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 small-scale, targeted mentoring can address organizational frictions that influence innovation and productivity. 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 innovation outcomes 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 confidence in identifying patentable ideas. The mentorship program is also associated with broader organizational improvements: engineers in the firm perceive the culture to be better, report higher job satisfaction, and are less likely to leave than engineers at peer firms. Accounting for these spillovers implies that small-scale mentoring programs can generate a meaningful return on investment and provide a cost-effective way to foster innovation and inclusivity.

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

*University of California Berkeley School of Law and Diversity Pilots Initiative
†Emory University and Diversity Pilots Initiative
‡University of Arizona, Eller College of Management and Diversity Pilots Initiative

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 grow, they typically transform from agile innovators into bureaucratic organizations producing more incremental innovations (Akcigit and Kerr, 2018), with organizational frictions impeding knowledge flows and reducing 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 mechanisms might help these firms maintain innovation efficiency at scale? Our analysis focuses on mentorship 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 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), and mentorship represents a natural extension of firms' existing human capital investments (Sandvik et al., 2025).

Specifically, we test whether a targeted patent mentoring program can increase innovation productivity 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 rel-

[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'm not confident about 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."

ative to their presence in STEM professions), are substantial. 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). Evidence suggests that these disparities represent structural 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 that was designed to promote innovation efficiency and overcome inclusion challenges, we contribute to policy debates about improving U.S. technological competitiveness (Bloom et al., 2019b), including those of the U.S. Patent and Trademark Office (USPTO), which identified the "innovator-inventor gap" as a priority (Chien, 2024). The mentoring program that we study aims to transform STEM workers who lack patenting experience ("innovators") into contributors who actively submit ideas for patents ("inventors") by increasing their exposure to and understanding of the invention process. By measuring both direct channels and spillovers to non-mentored peers, we provide comprehensive evidence on returns to human capital investments in R&D-intensive firms, overcoming measurement challenges that have historically complicated corporate training program evaluation (Black and Lynch, 1996).[3]

Specifically, we analyze the staggered roll-out of a mentorship program in a superstar technology firm, which involved 633 innovators participating in seven cohorts of the program, 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 program's curriculum paired experienced inventors with proven track records of patenting success (i.e., mentors) with small groups of mentees. The mentors guided their mentees through the inventive idea development and patenting process. The modal group had 5

[3]Note that human capital development investments are substantial, as U.S. firms spent over $100 billion on training in 2023 (Statista, 2024).

mentees and 1–2 mentors who met regularly over a calendar quarter.

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 estimator. The latter addresses potential biases arising from staggered treatment timing and evolving treatment effects (Baker et al., 2022; Roth et al., 2023).[4] We construct our stacked difference-in-differences estimator using two control groups: never-mentored innovators and not-yet-mentored innovators (i.e., those who were not selected to participate in an earlier cohort but who participated in a later cohort). 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 difference-in-differences approach using event study analyses and visualizations (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 multiple 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), allowing us to measure both the direct and indirect way organizational interventions may influence innovative ideas.

[4]Our inferences are also robust to the use of the estimator proposed by Callaway and Sant'Anna (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 future non-mentored collaborators from our analysis, restricting our control group to only those eventually mentored, focusing exclusively on high-quality disclosures that progress to approved patent applications, removing any idea 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.

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 gains from mentorship is economically meaningful, given a baseline disclosure rate of 0.5 ideas 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. Next, we show that these productivity gains are not driven by changes in team size and composition, helping to rule out concerns that the firm may be engaging in "diversity-washing."[5] Finally, triple difference-in-differences estimates suggest larger benefits for female and URM innovators, though standard errors are

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

large for these subgroup analyses resulting in significance only at the 5% level.[6]

Second, we present survey evidence, which reveals three key mechanisms driving the innovation productivity gains: knowledge transfer about patent processes, expanded professional networks, and increased confidence in identifying patentable innovations.[7] By surveying participants and non-participants, we distinguish between alternative explanations for organizational frictions suppressing innovative productivity. While underrepresented groups are less likely to anticipate benefits from patent mentorship, this hesitation stems from bureaucratic impediments to knowledge flows and collaboration costs rather than underconfidence or concerns about revealing weaknesses to leadership. These mechanisms align with existing theoretical work on barriers to innovation (Manso, 2011; Bohren et al., 2019).

Third, we document significant positive spillovers to non-participants. Our unique dataset allows us to track the formation of new collaborating teams and the submission patterns for both mentored individuals and their non-mentored peers. To test for such 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 (List, 2022).

Additionally, to validate the suggested social connectedness mechanism in the survey, we

[6]We use the race/ethnicity data to identify underrepresented racial minority (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 (C:2018).

[7]The externally validated June 2023 survey was conducted as part of a multi-firm study on innovation pipelines (Chien and Grennan, 2024).

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, job tenure, and previous patenting success. This improved team diversity is particularly important given evidence linking diverse teams to higher-value patents. Hence, these findings suggest mentorship can help inventors overcome the "homophily trap" that many teams find themselves subject to (Heath et al., 2023).

Fourth, we find meaningful empirical relationships between mentorship, corporate culture, and employee retention. Using 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—we examine changes in workplace perceptions and retention patterns. Engineers at the focal firm, which implemented the mentorship program, are 4% less likely to leave following program implementation compared to peers at similar firms. This retention pattern represents substantial savings given replacement costs for technical talent (Benson et al., 2019). Changes in workplace perceptions appear connected to this outcome: both mentored and non-mentored engineers report improved satisfaction, career opportunities, and stronger cultural norms relative to control groups of (i) non-engineering employees at the focal firm and (ii) engineers at peer firms. These culture patterns are intriguing, given prior work showing how persistent devaluation and lack of sponsorship often drive members of underrepresented groups to exit STEM roles (Spoon et al., 2023).

Finally, we 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 from conservative to stock-based ones for superstar firms, and even our most conservative estimate suggests direct benefits of approximately $5 million. We also document potential indirect benefits through improved retention and corporate culture, though

these effects are more challenging to quantify precisely and their magnitude is much larger. Since the mentoring program had modest implementation costs—primarily staff time for coordination—this indicates the program's ROI is high. Therefore, we evaluate mentor quality to see if that is a constraint. 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. Our analysis of mentor quality reveals that even mentees with below-median quality mentors experience significant productivity gains, a pattern suggesting the program's effectiveness scales across mentor quality levels.

Thus, this raises the question of why high-performing firms have not widely adopted such seemingly valuable programs. Our interviews suggest information 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. These 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.

The remainder of the paper is organized as follows: Section 2 describes the related literature and discusses our contributions. Section 3 describes the institutional setting and the design of the mentorship program. Section 4 describes our data and empirical approach. Section 5 presents the main results of the mentoring program's direct and indirect influence. Section 6 provides a cost-benefit analysis of the mentoring program. Section 7 discusses the implications of our findings for scalability and public policy. Section 8 concludes with directions for future research. There are also internet appendices detailing the survey instrument and our robustness and falsification tests.

2 Related Literature and Contributions

This interdisciplinary study integrates insights from personnel and organizational economics and innovation economics. First, our work contributes to personnel economics by examining how organizational practices beyond compensation relate to innovative productivity (Lazear and Shaw, 2007; Bloom et al., 2019c; Hoffman and Stanton, 2024). While previous work suggests that human capital investments enhance worker productivity, less is known about their relation with creativity or innovation outputs. Some even argue that it is impossible to force creativity into a rigid, process-driven framework.

We address this surprising gap in the literature by quantifying how mentoring (Porter and Serra, 2020; Sandvik et al., 2025) relates to the quantity, quality, and diversity of inventive ideas generated within a firm. By measuring the direct productivity gains from mentoring alongside spillover effects, we provide evidence of how investments in human capital translate into measurable innovation outputs. Our findings suggest that organizational frictions create substantial inefficiencies in the innovation process that can be mitigated through structured interventions. This is consistent with evidence that suggests variation in management practices explains meaningful productivity differences within firms (Benson et al., 2019; Bloom et al., 2019a). It is also consistent with research showing that corporate culture is a key driver of long-term value creation (Guiso et al., 2015; Li et al., 2020; Graham et al., 2022a,b, 2023; Gorton et al., 2022; Grennan and Li, 2023).

