

Task Assignment and Career Advancement

Fiona Chen

James Stratton

Harvard University

Harvard University

October 7, 2026

Preliminary

Abstract

How do the tasks workers are assigned shape their career advancement? We examine this question using data from an engineering analytics platform on tasks assigned to software engineers, merged with LinkedIn data on workers' career trajectories. We exploit variation in workers' task assignment arising from fluctuations in their teams' task portfolios at the time they join. Assignment to more complex tasks in the first quarter of employment significantly increases the likelihood of promotion within two years. We provide evidence consistent with a human capital accumulation channel: workers exhibit steep learning curves in their first year at the firm, with task completion times falling sharply with accumulated experience. We also examine the implications for gender disparities, documenting gender differences in task assignment and promotion rates.

Contact: fionachen@g.harvard.edu and jstratton@g.harvard.edu. We are indebted to Raj Chetty, Larry Katz, Mandy Pallais, and Jesse Shapiro for their guidance and support throughout this project. Additionally, we would like to thank participants at the Harvard Labor/Public Workshop and Stone Program on Inequality. The results are based on data that has been aggregated and anonymized. The manuscript was reviewed by the data partner solely to verify that it does not disclose proprietary information, as required by the research agreement; it was not reviewed for its scientific content.

1 Introduction

Promotions inside firms account for a meaningful share of workers’ life-cycle wage growth (Baker, Gibbs, & Holmstrom, 1994; Baker, Gibbs, & Holmström, 1994; McCue, 1996). In standard models of firm production, tasks are assigned to workers on the basis of comparative advantage, and jobs are bundles of tasks, chosen to balance the gains from specialization against hiring and coordination costs (Becker & Murphy, 1992; Lazear, 1992). In turn, workers are promoted when moving into a higher-level job better matches their comparative advantage to what the firm demands. In practice, however, firms’ task assignment may deviate from the benchmark for several reasons: firms may face short-run fluctuations in the demand for particular types of labor (Adhvaryu et al., 2022); managers may vary task assignment to learn about workers’ abilities (Waldman, 1984; Meyer, 1994); and managers may discriminate on characteristics unrelated to workers’ skills or experience (e.g., Babcock et al., 2017; Chu et al., 2022). How does variation in the tasks that workers are assigned affect their career advancement?

We study this question by leveraging a novel dataset from an analytics platform, which firms use to understand the activity of their software engineering teams. Our data provides information on the activity of 52,659 workers at 303 firms. We primarily make use of the platform’s integration with Jira, a project management tool, through which we can observe workers’ assignment of “issues” — i.e., tasks. Individual issues represent blocks of engineering or engineering-adjacent work, that generally takes several days to two weeks to complete, such as updates to website features or bug fixes. Data on each issue is highly detailed: we observe its assignee and supervisor, text description, and priority level. We merge this with LinkedIn data to on workers’ education and work histories, to assess their promotion trajectories.

Task assignment may play a particularly important role in this context, in which knowledge workers perform non-routine tasks. The non-routine nature of work provides scope for variation in task assignment, and also increases the importance of on-the-job learning. However, it also poses unique challenges for measurement of learning and productivity, which may explain why existing literature on firm-specific human capital has focused on more routine occupations such as production. Our data allows us to overcome such challenges, by providing us detailed insights into the outputs of knowledge workers.

We begin by documenting the importance of task assignment for career advancement. To do this, we use machine learning tools to classify Jira issues' complexity and category according to their text descriptions. We find substantial variation in task assignment across workers, even within job titles. We also provide evidence on sources of variation of task assignment within-occupation, documenting differences in task assignment by workers' educational credentials and gender. Second, we document that task assignment plays an important role in shaping workers' career advancement. We use an instrumental variables approach, exploiting quasi-random fluctuations in teams' task compositions at the time that workers join the firm. A 1 p.p. increase in a worker's share of complex tasks in her first quarter at the firm leads to a 1.8 p.p. increase in her likelihood of promotion.

We then turn to exploring one channel through which task assignment may matter for workers: opportunities to accumulate human capital on the job. We document workers' "learning curves" through their time at the firm, by grouping tasks by their complexity, then estimating the relationship between work experience and time to complete tasks of different categories. We find a sharp learning gradient within workers' first year at the firm, which then plateaus. These patterns are consistent with assignment to more complex tasks providing workers with greater scope for skill development.

Finally, we explore the role of task assignment in gender disparities in career advancement. Despite progress over the past several decades, significant gender gaps persist in high-wage occupations: within-occupation differences account for roughly two-thirds of the gender wage gap, and the gap is largest at the top of the wage distribution (Goldin, 2014; Blau & Kahn, 2017). We document that women represent only 24 percent of recent graduate roles in our sample, and that persistent promotion and salary gaps emerge even among workers in the same job titles. We show that women are assigned fewer "high priority" tasks than men, and that this gap grows over time within a firm. Furthermore, women are assigned project management and design tasks at a higher rate than men; and data analytics and bug fix tasks at a lower rate than men. Notably, women complete tasks at higher rates than men across all priority levels, suggesting the gap in task assignment is not driven by performance differences.

Related literature. This project relates to several bodies of literature. First, it relates to a large literature on internal labor markets, which uses detailed personnel records to understand how workers are evaluated, paid, and promoted within firms (e.g., Baker, Gibbs, & Holmstrom, 1994; Baker, Gibbs, & Holmström, 1994; Lazear & Oyer, 2003; Gibbons & Waldman, 2006; Kohlhepp, 2024). A

subset of this literature emphasizes the importance of firm experience and training for human capital accumulation (e.g., Lazear, 2009; Levitt et al., 2013; Frederiksen & Kato, 2018; Jedwab et al., 2023), and of task-specific human capital and learning through job assignment (Waldman, 1984; Meyer, 1994; Gibbons & Waldman, 2004; Gathmann & Schönberg, 2010; Lise & Postel-Vinay, 2020). We contribute to this literature by providing evidence on how the tasks workers perform within a job affect their career advancement, using variation in which tasks are available to a worker rather than variation in worker or manager choices.

