventually get mentored, focusing exclusively on high-quality disclosures that progress to approved patent applications, removing any disclosures submitted with the mentor as a collaborator, and clustering the standard errors by mentor rather than by innovator. In addition, when normalizing submissions by team size, the positive effect remains apparent and statistically significant, indicating that mentorship genuinely increases innovative productivity, rather than simply teaching participants to join existing teams. We conduct additional tests to further show that these productivity gains are not driven by inflated team sizes, helping to rule out concerns that the firm may be engaging in "diversity-washing."[6]

Our analysis of long-term outcomes reveals that mentorship delivers sustained benefits beyond the program period itself. When we exclude submissions made during and immediately after mentorship, we find mentored innovators continue producing approximately 0.33 additional disclosures per year. The combination of the short- and long-run productivity

[6]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.

gains from mentorship is economically meaningful, given a baseline disclosure rate of 0.5 disclosures per year. Our estimated productivity gains from mentorship are 2.7 times larger than the average productivity gains from high-performing coworker exposure documented in the meta-analysis by Herbst and Mas (2015), suggesting mentorship creates stronger knowledge diffusion than typical workplace peer effects.

Heterogeneity tests for mentor and mentee characteristics suggest more productivity gains for female and underrepresented racial minority (URM) innovators (submitting 1.33 vs. 1.00 additional disclosures annually).[7] Importantly, the baseline relation for mentorship remains statistically significant and positive in these heterogeneity tests, suggesting male and non-URM innovators still see productivity gains from mentoring. We see little evidence of heterogeneity across geography, age, or cohort for mentees, indicating the program's scalability. On the mentor heterogeneity side, we see that being paired with a high-quality mentor doubles the mentee's productivity gains.[8] Importantly, though, our analysis of mentor quality reveals that even mentees with below-median quality mentors experience significant productivity gains, again suggesting the program is scalable.

Second, we present survey evidence, which reveals several mechanisms underlying the innovation productivity gains: knowledge transfer about patent processes, expanded professional networks, increased satisfaction at work, and to a lesser extent, greater confidence in identifying patentable innovations. To validate the expanded professional networks mechanism, we use our unique dataset to track the composition of new teams of collaborating innovators. We find that mentored innovators form more diverse inventor teams after participating in the mentorship program, with diversity measured relative to pre-program team composition baselines. This diversity manifests across many dimensions, including gender, race, and previous patenting success. This improved team diversity is particularly important

[7] We use the race/ethnicity data to identify URM innovators as those who are neither Asian nor White, as Asian and White individuals are highly represented among the innovators in this firm and its peer firms (Cook, 2018).

[8] We measure mentor quality through fixed effects estimated from mentee performance residuals after controlling for demographic characteristics, time trends, cohort effects, and location, allowing us to isolate the mentor-specific contribution to mentee innovation outcomes.

given evidence linking diverse teams to higher-value patents and the importance of overcoming the "homophily trap" (Bernstein et al., 2022; Heath et al., 2023).

Third, our unique dataset allows us to track the submission patterns for non-mentored innovators and account for additional spillovers. We restrict our sample to innovators who are never directly associated with the mentorship program and construct a measure of mentor exposure using the composition of the innovator team. We find evidence of positive spillovers in innovative productivity among non-mentored innovators who collaborate with mentored innovators. These spillovers appear to be larger for individuals from well-represented demographic groups (e.g., white men), suggesting that organic knowledge flows may spread to certain types of employees more readily than others. Finding a positive indirect influence on non-mentored innovators is particularly useful for estimating ROI, because it indicates these programs increase innovativeness within the firm, rather than merely shifting the focus from one group of innovators to another. This points to the scalability of mentorship programs within large organizations.

Fourth, we validate the increased workplace satisfaction mechanism and consider broader workplace spillovers by documenting a positive relation between mentorship and engineers' perceptions of the corporate culture as well as their retention.[9] 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 patterns represent substantial savings given replacement costs for technical talent (Benson et al., 2019) and the role of culture in driving long-term value creation (Guiso et al., 2015; Graham et al., 2022a). Additional tests using employees' "Culture & values" star rating and culture measures derived from textual analysis of open-ended responses show the cultural improvements at the focal firm are more pronounced among engineers and driven by the perceived strengthening of norms around community orientation, collaboration, and

[9]We use data from Glassdoor—a career intelligence website that provides transparency about jobs, salaries, and company reviews—and Revelio—a labor market intelligence platform that tracks employee movements across firms—to examine changes in workplace perceptions and retention patterns.

integrity.

We use these four main findings to update traditional estimates of mentoring's ROI to account for the multiple channels through which mentorship creates positive returns. We use a range of patent valuation metrics, and even our most conservative estimate suggests direct benefits of approximately $5 million. Other indirect innovation benefits from higher-quality patents and spillovers to peers increase this by 20%. Finally, adding the indirect benefits from improved retention and corporate culture, though more challenging to quantify precisely, substantially increases the return. Since the mentoring program had modest implementation costs—primarily staff time for coordination—this indicates the program's ROI is high. We conclude by discussing the potential for scaling the program, given the growing recognition among experimentalists of its importance (Al-Ubaydli et al., 2020; List, 2022).

Finally, the positive relation between mentorship and productivity that we observe raises the question of why high-performing firms have not widely adopted such seemingly valuable programs. Our interviews with the firm's leaders suggest managerial frictions within the organization impeded adoption, which is consistent with recent empirical work showing that behavioral and cultural barriers often prevent managers from implementing efficiency-enhancing practices (Mishra et al., 2022; Gertler et al., 2025). Moreover, Graham et al. (2022a) document that corporate leaders frequently perceive a lack of investor support for such human capital and culture investments, even when they believe such programs would enhance firm value. Our findings suggest that, despite organizational barriers that may have historically limited adoption, structured mentoring programs represent a cost-effective intervention for improving innovation outcomes at large technology firms.

2 Related Literature and Contributions

This interdisciplinary study integrates personnel, organizational, and innovation economics insights to make several contributions. First, we contribute to the personnel economics literature by quantifying how organizational practices beyond compensation enhance innovative

productivity (Lazear and Shaw, 2007; Bloom et al., 2019c). While prior research shows innovation skills can develop over time (Azoulay et al., 2010) and respond to environmental factors (Coffman, 2014; Bol et al., 2023), our study is the first to systematically measure how mentorship transforms potential into concrete inventive outputs within a corporate setting.

The economic literature on corporate mentorship remains surprisingly limited given that mentorship is a ubiquitous practice (Athey et al., 2000; Sandvik et al., 2025). We bridge this gap by providing evidence from an intervention in a superstar firm, complementing findings from education (Ginther et al., 2020; Adams and Lowry, 2022) and military contexts (Lyle and Smith, 2014). Importantly, our results demonstrate that mentorship outperforms traditional training alone for innovation outcomes, aligning with microentrepreneurship research (Brooks et al., 2018) and studies showing human-AI complementarity (Tong et al., 2021; Otis et al., 2024). Similar to executive coaching studies that document performance gains through expanded networks (Cai and Szeidl, 2018), we document how mentorship facilitates connections, even in traditionally male-dominated settings, that reduce gender disparities (Cullen and Perez-Truglia, 2023; Benson et al., 2024). Unlike previous work focused on middle managers as teachers (Friebel et al., 2022; Hoffman and Tadelis, 2021), we show that mentorship allows peers to teach one another, which may fill critical gaps when employees have weak managers (Lazear et al., 2015; Benson et al., 2019).

A novel finding in our study is that mentoring creates organizational synergies beyond individual outcomes—i.e., the aggregate impact exceeds the sum of individual gains. These spillovers to corporate culture amplify productivity throughout the organization, supporting research on culture as a driver of long-term value creation (Guiso et al., 2015; Li et al., 2020; Graham et al., 2022b, 2023; Gorton et al., 2022; Grennan and Li, 2023) and management practices as determinants of productivity differences (Bertrand and Schoar, 2003; Bloom et al., 2016, 2019a).

We also provide new insights into innovation economics by analyzing the complete pre-patent pipeline, including rejected ideas (Serrano, 2010; Hegde and Sampat, 2009; Pairolero

et al., 2022; Aneja et al., 2024). Our unique contribution lies in linking comprehensive idea-level data to personnel records—the first study to do so. While others have examined academic idea generation (Fleming et al., 2019; Masclans et al., 2024), we examine commercial innovation screening processes in a superstar technology firm, demonstrating that patenting is process-driven and responsive to targeted interventions. This new understanding helps explain differences in inventor matching across contexts (Bhaskarabhatla et al., 2021; Gupta et al., 2024).

Finally, we advance the literature on team-based innovation by addressing a key friction in knowledge production. As technological advances increasingly require larger teams with diverse expertise (Wuchty et al., 2007; Jones, 2009), our mentoring intervention reduces collaboration barriers by facilitating knowledge transfer between experienced and novice innovators. Our empirical evidence on innovation team formation provides concrete mechanisms to enhance team productivity, connecting to broader work on team complementarities (Hamilton et al., 2003) and the critical role of diversity in strengthening innovative output (Hong and Page, 2004; Heath et al., 2023).

3 Firm Setting

This study examines a mentorship program at a publicly-traded, S&P 500 technology firm that develops high-technology services and products, including cloud management, data storage, networking, security software, and telecommunications equipment. With over 50,000 employees globally (primarily engineers and software developers) across more than 100 offices worldwide, this R&D-intensive firm implemented a structured mentorship program to enhance innovation capabilities. The mentorship program had 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." While open to all employees, the program particularly encouraged participation from underrepresented groups (URGs), including women and racial/ethnic minorities.

The mentorship program occurred in seven cohorts from Q4 of 2020 to Q3 of 2023, evolving from two-month to four-month durations, while maintaining a consistent curriculum. Participation required managerial approval and a commitment of several hours weekly. Cohort sizes were primarily constrained by mentor availability, with additional limitations from time zone compatibility and shared invention interests in areas such as machine learning, AI applications, and cloud services.[10]

The program structure grouped approximately four to five mentees with one or two mentors, who had ample experience developing and patenting their own ideas. Figure A.1 illustrates the typical timeline, with teams beginning R&D solutions about a month after team formation, drafting disclosures at the two-month mark, and with submissions occurring shortly after the three-month mark. Each team's invention submission culminated in a showcase event where mentees presented their work to senior technical talent teams.

While teams typically produced a disclosure submission within four months, the program's primary purpose was developing participants' independent innovation capabilities outside the program setting. In subsequent sections, we quantify the direct productivity gains and indirect spillovers to non-participants throughout the organization.

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,

[10]It is our understanding that receiving such approval was routine, and in the rare case that there was a product launch or another competing event, the applicant was encouraged to apply for the next cohort and was reminded to do so. The start dates of each cohort of the program were as follows: 2020-Q4, 2021-Q1, 2021-Q3, 2021-Q4, 2022-Q3, 2023-Q1, and 2023-Q3.

marital status, race/ethnicity, job role, and the country where they work. We use the race/ethnicity data to identify URM innovators and define them as those who are neither Asian nor White.

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 data includes details on each disclosure's status as of March 2024. Our main innovation productivity tests consider all IDF submissions, regardless of their status in the process, and we also conduct additional tests that focus only on IDFs that become (or are very likely to become) approved patent applications.[11] 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.[12]

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.

[11] We refer to these IDFs as Status 2 disclosures, which have any of the following statuses: Approved Patent Application; Prior Art Search Requested; Returned for Comments. At times we also discuss Status 1 IDFs, which have any of the following statuses: Declined; Auto Declined; Defensive Publication Candidate; Approved Defensive Publication; Approved Trade Secret; Deleted; Merged. In Appendix B, we provide evidence from a series of robustness tests that show increased productivity based on other status classification approaches.

[12] 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.).

4.2 Innovator Characteristics

In Table 1, we report the demographic characteristics of the mentorship program participants across all seven cohorts and for non-mentored innovators. There were 101 innovators in the first cohort of the mentorship program. The average participant was 35 years old, and they had about 3.5 years of tenure with the firm and 6.5 years of tenure in the industry. Ninety percent of the first cohort participants were female, 71% were Asian, 2% were Black, 13% were Latinx, 13% were White, and 16% were URMs.[13] Eighty-one percent of the first cohort participants were engineers, 4% were "top" engineers (a distinction signifying their rank within the firm), and 3% were data scientists. The rest of the innovators have alternative job titles, such as "management" and "executive." The participants in the second cohort (N = 73) and third cohort (N = 171), are quite similar to cohort-one participants in their demographic characteristics. The participants in cohorts four through seven are fairly similar as well, though the fraction of male participants increases substantially in the later cohorts (37%–48%, up from 10%–22% in the first three cohorts).[14]

To understand how representative the mentorship program participants are of the broader innovator population, we compare the characteristics of the mentored innovators to those of the over 28,500 innovators who were not involved in the program. We find that the non-participants are, on average, older and more tenured than the mentored innovators. They are also more likely to be male and/or White. These differences come as little surprise, given that the firm specifically targeted the mentorship program towards individuals from under-represented demographic groups, and as it is expected that innovators who are relatively new to the firm would be more likely to seek out mentorship resources to help them get up to speed on the firm's invention development and disclosure process. Interestingly, though, we see in the bottom rows of Table 1 that the pre-mentorship quarterly disclosure outcomes

[13] Individuals can be members of multiple different racial/ethnic groups, which is why these percentages sum to a number greater than 100%.