Second, this study contributes to the literature on innovation economics, particularly the emerging work on pre-patent dynamics (Serrano, 2010; Hegde and Sampat, 2009; Pairolero et al., 2022; Aneja et al., 2024). By analyzing the full innovation pipeline, including ideas rejected before even reaching the patent application stage, we meaningfully advance the literature, as no prior research has linked such detailed idea-level data to patenting and administrative personnel records. While prior research has examined idea generation and selection from academic publishing to patents (Fleming et al., 2019; Masclans et al., 2024), our setting differs substantially: we analyze the internal screening of commercial innovations

at a superstar technology firm, where economic incentives, organizational structures, and innovation objectives fundamentally differ from those in academic research. Interestingly, our data reveals that patent attorneys' internal screening processes filter out many innovative ideas before they reach patent application status, consistent with patenting being more process-driven and, therefore, capable of being influenced by targeted interventions.

Our research also contributes to the growing literature on team-based innovation and knowledge production. Jones (2009) demonstrates that scientific and technological advances increasingly require larger teams with diverse expertise due to the expanding knowledge frontier and growing burden of knowledge. Building on this work, Wuchty et al. (2007) show that team-based approaches typically generate more impactful innovations than solo efforts. Our mentoring intervention addresses a key friction in this process by facilitating knowledge transfer between experienced and novice innovators in a team setting, thereby reducing the costs of networking and collaboration. By examining how structured mentoring influences submission rates, team formation, and patenting outcomes, we provide empirical evidence on specific mechanisms to enhance team productivity in innovation-intensive environments. This connects to broader work on team complementarities (Hamilton et al., 2003) and the role of diversity in strengthening innovative output (Hong and Page, 2004; Bernstein et al., 2022; Heath et al., 2023).

3 Firm Setting

The mentorship program that we study occurred within a firm that develops and manufactures high-technology services and products. For example, services include automation and cloud management, data storage and collaboration, networking, and security capabilities, and products include business and personal software, hardware, and telecommunications equipment. The firm employs over 50,000 individuals, many of whom are engineers and software developers. The firm is headquartered and incorporated in the United States. It has over 20 offices in the United States and operates over 100 offices in more than 80 countries

outside the United States. The firm offers employees flexible work arrangements, including in-office, work-from-home, and hybrid options.

The firm is publicly traded on Nasdaq, and it is a part of the S&P 500. The firm relies heavily on research and development. As such, investments in human capital are an especially important aspect of the firm's business strategy. To this end, the firm instituted a mentorship program aimed at helping its employees, specifically its engineers, gain the skills and knowledge necessary to develop and patent their inventions successfully. While the mentorship program was open to all employees, those from underrepresented demographic groups (e.g., women and racial/ethnic minorities) were especially encouraged to participate. The advertisements for the mentorship program were sent to all employees involved or potentially involved in the innovation process, including those in upper management positions, so they could encourage their direct reports to apply. In this regard, two of the explicit goals of the mentorship program were (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."

The mentorship program occurred in seven cohorts from Q4 of 2020 to Q3 of 2023. For the initial cohorts, the program lasted approximately two months, and for later cohorts, the program lasted for four months. The content and curriculum remained largely the same across cohorts. As this program was voluntary and not part of regular work assignments, participants were credited with the time required to complete the program, to reduce the chance that they fell behind on their regular tasks and responsibilities. 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. Employees who were interested in participating in the program had to submit an application, wherein they had to acknowledge their commitment to engage with the program for several hours per week for the duration of the program. There was limited space within each cohort, with the main constraint on cohort size being the number of employees willing and able to mentor the program participants. Other constraints included

a lack of overlap in geographical locations (e.g., if no one was located in the preferred time zones of a particular applicant) and a lack of overlap in common invention interests (e.g., if only one person expressed interest in a particular research topic). The topics of common interest included (but were not limited to) the following: data analytics (e.g., machine learning); A.I. in security; A.I. in customer experience; and automation and cloud services. The mentors were all experienced inventors with proven track records of invention development and patenting success. They were instructed to teach research principles, ensure a fair division of labor among the team, connect mentees with resources, and coach mentees on writing effective invention disclosures.

At the beginning of the mentorship program, participants within a cohort were placed into small groups of approximately four to five mentees. Each group was randomly assigned one or two mentors. The program's curriculum had mentors guide their group of mentees through the invention development and patenting process. Figure A.1 illustrates the typical timeline by which these events took place, with teams beginning to work on R&D solutions to their chosen problem statement about a month after their first interactions and then with teams starting to draft up their solution disclosures at the two-month mark, with submissions occurring shortly after the three-month mark. Each team's invention submission culminated in a showcase event, during which the mentees presented their work to senior technical talent teams. In most cases, groups were able to successfully generate and submit an invention disclosure within four months of starting the mentorship program. While these invention submissions constitute successful R&D outcomes in and of themselves, the main purpose behind the mentorship program was to cultivate within the participants the skills and knowledge necessary to develop and patent their own inventions outside of the mentorship program setting. In the following sections, we quantify the direct benefits of the mentorship program such as invention submissions, and we estimate other, possibly unintended, benefits of the program, such as spillovers from program participants to those who did not participate in the program.

4 Data Sources, Innovator Characteristics, and Empirical Specifications

4.1 Data Sources

Estimating the relation between mentorship and the quantity and quality of ideas disclosed as invention disclosures and patent applications requires a detailed panel of innovator-level data. The firm provided us with this data, which covers all invention disclosures submitted by the firm's employees from 1998–2023. In total, 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 of the individuals, including their gender, age, marital status, race/ethnicity, job role, and the country in which 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 the status (as of March 2024) of each disclosure. We categorize disclosures based on their status after submission through the internal review process. Status 2 disclosures are those with any of the following statuses: Approved - Patent Application; Prior Art Search Requested; RFC (return for comments). Status 1 disclosures are those with any of the following statuses: Declined; Auto Declined; Defensive Publication Candidate; Approved - Defensive Publication; Approved - Trade Secret; Deleted; Merged.[8] 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 a filed or an approved patent. Submissions can also be returned for comments (RFC) by the review

[8]For our main analyses, we do not group "Approved - Defensive Publication" and "Approved - Trade Secret" in with the other Status 2 disclosures because defensive publications and trade secrets are generally less valuable, relative to patent applications. In addition, we omit disclosures with the statuses of "Submitted" or "Filing Candidate" from the Status 1 group because those disclosures have not yet been reviewed. They do, however, get included in the *Any Status* group. In Appendix B, we provide evidence from a series of robustness tests that support the conclusion that our inferences are not sensitive to these classification decisions.

committee.

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

4.2 Innovator Characteristics

In Table 1, we report the demographic characteristics of the participants in each cohort of the mentorship program. We report innovator characteristics separately for each of the seven cohorts. We also report the demographic characteristics of the 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, 15% were Latinx, 13% were White, and 16% were URMs.[9] 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).[10]

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

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

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 of mentored innovators are quite similar to those of non-participants. This indicates that innovators from the two groups exhibited similar levels of productivity at baseline, suggesting 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 the mentorship program and the participants' invention outcomes, we compare their quarterly disclosure outcomes before and after they begin the program. We use data from Q1 of 2020 onwards, as the first cohort of the mentorship program was in Q4 of 2020 and as this ensures all of our data is in the during-COVID or post-COVID period, mitigating concerns that using pre-COVID data may bias estimates (though all our inferences hold if we include pre-COVID data). We use a multi-pronged empirical strategy to mitigate bias in our estimates and ensure the robustness of our statistical inferences. Specifically, we use both a two-way fixed effects (TWFE) estimator and a stacked difference-in-differences estimator. Further, we also explore the robustness of our results using different control groups as a basis for comparison, including never-mentored innovators, not-yet-mentored innovators, and propensity-score-matched innovators.

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. To implement this estimator, we regress an innovator's quarterly disclosure outcome values on the binary indicator *Mentored*, which equals one in the quarter that their mentorship program begins and each subsequent quarter, and zero otherwise. 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). We estimate the following model using ordinary least squares:

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

where $y_{i,t}$ is the number of disclosure submissions of a particular status that innovator i has in quarter t, λ_i is an innovator fixed effect, γ_t is a year-quarter fixed effect, and ε_i is an idiosyncratic error. We double cluster standard errors by innovator and year-quarter.[11]

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 (Barrios, 2021). This bias can be especially problematic if treatment effects are dynamic and evolve over time (Baker et al., 2022). To circumvent these concerns, we use a stacked difference-in-differences estimator, which avoids the potential aforementioned bias (Cengiz et al., 2019). While Gormley and Matsa (2011) first applied stacked difference-in-differences estimation in finance research, the formal econometric properties of this approach were developed and established more

[11] 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.

recently.