Second, a growing literature documents variation in how firms assign tasks to workers. Several papers show that women are more likely to take on non-promotable tasks (Babcock et al., 2017; Chu et al., 2022; Lehmann, 2013; Bol et al., 2026; Murphy-Hill et al., 2023). Beyond gender, managers shape assignment by moving workers across jobs and task domains (Minni, 2026) and by reassigning workers when productivity shocks hit (Adhvaryu et al., 2022). Cullen et al. (2025) similarly use quasi-random assignment of Air Force jobs, but to estimate the career effects of job-skill mismatch. Bircan et al. (2025) is closest to our paper: using manager rotations, they show that supervisors causally shift who receives visible leadership roles, and relate this to the gender promotion gap. Their design identifies the effect of managers on assignment, whereas ours isolates variation in the work available to the worker, which lets us estimate the effect of task assignment on promotion directly.

The rest of the paper proceeds as follows. In Section 2, we describe our data and institutional details. In Section 3, we describe our main results on the effects of task assignment on workers’ career advancement. In Section 4, we provide evidence on the human capital accumulation channel. In Section 5, we document gender disparities in task assignment and career outcomes. Section 6 concludes.

2 Data and institutional context

This section summarizes the data used in our analysis.

2.1 Data sources

Engineering analytics platform. Our primary source of data is an engineering analytics platform. Firms use the platform to monitor software activity and project progress, and to understand the

engineers’ productivity, including in relation to their use of AI coding tools. To provide these analytics, the platform integrates data on engineers’ activity from several sources. We primarily make use of the data from the issue tracking system integration.

Analysis sample. We use data from 2015 to 2024. We construct an analysis sample of engineers who are active at a firm monitored by the platform for a period of at least three months in this time period. There are 52,659 workers in this sample from 303 firms.

HR data. Firms directly enter data into the platform on workers’ job titles; a subset of firms also enter workers’ salaries, and the names of their supervisors. Firms have an incentive to record this information accurately: workers are sorted by job title when managers review worker performance, and firms enter salary information in part so that they can calculate R&D tax credits.

Issue tracking data. Software teams employ “issue tracking” systems to organize, assign, and monitor engineering work.¹ Issue management systems are industry-standard for software engineering teams.

Atlassian, Jira’s owner, claims that four-fifths of Fortune 500 firms use Jira (Farquhar, 2019). An “issue” is a discrete unit of engineering work that can be assigned to a worker, such as a feature update or a bug fix. Our analysis sample includes 9,703,637 issues, for an average of 7.45 per worker-month. Figure 1 presents an example of Jira’s interface when assigning issues.

Each issue includes rich information, including the assignee and supervisor, as well as timestamps associated with its opening and resolution. Each issue also includes a text description, which we use to classify the content of the tasks, as described in Section 2.4. The text description of each issue is typically rich: in our analysis sample, the median length of an issue description is around 65 words. Each is also assigned a priority level by the assigner, ranging from low to critical, which captures the importance of the issue to the project, as well as its time urgency. Finally, it also includes information with its resolution status — i.e., whether it was successfully or unsuccessfully completed — which we use to assess the quality of work completed.

LinkedIn. We obtain supplemental data on workers’ characteristics from Coresignal, a data provider which scrapes LinkedIn profiles as of the end of 2022. We combine this with the HR-inputted data on the analytics platform to obtain a more comprehensive view of workers’ education and work

¹The platform integrates data on several issue management systems, including Jira and Microsoft Azure DevOps. For simplicity, we adopt the language of the most common issue management system globally, Jira.

histories. We first hand-match the analysis sample of firms with the list of firms whose data are collected from LinkedIn by Coresignal; we are able to match essentially every firm in the sample. Next, we merge worker names from data to profile names in the LinkedIn data, restricting attention to profiles that have, at some point, worked at a firm. We merge the names of educational institutions to the classification of colleges in Chetty et al. (2023) and define a worker as a graduate from a highly-selective college if any of the worker’s educational qualifications were from a public or private highly-selective college.

Dun & Bradstreet. Dun & Bradstreet (D&B) is a commercial provider of business data and analytics. D&B maintains records on hundreds of millions of businesses worldwide, compiling information across numerous sources, including public business registries, firm websites, and trade-payment records. We merge firms to D&B records by firm name, and are able to match nearly all firms.

2.2 Descriptive statistics from analysis sample

Worker-level summary statistics. Appendix Table 1 presents summary statistics on the workers, including income, gender, and education history. We have salary data on a minority of workers. Around five-sixths of the workers whom we are able to assign a gender are classified as male. Around four-fifths of workers for whom we have access to educational data are college-educated; most of the remaining workers completed coding courses offered by private platforms that are not associated with colleges (*e.g.*, Codecademy, Udemy). Around two-fifths of workers are in “senior” roles, which we classify as job titles indicating that the worker has been promoted relative to the typical positions new graduates would receive (*e.g.*, “senior developer”, “team leader”, “engineering manager”). Appendix Figure 1 presents the distribution of role types in the analysis sample. Around four-fifths of workers are engineers; the remainder are in other roles involved in the software development process, including managers.

Firm-level summary statistics. Appendix Table 1 presents firm-level summary statistics, which we obtain from Dun & Bradstreet. Most firms are in the professional services industry, which includes the software industry — *i.e.*, these are firms that are creating a software product. There is some representation of firms in other industries, including retail, wholesale, and manufacturing — these firms are largely using software for an internal process.

2.3 Gender

We use the names of engineers to assign genders, using the method in Day et al. (2020), which computes the empirical frequency of first names by gender in a large dataset of names from the U.S. and U.K. We are unable to assign genders to a subset of approximately 35% of workers, which includes workers whose names are not available in the data, or for whom we cannot reliably impute gender based on the algorithm.