[14] The change in cohort composition was partially in response to requests from human resources to be more gender inclusive.

of mentored innovators are quite similar to those of non-participants. This indicates that innovators from the two groups exhibited similar productivity levels at baseline. This suggests that non-mentored innovators can be used as a helpful comparison group as we estimate the relation between mentorship participation and invention disclosure outcomes.

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). As described below, we use both a TWFE estimator and a stacked DID estimator.

4.3.1 Two-Way Fixed Effects Estimator

To test for changes in disclosure outcomes after mentorship, we use a TWFE estimator, exploiting the staggered timing of the different mentorship cohorts and the presence of never-mentored innovators. We include innovator and year-quarter fixed effects, allowing us to estimate within-innovator level changes in productivity, using as a benchmark for comparison those not participating in the mentorship program at that time (or ever). Our exact specification is as follows:

$$\mathrm{y}_{i,t} = \alpha_i + \beta_1 \mathrm{Mentored}_{i,t} + \lambda_i + \gamma_t + \varepsilon_{i,t}, \tag{1}$$

where $y_{i,t}$ is the number of disclosure submissions that innovator i has in quarter t, λ_i is an innovator fixed effect, γ_t is a year-quarter fixed effect, and $\varepsilon_{i,t}$ is an idiosyncratic error. *Mentored* equals one in the quarter that innovator i's mentorship program begins and each subsequent quarter, and zero otherwise. We double cluster standard errors by innovator and

year-quarter.[15]

4.3.2 Stacked Difference-in-Differences Estimator

A critique of TWFE estimations with staggered treatment timing is that bias can stem from comparing already treated subjects with just treated subjects (Baker et al., 2022, 2025). We use a stacked DID estimator to circumvent these concerns, avoiding the potential bias.

To implement the stacked DID estimator, we create a new data set that compares the disclosure outcomes of innovators who participate in a particular cohort of the mentorship program to those of a set of "clean" control innovators. Specifically, for each cohort of treated innovators, we utilize two control groups: (1) those who never participate in the mentorship program (i.e., never-mentored innovators), and (2) those who first participate in the mentorship program in a later cohort (i.e., not-yet-mentored innovators). The latter group provides a particularly compelling counterfactual, as these individuals eventually select treatment but differ only in the timing of their participation. One benefit of this approach is that it allows us to isolate the role of mentorship as a distinct phenomenon from selection challenges because the mentored and not-yet-mentored innovators eventually demonstrate the same self-selection decision. In other words, not-yet-mentored innovators are likely to be selected on the same unobservable characteristics that lead currently-mentored innovators to participate in the mentorship program. The identifying assumption is thus milder, namely, that conditional on observable characteristics, the timing of initial program participation is as good as random with respect to potential outcomes.[16] While heterogeneity in treatment across cohorts remains possible, the consistent curriculum and mentor training means this design minimizes biases from unobserved factors.

We aggregate all cohort-specific data. Namely, the innovator-quarter observations for

[15]This clustering approach is conservative. If we, instead, single cluster standard errors by mentor, double cluster by mentor and year-quarter, or triple cluster by mentor, year-quarter, and innovator, then the standard errors shrink considerably, and our estimates become even more precise.

[16]Specifically, there were instances in which innovators wanted to participate in earlier mentorship cohorts but were unable to do so due to mentorship cohort capacity limits, so for some mentored innovators, the timing of their participation was quasi-random.

innovators treated in a specific cohort and the observations for that cohort's control group innovators. Then, we stack the aggregated data for all seven of the mentorship program cohorts and again regress disclosure outcomes on the *Mentored* indicator, with the inclusion of innovator-by-cohort fixed effects and year-quarter-by-cohort fixed effects. Our exact specification is as follows:

$$y_{i,t,g} = \beta_1 \text{Mentored}_{i,t,g} + \alpha_{i,g} + \lambda_{t,g} + \varepsilon_{i,t,g}, \tag{2}$$

where, $y_{i,t,g}$ is the disclosure outcome variable for innovator i in 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. We double cluster standard errors by innovator-by-cohort and year-quarter-by-cohort to account for repeated innovator observations across the cohorts (i.e., as the same innovator can be used in the control group for multiple cohorts).[17] This estimation approach is analogous to estimating separate DID treatment coefficients for each mentorship cohort group, using never-mentored and not-yet-mentored innovators as controls, and then taking the average of the treatment coefficients across all seven cohorts.

4.3.3 Parallel Trends Tests

Before we discuss our main findings, we first check to see if the assumption of parallel trends appears to hold, which is necessary for identification. To evaluate the credibility of this assumption, we compare the pre-mentorship trends in the quarterly disclosure submission outcomes between mentored and not-yet-mentored innovators. While it is not possible to definitively establish if the assumption is met, these tests can provide suggestive evidence in support (Rambachan and Roth, 2023).

Figure 1a plots quarterly disclosure submission patterns around the mentorship initiation quarter. The point estimates reflect the relative differences in disclosure submissions

[17] The inclusion of time-related fixed effects (year-quarter and year-quarter-by-cohort) and the clustering of standard errors by either year-quarter or year-quarter-by-cohort accounts for the serial correlations in disclosure outcomes that occur across time.

between mentored innovators and not-yet-mentored innovators in each quarter. Navy dots with light blue rectangles show the point estimates and 95% confidence intervals for quarters preceding mentorship participation, while maroon dots with light red rectangles show the corresponding estimates for subsequent quarters. We generate these estimates using the Callaway and Sant'Anna (2021) methodology. In Figure 1b, we repeat this procedure using an alternative dependent variable that scales each disclosure submission by the size of the team that produces the invention. We find evidence of common pre-trends in disclosure activity before the mentorship program, followed by significant increases in disclosure submissions post-mentorship.[18] Across all parallel trend tests, whether examining combined or separate disclosure types, the figures show similar submission patterns between mentored and not-yet-mentored groups before the implementation of the mentorship program. 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 0.381 to 0.960 (see figure notes). These patterns suggest comparable submission trajectories between mentored innovators and not-yet-mentored innovators in the quarters preceding program participation.

4.3.4 Survey Evidence on Non-Participation and Self-Selection

To identify potential self-selection patterns, we surveyed innovators at the firm about their participation decisions. This survey was part of a larger, multi-firm interview and survey process conducted by Chien and Grennan (2024) that is consistent and reliable, with response rates in line with other corporate surveys (Graham et al., 2005). This particular survey was administered among a small, demographically representative sample of engineers and technical staff in positions deemed by the focal firm as capable of contributing to IP development, along with an oversampling of engineers participating in mentorship.

The results are presented in Table A.1. Panel A shows that 43% of the survey respon-

[18] In addition, Appendix Figures A.3a and A.3b present corresponding analyses for Status 2 and Status 1 disclosures, respectively.

dents had participated in the mentorship program and confirms that female engineers were significantly more likely to participate, consistent with our administrative data. Next, we explore selection issues associated with participation. In particular, for non-participants, we asked: "What are the reasons you have not sought patent mentorship? (check all that apply)." To minimize response bias, we randomized the presentation order of answer choices using Qualtrics. Panel B of Table A.1 reveals that lack of awareness was the primary reason for non-participation across the sample. However, women were significantly more likely than men to believe patent mentorship would not benefit them (11% versus 3%). This pattern is particularly noteworthy given that women are more likely to participate overall, suggesting a negative relation between self-selection into program participation and ex-ante perceptions of the program's efficacy among some participants. Prior research on mentoring suggests that such self-selection would work against us finding significant positive benefits from mentoring (Sandvik et al., 2025).

This gender difference in perceived benefits admits two possible interpretations: it could reflect (i) a lack of confidence among women (a personal trait), which has been shown to generate gender gaps in STEM professions (Gallus and Heikensten, 2020), or (ii) bureaucratic barriers to knowledge flows (an organizational trait) that the program aimed to address. The survey responses help distinguish between these mechanisms. Women did not disproportionately cite avoiding activities that highlight weaknesses, suggesting the gap is not driven by confidence issues (Coffman, 2014; Bordalo et al., 2016). Moreover, virtually no respondents indicated they already received adequate IP support from colleagues, confirming the program was needed and not redundant. Time constraints were cited by only 6% of respondents, with no significant gender differences.

5 Results

This section presents our empirical findings in four parts. First, we examine the direct relation between mentorship participation and subsequent innovation activities, documenting

both immediate and long-term patterns in disclosure submissions. We then analyze how these patterns differ systematically across mentee and mentor characteristics such as gender, ethnicity, age, and geography (Section 5.1). Second, we examine survey evidence to identify the specific mechanisms through which mentorship generates productivity gains. We then leverage our unique administrative data to empirically validate these mechanisms, with particular attention to the role of expanded professional networks in facilitating innovation outcomes (Section 5.2). Third, we investigate potential indirect spillovers along several dimensions: the relation between mentored and non-mentored collaborators' productivity, engineers' perceptions of the culture and workplace, and talent retention rates (Section 5.3). Finally, we address possible alternative explanations for the observed relation between mentorship and innovation (Section 5.4). Throughout our analysis, we acknowledge potential limitations and conduct multiple specification tests to examine the robustness of the observed relations.

5.1 The Mentorship Program and Innovator Productivity

We begin by focusing on the direct relation between the mentorship program and mentored participants' invention disclosure submissions, our primary observable productivity outcome measure. We estimate the overall and long-term relation and discuss the robustness of the observed relation. We also consider how heterogeneity across mentee and mentor characteristics correlates with potential gains in innovative productivity.

5.1.1 Mentorship Program Participation and Overall Productivity

We report the relation between mentorship and innovator productivity in Table 2. In Panel A, the dependent variable in Columns (1)–(6) is the number of unique ideas an innovator submitted in a given quarter, unconditional on the idea's eventual status as a patent application. In the odd-numbered columns, we use the TWFE estimator, and we use the stacked DID estimator in the even-numbered columns. We use both not-yet-mentored innovators

and never-mentored innovators as the basis for comparison.[19] The coefficient estimates in Columns (1) and (2) suggest that mentorship is associated with an increase in disclosure submissions of 0.25 each quarter, or one additional disclosure per innovator per year. This positive association between mentorship and innovator productivity is economically meaningful and statistically significant in each specification, p-values < 0.01.

In Columns (3)–(8), we report the results of our primary robustness tests (we discuss the results of many additional robustness tests in Appendix B). First, in Columns (3) and (4), to reduce the concern that our main findings are driven by violations of the SUTVA due to spillovers to the innovators who collaborate with mentored individuals in the future, we remove the observations from these innovators. The coefficients on *Mentored* are essentially unchanged when making this restriction. Second, in Columns (5) and (6), we demonstrate the robustness of our results when we only use not-yet-mentored innovators in the control group (i.e., we remove never-mentored innovators from the analysis). Finally, in Columns (7) and (8), we use an alternative dependent variable, the ideas an innovator submitted in a given quarter that become approved patent applications (i.e., high-quality, Status 2 ideas). We continue to estimate a positive and statistically significant coefficient on *Mentored*, with the economic significance suggesting that, after mentorship, each mentored innovator produces approximately 0.2 additional high-quality inventions each year.

In Panel B of Table 2, we repeat the same tests from Panel A, but we use an alternative dependent variable that scales each disclosure submission by the size of the team that produces the invention. By doing this, we mitigate the concern that mentorship might simply teach participants how to join existing invention teams, which would increase their personal productivity outcomes without necessarily improving firm-level outcomes. The coefficients on *Mentored* are all positive and statistically significant (p-values < 0.01). This further suggests that a positive relation exists between mentorship program participation and subsequent innovative outcomes.

[19]The sample size increases when using the stacked DID estimator because we form control groups by drawing from the sample of control innovators with replacement.

5.1.2 Mentorship Program Participation and Long-Term Productivity

As mentioned, the mentorship program was designed to help innovators gain the skills, confidence, and knowledge necessary to develop and patent their inventions successfully. To this end, the program's curriculum required mentors to guide their mentees through the invention development and patenting process, with disclosure submission being an expected outcome. Given this expectation, it might be unsurprising that participation in the mentorship program is associated with more disclosures in the near term. What matters, from both the firm's perspective and that of the innovators, is that mentored participants could successfully develop and patent their inventions independently after the program.