To implement the stacked difference-in-differences 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 the outcomes of a set of "clean" control innovators. For each cohort of treated innovators, we use two control groups of innovators: (1) those that never participate in the mentorship program (i.e., never-mentored innovators), and (2) those that first participate in the mentorship program in a later cohort (i.e., not-yet-mentored innovators). One benefit of using not-yet-mentored innovators (i.e., eventually treated subjects) as a control group is that it allows us to exploit variation in the timing of mentorship program participation, rather than variation in the decision to participate at all (Deshpande and Li, 2019). In addition, 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. Furthermore, 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.

We aggregate all cohort-specific data—i.e., the innovator-quarter observations for 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. This allows us to estimate the following event-study regression using ordinary least squares (OLS):

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

where, $y_{i,t,g}$ is the disclosure outcome variable of interest 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. Finally, standard errors are double clustered 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).[12] This estimation approach is analogous to estimating separate difference-in-differences treatment effects for each mentorship cohort group, using never-mentored or not-yet-mentored innovators as controls, and then taking the average of the treatment effects across all seven cohorts.

4.3.3 Common Trends Tests

Before we tabulate and discuss our main findings, we first analyze the evidence to see if the common trends assumption, which is necessary for identification, appears to hold. While it is not possible to definitively establish if the assumption is met, these tests can provide suggestive evidence in support.

Our difference-in-differences estimations rely on the assumption that the disclosure of inventive ideas by mentored and non-mentored innovators would follow a common parallel trend in the absence of the mentorship program. 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 (i.e., those who are eventually mentored in a later period).

Figure 1a plots quarterly disclosure submission patterns around the mentorship initiation quarter. The points reflect the relative differences in disclosure submissions between mentored innovators and non-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, which addresses potential limitations of the TWFE approach by allowing for more flexible dynamics and reducing potential bias in settings with staggered program adoption. In Figure 1b, we repeat this procedure using an alternative dependent

[12] 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 disclosures outcomes that occur across time.

variable that scales each disclosure submission by the size of the team that produces the invention. In both figures, we find evidence of common pre-trends in disclosure activity before the mentorship program, followed up significant increases in disclosure submissions post-mentorship.

Following the same procedure, Appendix Figures A.3a and A.3b present corresponding analyses for Status 2 and Status 1 disclosures, respectively. Across all specifications, whether examining combined or separate disclosure types, the figures show consistent patterns in submission rates between mentored and non-mentored groups before the implementation of the mentorship program. This is consistent with no diverging pre-trends in disclosure outcomes. 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 0.321 to 0.960. These patterns suggest comparable submission trajectories between future mentored innovators and non-mentored innovators in the quarters preceding program participation.

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 participant characteristics such as gender and ethnicity (Section 5.1). Second, we analyze survey evidence to explore potential mechanisms associated with mentorship and innovative behavior, examining patterns in knowledge sharing, network formation, and confidence building (Section 5.2). Third, we examine potential indirect spillovers along four dimensions: the relation between mentored and non-mentored collaborators' productivity, patterns in team formation and demographic composition, workplace perceptions and corporate culture, and talent retention rates (Section 5.3). Finally, we address possible alternative explanations for the observed relation between mentorship

and innovation output (Section 5.4). Throughout our analysis, we acknowledge potential confounding factors and conduct multiple specification tests to examine the robustness of the observed relations.

5.1 The Mentorship Program and Innovator Productivity

We begin discussing our results 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 to alternative sampling restrictions and econometric specifications. We also consider how heterogeneity across participant demographic characteristics correlates with the propensity to submit invention disclosures.

5.1.1 Mentorship Program Participation and Overall Productivity

We report the relation between mentorship and various disclosure submission outcomes in Table 2. In Panel A, the dependent variable in Columns (1)–(6) is the number of unique submissions of any status that an innovator submitted in a given quarter. In the odd-numbered columns, we use the TWFE estimator, which includes innovator and time fixed effects. We use the stacked difference-in-differences estimator in the even-numbered columns. We use both not-yet-mentored innovators and never-mentored innovators as the basis for comparison.[13] The coefficient estimates in Columns (1) and (2) suggest that mentorship is associated with an increase in disclosure submissions of 0.25 each quarter, which translates to approximately one additional disclosure per innovator each 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

[13]The sample size increases when using the stacked difference-in-differences estimator because we form control groups by drawing from the sample of control innovators with replacement.

reduce the concern that our main findings are driven by violations of the Stable Unit Value Treatment Assumption (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 number of Status 2 submissions that an innovator submitted in a given quarter (i.e., the high-quality disclosures that become approved patent applications). We continue to estimate a positive and statistically significant estimate on *Mentored*, with the economic significance suggesting that, after mentorship, each mentored innovator produces approximately 0.2 additional inventions that progress well into the disclosure review process each year. These findings echo the patterns observed in Figure 1 and Figures 1b, which show visual evidence of a significant, positive relation between mentorship and innovator productivity in the quarter of the mentorship program and the quarters following the program.

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.

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 caused innovators to realize gains in disclosure in the near-term. What really matters, from both the firm's perspective and that of the innovators, is that participants finished the program with the skills, confidence, and knowledge necessary to successfully develop and patent their inventions independently.

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. 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 Status 2 disclosure counts in Columns (7)–(8) are also positive and significant at the 1% level. The economic magnitude is 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. Thus, these results suggest that mentored innovators gained the skills and knowledge necessary to independently develop and disclose inventive ideas that would lead to 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), and when using different disclosure status classifications (e.g., considering only submissions that have become approved patents). We also discuss the results of several other tests that bolster our main inferences.

5.1.3 Heterogeneity Tests for Mentorship and Innovation Productivity

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 major goal from the firm leadership's perspective was to assist employees from demographic groups that have historically been underrepresented in the innovation and patenting process (Bol et al., 2023). 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 lower confidence or less knowledge of the patenting process than their peers.

In Table 4, we report the results from these heterogeneity tests, wherein we continue to employ the TWFE estimator in the odd-numbered columns and the stacked difference-in-differences estimator in the even-numbered columns. In Panel A, we interact *Mentored* with the binary indicator *Female*, which allows us to estimate whether female mentees realize more or less mentorship program benefits than do male mentees. The coefficients on *Mentored* $\times$ *Female* are positive in all eight columns, but they are never statistically significant. Similarly, in Panel B, where we interact *Mentored* with the binary indicator *URM*, we estimate positive interaction coefficients in all columns, but none of the estimates are significant. In Panel C, we consider gender and URM status in tandem by interacting *Mentored* with the binary indicator *Female/URM*, which equals one if the innovator is either female or a URM, and zero otherwise. Again, we estimate positive coefficients on the interaction term in all eight

columns, and the estimates are statistically significant at the 5% or 10% levels in seven of the eight specifications.[14]

Taken together, the coefficient estimates on the interaction terms in Table 4 provide some evidence (albeit noisy evidence) that female and URM innovators benefited more from the mentorship program than did male and non-URM innovators. Importantly, though, the average treatment effects, which include male and non-URM innovators, continue to be positive and statistically significant, as indicated by the positive and statistically significant coefficients on *Mentored* in all columns and panels of Table 4. This suggests that all participants benefited from mentorship program participation regardless of their demographic characteristics.[15]

Next, we explore whether heterogeneity in mentor quality is meaningfully associated with 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 from 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

[14]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.

[15]In Table A.1, we consider whether more pronounced productivity gains are present based on the country in which the innovator works (Panel A) or based on whether the innovator participated in one of the first three cohorts or the latter cohorts (Panel B). The coefficient estimates on the interaction term in Panel A are never statistically significant. The coefficient estimates on the interaction term in Panel B are not significant in Columns (1)–(6), but they are negative and significant at the 10% level in Columns (7) and (8), which focus on Status 2 disclosure counts. This suggests that mentees in early and late cohorts realized similar productivity benefits due to 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 geographic or timing differences matter.