2.4 Measuring task content

We use three measures of task content in our analysis. The first, issue “priority”, is inputted by the workers into Jira. Then, we use machine learning tools to classify issues by their complexity and content.

Complexity prediction. We begin by using machine learning tools to classify issues by their complexity. To do this, we vectorize the Jira issue text descriptions, and train a model on the outcome of the seniority of the assigned worker. Thus, this model assigns each issue with a probability that it is assigned to a senior worker rather than a junior worker. We then define a task as a “complex task” if it has a complexity level greater than or equal to 0.5 (*i.e.*, it is more likely than not to be assigned to a senior worker), and we define a task as a “simple task” if it has a complexity level less than 0.5.

To train the model, we begin by restricting attention to the sample of tasks assigned prior to the introduction of GitHub Copilot for Business in 2023. We randomly select 25% of those issues to use in training. We first represent the raw text description of each task t , $\mathcal{T}_t$, as a numerical array $\mathbf{D}_t$. To do so, we stem words, remove especially common and rare words, and then convert each text description into a TF-IDF representation: that is, for each task t , $\mathbf{D}_t$ collects the frequencies with which each word in the corpus of words in our task documents is used. Next, we flexibly estimate the function $\mathbb{P}(\text{Seniority}_{w(t)}|\mathbf{D}_t)$ through a logistic regression in the training sample, where $\text{Seniority}_{w(t)}$ is a binary variable indicating whether the worker to whom task t was assigned, $w(t)$, is a senior worker. We then construct the complexity of each task, $\hat{s}_t$, as the fitted value of this logistic regression.

Appendix Table 2 presents some examples of tasks and the complexity assigned to those tasks by our classification algorithm.

Task categories. We group tasks into categories using an unsupervised topic model. Specifically, we fit a latent Dirichlet allocation (LDA) model to the text of each Jira issue, which “discovers” latent topics, according to words that tend to appear together, and then assigns each task to its dominant topic. We fit topic categories, including infrastructure, review, DevOps, project management, and UI/UX.

Validation of classifications. Finally, we assess the predictive accuracy of our classification of issues in a test set of tasks that was excluded from the initial training sample. Appendix Figure 2A shows the distribution of task complexity among tasks that were assigned to senior vs. junior workers, restricting attention to tasks in the test set. The distribution of predicted task complexity is systematically lower among tasks assigned to junior workers relative to senior workers.

Then, we compare the complexity labels with the category labels. The complexity classification algorithm predicts both variation in complexity levels *across* these categories (Appendix Figure 2B), and *within* these categories (Appendix Figure 2C).

3 Task assignment and career advancement

We now move onto our analysis of the structure of work and workers’ career advancement.

3.1 Variation in task assignment

We begin by providing descriptive evidence of variation in task assignment, among workers with the same job title. We provide evidence of variability across three dimensions: years of experience, college quality, and gender. The results are displayed in Figure 2, where analysis is restricted to workers in junior job positions, and estimates come from the univariate regressions of task complexity on worker characteristics.

Figure 2A shows that workers with more years of experience receive more complex tasks: 1 additional year of work experience is associated with a 1.9 p.p. increase in share of complex tasks in the first two years. Figure 2B shows that workers who attended a highly selective college receive 7.5 p.p. more complex tasks than workers who attended a non-selective college. Finally, Figure 2C shows that female workers receive 0.7 p.p. more complex tasks than male workers upon entry.

3.2 Descriptive evidence on task assignment and career advancement

Next, we assess the predictive relationship between task complexity and a worker's advancement through the firm. We restrict attention to the subset of our analysis sample consisting of workers who are new entrants to pre-existing firms. We find that more complex tasks predict promotion, conditional on pre-entry characteristics. Conversely, holding fixed pre-entry characteristics, a worker's task portfolio is predictive of promotion. Figure 3 shows that new hires given more complex tasks upon entry into the firm are more likely to be promoted, conditional on years of experience and whether the worker attended a highly-selective school. The point estimate implies that a 10 percentage point increase in a worker's complex task share predicts a 2.4 percentage point increase in the likelihood of the worker being promoted.

3.3 Causal effect of task assignment on career advancement

One interpretation of the positive cross-sectional relationship between the complexity of a worker's tasks and her promotion likelihood is that task assignment causally affects promotion likelihood, e.g., through opportunities for human capital accumulation. An alternative possibility is that this positive cross-sectional relationship simply reflects unobserved differences in workers' capabilities upon entering the firm: managers assign more complex tasks to some higher-skilled workers, who are also more likely to be promoted.

To distinguish between these possibilities, we develop an instrument for the causal effect of task complexity on a worker's likelihood of promotion. The instrument isolates variation in task assignment that arises due to changes in the environment facing the worker's manager, and are plausibly unrelated to a worker's latent productivity. The intuition for the instrument is as follows. There are fluctuations over time in the complexity of the full set of tasks which a manager can assign to workers: some workers join a firm at a point in time at which her team is embarking on a new project, and there are more complex tasks available to be assigned to workers; other workers join a firm at a point in time at which only simpler tasks are available.

To fix ideas, it is helpful to consider by way of example the path of task complexity of one particular team in the analysis sample, provided in Figure 4A. This team experiences a substantial increase in task complexity in the third quarter of 2019; after this spike, task complexity returns to its standard level. One worker joins the team immediately prior to this spike; the other joins the

team after this spike. Under an assumption that the precise timing of a worker joining a team is orthogonal to her latent likelihood of promotion, we can use the comparison between these workers to determine the causal effects of task assignment on promotion.

To generalize this idea, we measure “shocks” to the composition of tasks within a team by constructing the mean predicted complexity of issues within each team $\times$ quarter, and then residualizing this mean complexity with respect to team fixed effects. We then use the measured shock to the complexity of a team’s tasks in the quarter prior to a new worker joining the team as an instrument for the mean complexity of tasks assigned to the newly-entering worker.