To test whether this benefit of mentorship occurred, we estimate the long-term relation between mentorship and disclosure outcomes. We use the same specifications as we did in the analyses in Table 2, but we remove the observations of mentored workers in the quarter of their mentorship program and the first quarter after the program. In doing so, we exclude the disclosure submissions that most likely occurred directly as part of the mentorship program curriculum, allowing us to isolate post-mentorship influence. In the stacked DID specifications, we also remove the corresponding quarters of data from each cohort's control group observations. We tabulate these results in Table 3.

The estimated relation between innovator mentoring and aggregate disclosure counts in Columns (1)–(6) of Panel A continues to be positive and statistically significant at the 1% or 5% level. The estimated relation for high-quality ideas in Columns (7)–(8) is also positive and significant at the 1% level. The economic magnitudes are approximately one-third of the corresponding estimates from Table 2, but they still constitute a meaningful amount of additional innovation in the long run from participation in the program (approximately 0.33 additional disclosures per innovator each year).

In Panel B, where we normalize each submission by the innovator team size before aggregating to the innovator-year-quarter level, we continue to find evidence of a positive and significant relation between mentorship and long-term innovator productivity. The effects

are significant at the 1% level when considering high-quality disclosure outcomes (Columns (7)–(8)). Thus, the long-term productivity gain results suggest that mentored innovators acquired the skills and knowledge necessary to independently develop and disclose inventive ideas that generate successful patent applications, further highlighting the merits of the mentorship program.

We show in Appendix B that the estimated relation between mentorship and productivity is robust when using alternative control groups (e.g., only never-mentored innovators), when employing different estimators (e.g., the Callaway and Sant'Anna (2021) estimator or a Poisson estimator), when using different disclosure status classifications (e.g., considering defensive publications and other parts of a firm's IP strategy), and when using different patent application quality metrics (e.g., forward citations, claims, and scope). We also discuss the results of other tests that help us understand the variation in our data and what is driving our inferences.

5.1.3 Heterogeneity Tests Based on Mentee and Mentor Characteristics

Next, we consider whether certain mentees, based on their demographic characteristics, are more or less likely to have positive productivity gains from participating in the mentorship program. While the mentorship program was available to all employees, regardless of their gender or race/ethnicity, a goal from the firm leadership's perspective was to assist employees from demographic groups that have historically been underrepresented to be more successful in the innovation and patenting process. As such, it may have been the case that female and URM innovators benefited more from the program if, for instance, they were especially engaged in the curriculum or if they started with less knowledge of the patenting process than their peers.

In Table 4, we report the results from this heterogeneity test, wherein we continue to employ the TWFE estimator in the odd-numbered columns and the stacked DID estimator in the even-numbered columns. For these tests, we interact *Mentored* with the binary indicator

Female/URM, which equals one if the innovator is either female or a URM (neither Asian nor White), or both, and zero otherwise. The sample size is 2/3 smaller than that of our main tests due to incomplete data on gender and ethnicity for all innovators.

The heterogeneity analysis reported in Table 4 reveals evidence that mentorship yields larger benefits to mentees from URGs. The coefficient estimates on the interaction terms are positive across all specifications and statistically significant at the 5% or 10% levels in seven of the eight specifications. To interpret the economic magnitude, consider that our baseline estimates suggest mentorship increases idea submission by 0.25 ideas per innovator per quarter, on average. For innovators from URGs, this relation increases by 0.08 ideas per quarter, translating into 1.33 additional ideas submitted annually compared to non-mentored peers. This pattern suggests mentorship helps innovators from URGs overcome organizational impediments that may prevent them from fully realizing their innovative potential. Importantly, though, the coefficient estimate on *Mentored* remains positive and statistically significant in all columns, indicating that mentorship program participation is linked to productivity gains for all innovators, including White and Asian males.[20]

We also consider other variants of mentee heterogeneity, including age, geographic location, and timing of participation. These results are included in Appendix Table A.3 and Table A.4. We see that mentees across all age groups benefit from mentorship. Similarly, U.S.-based mentees and those in early mentorship cohorts realize productivity gains that are similar to non-U.S.-based mentees and those in later mentorship cohorts, respectively. This helps strengthen potential causal interpretations by ruling out two important alternative explanations. First, if benefits were concentrated among younger innovators, one might worry that mentorship accelerates learning that would naturally occur with experience rather than unlocking additional innovative potential. Second, if the relation varied substantially by region, local unobservable factors might confound our estimates. The uniformity of the estimated relation suggests instead that mentorship addresses organizational-specific con-

[20] Appendix Table A.2 shows qualitatively similar, albeit noisier, estimates when gender and ethnicity are evaluated individually and intersectionally.

straints in the innovation process such as knowledge barriers, rather than context-specific factors. This pattern of results is consistent with mentorship providing fundamental tools that enhance innovative capacity regardless of an innovator's career stage or location.[21]

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 estimate mentor fixed effects using the following approach. First, for the mentees in the sample, we regress their quarterly number of submissions during and after their first quarter in the mentorship program on their demographic characteristics (age, tenure with firm, tenure in industry, gender, underrepresented minority status, and job role), and we include fixed effects for year-quarter, their initial mentorship cohort, and their work country. We save the residuals from this estimation, and then, for a given mentor, we take the average residuals across all the mentees that the mentor supervised. These averages become the mentor fixed effects. We then split mentors into above versus below the median value of the mentor fixed effects, creating a *High FE* indicator, and we interact this indicator with the *Mentored* indicator and re-estimate our main regression models.

The results, which are reported in Panel A of Table 5, show heterogeneity in the quality of mentoring, but importantly, all innovators, even those with the lowest-quality mentors, benefit. The positive and statistically significant coefficients on *Mentored* in all eight columns indicate a positive relation between mentorship and innovation outcomes among those mentees whose mentors had relatively low fixed effects. The positive and significant coefficients on the interaction term indicate that mentees with high-fixed-effect mentors realized an even greater (i.e., more positive) association between their involvement in the mentorship program and their innovative output. We arrive at similar conclusions when we view the results

[21]The coefficient estimates on cohort timing are insignificant in 6 out of 8 specifications, but they are negative and significant at the 10% level for high-quality ideas. This suggests that mentees in early and late cohorts realized similar productivity gains from mentorship program participation, with a possibility that later participants benefited slightly more. This could be due to improvements in mentorship program quality as mentors gained experience over time. However, given the lack of precision across most of the specifications and best practices under multiple hypothesis testing (Harvey et al., 2015), it is best to conclude that we fail to reject the null hypothesis that timing differences matter.

in Panel B, which interact *Mentored* with indicators for mentor fixed effects quartiles, with *Q4 FE* capturing mentors in the top quartile of the distribution. We find that mentees with mentors in the bottom quartile of the fixed effects distribution realize gains to mentorship, and these gains are similar to those with mentors in the second fixed effects quartile. We see greater productivity gains for those with mentors in the third fixed effects quartile, and the gains are greatest for those with the best mentors.[22]

These findings have important implications for the scalability of the intervention, an issue gaining increased traction in experimental settings (List, 2022), and which we discuss in detail in Section 7. While the highest quality mentors appear to be associated with the greatest program returns, it is evident that even the lower quality mentors can help participants gain the skills and knowledge they need to successfully develop inventive ideas independently. We also consider heterogeneity based on alternative mentor characteristics in Appendix Table A.6. In general, we find little evidence of significant differences, suggesting that mentors were, on average, quite effective regardless of their gender, job status (e.g., awarded "top engineer"), or prior involvement in the mentorship program.

5.2 Exploring the Mechanisms Driving the Productivity Gains

To better understand the mechanisms through which mentorship works and how it enables productivity gains, we surveyed innovators involved in the mentorship program, using the same survey instrument capturing non-participant behavior described in Section 4.3.4. Table 6 summarizes the perceived experiences of those with a mentor and shows the program was positively received; so much so that mentees continued to meet with their mentors after the formal mentorship program ended. Mentees state that their mentors provide advice, suggestions, or support around twice a month, and this finding is the same across demographic groups.

When looking at the perceived benefits, mentored innovators reported benefits from the

[22]We report the demographic characteristics of the mentors, split by mentor fixed effect quartiles, in Table A.5.

mentorship received across all response choices.[23] All mentees "strongly agree" that "they are planning to work on more invention disclosures at work" (mean of 1.7 on a scale of -2 = strongly disagree to 2 = strongly agree). In addition, mentees strongly agree that they have: (i) learned a lot about patenting and becoming a better inventor, (ii) are more satisfied at work, (iii) have developed professionally, gained leadership skills, and think more broadly, (iv) have benefited from their mentoring relationship, and (v) their mentor introduced them to new inventors that they anticipate connecting with in the future. Two other response choices only received aggregate levels of "agreement." Both related to confidence: (i) confidence that the mentee could incorporate their mentor's tips into their work process and (ii) confidence that the mentee had the social connections necessary to get their invention ideas accepted by IP professionals.

While most responses were statistically indistinguishable by gender or URM status, there were some differences. Notably, females' perceptions of mentorship are more favorable. Females are in stronger agreement that mentorship increases their job satisfaction (average response of 1.4 compared to 1.0 for males, and the difference is statistically significant). Females also perceive that they have benefited more in terms of learning about patenting (1.5 compared to 0.9 for males, and the difference is statistically significant). Overall, the data suggests that targeted mentorship helps transform non-patenting STEM workers into active patent contributors through a few key mechanisms: (i) enhanced understanding of the invention process, (ii) expanded social networks, and (iii) increased satisfaction at work.

5.2.1 Validating the Expanded Social Networks Mechanisms

Next, to help validate the survey mechanism of expanded social networks, we explore whether the mentorship program is related to the demographic composition of mentored innovators' subsequent inventive teams. We may expect increased team diversity due to the mentorship program for multiple reasons. First, the mentorship program may have exposed innovators to

[23] We randomize the display of response choices to mitigate survey bias.

a more diverse set of possible collaborators with whom they could work. Second, the program may have highlighted the possible benefits of increasing the diversity of their collaborator teams. Third, the protocol materials for the mentorship program specifically emphasized creating space for "diverse perspectives and inclusivity." Our empirical analysis investigates whether these program features translate into systematic differences in the demographic composition of mentored innovators' research teams following program completion.

To do this, we employ the same two empirical approaches from our earlier tests: the TWFE estimator and the stacked DID estimator. To study disclosure team diversity, 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)–(6) of Table 7. 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.[24] This HHI metric becomes the dependent variable in Columns (7) and (8) of Table 7. 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 (9) and (10) of Table 7.

[24]Smaller HHI values capture larger amounts of diversity (i.e., less concentration in representation among a single group).

In Panel A, we use the full sample of innovator-quarter disclosure submission data, whereas in Panel B we remove near-term observations to estimate long-term effects. We find evidence that participation in the mentorship program is positively related to the level of female representation, gender-by-racial/ethnic diversity, and patenting experience diversity of an innovator's subsequent collaborator teams. The positive coefficients in Columns (1)–(6) of Panel A suggest that the female representation of mentored innovators' collaborator teams increases by 15–24 percentage points after their involvement in the mentorship program. The negative coefficients in Columns (7)–(10) suggest that gender-by-racial/ethnic diversity increases, as does patenting experience diversity. The results in Panel B suggest that the increases in the gender-by-racial/ethnic and patenting experience diversity of the teams persists in the long-term, whereas the increases in female representation do not.

These findings are important economically because the indirect influence of the mentorship program on the diversity of research teams post-mentorship could generate substantial material gains for the firm. For instance, Heath et al. (2023) measures the diversity and productivity of over 40,000 teams from open-source software projects, and they find that team diversity is positively associated with the survival, development, and popularity of the projects. They attribute this increase in performance to teams successfully avoiding the "homophily trap," whereby individuals' natural preference to associate with similar peers may cause research teams to miss out on the opportunity to collaborate with talented, albeit less similar, peers.

5.3 Spillovers from the Mentorship Program

We find evidence consistent with mentorship having a direct, positive influence on the near-term and long-term productivity of the program participants. In this subsection, we extend our analyses to investigate whether innovators who did *not participate* in the program realized any potential benefits by interacting with mentored innovators.

Several hypotheses guide these analyses. First, it is possible that program participants

shared their newfound skills and knowledge with their non-mentored peers, leading to indirect benefits of the program. The program's implementation may also affect employees' overall perceptions of the firm, potentially influencing factors such as job satisfaction and organizational commitment. Furthermore, the presence of mentored colleagues may impact the likelihood of non-participants remaining with the firm, either positively through knowledge spillovers or negatively through perceived disparities in career development opportunities. We also consider that non-mentored employees may experience relative declines in productivity or career progression compared to their mentored peers, potentially due to resource allocation or attention disparities. Lastly, the program's selective nature may influence non-participants' motivation and engagement, potentially leading to either increased effort to match perceived advantages or decreased morale due to exclusion.