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 are reported in Panel A of Table 5. 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 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. Then the gains increase for those with mentors in the third fixed effects quartile, and the gains are greatest for those with the best mentors (as proxied by the mentor fixed effects).[16]

These findings have important implications for the firm's scale-ability of the mentorship program. 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 their innovative ideas.[17]

5.2 Exploring the Mechanisms via Survey Response Data

To better understand the mechanisms through which mentorship works and the barriers innovators face when trying to patent their inventive ideas, we surveyed innovators involved

[16] We report the demographic characteristics of the mentors, split by mentor fixed effect quartiles, in Table A.2.

[17] We also consider heterogeneity based on alternative mentor characteristics in Table A.3. In general, we find little evidence of significant differences, suggesting that mentors were, on average, quite effective regardless of their gender, job status (i.e., "top engineer"), or prior involvement in the mentorship program. In Table A.4, we show that mentees across the full age distribution benefit from mentorship, which further speaks to the scale-ability of the program, as both old and young innovators stand to gain from program participation.

in the mentorship program. This survey was part of a larger, multi-firm interview and survey process conducted by Chien and Grennan (2024). The study began with interviews of patent professionals, mostly general counsels or heads of IP. This particular survey was a small, demographically representative sample of engineers and technical staff in positions deemed by the focal firm as capable of contributing to IP development but who had not necessarily already developed a patent. For our purposes, the firm oversampled from engineers who had participated in the program. Overall, the response rate was about 17%, which is comparable to other corporate surveys (Graham et al., 2005). Consistent with the oversampling, Panel A of Table 6 shows that 43% of the survey respondents had participated in the mentorship program.

Panel B of Table 6 begins by summarizing the perceived experiences of those without a mentor. It shows that the engineers who did not participate mostly did not participate because they did not know about the program. However, women are significantly more likely than men to believe that patent mentorship would not benefit them (11% vs. 3%). This works against the sample selection, as women are more likely to participate.[18] We observe statistically indistinguishable rates of inability to participate across demographic groups. This again supports our identification approach that the capacity constraints on the mentorship program introduced quasi-random variation in participation timing.

Panel C of Table 6 summarizes the perceived experiences of those with a mentor. The frequency of meetings with their mentors after the formal mentorship program ends is consistent across demographic groups, with mentees suggesting their mentor provides advice, suggestions, or support around twice a month. When looking at the perceived benefits, mentored innovators reported benefits from the mentorship received. Mentees indicate a desire to work on more inventions, greater confidence in their ability to incorporate their mentor's tips into their work process, a greater likelihood of having the social connections necessary to get their invention ideas accepted by patent professionals, and having more social connections

[18]These survey findings relate to recent work on the association between self-selection and treatment effects in human capital development programs (Sandvik et al., 2025).

from introductions to innovators that they anticipate working with. Females' perceptions of mentorship are more positive. Females indicate a statistically significant increase in job satisfaction of 1.4 units, compared to 1.0 units for their male counterparts. Females also perceive that they have benefited more in terms of learning about patenting, at 1.5 units compared to 0.9 units for their male counterparts. Overall, the data suggests that mentorship, particularly when formalized, could be a key mechanism for fostering greater inclusivity in the invention process among engineers.

5.3 Spillovers from the Mentorship Program

We find convincing evidence that the mentorship program had 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. Additionally, the program may influence perceptions about the benefits of diverse innovator teams, potentially leading to changes in subsequent team formation and composition. 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, collaboration patterns, 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 innovator 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 difference-in-differences estimator. We report the results in Table 7. 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.5, 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* × *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 Innovator Team Composition

Next, we explore whether the mentorship program is related to the demographic composition of mentored innovators' subsequent inventive teams. We may expect to see increases in team diversity due to the mentorship program for multiple reasons. First, the mentorship program

may have exposed innovators to 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," with an explicit goal being to "diversify invention submissions and subsequent patentable inventions from all parts of the engineering talent pool." Our empirical analysis investigates whether these program features correspond to 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 difference-in-differences 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 8. 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.[19] This HHI metric becomes the dependent variable in Columns (7) and (8) of Table 8. 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

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

collaborator teams for a given innovator-quarter.[20] This HHI metric becomes the dependent variable in Columns (9) and (10) of Table 8.

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 and gender-by-racial/ethnic diversity of an innovator's 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.

This is 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.

5.3.3 Perceptions of the Workplace

The survey results, discussed in Section 5.2, provide useful information about participants' perceptions of the mentorship program, but it is hard to determine if there is appreciable bias given the anonymous, one-shot nature of the survey. Therefore, to explore the work-

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

place implications of the program on a broader sample of employees, we turn to data from crowd-sourced employee reviews on Glassdoor, a career intelligence website that attempts to provide transparency about jobs, salaries, and companies. Specifically, we consider the relation between the mentorship program and employees' perceptions of their workplace environment. 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 impact their sentiment towards the firm.

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 (N = 294,277) during the period of study. The industry peer group is defined based on market capitalization, R&D expenditures, and industrial classification. 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 November 2023, the last month in our dataset (N = 242,066).

First, we estimate the firm's overall star rating change on Glassdoor after the implementation of the mentorship program. Restricting to the 12,019 reviews left by employees of the focal firm, we regress the overall firm rating on the indicator *Post*, which equals one in and after October 2020. The results in Column (1) of Table 9 show that the ratings of the focal firm are significantly greater after the start of the mentorship program. However, this result could be due to general positive trends in employee ratings independent of the mentorship program. To account for this and other possible confounding factors (e.g., work-from-home initiatives increasing employee satisfaction), we consider whether the correlation is more pronounced among engineers, who were the main program participants, than

among the other firm employees. To do this, we interact the binary indicator *Engineer* with *Post*. The positive and significant coefficient on *Post* × *Engineer* in Column (2) shows that engineers experience a significantly larger increase in ratings than non-engineers. The fact that engineers are more likely to perceive the workplace as better than their non-engineering peers suggests that the increase in ratings is likely due, in part, to the mentorship program, which impacted engineers much more than the firm's other employees. We spoke to those originally administering the patent mentorship program to understand if there were other confounds present in the firm that may lead engineers to have a different perception of the firm for another reason while the mentorship program was first being implemented. The program administrators indicated that there were no other meaningful changes in the workplace environment for engineers at the time.

One potential concern is that the engineers' improved perceptions of the workplace are driven by broader macro-trends in employee satisfaction, possibly trends that impact engineers differently than non-engineers. Therefore, we also consider whether we observe similar patterns in sentiment among the employees of the focal firm's industry peers. In Columns (3) and (4) of Table 9, we replicate the analyses of Columns (1) and (2) using only the overall firm ratings of current employees from the focal firm's industry peers. In Column (3), we find evidence that peer firms also realize an increase in ratings after the beginning of the mentorship program in the focal firm, suggesting that ratings may tend to drift upward over time. However, in Column (4) we find that the coefficient on *Post* × *Engineer* is small (about a tenth the size of the coefficient in Column (2)) and 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. Specifically, we show that engineers' perceptions of the focal firm improve more than do the perceptions of engineers in peer firms by estimating the triple-interaction model in Column (5), wherein we find a positive and significant coefficient on *Post* × *Engineer* × *Focal Firm*. We estimate similar effects when

we use employees' ratings of the firms' culture and values.[21]

5.3.4 Innovator Departures

To finish, we explore potential associations between the timing of the mentorship program implementation and patterns in innovator turnover, focusing on differences between engineers and non-engineers, and 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. 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, *Outflows*, 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 use, as an alternative dependent variable, the inverse hyperbolic sine transformation of *Outflows*, which accounts for skewness in the data while accommodating year-months in which no turnover is observed within a particular group of individuals. 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 then estimate OLS regressions that mimic those from the previous workplace perception analysis (in Table 9). Specifically, we report the results in Table 10. The estimate displayed in Column (1) shows that the 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

[21] If we re-estimate these models for the other ratings measures provided in Glassdoor, we find that engineers in the focal firm, relative to non-engineers, provided increased ratings of their career opportunities, their compensation and benefits, and the firm's culture and values. 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.