Formally, we estimate the following two-stage least-squares model:

$$s_{i,1} = \alpha_{t(i,1)} + \beta \bar{s}_{t(i,1),0} + \varepsilon_{i,1} \quad (1)$$

$$\text{Promotion}_i = \gamma_{t(i,1)} + \delta \hat{s}_{i,1} + \eta_i, \quad (2)$$

where $s_{i,q}$ is the mean complexity of tasks assigned to a worker at quarter q relative to the worker’s first quarter in the firm ($q = 1$), $t(i, 1)$ is the team to which worker i is assigned during her first quarter, and $\bar{s}_{t(i,1),0} \equiv \frac{1}{|\{j: t(j,0)=t(i,1)\}|} \sum_{j: t(j,0)=t(i,1)} s_{j,0}$ is the mean complexity of issues among workers assigned to i ’s team in the quarter prior to her joining the firm, and Promotion_i is an indicator variable equal to one if the worker is promoted to a senior or managerial role during her first two years at the firm.

Identification assumptions. The instrument recovers the local average treatment effect of being assigned more complex issues under four identifying assumptions, which we discuss in detail below.

1. *Relevance.* The relevance condition requires that $\beta \neq 0$. Figure 4B presents a binned scatterplot of our estimates of equation (1). New entrants to a firm tend to be assigned more complex issues when they are assigned to teams with more complex issues in the quarter prior to their arrival. The magnitude of the coefficient is reasonably large: a 10 p.p. increase in the complexity of a team’s issues increases the complexity of a new worker’s issues by roughly 1.5 p.p. (a pass-through rate of roughly 15%), after removing team fixed effects. We are able to reject the null that there is no effect of the instrument on the endogenous variable at all standard levels of statistical significance.

To interpret this pass-through rate, it is helpful to consider a natural benchmark in which shocks to a team’s task composition at quarter $q - 1$ are fully passed through to an individual worker’s composition at quarter q , in which case pass-through would be 100% ($\beta = 1$). Two factors drive

attenuation relative to this full pass-through: first, the shock to a team’s tasks in quarter q may be less than fully passed through to a team’s tasks at $q + 1$ (that is, the time structure of tasks may not feature full pass-through due to mean reversion); and second, a team’s tasks at $q + 1$ may not pass through one-to-one into an individual worker’s tasks at $q + 1$ (that is, the “vertical” pass-through of shocks may be incomplete – for instance, because managers avoid giving difficult tasks to new team members).² Appendix Figure 3 presents this decomposition. The attenuation of the pass-through rate relative to full pass-through is almost entirely driven by mean-reversion in a team’s task composition.

2. *Exogeneity.* In our context, the key challenge to exogeneity is the possibility that managers respond to short-term increases in the difficulty of their task portfolio by hiring more productive workers, who are subsequently more likely to be promoted. To assess the plausibility of this exogeneity concern, we examine whether teams experiencing a shock inducing a higher share of complex tasks hire workers with stronger pre-employment characteristics. We find that teams experiencing these shocks do not hire workers with more experience (Appendix Figure 5A), and are no more likely to hire workers who have attended highly-selective colleges (Appendix Figure 5B). Our interpretation of this evidence is that firms’ hiring is not responsive to fluctuations in task difficulty that occur immediately prior to the worker joining the firm.

One possible reason for this unresponsiveness is that firms have already hired workers more than one quarter prior to a worker’s entry to the firm, and so, for this reason, their hiring does not respond to essentially contemporaneous shocks to task difficulty. In the future, we hope to directly assess this point by gathering data on the hiring cycles of technology firms, perhaps through observing the timing of firms’ job postings.

3. *Exclusion restriction.* The exclusion restriction requires that a shock to a team’s task composition prior to a worker’s entry into the firm affects the worker’s promotion likelihood only through its effect on that worker’s task composition. The exclusion restriction may be violated if a firm’s promotion policy is such that a worker’s likelihood of being promoted is a function of not just *her* performance,

²Formally, given three random variables A , B , and C , denote by $\beta_{A \rightarrow B|C} \equiv \frac{\text{Cov}[A, B|C]}{\text{Var}[A|C]}$ the coefficient from a regression of B on A , residualized with respect to C . In a statistical sense, the regression coefficient for the effect of team task complexity in the quarter prior to entry on a worker’s task assignment on entry can be written as

$$\beta_{\bar{s}_{t(i,1),0} \rightarrow s_{i,1}|t(i,1)} = \beta_{\bar{s}_{t(i,1),0} \rightarrow \bar{s}_{t(i,1),1}|t(i,1)} \times \beta_{\bar{s}_{t(i,1),1} \rightarrow s_{i,1}|t(i,1)} + \beta_{\bar{s}_{t(i,1),0} \rightarrow s_{i,1}|t(i,1), \bar{s}_{t(i,1),1}}.$$

That is, the observed effect of the instrument on tasks assigned to a new entrant can be decomposed into two terms: (1) the direct effect of tasks assigned prior to entry on a worker’s tasks assigned on entry; and (2) the product of the effect of team tasks prior to entry on team tasks at entry, and the effect of team tasks at entry on an individual worker’s tasks at entry. We find that the first term is essentially zero.

but also of the performance of other members of her team. For instance, if a firm’s likelihood of promoting any given worker was decreasing in its willingness to promote other workers within the same team – as would be the case if the firm had a policy of promoting one worker within each team – then our estimated effects of task complexity would be attenuated. For this reason, we interpret our results as a lower bound for the effect of task complexity on promotion likelihood.

4. *Monotonicity.* Finally, the monotonicity assumption requires that harder tasks within a team uniformly increase each worker’s task difficulty.