The mentorship program's influence on non-participating employees likely involves multiple interconnected relationships that could be either complementary or offsetting. To systematically examine these relationships, we implement several econometric specifications that measure different types of spillover patterns. This comprehensive empirical approach allows us to document the intended and unintended relations between the mentorship program and the broader organizational environment. Specifically, we examine how the program correlates with non-participants' productivity, perceptions of workplace culture, and talent retention.

5.3.1 Productivity of Non-Participants

To explore potential productivity pattern changes among non-participants following their exposure to mentored colleagues, we restrict our sample to only include innovators who were not directly associated with the mentorship program at any time, either as a mentor or mentee. 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 create an indicator variable, *Post First Exposure*, which equals one in the first quarter a non-participating inno-

vator is observed to collaborate with a mentored innovator and each subsequent quarter, and zero otherwise. We observe 520 non-participating innovators in our sample who collaborated with a mentored innovator, and we collect their disclosure submission data from four quarters before to four quarters after their first collaboration with a mentored teammate. We then regress disclosure submission outcomes on the *Post First Exposure* indicator using both the TWFE estimator and the stacked DID estimator. We report the results in Table 8. We estimate positive coefficients in all six columns, and the inferences about mentored collaborator exposure are statistically significant in Columns (2), (3), (4), and (6). These results provide evidence consistent with the mentorship program having some positive benefits for non-participating innovators' productivity.

In Table A.7, we consider the extent to which there is heterogeneity in the productivity gains based on the non-participating innovator's demographic characteristics. In Panel A, we interact *Post First Exposure* with the binary indicator *Female*, which allows us to estimate whether female non-participating innovators realize different productivity patterns than male non-participating innovators when they are exposed to mentored collaborators. The coefficients on *Post First Exposure* $\times$ *Female* are negative in all six columns, and they are statistically significant in Column (4), p-value < 0.10, which suggests that the non-mentored female innovators in our setting may benefit less from exposure to mentored collaborators than non-mentored male innovators. One potential reason for this pattern may be related to information sharing. Survey evidence from engineers suggests that women engineers are less likely to be privy to the information necessary for inventive purposes and are less likely to have a strong network of engineers to work with (Chien and Grennan, 2024). Both of the patterns evident in the survey are consistent with our observation of women seeing fewer indirect productivity gains from having a mentored colleague.

Similarly, in Panel B, where we interact *Post First Exposure* with the binary indicator *URM*, we estimate negative interaction effects in all six specifications and statistically significant interaction terms in four of the six specifications. In Panel C, where we interact *Post*

First Exposure with the binary indicator *Female/URM*, which equals one if the innovator is either female or a URM, and zero otherwise, we estimate negative and statistically significant interaction terms in all six models. These findings suggest that women and URM innovators realize fewer spillover benefits from the mentorship program, relative to their male and non-URM peers.

5.3.2 Perceptions of Corporate Culture and Engineer Turnover

The survey results, discussed in Section 5.2, indicate that mentorship substantially increases workplace satisfaction and serves as a key mechanism enabling innovators to develop into productive inventors. Recognizing potential self-reporting biases in the anonymous survey data, we validate that finding in this section using objective, external measures from a broader employee sample. This approach parallels our earlier validation of the expanded social network mechanism, allowing us to differentiate among the channels through which mentorship works by using complementary data sources with different strengths and limitations. Specifically, in this section, we examine engineers' perceptions of the focal firm's corporate culture and their turnover decisions.

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.[25] While only 633 innovators participated in the mentorship program (a small number of innovators compared to the total number of innovators within the firm), we have already documented multiple examples of spillovers that may have influenced the perceptions of innovators more broadly. Furthermore, an employee's knowledge that the firm is willing to invest in them by offering a mentorship program, even if they don't participate, could also influence their sentiment towards the firm.

[25]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 (Tambe et al., 2020; Graham et al., 2022a; Grennan, 2023; Liu et al., 2023; Martellini et al., 2024; Gornall et al., 2025).

Our empirical approach is a triple difference estimator (DDD). Our first difference is between engineers and non-engineers at the focal firm. Our second difference is between our focal firm and its peer firms. The industry peer group is defined based on market capitalization, R&D expenditures, and industrial classification. For a causal interpretation, only one parallel trend assumption must hold (Olden and Møen, 2022). Figure B.1a shows evidence consistent with the assumption of parallel trends being satisfied. Our DDD specification is as follows:

$$\begin{aligned} \mathrm{y}_{i,j,t,c} = \beta_1 \mathrm{Eng}_{i,j,t,c} + \beta_2 \mathrm{Post}_{i,j,t,c} + \beta_3 \text{Focal Firm} \times \mathrm{Eng}_{i,j,t,c} + \beta_4 \mathrm{Post} \times \text{Focal Firm}_{i,j,t,c} + \\ \beta_5 \mathrm{Post} \times \mathrm{Eng}_{i,j,t,c} + \beta_6 \mathrm{Post} \times \mathrm{Eng} \times \text{Focal Firm}_{i,j,t,c} + \alpha_i + \lambda_c + \varepsilon_{i,j,t,c}, \end{aligned} \tag{3}$$

where, $y_{i,j,t,c}$ is the overall rating for firm i by employee j in quarter t in country c. α_i is a firm fixed effect, λ_c is a country fixed effect, $\varepsilon_{i,t,j,c}$ is an idiosyncratic error term, and we use robust standard errors. *Post* equals one after the beginning of the mentorship program in the focal firm, and zero otherwise. *Eng* equals one for jobs titles containing engineer and zero otherwise. An indicator for just the *Focal Firm* is excluded from the specification as it is collinear with the firm fixed effect (α_i).

We begin with the full sample of over 13 million individual Glassdoor reviews. Glassdoor has a star rating system, and a five-star rating is considered the top rating. We keep the reviews from current employees of the focal firm and its industry peers, restricting them to those with non-missing job titles during the period of study. To create a relatively balanced panel around the first cohort of the mentorship program, which began in October 2020, we keep reviews from January 2018 to June 2024, the last month in our dataset (N = 269,437).

We present our estimates in three ways. First, we report the DID based on job title, comparing engineers and non-engineers. Second, we repeat this exercise but for the peer firms, which serves as a placebo test. Third, we combine the two samples to run the DDD estimator.

Panel A of Table 9 examines engineers' workplace perceptions in relation to the mentor-

ship program. The dependent variable is the firm's overall star rating given by an employee on Glassdoor. The results in Columns (1) and (2) show that the ratings of the focal firm are significantly larger after the start of the mentorship program, and this correlation is more pronounced among engineers, who are the main program participants. In Columns (3) and (4), we replicate the analyses for the focal firm's industry peers. In Column (4), we find that the coefficient on *Post* × *Engineer* is almost a precise zero (0.01 relative to the 0.13 in the focal firm) and it is statistically insignificant. The negligible coefficient on the interaction term for industry peers suggests that the improvement in overall Glassdoor ratings among engineers in the focal firm is not likely due to broader macro-trends in employee satisfaction. Finally, we run the DDD analysis around the implementation of the mentorship program using the engineers in peer firms as the control group. It reveals that engineers' perceptions of the focal firm improve more than the perceptions of engineers in peer firms. The estimate reported in Column (5) is statistically significant and economically meaningful.

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. 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. 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, as we did in the analysis of engineers' perceptions of their workplace.

We report the results in Panel B of Table 9. The estimate in Column (1) shows that the focal firm experienced fewer departures in the months following the implementation

of the mentorship program relative to the months before. In Column (2), we find that this reduction in turnover is stronger among the firm's engineers (i.e., they experience even *less* turnover), relative to employees in the other occupational roles. These findings are statistically significant, and the economic interpretation of the result is that engineers in the focal firm experience about 4% less attrition after the beginning of the mentorship program, relative to non-engineers.[26] This seemingly modest decrease can translate into substantial economic benefits for the firm when considering the costs associated with employee turnover, particularly for highly skilled positions such as engineers, which we explore further in our ROI calculation presented in Section 6.

In Column (3), where we consider the focal firm's peer firms, we find evidence of *increased* turnover after the implementation of the mentorship program in the focal firm. In Column (4), we find mixed evidence as to whether engineers experience a differential average treatment effect, relative to non-engineers, as the coefficient on *Post* × *Engineer* is not statistically significant when using the raw departure count outcome, but it is significant when using the transformed outcome in Table A.8. However, Column (5) of both tables provides evidence that the reduced attrition among engineers is greater in the focal firm, relative to its peer firms. We conduct falsification tests for the workplace perceptions and employee departure tests in Appendix B, which support the robustness of our inferences.

Next, in Table 10, we examine workers' perceptions of the firm's culture specifically. We arrive at similar inferences when we use employees' ratings of the firms' "Culture & values."[27] As with the overall firm rating, we see significant improvements at the focal firm and insignificant improvements in the peer firms, and the DDD specification reveals significant improvements among engineers at the focal firm relative to peer controls. The

[26] In Appendix Table A.8 we use, as an alternative dependent variable, the inverse hyperbolic sine transformation of departures, which accounts for skewness in the data while accommodating year-months in which no turnover is observed within a particular group of individuals. This measure also provides a straight-forward economic interpretation, leading to the inference of a greater reduction in attrition of 4%.

[27] Note, the triple-interaction regression model includes the *Post* × *Focal Firm* and *Focal Firm* × *Engineer* interaction terms, but, for brevity, we do not report these in the table. Similarly, all models include a constant, which we do not report.

culture patterns are also intriguing, given prior work showing how persistent devaluation often drives members of underrepresented groups to exit STEM roles (Spoon et al., 2023).

In Panel B, 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.[28] 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). The seven cultural values include adaptability, collaboration, community-orientation, customer-orientation, detail-orientation, integrity, and results-orientation. Across the seven cultural values, the largest change occurs in the community-oriented value, and it is statistically significant. The cultural norms underlying the community-oriented cultural value include personal growth, embracing diversity, career advancement opportunities, and managerial guidance. These norms are consistent with mentorship catalyzing culture change. To a lesser extent (about 40% of the economic magnitude), we see increases in collaboration, which is consistent with the patterns we observe in the administrative data. Finally, we also observe a significant increase in integrity, which seems important given its role in value creation at public firms (Guiso et al., 2015).

Taken together, the patterns that emerge across multiple specifications suggest the culture improved and engineers' departures decreased at the focal firm. From a cost-benefit analysis perspective, this points toward spillover gains from retention and a workplace perception coinciding with the mentorship intervention. The robustness of this relation across various model specifications helps support the notion that improved satisfaction at work is a key mechanism driving innovative productivity gains.

[28]These text-based measures are described in more detail in Internet Appendix D. 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 information (i.e., pros vs. cons) and oversampling (to account for imbalance in categories).

5.4 Addressing Possible Alternative Explanations

One potential alternative explanation for our findings is that existing teams incorporated mentored individuals into ongoing projects rather than mentorship generating genuinely 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 our main results are robust to scaling by team size (see Panel B of Table 2), in this subsection we conduct more 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 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 Appendix Table A.9). 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 Appendix Table A.10). Finally, we examine the time to the next idea submission to see if innovators can submit a new idea independent of their original team, and find that mentored innovators submitted their second innovation in 1.0 year vs. 1.25 years for non-mentored innovators. Importantly, quality remained high in these subsequent disclosures, with comparable approval rates (38% vs. 35%) and slightly lower rejection rates for mentored teams. This battery of tests confirms mentorship truly enhanced innovation efficiency rather than merely reshuffling mentored innovators into existing projects.

6 Return on Investment to the Firm

This section evaluates the financial returns from the mentorship program using a framework that builds on established methodologies for assessing human capital investments (Black and Lynch, 1996; Bell et al., 2019; Lazear et al., 2015). We present conservative estimates based on directly quantifiable outcomes, while acknowledging the additional value of harder-to-

measure organizational benefits.

First, we evaluate direct innovation productivity gains. Our most conservative estimates focus on patent productivity gains. Mentored innovators produced one additional invention disclosure annually, with approximately 11% converting to granted patents. Valuing patents conservatively at $38,460 per patent based on application and attorney fees, this translates to approximately $4.017 million (633 mentees $\times$ 1 submissions $\times$ 0.11 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 $1.354 million. Detailed calculations and alternative valuation approaches for patents (e.g., Kogan et al. (2017)), are presented in Appendix Subsection B.7.

Second, we evaluate indirect innovation productivity gains. In our analysis, we documented two measurable spillover channels. First, 520 non-mentored innovators who collaborated with mentored participants showed a 0.5 increase in annual submissions, generating an estimated $1.1 million in additional value. Second, increased team diversity post-program added approximately $0.26 million in value. These estimates employ the same conservative valuation approach described above.

Third, we evaluate the indirect value gains from improved retention and culture. The program generated substantial value through improved retention, with engineer turnover decreasing by 4 percentage points. Even attributing just 1% of this reduction to mentorship yields savings of approximately $45.9 million in replacement costs. Additionally, survey evidence indicates improved workplace satisfaction and knowledge sharing, enhancing organizational capital.