(i.e., they experience even *less* turnover), relative to employees in the other occupational roles. These findings are statistically significant both when using the count *Outflows* measure in Panel A, and when using the transformed measure in Panel B. The transformed measure provides a tractable economic interpretation of the result, with engineers in the focal firm experiencing about 4% ($= e^{-0.040} - 1$) less attrition after the beginning of the mentorship program, relative to non-engineers. 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 in Panel A, but it is significant in Panel B. However, Column (5) of both panels provides evidence that the reduced attrition among engineers is greater in the focal firm, relative to its peer firms.[22]

Overall, the empirical analyses of engineers' departures reveal a consistent pattern across multiple specifications, indicating a negative association between the implementation of the mentorship program and engineer departures from the focal firm. While inferring a direct causal relation is challenging, the robustness of this relation across various model specifications and derived from a quasi-experimental design supports a link between the presence of the mentorship program and lower turnover rates among engineers. This departure pattern, which was observed consistently in our analyses, points to a potential retention benefit and a workplace perception benefit coinciding with the mentorship initiative.

[22] We conduct falsification tests for the workplace perceptions and employee departure tests in Appendix B, which bolster the inferences we draw from the results in Table 9 and Table 10. We also discuss the robustness of the results when different iterations of *Post* are used.

5.4 Addressing Possible Alternative Explanations

One alternative explanation for the observed relationship between mentorship and innovator productivity relates to team composition patterns. Specifically, existing innovation teams might have simply incorporated newly mentored individuals into their ongoing projects, rather than mentorship leading to genuinely new innovative activities. Under this scenario, mentored individuals would show higher submission rates following program participation, as they join more project teams. However, these increased individual-level submissions would not translate into additional firm-wide innovation output since the underlying projects would have occurred regardless of the mentorship program. This distinction between individual and organizational outcomes is important for understanding the program's broader impact on innovation productivity.

However, the patterns we observe in disclosure submissions suggest genuine increases in firm-wide innovation productivity, rather than merely reflecting larger team sizes on existing projects. For instance, before the mentorship program, the average innovation team size for accepted patent applications that eventually-mentored innovators were a part of was 4.62. Whereas, during the mentorship program, it was 4.59, and after the program it was 4.31. None of these team size values are significantly different from each other (p-values range from 0.317 to 0.920). This provides strong evidence against the team-size-boosting concern for accepted patent application disclosures.[23]

To further mitigate the concern about inorganic increases in innovation team size, we use data from the U.S. PTO to compare patent grant team sizes in the focal firm to those in peer firms. In doing so, we show that the team size dynamics within the focal firm mirror those of peer firms, suggesting that the mentorship program did not prompt artificially larger patent

[23]It does appear to be the case that individuals who participated in the mentorship program tend to be on larger teams in general (both before, during, and after the program). Among innovator teams with no individuals who participate in the mentorship program (past, present, or future), the team size for accepted patent applications is 3.08. Importantly, though, the team sizes of innovator teams with no mentored individuals has increased since the start of the mentorship program, to 3.44. This suggests that innovation team sizes do increase over time, regardless of mentorship associations.

grant team sizes. These results are reported in Appendix Table A.6, where we use patent data from the U.S. PTO from 2018 to 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 OLS regressions to analyze how these patent team size measures change in relation to the binary indicator *Post*, which equals one from when the first cohort of the mentorship program began in the focal firm and onward (in Panel A).[24] Some models include interactions between *Post* and the binary indicator *Focal Firm*. The positive and significant coefficients on *Post* in Columns (1), (2), (4), and (5) suggest that patent grant team sizes are larger in the post-period than they were in the pre-period for both the focal firm and its peers, and the small, insignificant estimates on *Post* × *Focal Firm* in Columns (3) and (6) suggest that these changes are no different for the focal firm than for peer firms.

As additional evidence of the mentorship program's effectiveness in boosting overall innovative efficiency at the firm, we also estimate the time it takes for mentored innovators to achieve their second disclosure after submitting their first disclosure. Across all innovators who are not part of the mentorship program, and who we observed submitting at least two disclosures, the time-to-second-submission is about 1.25 years (i.e., five quarters). The average value of time-to-second-submission is very similar for individuals who participated in the mentorship program, but who achieved their second submission before the beginning of the program. In contrast, the average value of time-to-second-submission is less than one year (i.e., less than four quarters) among individuals who achieved their second submission after they enrolled in the mentorship program. Furthermore, innovators who had their first submission experience during or after the mentorship program had an average time-to-second-submission of about one quarter. These findings suggest that the mentorship program substantially expedited the idea-development learning curve among mentored

[24] In Panel B (Panel C), *Post* equals one from when the second (third) mentorship program began and onward.

innovators.

Finally, we do not find evidence of meaningful differences in the quality of the disclosure submissions made by teams with mentored innovators, relative to those made by teams with no mentored innovators. Among all disclosure submissions made by teams with no mentored innovators, approximately 65% have the status of "Declined," whereas the declined rate is only 62% for teams with currently or previously mentored innovators. Furthermore, 35% of the submissions from teams with no mentored innovators have the status of "Approved - Defensive Publication/Patent Application/Trade Secret," "RFC," or "Submitted," and 35% of the disclosure submissions from teams with currently or previously mentored innovators do.[25] These patterns in disclosure status indicate that, during and after the mentorship program, mentored innovators were a part of valuable innovation projects.

6 Return on Investment to the Firm

This section develops an analytical framework to evaluate the returns from the mentorship program, building on established methodologies for assessing human capital investments and organizational productivity (Black and Lynch, 1996; Bell et al., 2019; Lazear et al., 2015). Our approach begins with directly quantifiable outcomes to establish a baseline estimate of the program's value, then extends to consider broader organizational benefits that are more challenging to measure precisely. For each component of our analysis, we explicitly state our underlying assumptions and detail our calculation methodology. While acknowledging the inherent complexity in measuring returns to human capital programs, we construct careful estimates of the program's ROI by systematically examining both direct innovation outcomes and indirect organizational benefits.

Our exercise incorporates the many value creation streams described in our empirical tests. Specifically, we begin by analyzing the direct channel through which patent mentorship creates value for the firm. Specifically, we observed an increase of 1.5 invention disclosures

[25]The other 3% of submissions from teams with currently or previously mentored innovators have a status such as "Deleted," "Merged," or "Filing Candidate."

per year per mentored innovator. While internal invention disclosure rejection rates vary by firm, and external U.S. PTO patent application rejection rates vary by art unit and by gender, a blended estimate based on statistics reported in Aneja et al. (2024) and Chien and Grennan (2024) suggest that the overall conversion rate from invention disclosure submitted internally to a patent attorney to a granted patent is 11%. To value these patents, we use a highly conservative approach based solely on actual costs: patent attorney fees ($25,000) plus mandatory renewal fees ($13,460), totaling $38,460 per patent. This approach likely understates patent value but provides a concrete lower bound. Thus, the direct productivity gains are approximately $4.017 million (633 mentees $\times$ 1.5 submissions $\times$ 0.11 conversion $\times$ $38,460 per patent). Given that team size for accepted patents is 3.08 named inventors, an even more conservative lower bound estimate of the direct benefit is $1.354 million. Conversely, if we were to measure value using stock market indicators, as proposed by Kogan et al. (2017), our estimated impact would dwarf this estimate and be more than ten times larger since any movement in this superstar technology firm's stock price implies a much larger value for its patents.

We observe limited but measurable spillover effects through two channels. First, the 520 non-mentored innovators who directly collaborated with mentored participants showed a 0.5 increase in annual submissions, generating an estimated $1.1 million in additional patent value using the same conservative valuation approach. Second, we document increased team diversity post-program. Based on prior literature linking team diversity to patent value, we conservatively estimate a 5% increase in value for patents from diverse teams, adding approximately $0.256 million.[26]

The program likely generates additional value through improved retention and cultural

[26]To value the benefits of this type of diversity, we are implicitly assuming 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-valued 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).

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 in the technology sector. Similarly, survey evidence suggests improved workplace satisfaction and knowledge sharing, which may enhance long-term organizational capital, though we acknowledge the challenges in precisely quantifying these cultural benefits.

For instance, one could estimate the value implications of the mentorship program's link to more efficient recruiting from improved employee perceptions of the workplace and higher retention rates among engineers. 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, requiring extensive training to get new hires up to speed and beyond their (typically) 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 $\times$ 29,212 inventors $\times$ replacement cost (1 $\times$ $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 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 new recruits (29,212 inventors $\times$ 0.19 turnover rate $\times$ $1,159 lower salary per year $\times$ 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., 3 p.p.) of that increase in Tobin's Q to culture improvements would suggest benefits in excess of $2 billion. Yet this seems 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.[27] 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.

