Big Data and Bigger Firms:
A Labor Market Channel *

Abhinav Gupta† **Naman Nishesh**‡ **Elena Simintzi**§

September 5, 2024

Abstract

This paper studies the impact of employee output information disclosure through GitHub on labor reallocation towards large firms. Github, which is the world's largest software management platform, tracks and publicly displays real-time individual contributions. In 2016, a policy change enabled GitHub users to display their contributions more accurately on their profiles. Following this update, employees with 1 standard deviation higher GitHub contributions witnessed a 5% increase in job transitions to large firms, predominantly at the expense of smaller companies. While productive individuals left small firms for senior roles in larger companies, the latter retained them through internal promotions. The departure of productive workers led to an overall reduction in employment growth and productivity for small firms with more productive employees prior to the shock. Our findings highlight the role of labor-related big data in amplifying the dominance of large firms in recent years.

1 Introduction

The past two decades have witnessed an unprecedented surge in the prominence of large firms (Autor et al. (2020)), sparking a debate about the economic implications of this phenomenon (De Loecker et al., 2020). Several studies have attributed this trend to changes in technology (Crouzet and Eberly, 2018; Lashkari et al., 2018; Autor et al., 2020; Firooz et al., 2022; Hsieh and Rossi-Hansberg, 2023) and the growing importance of online platforms and big data (Veldkamp and Chung, 2024) in enhancing the product market power of large firms. Our paper instead examines how large firms may benefit from the implications of these technology-related trends on the labor market. Specifically, we argue that the abundance of data made available via web-based platforms, which enable employees to provide real-time output information, has implications for labor reallocation and, consequently, the growth of large versus small firms.

The rise of internet-based platforms has revolutionized employee recruitment, with an overwhelming majority of job seekers (over 80%) and employers (over 90%) now utilizing them (Smith (2015); Bradshaw (2024)). Information on individual output available on such platforms provides a signal on employee productivity. The impact of these signals on labor reallocation is not clear ex-ante. On one hand, in the absence of other employee-level information, employment at larger firms may serve as a signal of higher employee quality, as these firms typically offer higher wages due to better worker quality and productivity (Brown and Medoff (1989); Oi and Idson (1999); Troske (1999)). The introduction of new productivity information may diminish the significance of job assignment and employer brand as indicators of quality (Waldman (1984, 1990); Bernhardt and Scoones (1993)), enabling productive employees to work at smaller firms. On the other hand, detailed signals may help larger firms overcome adverse selection and recruit talented employees from smaller firms (Greenwald (2018); Gibbons and Katz (1991)). Additionally, larger firms can invest more significantly in obtaining such data, and better utilize them, due to economies of scale (Brynjolfsson and McElheran (2016)).

To empirically study this question, we focus on GitHub, a software management and version control platform, used by over 100 million worldwide users and 200,000 firms in the United States alone.[1] Users and firms leverage GitHub to host their code and data in repositories, fostering collaborative work through contributions from multiple users. GitHub provides contributors the flexibility to choose the visibility of their repositories — whether public, allowing anyone to view and download the code, or private, restricting access solely to users within the same firm. GitHub's unique feature is its public display of users' contributions, including time stamps, turning it into a hub for recruiters seeking potential employees.[2] The large dimensionality and extensive public use make GitHub an ideal starting point for studying how big data[3] hosted by technology platforms impacts labor markets.

We identify the impact of information on individual output available on GitHub by exploiting a quasi-natural experiment that changed users' ability to signal their quality. Prior to May 2016, the GitHub contribution calendar only displayed public contributions. However, on May 19, 2016, GitHub introduced the option for users to include anonymized private contributions on their profiles. Given these private contributions typically refer to projects managed on Github by employers, the quantity of individuals' private contributions capture their productivity in their current firm. With this update, Github users could instantaneously reveal the number of previously hidden private contributions, without incurring any additional costs. Since revealing the number of private contributions could only improve the profile of an individual at no cost, increasing their outside option, individuals' rational response was to opt in and display their total contributions on their profiles.