Results from instrumental variables design. Figure 4C presents a binned scatterplot corresponding to the reduced-form regression in our instrumental variables design:

$$\text{Promotion}_i = \kappa_{t(i,1)} + \lambda \bar{s}_{t(i,1),0} + \iota_i.$$

Under the exogeneity restriction, the estimated coefficient $\hat{\lambda} = 0.28$ indicates that a 10 percentage point increase in the mean complexity of a team’s tasks results in a 2.8 percentage point increase in a worker’s likelihood of promotion over the following two years. Under the exclusion and monotonicity restrictions, this implies that a one percentage point increase in a worker’s own complexity share increases her likelihood of promotion over the following two years by 1.81 percentage points (Figure 4D).

4 On-the-job learning

In this section, we estimate the effects of work experience and task assignment on human capital accumulation.

4.1 Estimation of learning curve

We plot the learning curves associated with two categories of task, displayed in Figure 5. Figure 5A displays the main results, with the average time to complete tasks assigned each month since the start of employment, plotted separately for simple vs. complex tasks. We observe a steep decline in completion time in the first year at which workers are at the firm, followed by a plateau for the two years after. Appendix Figure 4 replicates these results, but with a balanced panel.

Next, we turn to more formally estimating the relationship between productivity and work experience. As in Levitt et al. (2013), we begin with a standard power law specification of the relationship: $S_t = AE_t^\beta$. Here, S_t is productivity in period t (average number of days to complete each task assigned in month t), E_t is work experience (cumulative number of tasks completed), and A and β are parameters. Formally, $E_t = \sum_{i=1}^{t-1} Task s_i$. Taking logs of this specification, we can get an empirically estimatable expression:

$$\log S_t = \alpha + \beta \log E_t + \varepsilon_t$$

where $\alpha \equiv \log A$ and ε_t is an error term. Thus, we can interpret β as a “learning rate”: i.e., the relationship between experience on productivity. To separate out effects associated with simply staying in the firm, we also include a time control.

Table 1 displays the results from estimating this expression. Panel A displays the estimated learning rates for all engineers in the first year as well as in the first 3 years at the firm. In the first year at a firm, a 1% increase in number of simple tasks completed in the first $t - 1$ months is associated with a 0.15% decrease in average time to complete 1 task in month t . Looking at complex tasks, the estimate is similar but of a smaller magnitude: a 1% increase in number of complex tasks completed is associated with a 0.11% decrease in average time to complete one task. Intuitively, the learning rates are lower in the next two years than in the first year: in the first 3 years, a 1% increase in simple tasks is associated with just a 0.14% decrease in average time to complete 1 task. We are also able to separately examine these effects for junior vs. senior engineers, with results in Panel B. We see that the effects of experience on simple tasks are more pronounced for junior than senior engineers; vice versa, the effects for complex tasks are more pronounced for senior engineers.

Taken together, these results are consistent with assignment to more complex tasks providing workers with greater opportunities to accumulate task-specific human capital, which in turn contributes to their advancement within the firm.

5 Gender disparities in task assignment and promotion

In this section, we document gender disparities in career outcomes and explore the role of task assignment in generating these gaps.

5.1 Gender gaps in career outcomes

We observe gender gaps in several dimensions. First, we observe a gender gap in entry into the field: women represent only 24 percent of recent graduate roles in our data. Second, we document a gender gap in promotion rates in Figure 6. Panel A displays the share of men and women who have been promoted to senior positions (*e.g.*, “senior software engineer” or “senior data scientist”) each month since graduation; Panel B displays the equivalent for managerial positions (*e.g.*, “manager” or “director”). Two years after graduating, around 12 percent of men have been promoted to senior roles, whereas only 9 percent of women have. In the same timeframe, around 3.6 percent of men have been promoted to managerial positions, compared with around 3.3 percent of women. Finally, Panel C shows a gender gap in salaries: on average, women earn 86 percent of what men earn, unconditional on any factors.

5.2 Gender differences in task assignment

We next document gender differences in the types of tasks assigned to workers. As shown in Figure 2C, women and men receive a similar share of complex tasks upon entry to the firm. However, a gap in task assignment emerges over time. Figure 7 displays the share of tasks that are high priority over time, broken down by gender. At the point of entry into a firm, there is no gender gap in task importance. However, a gap emerges approximately 1.5 years in: three years into employment, men are assigned 3 percentage points more high priority tasks than women.

We assess whether this gap might be driven by differences in performance by gender. Figure 8 presents task completion rates by gender and priority level. For every priority level, women complete tasks at a higher rate than men, suggesting that the emerging gap in task assignment is not attributable to differential performance.

Figure 9 documents gender differences in the types of tasks assigned, using an unsupervised clustering approach to group task descriptions into categories. Women are assigned a larger share of UI-related tasks, while men receive a larger share of data analysis and debugging tasks. These differences in task type may compound the effects of differences in task complexity on workers’ human capital accumulation and promotion prospects.

6 Conclusion

This paper studies how task assignment shapes career advancement for knowledge workers. Using data on 52,659 engineers at 303 firms, we exploit quasi-random variation in the complexity of tasks available to new workers — arising from fluctuations in a team’s task portfolio in the quarter before they join — to identify the causal effect of task assignment on promotion. A one percentage point increase in a worker’s share of complex tasks in her first quarter raises her likelihood of promotion within two years by 1.8 percentage points.

We provide evidence consistent with human capital accumulation as a key channel for this effect. Workers exhibit steep learning curves in their first year at the firm: completion times fall sharply with experience on both simple and complex tasks, and then plateau. These patterns suggest that early assignment to more complex work expands the scope for on-the-job learning and skill development.

We also provide evidence of gender disparities in task assignment and career outcomes. While women and men enter firms with similar task portfolios, a gap in task complexity and priority opens up over the first few years of employment, and this gap is not explained by differential performance. Women complete tasks at higher rates than men across all priority levels, yet receive fewer high-priority tasks as their tenure grows. These findings suggest that differences in task assignment may contribute to persistent promotion and wage gaps between men and women in high-wage occupations.