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 large labor market gains as well. These calculations, while systematic and grounded in observable data, require several important qualifications. Our patent valuation methodology relies on cost-based measures rather than market prices, which likely understates the true economic value that patents generate for the firm. When measuring collaboration effects, our approach captures only direct working relationships between

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

mentored and non-mentored employees, potentially missing broader ripple effects through extended professional networks. The relationship between the mentorship program and employee retention presents particular analytical challenges, as other organizational changes could have occurred during our study period that contribute to changes in retention patterns. Although our interviews with program administrators indicated no such changes. Finally, while our survey data and Glassdoor analysis reveal positive shifts in workplace culture following the mentorship program implementation, converting these harder-to-quantify improvements in the organizational environment into quantitative economic measures remains methodologically complex.

Nevertheless, even using these conservative assumptions and focusing primarily on direct benefits, the program's returns appear to exceed its modest implementation costs. The program required minimal infrastructure—just two part-time administrators and voluntary participation from mentors whose productivity did not decrease during their mentorship period. While more speculative benefits could increase the ROI substantially, even our most conservative estimates imply that the mentorship program generates meaningful positive returns.

7 Scaling and Public Policy Considerations

Innovation has been a key driver of U.S. economic growth and technological leadership (Hall et al., 2005; Moser, 2013). Yet history is rife with examples of barriers to reaching our full innovation potential.[28] These historical lessons emphasize that addressing organizational frictions, which limit participation in patenting, represents not only an equity concern but an economic imperative for maintaining technological competitiveness. Despite progress in broadening participation in the patent system, significant gaps persist. While the literature has focused on various painpoints ranging from socioeconomic status (Bell et al., 2019) to

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

geographic talent allocation (Hsieh et al., 2019) to assertiveness after rejections by patent examiners (Aneja et al., 2024), there is surprisingly little evidence examining what firms can do specifically.

Yet looking within firms is a critical imperative now. To provide context, the USPTO's empirical examination of inclusive patenting began during the 1990 bicentennial of the Patent Act, with a hand-coded analysis of gender representation among biotech patent inventors (PTO, 1999). While women represented only 2% of inventors in 1977, this increased to 10% by 1996, but recent analyses show this progress has plateaued (Toole et al., 2020). Subsequently, Chien (2024) documented substantial heterogeneity across firms in the gap between women in technical roles and women inventors, with some firms showing gaps exceeding 25 percentage points while others approached parity. This firm-level heterogeneity mirrors broader productivity differences across seemingly similar firms (Syverson, 2004; Gibbons and Henderson, 2012; Backus, 2020) and suggests the need to examine internal firm practices.

In fact, the stalled progress for women inventors prompted researchers to investigate factors affecting the conversion of innovative ideas into patents. Chien and Grennan (2024) surveyed thousands of engineers and discovered that management practices and corporate culture are the most important barriers to submitting inventive ideas, more so than personal characteristics or the submission process itself. For example, women engineers were significantly less likely to report managers sharing information critical for developing patentable inventions. Then, beginning in 2021, over 40 major patent-filing firms signed the "Diversity Pledge," committing to take actionable and measurable steps to increase diversity in innovation (Harrison, 2021). Such firm-focused initiatives recognize that reducing barriers to innovation participation has implications beyond firm productivity to aggregate economic growth and inclusion.

Methodologically, our research contributes to the experimental economics literature by implementing and evaluating a specific intervention, that of patent mentorship, that addresses these frictions. A growing concern in this literature pertains to external validity,

specifically how accurately experimental estimates predict the impact of subsequent interventions (Muralidharan and Niehaus, 2017; Al-Ubaydli et al., 2017, 2020). Given that our approach connects directly to work on the "voltage effect," we consider whether this program could maintain effectiveness at scale by systematically addressing five threats 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 effects 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 utilized the rich demographic data available to understand potential differences in 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.[29] One interesting finding from our study is that the white male non-mentored peers appeared to benefit more than the underrepresented peers. This contrasts with the underrepresented mentors benefitting 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 effectiveness of nudges at scale (DellaVigna and Linos, 2022).

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

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.

Finally, it is important to acknowledge the limitations of our study. Namely, examining a superstar technology firm provides immense scale 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. 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. Future research should examine how similar mentoring approaches perform in firms with different resource constraints, organizational cultures, and worker skill distributions to understand better when and where such interventions remain effective.

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 1.3 additional invention disclosures annually compared to control groups, with gains persisting well beyond the pro-

gram 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 confidence in identifying patentable ideas. 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 employee satisfaction ratings and engineering talent retention relative to multiple control groups. These patterns suggest the mentorship program's benefits ripple throughout the organization's culture.

We synthesize our findings by presenting detailed estimates of the mentorship program's ROI. Direct innovation benefits 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 (List, 2022), by considering common threats in field experiments and 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 both innovation efficiency and inclusion challenges within technology firms.

While our single-firm setting provides 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 particularly 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, our findings indicate that well-designed mentorship programs can meaningfully enhance innovation productivity while improving organizational knowledge flows. For leaders seeking to develop their innovation capabilities and policymakers aiming to increase returns on R&D investment, structured 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.

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* .

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.

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.

Barrios, John Manuel. 2021. Staggeringly problematic: A primer on staggered did for accounting researchers. *Available at SSRN 3794859* .

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.

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.

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, 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.

Bohren, J. Aislinn, Alex Imas, Michael Rosenberg. 2019. The dynamics of discrimination: Theory and evidence. *American Economic Review* **109**(10) 3395–3436.

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

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.

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

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, Ben Zipperer. 2019. The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics* **134**(3) 1405–1454.

Chien, Colleen. 2014. Startups and patent trolls. *Stanford Technology Law Review* **17** 451–506.

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* .

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

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

Davis, Jonathan M.V., Jonathan Guryan, Kelly Hallberg, Jens Ludwig. 2017. The economics of scale-up. Working Paper 23925, National Bureau of Economic Research. doi:10.3386/w23925. URL http://www.nber.org/papers/w23925.

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

Deshpande, Manasi, Yue Li. 2019. Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy* **11**(4) 213–248.

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.

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.

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. URL http://www.nber.org/papers/w33387.

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

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.

Gormley, Todd A., David A. Matsa. 2011. Growing out of trouble? corporate responses to liability risk. *The Review of Financial Studies* **24**(8) 2781–2821.

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. doi:10.1111/jacf.12582.

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. doi:10.1111/jacf.12528.

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

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

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

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.

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.

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.

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.

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.

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

Manso, Gustavo. 2011. Motivating innovation. *Journal of Finance* **66**(5) 1823–1860.

Masclans, Roger, Sharique Hasan, Wesley M Cohen. 2024. Measuring the commercial potential of science. Working Paper 32262, National Bureau of Economic Research. doi:10.3386/w32262.

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. 2013. Patents and innovation: Evidence from economic history. *Journal of Economic Perspectives* **27**(1) 23–44.

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) 103–24. doi:10.1257/jep.31.4.103. URL https://www.aeaweb.org/articles?id=10.1257/jep.31.4.103.

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

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* .

Porter, Catherine, Danila Serra. 2020. Gender differences in the choice of major: The importance of female role models. *American Economic Journal: Applied Economics* **12**(3) 226–54.

PTO, U.S. 1999. Buttons to biotech - u.s. patenting by women, 1977 to 1996. *Technology Assessment and Forecast Program Report* .

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.

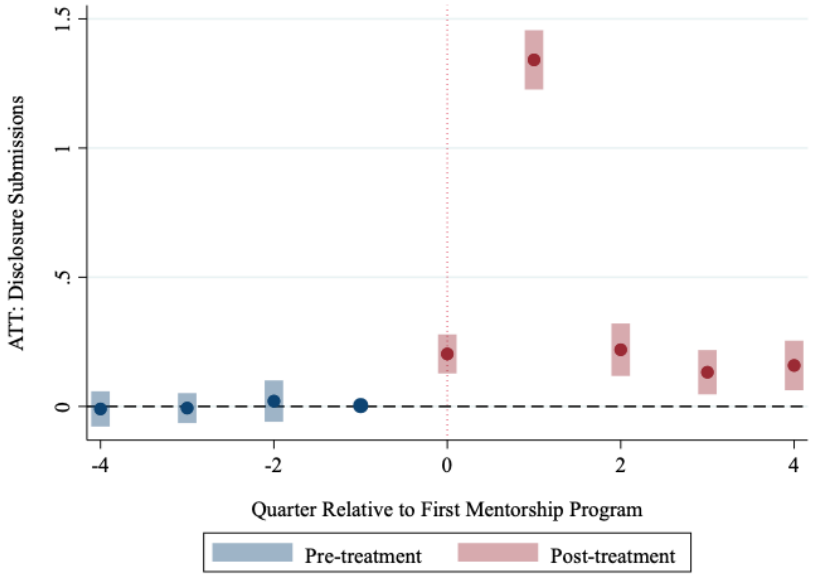
Syverson, Chad. 2004. Market structure and productivity: A concrete example. *Journal of Political Economy* **112**(6) 1181–1222. doi:10.1086/424743.