Importantly, the program required minimal infrastructure, just two part-time administrators, and voluntary participation from mentors. 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 ended. This suggests minimal opportunity costs from mentor time allocation.

Nevertheless, to present the most conservative estimate possible, we instead assume that participation replaced work hours and count that portion of the mentor's salary as a cost, which totals $0.74 million.

Thus, even using conservative assumptions and focusing primarily on direct benefits, the program's returns substantially exceed its implementation costs. While including harder-to-measure benefits like improvements in culture and retention could increase the ROI further, our baseline estimates imply that mentorship generates meaningful positive returns through multiple value-creation channels.

7 Scaling and Public Policy Considerations

Innovation drives economic growth, yet history is rife with examples of barriers to reaching maximum innovation potential.[29] These historical lessons emphasize that addressing frictions, which limit participation in patenting, is imperative for technological competitiveness. Despite progress in broadening participation in the patent system, significant gaps persist due to a variety of challenges: socioeconomic (Bell et al., 2019), geographic (Hsieh et al., 2019), institutional (Aneja et al., 2024), and organizational (Chien, 2019; Toole et al., 2020). The disparities at the organizational level mirror broader productivity differences observed across seemingly similar firms (Syverson, 2004; Gibbons and Henderson, 2012; Backus, 2020), suggesting that interventions could catalyze greater innovative productivity.

Our study contributes to research on inclusive innovation by examining whether targeted mentoring can help improve innovative productivity and smooth organizational frictions. Replicating this type of mentorship program at other firms is possible. For example, beginning in 2021, over 40 major patent-filing firms signed the "Diversity Pledge," committing to taking actionable and measurable steps to increase diversity in innovation (Harrison, 2021). However, before encouraging other firms to adopt such a program, it is important to con-

[29] Experimental settings ranging from women's entry into STEM during WWII (Moser and Lubczyk, 2024), to racist lynchings in the South (Cook, 2014), to the impact of immigration restrictions (Moser et al., 2014) document the substantial economic benefits of inclusive innovation.

sider the external validity, or specifically how accurately our experimental estimates predict the impact of subsequent interventions (Muralidharan and Niehaus, 2017; Al-Ubaydli et al., 2017, 2020). Examining external validity is one part of the broader "voltage effect," which considers whether an intervention can maintain effectiveness at scale by systematically addressing a number of threats to scaling identified in the literature (List, 2022, 2024).

First, regarding statistical inference, our results remain robust across multiple estimation approaches and sample restrictions, reducing concerns about false positives that often plague smaller studies (Davis et al., 2017). The persistence of productivity estimates over longer horizons further supports the program's efficacy.

Second, we address population representativeness by documenting the program's effectiveness across demographic groups. Unlike many pilot interventions studying selected samples, our program was open to anyone working in the firm. Thus, we had the entire universe of workers to consider, and we utilized the rich demographic data available to understand potential differences across various dimensions of diversity including gender, race, job tenure, previous productivity, and geographic location.

Third, our explicit measurement of spillover effects quantifies how mentored innovators influence non-mentored peers. These positive spillovers increase productivity by approximately 25 percent for non-mentored peers, suggesting scaling would enhance rather than diminish effectiveness.[30] One interesting finding from our study is that the male-non-URM non-mentored peers appeared to benefit more than their underrepresented peers. This contrasts with the underrepresented mentees gaining marginally more from the intervention. But this also speaks to a way to scale the program, by continuing the targeted intervention (i.e., overselecting underrepresented groups or especially encouraging them to apply).

Fourth, our firm-wide analysis captures potential general equilibrium effects. The sustained productivity increases even as the program expanded across cohorts indicate minimal diminishing returns that would typically reduce the effectiveness of nudges at scale (DellaV-

[30]This spillover magnitude is 2.7 times larger than the average peer effect identified in previous literature (Herbst and Mas, 2015).

igna and Linos, 2022).

Fifth, our analysis of mentor-fixed effects addresses implementation fidelity. While the highest-quality mentors produced the largest gains, even lower-quality mentors generated significant productivity improvements of approximately 15 percent. The fact that even lower-quality mentors generated significant productivity improvements indicates that the program design itself, rather than individual mentor characteristics, drives the core benefits. This suggests the program's core benefits derive from its design, which was consistent across cohorts and mentors, rather than depending entirely on exceptional implementation.

Sixth, to assess scalability, we replace our individual productivity measures used throughout the paper with firm-level measures of innovative productivity. We find mentorship is positively and statistically significantly related to patent quantity and quality at the focal firm. We report these results in Appendix Table A.10.[31] The findings are consistent with the firm in our study producing more total patents per quarter after program implementation than its peers. Moreover, about 10% of the patents being produced at the time of the mentorship program turned out to be of the highest quality as measured by forwarded citation concentration. These firm-level findings reinforce our individual-level results, suggesting that mentorship benefits aggregate to meaningful organizational improvements in innovation outputs rather than merely redistributing innovative activity among employees. They also support the conclusion that a small-scale targeted mentoring intervention can have a large-scale influence on innovation.

Finally, we acknowledge the limitations of our study. Namely, we examine a superstar technology firm, which engages in immense levels of innovation and R&D investment, but that may be atypical compared to the standard firm. This firm is a technology leader with over 50,000 technical employees, substantial capital reserves, and access to top-tier talent.

[31]The dependent variable in Column (1) is the number of patent applications produced per quarter that eventually become granted, and in Column (2), it is the number of high-quality patents as proxied by those patents in the upper quartile of forward citations based on the total forward citations in the first three years. The dependent variables in Columns (3) and (4) reflect nominal and real patent value using the methodology pioneered by Kogan et al. (2017).

Treatment responses could differ in resource-constrained organizations or firms with different worker skill distributions. Consequently, our findings are most directly applicable to similar high-tech, knowledge-intensive workplaces.

8 Conclusion

This study examines the relationship between targeted mentoring and innovation outcomes within a superstar technology firm. Using proprietary data on the innovation pipeline from initial idea disclosure through patent application and exploiting the staggered implementation of a mentoring program for identification purposes, we document several findings that contribute to our understanding of human capital development and innovation within knowledge-intensive organizations.

First, we find that mentored innovators submit approximately one additional invention disclosure annually compared to control groups, with gains persisting well beyond the program period end. Second, we document positive spillovers to non-participants collaborating with mentored innovators, contributing to the literature on peer effects and knowledge diffusion within firms. Third, we survey the innovators and find support for three potential mechanisms leading to short- and long-term innovative gains: knowledge transfer about patent processes, expanded professional networks, and increased satisfaction at work. The results align with theoretical work on barriers to innovation and provide insight into how mentoring might reduce these frictions. Fourth, we extend our analysis to document broader organizational changes following the mentorship program implementation. Examining Glassdoor reviews and employee retention data reveals meaningful improvements in workplace satisfaction, the perceived strength of cultural norms associated with community, collaboration, and integrity, and engineering talent retention relative to multiple control groups. These patterns suggest the mentorship program's benefits ripple throughout the organization and its culture.

We present detailed estimates of the mentorship program's ROI. Direct innovation bene-

fits alone, conservatively valued using patent renewal costs, exceed implementation costs by a substantial margin. When accounting for spillovers to non-mentored peers and organizational benefits such as improved retention and stronger cultural norms, the economic case for mentorship strengthens further. We also address concerns about external validity and scalability by considering common threats in field experiments (List, 2022), and we show that the program maintains its effectiveness when expanded, with persistent productivity gains observed even with lower-quality mentors and across diverse participant groups. These findings suggest that targeted mentoring represents a cost-effective, scalable intervention for addressing innovation efficiency and inclusion challenges within technology firms.

While our single-firm setting provides an important proof of concept, future research should examine how corporate governance structures influence firms' ability to implement similar programs and whether capital markets efficiently price the innovation option value unlocked by more effective knowledge sharing. These questions are important given our finding that mentorship programs can effectively tap into latent innovation potential across the organization, potentially creating substantial economic value that extends beyond the individual firm to broader technological progress and economic growth. In conclusion, for leaders seeking to develop their innovation capabilities or policymakers aiming to increase returns on R&D investment, targeted mentoring appears to be a promising and relatively low-cost tool worth further exploration and experimentation.

References

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

Acemoglu, Daron, Jörn-Steffen Pischke. 1998. Why do firms train? theory and evidence. *Quarterly Journal of Economics* **113**(1) 79–119.

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.

Akcigit, Ufuk, Sina T. Ates. 2021. Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics* **13**(1) 257—-298.

Akcigit, Ufuk, Nathan Goldschlag. 2022. Measuring the characteristics and employment dynamics of u.s. inventors. *Census Center for Economic Studies Working Paper* 1–41.

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 Suskind. 2020. The science of using science: Toward an understanding of the threats to scalability. *International Economic Review* **61**(4) 1387–1409. doi:https://doi.org/10.1111/iere.12476.

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.

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

Athey, Susan, Christopher Avery, Peter Zemsky. 2000. Mentoring and diversity. *American Economic Review* **90**(4) 765–786.

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.

Azoulay, Pierre, Waverly Ding, Toby Stuart. 2007. The determinants of faculty patenting behavior: Demographics or opportunities? *Journal of Economic Behavior & Organization* **63**(4) 599–623.

Azoulay, Pierre, Joshua S. Graff Zivin, Jialan Wang. 2010. Superstar extinction. *Quarterly Journal of Economics* **125**(2) 549–589.

Backus, Matthew. 2020. Why is productivity correlated with competition? *Econometrica* **88**(6) 2415–2444.

Baker, Andrew, Brantly Callaway, Scott Cunningham, Andrew Goodman-Bacon, Pedro H. C. Sant'Anna. 2025. Difference-in-differences designs: A practitioner's guide. *Working Paper* .

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.

Becker, Gary S. 1965. A Theory of the Allocation of Time. *The Economic Journal* **75**(299) 493–517.

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.

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. 2024. "potential" and the gender promotion gap. *Yale Working Paper* .

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

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

Bhaskarabhatla, Ajay, Luis Cabral, Deepak Hegde, Thomas Peeters. 2021. Are inventors or firms the engines of innovation? *Management Science* **67**(6) 3899–3920.

Black, Sandra, Lisa Lynch. 1996. Human-capital investments and productivity. *American Economic Review* **86**(2) 263–67.

Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, John Van Reenen. 2019a. What drives differences in management practices? *American Economic Review* **109**(5) 1648–83.

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

Bloom, Nicholas, John Van Reenen, Heidi Williams. 2019b. A toolkit of policies to promote innovation. *Journal of Economic Perspectives* **33**(3) 163—-184.

Bloom, Nicholas, Heidi Williams, John Van Reenen. 2019c. A toolkit of policies to promote innovation. *Journal of Economic Perspectives* **33**(3) 163–184.

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.

Brooks, Wyatt, Kevin Donovan, Terence R. Johnson. 2018. Mentors or teachers? microenterprise training in kenya. *American Economic Journal: Applied Economics* **10**(4) 196—-221.

Budish, Eric, Benjamin Roin, Heidi Williams. 2015. Do firms underinvest in long-term research? Evidence from cancer clinical trials. *American Economic Review* (105) 2044–2085.

Cai, Jing, Adam Szeidl. 2018. Interfirm relationships and business performance. *Quarterly Journal of Economics* **133**(3) 1229–1282.

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

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* .

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

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

Cook, Lisa. 2014. Violence and economic activity: Evidence from african american patents, 1870–1940. *Journal of Economic Growth* **19** 221–257.

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

Cullen, Zoë, Ricardo Perez-Truglia. 2023. The old boys' club: Schmoozing and the gender gap. *American Economic Review* **113**(7) 1703–1740.

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

Davis, Jonathan M.V., Jonathan Guryan, Kelly Hallberg, Jens Ludwig. 2017. The economics of scale-up. Working Paper 23925, National Bureau of Economic Research.

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

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. Jill Burstein, Christy Doran, Thamar Solorio, eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186.

Eisfeldt, Andrea, Dimitris Papanikolaou. 2013. Organization capital and the cross-section of expected returns. *Journal of Finance* **68**(4) 1365–1406. doi:10.1111/jofi.12034.

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.

Fleming, L., H. Greene, G. Li, M. Marx, D. Yao. 2019. Government-funded research increasingly fuels innovation. *Science* **364**(6446) 1139–1141. doi:10.1126/science.aaw2373.

Fleming, Lee, Santiago Mingo, David Chen. 2007. Collaborative brokerage, generative creativity, and creative success. *Administrative Science Quarterly* **52**(3) 443–475.

Friebel, Guido, Matthias Heinz, Nikolay Zubanov. 2022. Middle managers, personnel turnover, and performance: A long-term field experiment in a retail chain. *Management Science* **68**(1) 211–229.

Gallus, Jana, Emma Heikensten. 2020. Awards and the gender gap in knowledge contributions in stem. *AEA Papers and Proceedings* **110** 241–44.