Our findings point to task assignment as an understudied margin through which firm decisions shape both individual career trajectories and aggregate inequality. Future work could examine the determinants of task assignment — including the role of manager characteristics and firm policies — and extend the analysis to other outcomes such as wages and retention.

Figures

FIGURE 1. Example of Jira issue

Redesign users' account registration workflow to make account creation faster

Attach

Add a child issue

Link issue

▼

Create ▼

...

Description

The current account registration process asks for address details that we no longer regularly use. Please update the process to remove these questions. Make sure that the back-end data collection workflows fits smoothly with the revised account creation process.

Attachments 1

Project plan.pdf
08 Oct 2024, 06:41 PM

To Do ▼ ⚡ Actions ▼

Details

Assignee

JS James Stratton

Labels

registration user-facing

Parent ⓘ

SCRUM-2 User registration

Team

Front-end

Sprint

None

Story point estimate

None

Development


Reporter

JS James Stratton

Notes: This figure displays an illustrative example of a Jira issue, including the assignee and text description.


16

FIGURE 2. Relationship between pre-entry characteristics and task assignment




Notes: This figure displays estimates of the cross-sectional relationship between workers' pre-entry characteristics and task assignment. Panel A shows a binned scatterplot in which the sample is new entrants to a firm, the x -axis variable is the entrant's experience (measured in years since college graduation) at the time of entry, and the y -axis variable is the mean complexity of tasks assigned to the new hire. The estimated coefficient and standard error from a regression of mean complexity on experience are displayed in the top left corner. Panel B compares the mean complexity of tasks assigned to new entrants who graduated from non-selective vs. selective colleges, restricting attention further to the sample for which we have access to educational records from merging Coresignal-LinkedIn data. Panel C replicates Panel B, instead comparing the complexity of tasks assigned to new entrants who are male vs. female.

FIGURE 3. OLS relationship between task difficulty upon entry and probability of promotion



Notes: This figure displays OLS estimates of the relationship between task assignment and a worker’s likelihood of promotion within their first two years at the firm.


FIGURE 4. Two-stage least-squares estimates of the causal effects of task assignment



Notes: This figure displays two-stage least-squares estimates of the causal effect of task assignment on a worker's likelihood of promotion within their first two years at a firm. Panel A shows an example of one team's activity, as recorded in data, to display the intuition behind the instrument for task complexity. The blue line displays mean task complexity for the team in each quarter. Worker A joins the team in 2019q2, during a time with greater team activity. Worker B joins in 2020q3, during a time of slower activity. Panel B displays the first stage result, with the OLS regression of a worker's own mean task complexity at quarter t on their team's mean task complexity at quarter $t - 1$. Panel C displays the reduced form result, with the OLS regression of an indicator for whether a worker was promoted within two years on their team's mean task complexity at quarter $t - 1$. Panel D displays the main IV result, with the 2SLS regression of the promotion indicator on a worker's mean task complexity in their first quarter at the firm, instrumented by their team's mean task complexity from the quarter before they join.


FIGURE 5. Learning curves for simple vs. complex tasks

A. All workers




Notes: This figure displays the “learning curves” associated with simple and complex tasks — i.e., the average amount of time taken to complete tasks assigned in each month since a worker’s start at the firm.

FIGURE 6. Gender gaps among tech workers




Notes: This figure displays gender gaps in career outcomes among tech workers. Panel A displays the share of men and women promoted to senior positions each month since graduation. Panel B displays the equivalent for managerial positions. Panel C displays the distribution of salaries by gender.

FIGURE 7. Evolution of task importance by gender




Notes: This figure displays the share of tasks that are high priority (Jira priority levels 3 or 4) over time, broken down by gender. The sample is restricted to workers in junior positions.

FIGURE 8. Task completion rates by gender



Notes: This figure displays the fraction of tasks completed, broken down by priority level and gender.

FIGURE 9. Task type by gender



Notes: This figure displays the share of tasks falling into each category, broken down by gender. Task categories are derived from an unsupervised clustering of Jira issue text descriptions.

Tables

TABLE 1. Regression estimates of learning curves for simple vs. complex tasks

A. All Engineers

	First Year				First 3 Years			
	(1) Junior Tasks	(2) Junior Tasks	(3) Senior Tasks	(4) Senior Tasks	(5) Junior Tasks	(6) Junior Tasks	(7) Senior Tasks	(8) Senior Tasks
Experience	-0.146*** (0.00189)	-0.176*** (0.00218)	-0.112*** (0.00222)	-0.127*** (0.00251)	-0.137*** (0.00106)	-0.163*** (0.00125)	-0.105*** (0.00121)	-0.114*** (0.00140)
Months Since Start		0.0236*** (0.000864)		0.0127*** (0.000998)		0.00667*** (0.000173)		0.00265*** (0.000199)
Observations	178889	178889	141693	141693	466971	466971	376789	376789

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B. Junior vs. Senior Engineers in First Year

	Junior Engineers				Senior Engineers			
	(1) Junior Tasks	(2) Junior Tasks	(3) Senior Tasks	(4) Senior Tasks	(5) Junior Tasks	(6) Junior Tasks	(7) Senior Tasks	(8) Senior Tasks
Experience	-0.165*** (0.00230)	-0.205*** (0.00269)	-0.102*** (0.00363)	-0.110*** (0.00406)	-0.132*** (0.00341)	-0.160*** (0.00391)	-0.143*** (0.00290)	-0.175*** (0.00337)
Months Since Start		0.0299*** (0.00107)		0.00739*** (0.00154)		0.0215*** (0.00148)		0.0251*** (0.00132)
Observations	112218	112218	64893	64893	66671	66671	76800	76800


Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: This table displays estimates of learning rates for simple vs. complex tasks from regressions of the form $\log S_t = \alpha + \beta \log E_t + \varepsilon_t$, where S_t is the average number of days taken to complete tasks assigned in month t since the worker joined the firm, and E_t is the cumulative number of tasks of the corresponding type completed in months 1 through $t - 1$. Coefficients can be interpreted as elasticities: the percent change in average completion time associated with a one percent increase in prior tasks completed. Columns (2), (4), (6), and (8) of each panel additionally include a control for the number of months since the worker joined the firm; columns (1), (3), (5), and (7) do not. Panel A pools all engineers and reports estimates separately for workers' first year and first three years at the firm. Panel B restricts to workers' first year at the firm and reports estimates separately for engineers in junior and senior job titles. The unit of observation is the worker-month. Standard errors clustered at the worker level are reported in parentheses.