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* .

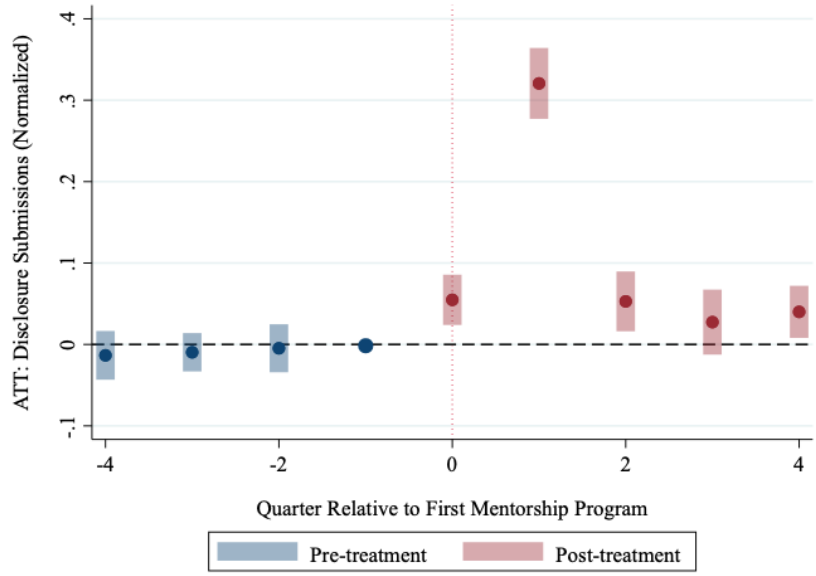
Wuchty, Stefan, Benjamin Jones, Brian Uzzi. 2007. The increasing dominance of teams in the production of knowledge. *Science* **316** 1036—-1039.

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. 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.

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 Step	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)			
Step 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 Program Participation 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 an analysis of the relationship between participation in the mentorship program and innovator productivity, as measured by invention disclosure form (IDF) submissions. The data comprises a panel of innovator-quarter submissions. In Columns (1)–(6), the dependent variable is the total number of disclosure submissions of any status that an innovator submitted in a given quarter. The dependent variable in Columns (7)–(8) is the number of submissions at Status 2. In Panel B, the submissions are first scaled by team size before they are aggregated 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 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 3: Mentorship Program Participation 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 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, but removes the observations of mentored workers in the quarter of first mentorship and the first quarter after mentorship. In the stacked difference-in-differences specifications, control group observations are removed if they occur in the quarter of mentorship, for the focal cohort, and the first quarter after mentorship. 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. In Panel B, the submissions are first scaled by team size before they are aggregated 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 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 4: Heterogeneity in the Association Between Mentorship and Productivity

Panel A: Gender

			Primary Robustness Considerations					
	Full Sample		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

			Primary Robustness Considerations					
	Full Sample		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: Gender or URM

			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 our analyses of potential heterogeneity in the relation between mentorship program participation and innovator productivity, focusing on gender and racial/ethnic diversity. 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 either or 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 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: Survey Evidence on the Mentorship Program

	Full sample		Specific Subgroups					
			RG		URG	Male	Female	URM
Panel A. Mentorship	Obs (1)	Mean (2)	Obs (3)	Mean (4)	Mean (5)	Mean (6)	Mean (7)	Mean (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%
Panel C. Perspective of those with a mentor								
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 (5), (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: 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 submissions at Status 2. 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 8: Changes in Innovator Team Composition After Mentorship

Panel A: Full Sample of Data

	Full Sample		Primary Robustness Considerations							
			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

	Full Sample		Primary Robustness Considerations							
			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 9: Engineers' Workplace Perceptions After the Mentorship Program

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	0.077***	0.024	0.048***	0.046***	0.045***
	(4.087)	(0.991)	(8.211)	(6.872)	(6.848)
Engineer		-0.198***		-0.001	-0.001
		(-5.866)		(-0.099)	(-0.072)
Post × Engineer		0.128***		0.013	0.013
		(3.343)		(0.904)	(0.917)
Post × Engineer × Focal Firm					0.112***
					(2.773)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.008	0.012	0.056	0.056	0.057
Observations	12,019	12,019	230,044	230,044	242,066

Notes. This table presents an analysis of potential changes in engineers' perceptions of their workplace coinciding with the implementation of the patent mentorship program. The data used is individual reviews posted from 2018 to 2024 on Glassdoor, which is a career intelligence website that attempts to provide transparency about jobs, salaries, and companies. We focus on the reviews provided by current employees of the focal firm and its industry peers. The analysis is restricted to reviews with non-missing job title data, which is necessary for identifying engineers. The dependent variable is the overall rating of the firm provided by the reviewer. We employ ordinary least squares regressions to analyze these ratings 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 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. The sum of observation counts between Columns (1) and (3), 242,063, is smaller than the value in Column (5) due to the omission of singleton observations. 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' Departures After the Mentorship Program

Panel A: Outflows

	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

Panel B: Inverse Hyperbolic Sine of Outflows

	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 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 scaled external outflows measure. 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.