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.

Gibbons, Robert, Rebecca Henderson. 2012. Relational contracts and organizational capabilities. *Organization Science* **23**(5) 1350–1364.

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.

Goodman-Bacon, Andrew. 2021. Difference-in-differences with variation in treatment timing. *Journal of Econometrics* **225**(2) 254–277. doi:10.1016/j.jeconom.2021.03.

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.

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 Grennan, Campbell R. Harvey, Shivaram Rajgopal. 2023. Corporate culture in a new era: Views from the c-suite. *Journal of Applied Corporate Finance* **35**(4) 7–21.

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

Grennan, Jillian, Kai Li. 2023. *Corporate Culture: A Review and Directions for Future Research*. Handbook of Financial Descision Making, UK: Edward Elgar Publishing.

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.

Gupta, Abhinav, Naman Nishesh, Elena Siminitzi. 2024. Big data and bigger firms: A labor market channel. *ECGI Working Paper* .

Hamilton, Barton H., Jack A. Nickerson, Hideo Owan. 2003. Team incentives and worker heterogeneity: An empirical analysis of the impact of teams on productivity and participation. *Journal of Political Economy* **111**(3) 465–497. doi:10.1086/374182.

Harrison, Suzanne. 2021. Increasing diversity in innovation. URL https://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.

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

Hegde, Deepak, Bhaven Sampat. 2009. Examiner Citations, Applicant Citations, and the Private Value of Patents. *Economics Letters* **105**(3) 287–289.

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

Hoffman, Mitchell, Christopher T Stanton. 2024. People, practices, and productivity: A review of new advances in personnel economics. Working Paper 32849, National Bureau of Economic Research. doi:10.3386/w32849.

Hoffman, Mitchell, Steven Tadelis. 2021. People management skills, employee attrition, and manager rewards: An empirical analysis. *Journal of Political Economy* **129**(1) 000–000.

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

Hong, Lu, Scott E. Page. 2004. Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* **101**(46) 16385–16389. doi:10.1073/pnas.0403723101.

Hoogendoorn, Sander, Hessel Oosterbeek, Mirjam van Praag. 2013. The impact of gender diversity on the performance of business teams: Evidence from a field experiment. *Management Science* **59**(7) 1514–1528.

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.

Ichniowski, Casey, Kathryn Shaw. 2013. *Insider Econometrics: Empirical Studies of How Management Matters*. Princeton University Press, Princeton, 263–312.

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.

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

Koffi, Marlene, Matthew Marx. 2023. Cassatts in the attic. *NBER working paper 31316* .

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

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.

Lazear, Edward P. 1999. Globalisation and the market for team-mates. *The Economic Journal* **109**(454) 15–40.

Lazear, Edward P., Kathryn L. Shaw. 2007. Personnel economics: The economist's view of human resources. *Journal of Economic Perspectives* **21**(4) 91—114.

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

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

Lerner, Josh, Julie Wulf. 2007. Innovation and incentives: Evidence from corporate r&d. *The Review of Economics and Statistics* **89**(4) 634–644.

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

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

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.

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

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

Masclans, Roger, Sharique Hasan, Wesley M Cohen. 2024. Measuring the commercial potential of science. Working Paper 32262, National Bureau of Economic Research.

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

Mkrtchyan, Anahit, Jason Sandvik, Da Xu. 2024. Employee responses to ceo activism. *Journal of Accounting and Economics* 101701.

Moser, Petra, Moritz Lubczyk. 2024. The ms. allocation of talent. *Working Paper* .

Moser, Petra, Alessandra Voena, Fabian Waldinger. 2014. German Jewish emigres and US invention. *American Economic Review* **104**(10) 3222–3255.

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* .

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.

Otis, Nick, Rowan Clarke, Solene Delecourt, David Holtz, Rembrand Koning. 2024. Does generative ai crowd out social learning? evidence from a field experiment with entrepreneurs. UC Berkeley Working Paper.

Pairolero, Nicholas A., Andrew A. Toole, Peter-Anthony Pappas, Charles A. deGrazia, Mike H. Teodorescu. 2022. Closing the gender gap in patenting: Evidence from a randomized control trial at the uspto. *USPTO White Paper* .

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. *The Review of Economic Studies* **90**(5) 2555–2591.

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski, John Poe. 2023. What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics* **235**(2) 2218–2244.

Sandvik, Jason J, Richard E Saouma, Nathan T Seegert, Christopher T Stanton. 2020. Workplace knowledge flows. *Quarterly Journal of Economics* **135**(3) 1635–1680. doi:10.1093/qje/qjaa013.

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.

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.

Statista. 2024. Training industry in the u.s. - statistics & facts. URL https://www.statista.com/topics/4896/training-industry-in-the-us/#statisticChapter.

Sun, Liyang, Jesse M. Shapiro. 2022. A linear panel model with heterogeneous coefficients and variation in exposure. *Journal of Economic Perspectives* **36**(4) 193—-204.

Syverson, Chad. 2004. Market structure and productivity: A concrete example. *Journal of Political Economy* **112**(6) 1181–1222. doi:10.1086/424743.

Tambe, Prasanna, Xuan Ye, Peter Cappelli. 2020. Paying to program? engineering brand and high-tech wages. *Management Science* **66**(7) 3010–3028.

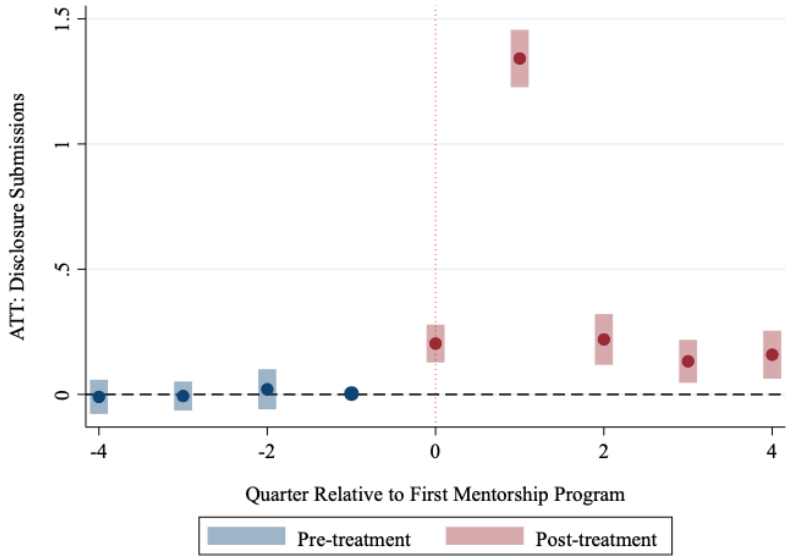
Tong, Siliang, Nan Jia, Xueming Luo, Zheng Fang. 2021. The janus face of artificial intelligence feedback: Deployment versus disclosure effects on employee performance. *Strategic Management Journal* **42**(9) 1600–1631. doi:10.1002/smj.3322.

Toole, Andrew, Michelle Saksena, Charles deGrazia, Katherine Black, Francesco Lissoni, Ernest Miguelez, Gianluca Tarasconi. 2020. Progress and potential: 2020 update on u.s. women invetor-patentees. *U.S. Patent and Trademark Office, Office of the Chief Economist* .

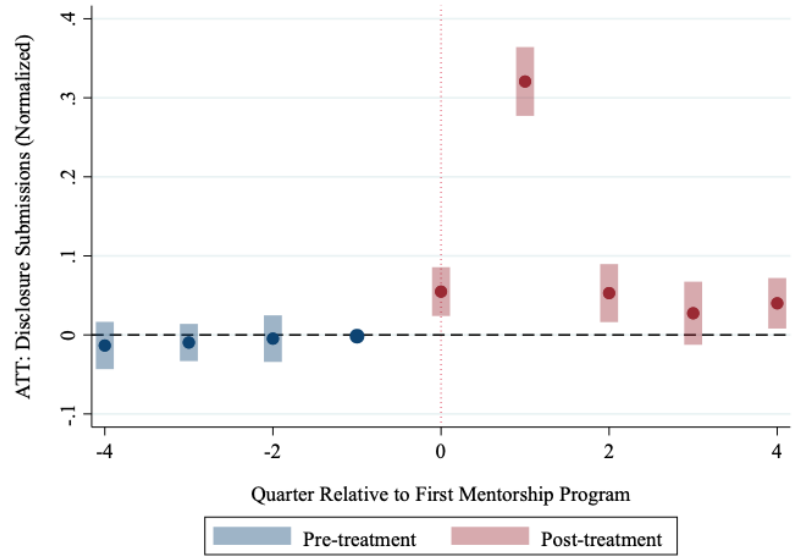
Wuchty, Stefan, Benjamin Jones, Brian Uzzi. 2007. The increasing dominance of teams in the production of knowledge. *Science* **316** 1036—-1039.

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

(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 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 from unscaled submission values. The average pre-period coefficient in figure (a) is 0.002 (p-value = 0.958), and the p-values on the individual pre-period coefficients range from 0.613–0.829. 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 -0.009 (p-value = 0.429), and the p-values on the individual pre-period coefficients range from 0.381–0.754.

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	0.10	0.05	0.11	0.11	0.08	0.06	0.08	0.08	0.481	0.873	0.760
	(0.39)	(0.12)	(0.40)	(0.33)	(0.22)	(0.10)	(0.28)	(0.51)			
Status 2	0.04	0.02	0.02	0.01	0.00	0.01	0.01	0.03	0.515	0.453	0.209
	(0.22)	(0.07)	(0.12)	(0.05)	(0.01)	(0.04)	(0.10)	(0.25)			
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

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***
	(3.372)	(3.753)	(3.324)	(3.732)	(4.269)	(3.983)	(4.023)	(3.932)
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

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***
	(3.061)	(3.521)	(3.024)	(3.505)	(4.304)	(3.696)	(3.896)	(3.853)
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

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. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table 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***
	(2.576)	(3.459)	(2.453)	(3.375)	(2.731)	(2.802)	(3.623)	(3.309)
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

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***
	(1.784)	(2.324)	(1.719)	(2.281)	(2.401)	(1.847)	(3.418)	(3.121)
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

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. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table 4: Mentorship and Innovator Productivity from Underrepresented Groups

			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

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. The key variable of interest is the binary indicator *Female/URM*, which equals one if the innovator is either female or a URM, 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. The empirical specification structure matches that used in Table 2. 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

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	0.138**	0.136**	0.138**	0.136**	0.425***	0.228***	0.021***	0.021***
	(2.427)	(2.360)	(2.394)	(2.344)	(3.702)	(2.882)	(3.387)	(2.913)
Mentored × High FE	0.197***	0.197***	0.197***	0.197***	0.208***	0.203***	0.051**	0.051***
	(3.558)	(4.801)	(3.557)	(4.804)	(3.768)	(4.840)	(2.527)	(2.936)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.210	0.228	0.308	0.309
Observations	473,939	3,272,940	465,778	3,215,813	9,549	22,210	473,939	3,272,940

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	0.143**	0.141**	0.143**	0.141**	0.420***	0.226***	0.017*	0.017
	(2.225)	(2.104)	(2.195)	(2.089)	(3.597)	(2.722)	(1.870)	(1.510)
Mentored × Q4 FE	0.263***	0.263***	0.263***	0.263***	0.275***	0.270***	0.066**	0.066***
	(3.720)	(4.849)	(3.723)	(4.853)	(3.716)	(4.664)	(2.482)	(2.858)
Mentored × Q3 FE	0.128**	0.127***	0.128**	0.127***	0.159***	0.144***	0.046*	0.046**
	(2.634)	(2.729)	(2.637)	(2.731)	(3.119)	(3.534)	(1.907)	(2.136)
Mentored × Q2 FE	-0.009	-0.009	-0.009	-0.010	0.012	0.006	0.008	0.008
	(-0.199)	(-0.202)	(-0.204)	(-0.208)	(0.217)	(0.129)	(0.536)	(0.552)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.402	0.392	0.397	0.210	0.228	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 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 *High FE*, 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), third, or second highest quartile of the mentor fixed effect distribution, *Q4 FE*, *Q3 FE*, and *Q2 FE*, 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 2. 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 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	86	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: 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***	-0.183***	-0.184***
	(5.372)	(4.924)	(5.416)	(4.895)	(7.137)	(2.957)	(-4.026)	(-4.008)	(-5.565)	(-5.164)
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	0.439	0.444
Observations	11,129	73,031	9,781	63,595	853	1,080	11,280	74,036	11,390	74,794

Panel B: Long-Term Productivity

			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**	-0.148**	-0.150***
	(1.255)	(1.347)	(1.251)	(1.330)	(2.245)	(0.930)	(-2.449)	(-2.604)	(-2.699)	(-2.643)
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	0.441	0.441
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 8: Productivity of Non-Participants After Exposure to Mentored Colleagues

			Primary Robustness Considerations			
	Full Sample		Alt. Control Group		Alt. Outcomes	
	(1)	(2)	(3)	(4)	(5)	(6)
Post First Exposure	0.517	0.615***	0.616***	0.312*	0.215	0.199*
	(1.744)	(2.764)	(3.001)	(1.850)	(1.554)	(1.920)
Two-Way Fixed Effects	✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓
Adj. R-Square	0.303	0.329	0.386	0.482	0.219	0.237
Observations	462,427	3,588,318	6,834	24,780	462,427	3,588,318

Notes. This table reports our analyses of potential productivity changes among individuals who did not directly participate in the mentorship program but 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. Our sample includes 520 non-participating innovators who we observe collaborating with mentored innovators. We then regress disclosure submission outcomes on this indicator using both the two-way fixed effects estimator (odd-numbered columns) and the stacked difference-in-differences estimator (even-numbered columns). In Columns (1)–(4), 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 (5)–(6) is the number of IDFs that become (or are likely to become) approved patent applications (i.e., high-quality ideas). The control group includes both not-yet-mentored innovators and never-mentored innovators in Columns (1)–(2) and Columns (5)–(6). In Columns (3)–(4), the control group is composed of not-yet mentored innovators only. Standard errors are clustered by innovator and quarter (innovator-by-cohort and quarter-by-cohort) when using the two-way fixed effects (stacked difference-in-differences) estimator, and t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table 9: 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 10: 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-oriented	Customer-oriented	Detail-oriented	Integ-rity	Results-oriented
	(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.