Appendix Figures


FIGURE 1. Role descriptions in analysis sample




Notes: This figure shows percent of workers in different job roles, broken down by junior versus senior workers.

FIGURE 2. Validation of complexity predictions


A. Predicted complexity, by worker seniority



B. Predicted vs. Actual complexity, across task categories




C. Predicted vs. Actual Complexity, Within Task Categories



Notes: This figure shows validation exercises for the complexity prediction. Panel A displays the distribution of predicted complexity among tasks assigned to senior workers and to junior workers. Panel B displays the relationship between actual and predicted complexity share for several categories of tasks. The x-axis displays the share of tasks in each category that are actually assigned to a senior worker; the y-axis displays the share of tasks in each category that are predicted to be assigned to a senior worker. Panel C displays the correlation between actual and predicted complexity within each task category. Task categories are derived from Jira labels that are manually inputted by workers.


FIGURE 3. Decomposition of sources of incomplete pass-through in first stage of instrumental-variables design



Notes: This figure presents a binned scatterplot in which the x -axis variable is the mean complexity of tasks within a team one quarter prior to a new worker's entry, and the y -axis variable is the mean complexity of tasks within a team during the quarter of the new worker's entry. The slope and standard error from a corresponding regression are displayed in the top left corner of the figure.

FIGURE 4. Learning curves — balanced panel


A. All workers




Notes: This figure displays the “learning curves” associated with simple and complex tasks — i.e., the average amount of time taken to complete tasks assigned in each month since a worker’s start at the firm. This figure uses a balanced panel, restricting attention to workers who are at the firm for 3+ years.

FIGURE 5. Balance check for instrumental variables design

A. Prior experience and task difficulty upon entry



B. Education selectivity and task difficulty upon entry



Notes: This figure presents placebo outcome tests for our instrumental variables design. Panel A displays a binned scatterplot in which the sample is new entrants to a firm, the x-axis is their team's mean complexity of issues in the quarter prior to entry, and the y-axis is the experience of the new hire, measured in years relative to college graduation. The coefficient and standard error from a regression of years of experience on team mean complexity are displayed in the top left corner. Panel B replicates Panel A, instead using the probability that the new hire attended a selective college as the y-axis variable.

Appendix Tables

TABLE 1. Summary statistics in analysis sample

<i>Panel A: Summary statistics for workers</i>		
	Mean	Standard Deviation
Male (%)	81.4	–
Managerial role (%)	17.0	–
Senior role (%)	40.5	–
Salary (\$)	129,275.2	62,241.55
<i>Panel B: Summary statistics for firms</i>		
	Mean	Standard Deviation
NAICS: Professional, scientific, technical services (%)	62.7	–
NAICS: Retail (%)	5.6	–
NAICS: Wholesale (%)	4.5	–
NAICS: Manufacturing (%)	4.1	–

Notes: Table displays summary statistics about the analysis sample. Panel A displays summary statistics about the workers, including the percent male, percent in managerial positions, percent in senior engineering roles, and salary distribution. Panel B displays summary statistics about the firms, including the number of employees and industries.

TABLE 2. Examples of complexity assigned to tasks

<i>Panel A: Examples of Tasks Predicted to be Junior</i>	
Task Description	Predicted Seniority
Fix misaligned Save button on settings modal in Safari	0.0874
Update error message text on password reset page	0.112
Bump lodash version to resolve Dependabot alert	0.143
Add unit tests for date parsing helper	0.228
Rename <code>usr_id</code> to <code>user_id</code> in CSV export	0.291
<i>Panel B: Examples of Tasks Predicted to be Senior</i>	
Task Description	Predicted Seniority
Define SLOs for payments API and propose alerting thresholds	0.702
Deprecation plan for v1 REST endpoints. Coordinate migration timeline with dependent teams.	0.768
Scope Q4 roadmap items in engineer-weeks and identify cross-team dependencies.	0.824
Evaluate build vs. buy for feature flag service. Present recommendation to Eng leadership.	0.869
Design sharding strategy for events table ahead of Q3 traffic growth. Write RFC and review with Platform team.	0.931

Notes: This table shows stylized illustrative examples of task descriptions alongside the complexity levels predicted by the classification algorithm described in Section 2.4. Predicted complexity is the fitted probability that a task is assigned to a senior rather than a junior worker. Panel A shows tasks predicted to be simple (predicted complexity < 0.5); Panel B shows tasks predicted to be complex (predicted complexity ≥ 0.5).