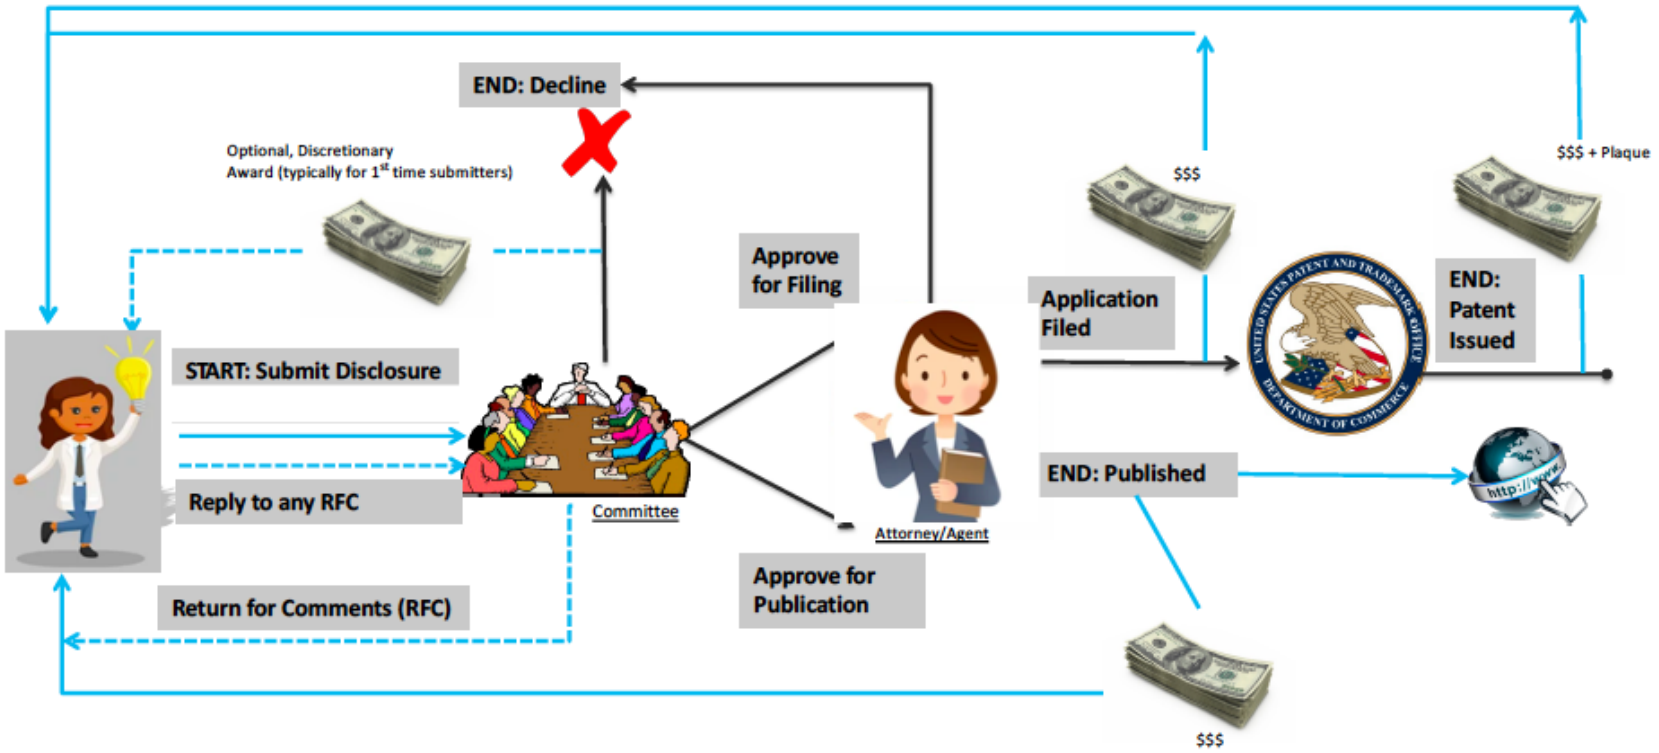
A Appendix Figures and Tables

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.

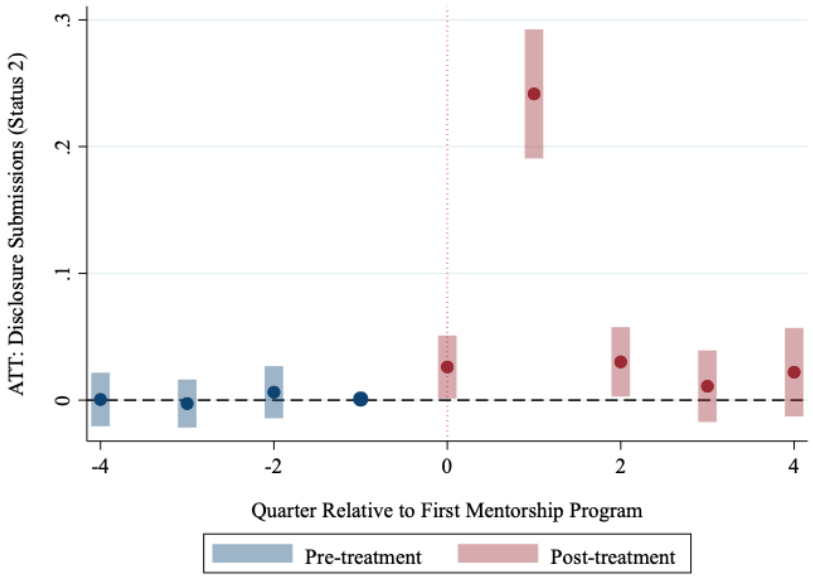
Figure A.2: Patent Process



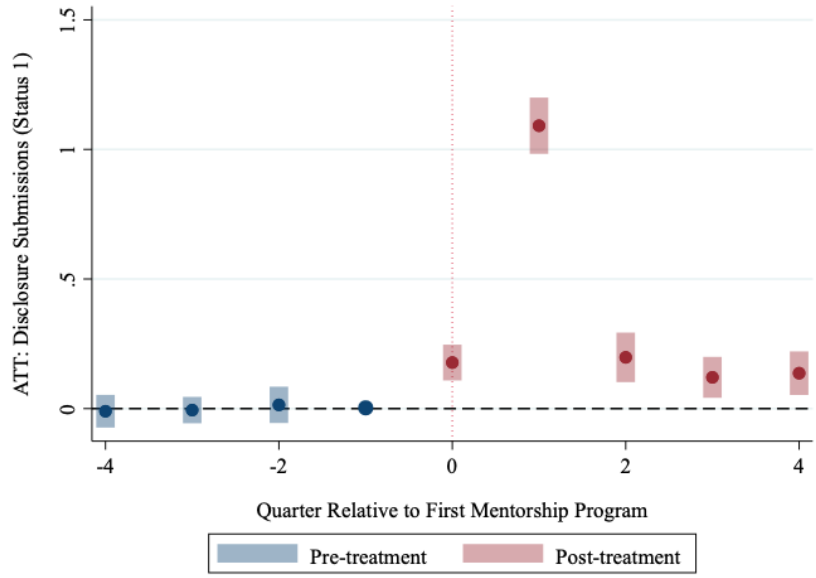
Notes. This figure highlights the patent process as taught to mentees by mentors as part of the mentorship program. The main outcomes that a disclosure submission can realize are (i) being declined, (ii) being approved for defensive publication, and (iii) being a filed or approved patent.

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.

Table A.1: 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.2: 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.3: 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.4: 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.5: 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 7. t-statistics are reported in parentheses. Statistical significance of 10%, 5%, and 1% is denoted by *, **, and ***, respectively.

Table A.6: 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 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.

B Robustness and Falsification Tests

This appendix demonstrates the robustness of our main findings on the relation between patent mentorship and innovation productivity. Specifically, we validate our results when using alternative control groups, when employing different estimators, and when using different disclosure status classifications. In addition, we conduct placebo tests for the employee satisfaction and turnover outcomes.

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 Callaway and Sant'Anna (2021) estimator and the Poisson pseudo-likelihood estimator, which accounts for the count-like nature of disclosure submissions. We use the Poisson pseudo-likelihood estimator in

Column (3), and we employ the Callaway and Sant'Anna (2021) estimator in Column (4).[30] 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).[31] 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 first mentorship and the first and second quarters after first 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 Additional Robustness Considerations

In Table B.3, 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

[30] 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.

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

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.4, 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.4 are very similar to our main results, mitigating the concern that additional mentorship is necessary for benefits to be realized.

B.5 Alternative Timing of Employee Sentiment and Retention Tests

The results in Table 9 suggested 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.5, 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 two 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 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.6. 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 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.7 show that the estimated relation between the mentorship program and engineers' departures from the focal firm is robust to alternative definitions of *Post*. In these tests, we note that the magnitudes of the coefficients are very similar to those in Table 10. 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.8, which report null in-

teraction 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 Table 10 are not driven by some alternative factors.

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 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 employe 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 first mentorship and the first and second quarters after first 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: 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.4: 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.5: 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	0.101***	0.060***	0.017***	0.012**	0.013**
	(5.732)	(2.584)	(3.195)	(2.079)	(2.092)
Engineer		-0.173***		-0.007	-0.007
		(-5.738)		(-0.644)	(-0.602)
Post × Engineer		0.102***		0.022*	0.022*
		(2.860)		(1.724)	(1.727)
Post × Engineer × Focal Firm					0.074**
					(1.973)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.010	0.014	0.055	0.055	0.056
Observations	12,019	12,019	230,044	230,044	242,066

Panel B: After the Third Mentorship Cohort

	Focal Firm Only		Peer Firms Only		Both
	(1)	(2)	(3)	(4)	(5)
Post	0.094***	0.069***	0.030***	0.030***	0.030***
	(5.726)	(3.162)	(6.618)	(5.804)	(5.853)
Engineer		-0.135***		0.010	0.010
		(-5.593)		(1.110)	(1.180)
Post × Engineer		0.060*		-0.000	-0.000
		(1.843)		(-0.005)	(-0.020)
Post × Engineer × Focal Firm					0.060*
					(1.740)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.010	0.013	0.056	0.056	0.056
Observations	12,019	12,019	230,044	230,044	242,066

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.6: 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	0.055	0.028	0.046***	0.041**	0.040**
	(1.242)	(0.491)	(2.987)	(2.434)	(2.352)
Engineer		-0.235***		0.002	0.005
		(-3.180)		(0.062)	(0.153)
Pseudo-Post × Engineer		0.063		0.023	0.022
		(0.713)		(0.609)	(0.581)
Pseudo-Post × Engineer × Focal Firm					0.036
					(0.383)
Country FE	✓	✓	✓	✓	✓
Firm FE			✓	✓	✓
Adj. R-Square	0.012	0.020	0.068	0.068	0.070
Observations	2,261	2,261	24,345	24,345	26,620

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.7: 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 scaled 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.8: 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 scaled 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.

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