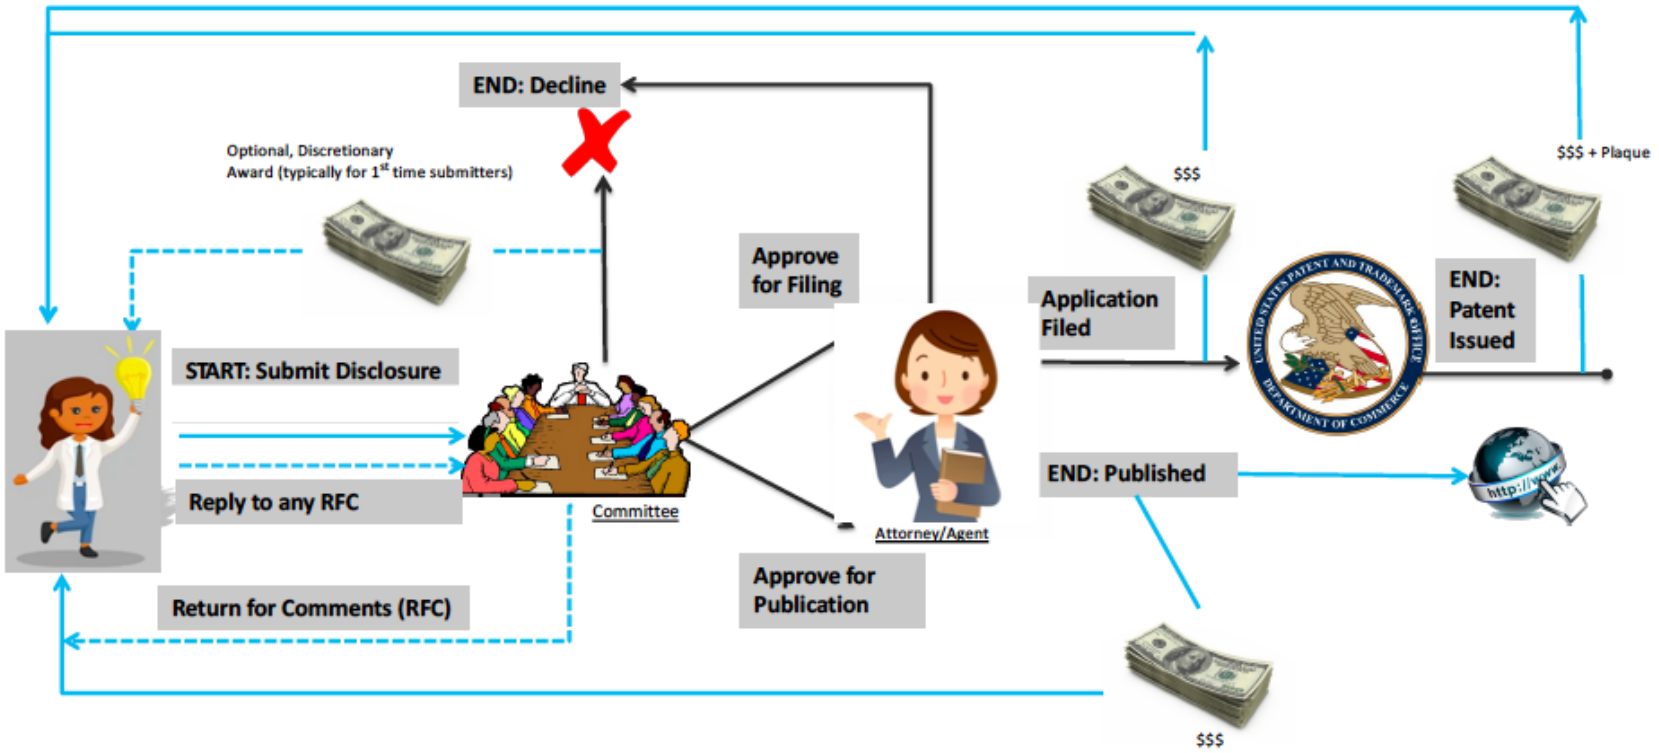
A Appendix Figures and Tables

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.

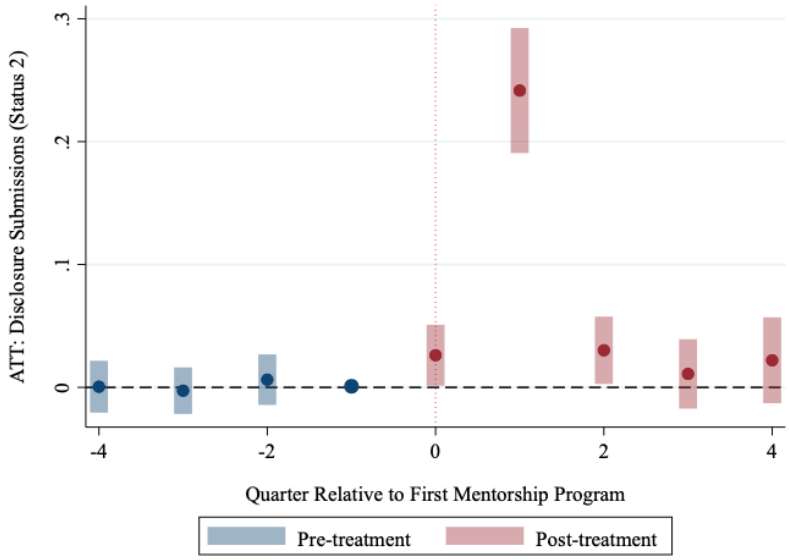
Figure A.2: Patent Process



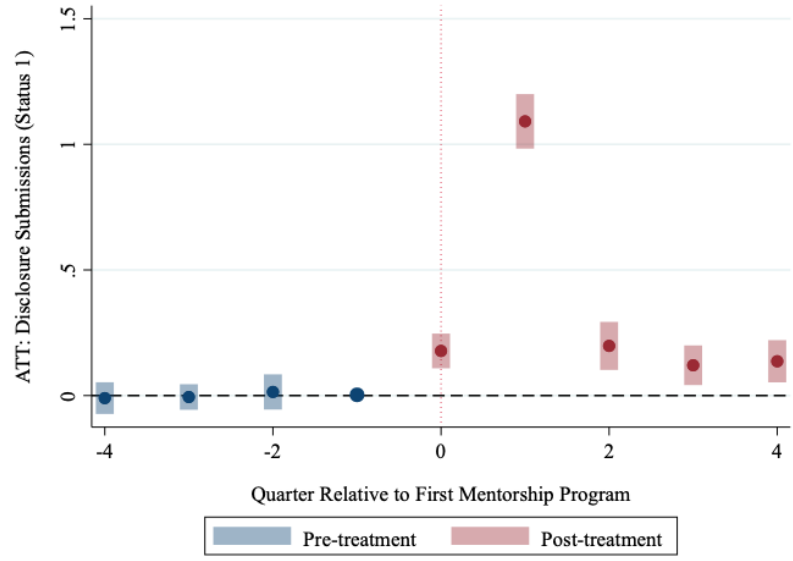
Notes. This figure highlights the patent process as taught to mentees by mentors as part of the mentorship program. The main outcomes that a disclosure submission can realize are (i) being declined, (ii) being approved for defensive publication, and (iii) being a filed or approved patent.

Figure A.3: Test of Common Pre-Trends by Submission Type

(a) Number of Submissions: Status 2



(b) Number of Submissions: Status 1



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. Status 2 submissions are those with any of the following statuses: Approved - Patent Application; Prior Art Search Requested; RFC. Status 1 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.

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	2%	0%	0%	2%	0%
I felt that the program would benefit my employer more than it would benefit me	125	1%	63	2%	0%	1%	0%	0%

Notes. This table provides descriptive statistics of the innovators' decisions to participate in patent mentorship and the perspective 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.

Table A.2: 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 2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.3: 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.4: 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	0.304***	0.299***	0.304***	0.299***	0.577***	0.366***	0.060***	0.059***
	(3.900)	(4.379)	(3.831)	(4.365)	(4.459)	(4.080)	(4.585)	(4.641)
Mentored × USA	-0.037	-0.035	-0.037	-0.036	-0.059	-0.041	-0.003	-0.003
	(-0.729)	(-0.776)	(-0.733)	(-0.784)	(-1.219)	(-0.898)	(-0.179)	(-0.200)
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 2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.5: Demographics Characteristics of Mentors

	Mentor Fixed Effect Quartile			
	1st (Lowest)	2nd	3rd	4th (Highest)
Age	47.22	46.04	46.08	46.17
Tenure (Firm)	8.58	10.47	9.84	10.28
Female	0.26	0.30	0.04	0.22
Asian	0.57	0.58	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	0.39	0.38	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	1.71	0.97	2.73	2.02
Number Mentors	27	29	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.6: 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 2. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.7: 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. The specification structure mirrors that in Table 8. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.8: 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 the main table. 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.9: 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	0.212***	0.187***	0.187***	0.061***	0.046***	0.046***
	(3.962)	(9.382)	(9.381)	(4.315)	(9.509)	(9.509)
Post × Focal Firm			0.025			0.015
			(0.446)			(0.977)
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.004	0.028	0.027	0.004	0.035	0.034
Observations	4,075	60,438	64,513	4,075	60,438	64,513

Panel B: After Second Mentorship Cohort

	Total Innovators			IHS(Total Innovators)		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.222***	0.197***	0.197***	0.063***	0.049***	0.049***
	(4.067)	(9.439)	(9.439)	(4.355)	(9.612)	(9.612)
Post × Focal Firm			0.025			0.014
			(0.425)			(0.907)
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.004	0.028	0.027	0.004	0.035	0.034
Observations	4,075	60,438	64,513	4,075	60,438	64,513

Panel C: After Third Mentorship Cohort

	Total Innovators			IHS(Total Innovators)		
	Focal	Peer	Both	Focal	Peer	Both
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.202***	0.213***	0.213***	0.054***	0.052***	0.052***
	(3.422)	(9.293)	(9.292)	(3.455)	(9.367)	(9.367)
Post × Focal Firm			-0.011			0.002
			(-0.173)			(0.093)
Firm FE		✓	✓		✓	✓
Adj. R-Square	0.003	0.028	0.027	0.003	0.035	0.033
Observations	4,075	60,438	64,513	4,075	60,438	64,513

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 began and onward (in Panel A). In Panel B (Panel C), *Post* equals one from when the second (third) mentorship program began 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.10: 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 our 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 variable 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 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 Tables 2 and 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 Table 2. 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-likelihood estimator in

Column (3), and we employ the Callaway and Sant'Anna (2021) estimator in Column (4).[32] 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.

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); "Submitted" (1% of the disclosures).[33] 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 "Submitted." The dependent variable in Columns (3)–(4) is the number of disclosure submissions with the status "Declined" or "Auto Declined." The dependent variable in Columns (5)–(6) is the number of disclosure submissions with the status "Approved - Patent Application." The dependent variable in Columns (7)–(8) 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.

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

[32] 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.

[33] The remaining 1% have statuses such as "RFC" and "Deleted."

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 on Collaborations with Mentors

In Table B.4, we show the robustness of our main findings when removing all IDFs for which a mentor (past, present, or future) is involved as a collaborator. By removing these IDFs, 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.5, 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.5 are very similar to our main results, mitigating the concern that additional mentorship is necessary for benefits to be realized.