References

- Adhvaryu, A., Kala, N., & Nyshadham, A. (2022). Management and shocks to worker productivity. *Journal of Political Economy*, 130(1), 1–47. doi: 10.1086/717046
- Babcock, L., Recalde, M. P., Vesterlund, L., & Weingart, L. (2017, March). Gender Differences in Accepting and Receiving Requests for Tasks with Low Promotability. *American Economic Review*, 107(3), 714–747. Retrieved 2024-04-17, from <https://www.aeaweb.org/articles?id=10.1257/aer.20141734> doi: 10.1257/aer.20141734
- Baker, G., Gibbs, M., & Holmstrom, B. (1994). The Internal Economics of the Firm: Evidence from Personnel Data. *The Quarterly Journal of Economics*, 109(4), 881–919. Retrieved 2024-10-01, from <https://www.jstor.org/stable/2118351> (Publisher: Oxford University Press) doi: 10.2307/2118351
- Baker, G., Gibbs, M., & Holmström, B. (1994). The wage policy of a firm. *Quarterly Journal of Economics*, 109(4), 921–955.
- Becker, G. S., & Murphy, K. M. (1992). The division of labor, coordination costs, and knowledge. *Quarterly Journal of Economics*, 107(4), 1137–1160.
- Bircan, Ç., Friebel, G., & Stahl, T. (2025). *Gender promotion gaps in knowledge work: The role of task assignment in teams* (RFBerlin Discussion Paper No. 2518). ROCKWOOL Foundation Berlin.
- Blau, F. D., & Kahn, L. M. (2017, September). The Gender Wage Gap: Extent, Trends, and Explanations. *Journal of Economic Literature*, 55(3), 789–865. Retrieved 2024-04-03, from <https://www.aeaweb.org/articles?id=10.1257/jel.20160995> doi: 10.1257/jel.20160995
- Bol, J. C., Fogel-Yaari, H., Grabner, I., & Sedatole, K. L. (2026). Promoting organizational citizenship behavior (OCB): Unintended gender disparities in the participation and evaluation of different types of OCB. *The Accounting Review*, 101(1), 27–55.
- Chetty, R., Deming, D. J., & Friedman, J. N. (2023). *Diversifying society's leaders? the causal effects of admission to highly selective private colleges* (Tech. Rep.). National Bureau of Economic Research.
- Chu, B., Handel, B., Kolstad, J., & Malmendier, U. (2022, April). Gender Differences in Non-Promotable Tasks: The Case of Clinical Note-Taking.
- Cullen, J. B., Dahl, G. B., & De Thorpe, R. (2025, September). *Job mismatch and early career success* (Working Paper No. 34215). Cambridge, MA: National Bureau of Economic Research. doi: 10.3386/w34215
- Day, A., Corbett, P., & Boyle, J. (2020). Is there a gender gap in chemical sciences scholarly communication? *Chemical Science*, 11(8), 2277–2301.
- Farquhar, S. (2019). *Reaching new heights in the cloud*. <https://www.atlassian.com/blog/platform/cloud-premium>.
- Frederiksen, A., & Kato, T. (2018, August). Human Capital and Career Success: Evidence from Linked Employer-Employee Data. *The Economic Journal*, 128(613), 1952–1982. Retrieved 2024-07-12, from <https://academic.oup.com/ej/article/128/613/1952/5088761> doi: 10.1111/eoj.12504
- Gathmann, C., & Schönberg, U. (2010, January). How general is human capital? A task-based approach. *Journal of Labor Economics*, 28(1), 1–49. doi: 10.1086/649786
- Gibbons, R., & Waldman, M. (2004, May). Task-specific human capital. *American Economic Review*, 94(2), 203–207. doi: 10.1257/0002828041301579

- Gibbons, R., & Waldman, M. (2006). Enriching a Theory of Wage and Promotion Dynamics inside Firms.
- Goldin, C. (2014, April). A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, 104(4), 1091–1119. Retrieved 2024-04-02, from <https://pubs.aeaweb.org/doi/10.1257/aer.104.4.1091> doi: 10.1257/aer.104.4.1091
- Jedwab, R., Romer, P., Islam, A. M., & Samaniego, R. (2023, July). Human Capital Accumulation at Work: Estimates for the World and Implications for Development. *American Economic Journal: Macroeconomics*, 15(3), 191–223. Retrieved 2024-07-16, from <https://www.aeaweb.org/articles?id=10.1257/mac.20210002> doi: 10.1257/mac.20210002
- Kohlhepp, J. (2024, August). The Inner Beauty of Firms.
- Lazear, E. P. (1992). The job as a concept. In W. Bruns (Ed.), *Performance measurement, evaluation, and incentives* (pp. 183–215). Boston: Harvard University Press.
- Lazear, E. P. (2009). Firm-Specific Human Capital: A Skill-Weights Approach. *Journal of Political Economy*.
- Lazear, E. P., & Oyer, P. (2003, December). *Internal and External Labor Markets: A Personnel Economics Approach* [Working Paper]. National Bureau of Economic Research. Retrieved 2024-10-01, from <https://www.nber.org/papers/w10192> doi: 10.3386/w10192
- Lehmann, J.-Y. K. (2013, October). Job Assignment and Promotion Under Statistical Discrimination: Evidence from the Early Careers of Lawyers.
- Levitt, S. D., List, J. A., & Syverson, C. (2013, August). Toward an Understanding of Learning by Doing: Evidence from an Automobile Assembly Plant. *Journal of Political Economy*, 121(4), 643–681.
- Lise, J., & Postel-Vinay, F. (2020, August). Multidimensional skills, sorting, and human capital accumulation. *American Economic Review*, 110(8), 2328–2376. doi: 10.1257/aer.20162002
- McCue, K. (1996). Promotions and wage growth. *Journal of Labor Economics*, 14(2), 175–209.
- Meyer, M. A. (1994). The dynamics of learning with team production: Implications for task assignment. *Quarterly Journal of Economics*, 109(4), 1157–1184.
- Minni, V. (2026). Making the invisible hand visible: Managers and the allocation of workers to jobs. *The Quarterly Journal of Economics*, 141(3). doi: 10.1093/qje/qjag017
- Murphy-Hill, E., Dicker, J., Horvath, A., Hodges, M., Egelman, C. D., Weingart, L. R., ... Chen, N. (2023, April). Systemic gender inequities in who reviews code. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1).
- Waldman, M. (1984). Job assignments, signalling, and efficiency. *RAND Journal of Economics*, 15(2), 255–267.