B.6 Alternative Employee Perception and Retention Tests

The results in Table 9 and Table 10 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.6, 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 results in Panel A of Table 9 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.7. 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 Panel A of Table 9, 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.8 show that the estimated relation between the mentor-

ship 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 are very similar to those 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.9, 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 Panel B of Table 9 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.7 Alternative ROI Calculation Assumptions

This appendix section provides detailed calculations and 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, then extends to broader organizational benefits. Given the inherent complexity in measuring returns to human capital investment, we present multiple estimates with varying degrees of conservativeness. The direct productivity channel values patents using a conservative approach based solely on application costs and akin to classic renewal fee methods of valuing patents (Pakes, 1986). Currently, patent attorney fees ($25,000) plus mandatory renewal fees ($13,460) total $38,460 per patent. This approach likely understates patent value but provides a concrete lower bound. Using patent value estimates based on equity price movements, as proposed by Kogan et al. (2017), the median value of a granted patent between 2018 and 2023 is $19.7 million. The mean value is $68.1 million. For comparison, our focal firm's median value over that period is 2.5 times larger. Thus, if we were to use our focal firm's assigned value it would be approximately 1,250 times larger. Thus, our estimated impact would be more than 1,000 times larger since any movement in this superstar technology firm's stock price implies a much larger value for its patents. Even if we look at a blended price across all the peer firms, this leads to a value

400 times higher than that derived from attorney and renewal fees.

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 (Lazear, 1999; 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; Fleming et al., 2007; Azoulay et al., 2007; Lerner and Wulf, 2007; Hoogendoorn et al., 2013; 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 benefits, though these effects are harder to quantify precisely. We observe a 4% reduction in engineer turnover following program implementation. While multiple factors may contribute to this improvement, even a fraction of this effect could generate meaningful savings given typical replacement costs for employees in the technology sector. Attracting and hiring new engineers often involves substantial expenditures on job postings, recruiters, and candidate interviews. New engineering hires also require specialized knowledge and project familiarity to reach full productivity, which requires extensive training to get new hires up to speed and beyond their low initial productivity levels. This period of reduced productivity and the resources allocated to training represent meaningful expenditures for firms. Research indicates that a new engineer can take up to six months to achieve full productivity. This gap between an engineer's departure and a new recruit reaching full productivity along with the recruiting costs are estimated to be as high as 1.5 times an engineer's annual salary (Berk and Kase, 2010).

Thus, if one were to follow an approach similar to the one used by Lazear et al. (2015) to estimate the portion of value of good bosses associated with reduced churn, one could come up with a value. In this case, we could base our cost savings on the reduction in departures (conservatively, we could assume only 1% of the observed 4% reduction in turnover of the firm's 29,212 inventors is truly from mentoring), and we could conservatively estimate a cost of 1 times the average annual salary of engineers, even though evidence suggests it is 1.5 times. Finally, we get our salary estimate from Revelio for the year before the mentorship program, weighted based on the representation of engineers in different geographic regions and seniority levels ($157,297). This is a superstar technology firm with the capacity and likelihood of paying much higher salaries, but we will go with this more conservative industry

average. Overall, this suggests indirect cost savings from retention of $45.9 million (0.01 × 29,212 inventors × replacement cost (1 × $157,297 annual salary)) based on conservative assumptions.

Further, we know that employees are willing to take a salary cut for better non-pecuniary benefits such as exceptional corporate culture and workplace environments. Based on estimates from Makridis (2018) about the willingness to forgo a salary for a better culture ($1,159 per year) and an average turnover rate in the IT industry of 19% according to the U.S. Bureau of Labor and Statistics, we estimate an additional cost savings of $19.2 million from lower initial salaries for recruits (29,212 inventors × 0.19 turnover rate × $1,159 lower salary per year × 3-year mentorship program).

Finally, to align with the recent literature emphasizing the importance of intangible assets, particularly organizational capital and corporate culture, in driving long-term firm value, we explore how the patent mentorship program contributes to the firm's long-term value creation through this lens. For instance, Graham et al. (2022a) find that improving corporate culture as proxied through Glassdoor culture ratings and other measures is associated with an increase in Tobin's Q. Our focal firm has seen improvements in its Tobin's Q ratio since implementing the patent mentorship program. Given the focal firm's mega market capitalization, even attributing a few percentage points (e.g., 1 p.p.) of that increase in Tobin's Q to culture improvements would suggest benefits in excess of $650 million. Yet this appears quite conservative and small when considering that engineers account for nearly 30% of the focal firm's workforce.

Alternatively, Eisfeldt and Papanikolaou (2013) suggest that organizational capital is embodied via key talent, and they show that firms with more organizational capital have higher and more variable risk premia. In the context of the patent mentorship program, we can view the program as an investment in organizational capital that enhances the firm's innovative capacity and knowledge base. Eisfeldt and Papanikolaou (2013) estimate a scaling parameter of 0.3 representing the fraction of SG&A expenditures contributing to organizational capital formation. Based on the salaries of the participating mentees and using their scaling value, this suggests that the firm added $89.6 million (633 mentees × $157,297 average salary × 0.3 scaling parameter × 3-year mentorship program) to its organizational capital through its mentorship program.

The firm anticipated that the main cost of the mentorship program was going to be the opportunity cost of mentors' time. The mentors were among the most prolific innovators in the company, and one might have been concerned that the time required of them to mentor others may have reduced their overall level of innovativeness. We do not find evidence in support of this concern. In fact, we find that mentors increased their innovativeness during

the program, and then their innovative output returned to their pre-mentorship baseline levels after the program ended. Specifically, we find that before they began their duties as mentors in the program, mentors submitted 2.6 disclosures per quarter, on average.[34] During their involvement in the program, mentors submitted 3.7 disclosures per quarter, on average, which is equivalent to their normal level of productivity plus the additional disclosure that stemmed from the program itself. Then, in the quarters after the program ended, mentors reverted back to their approximate pre-program levels of disclosure production.

The mentorship program was completely voluntary. But to be conservative, we assume it did 100% crowd out actual work hours for both mentees and mentors. When signing up for the mentorship program, mentees were told to anticipate spending 70 hours total on the program, and for mentors, only about 40 hours, including the training to ensure a consistent curriculum. Recall that the mentees had inventive tasks to work on in between check-ins, making their hours higher. Based on the salaries of participating mentees, mentors, and administrators, we estimated $4.4 million in labor costs (70 hours is 3.6% of annual working hours assuming 48 working weeks $\times$ 40 hours per week). At a salary of $157,297 $\times$ 3.6% $\times$ 633 mentees = $3.63 million. Similarly, for mentors, 40 hours is 2.1% of annual working hours. Assuming mentors have a salary that is double that of engineers, we get $314,594 $\times$ 2.1% $\times$ 113 mentors = $0.74 million. On the administrative side, it took about 40 hours per cohort to administer $\times$ 7 cohorts $\times$ $50 per hour for administrators at high-tech firms = $14,000.

In conclusion, the total benefits of the mentorship program include direct benefits from greater innovation productivity among those mentored ($4.017 million) and indirect benefits from greater innovation productivity among those working with the mentored innovators ($1.1 million), more valuable patents attributable to increased innovator team diversity ($0.256 million), and plausibly larger gains from retention and culture improvements. In contrast, labor costs under the worst-case scenario are still $1 million less than the innovation gains, suggesting the program was cost-effective and had a positive ROI.

These calculations, while systematic and grounded in observable data, require several important qualifications. The patent valuation methodology likely understates true economic value. Further, our approach to measuring innovation gains captures only direct working relationships, potentially missing broader ripple effects through extended professional networks. The relation between mentorship and employee retention presents particular analytical challenges, as other organizational changes could contribute to retention patterns, although our interviews with program administrators indicated no such changes. Finally, converting im-

[34] In comparison, non-mentor innovators who were also not part of the program submitted fewer than 1 disclosure *per year*, on average.

provements in the organizational environment into quantitative economic measures remains methodologically complex, although 92% of executives believe that improving their firm's culture would increase the value of their company (Graham et al., 2022a).

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)	(0.023)
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)	(0.024)
Two-Way Fixed Effects			✓	
Stacked Diff-in-Diff	✓	✓		
PS-Matched Sample		✓		
Adj./Pse. R-Square	0.396	0.412	0.408	–
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

	Submission		Declined		Approved Patent		Publication	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.212***	0.210***	0.128**	0.126***	0.045***	0.045***	0.025**	0.025***
	(3.298)	(4.027)	(2.823)	(3.770)	(3.824)	(3.781)	(2.828)	(3.276)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.399	0.403	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	473,939	3,272,940

Panel B: Long-Term Productivity

	Submission		Declined		Approved Patent		Publication	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mentored	0.075**	0.074***	0.048*	0.047***	0.016***	0.016***	0.008*	0.008**
	(2.348)	(3.195)	(1.844)	(2.627)	(3.091)	(2.870)	(1.901)	(2.344)
Two-Way Fixed Effects	✓		✓		✓		✓	
Stacked Diff-in-Diff		✓		✓		✓		✓
Adj. R-Square	0.402	0.395	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	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 "Submitted." The dependent variable in Columns (3)–(4) is the number of disclosure submissions with the status "Declined" or "Auto Declined." The dependent variable in Columns (5)–(6) is the number of disclosure submissions with the status "Approved - Patent Application." The dependent variable in Columns (7)–(8) 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 (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 submissions at Status 2. 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.5: 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 submissions at Status 2. 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: 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.7: Engineers' Workplace Perceptions: Falsification Tests

Panel A: 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.8: 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.9: 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

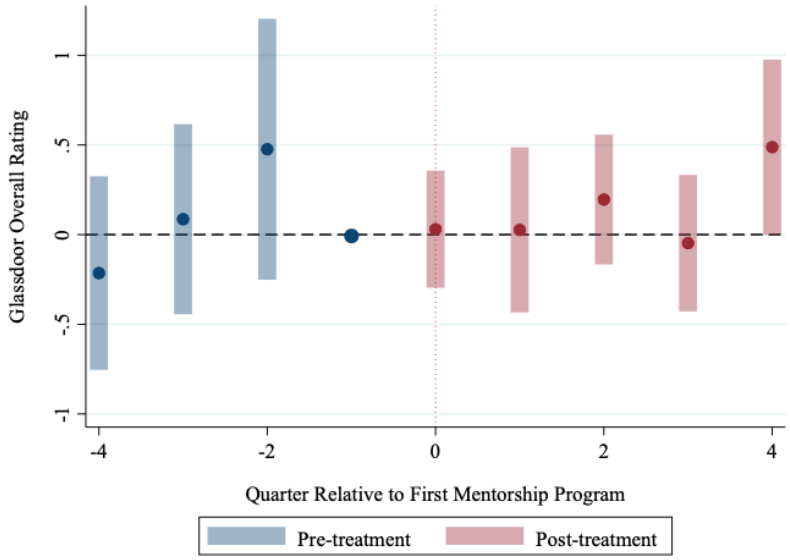
Panel B: Inverse Hyperbolic Sine of Outflows (Two-Year Lagged Timing)

	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

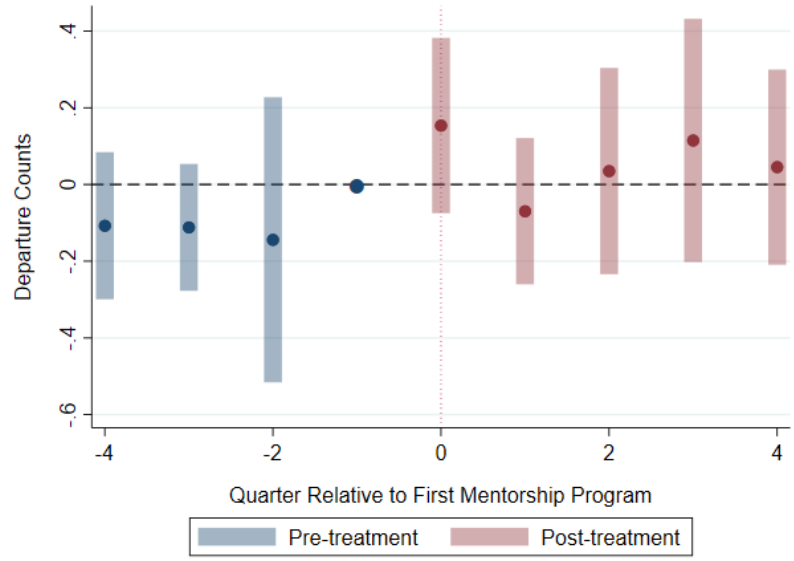
Notes. This table presents an analysis of potential changes in employee departure patterns coinciding with the implementation of the patent mentorship program. The data used is employee departure measures from Revelio, a labor market intelligence firm that tracks employee movements across firms. Panel 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 Percetions



(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 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

D 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 (Devlin et al., 2019).

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 B of Table 10 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,