We construct a novel database that integrates employee work history with individual output to implement our research design. Employee work history is sourced from LinkedIn, the largest online professional networking platform globally. We link this dataset to

---

[1]https://github.com/about and https://enlyft.com/tech/products/github.

[2]See https://arc.dev/hire-developers/github or https://www.toptal.com/github as examples of how numerous third-party providers assist recruiters in sourcing candidates through GitHub.

[3]GitHub data meet all three criteria for big data as defined by Goldstein et al. (2021): they are large in size (a compressed subset of GitHub repositories as of November 2020 was 21 TB in size https://archiveprogram.github.com/arctic-vault/), possess multiple dimensions, and have a complex layout.

employees' GitHub profiles, which provide records of their daily contributions. Our final sample covers approximately 300,000 individuals and more than 36,000 firms in the U.S. We develop a person-specific measure of output by employing an AKM model (Abowd et al., 1999) on total pre-shock private contributions on GitHub. This is achieved by regressing total private contributions on individual fixed effect, firm fixed effects and time-varying controls, and using the individual component of contributions as our measure of employee output.

At the individual level, we examine whether information on employee-specific output impacts the reallocation of employees between firms. We find that individuals with 1 standard deviation higher productivity increased their mobility to larger firms (those employing over 1000 individuals), by 5% over the sample mean in our most robust specification, after the GitHub policy change. The coefficients remain unchanged regardless of granular dynamic employee and market controls, alleviating concerns of concurrent trends influencing our results. Importantly, the lack of pre-trends in the eight quarters preceding the change strengthens our confidence in attributing the observed impact to this specific event. Moreover, the absence of any significant effects on a placebo sample, where individual output is based on public contributions unaffected by the shock, helps alleviate concerns regarding potential confounding factors. Another concern might be that only individuals looking to move jobs choose to reveal their private contributions, biasing our results. However, our results persist even within users who reveal their contributions and for whom we observe at least one pre-period private contribution, thereby minimizing the likelihood of such bias.

In contrast to the increasing flow of talent towards large firms after the shock, individuals with 1 standard deviation higher productivity are less likely to move to medium-sized firms (those employing between 50 and 1,000 employees) by 5% relative to the sample mean, and less likely to move to small firms (those employing 50 employees or less) by 6% relative to the sample mean. These results are also robust to a dynamic analysis, which shows there are no pre-trends prior to the GitHub policy change. Interestingly, when we estimate our baseline specification by firm size buckets, we find that talent flow increases almost monotonically

across size buckets from -8% at the 11-50 size bucket to a +7% at firms with more than 10,000 employees. As such, the largest gains are captured by the largest employers in the economy, the so-called "superstar" firms.

Moreover, we examine where the talent flow to large firm originates from. Do we observe reallocation of talent between large firms or are these gains to large firms at the expense of smaller firms? Our results suggest that reallocation of talent within large firms drops, while it is small and medium-sized firms who lose their talent to large firms. In economic terms, the larger effects are seen amongst mobility from small firms, and less so from medium-sized firms: individuals with 1 standard deviation higher productivity are more likely to move from small (medium-sized) firms to large firms by 5% (4%) relative to the sample mean, following the shock. Such moves from small/med-sized to large firms are associated with career upgrades for talented employees, who move to more senior or higher-paying jobs in large firms. In contrast, large firms do not lose their talented employees, but they promote them, consistent with the fact that the outside option of talented employees increases after the policy. Overall, treated employees experience career upgrades following the GitHub policy change either by moving towards large firms or by getting promotions within large firms.

We next test the idea that the reduction of information asymmetry regarding employee quality through GitHub should reduce the importance of traditional signals of employee quality. To this end, we consider three traditional signals: the work experience of the individual, having a degree from a highly-ranked university, and the network of universities large firms typically hire from. We find consistent results across all our measures that talented individuals with less experience, with degrees outside of the highly-ranked universities, and outside the traditional networks large firms typically hire from are more likely to move to large firms following the GitHub policy change. These results support the hypothesis that web-based platforms have led to a 'democratization' of screening in the labor markets, allowing talent to flow more freely across firms.

4

The flow of talent towards large firms can be a result of both demand and supply-side factors. On the one hand, employees may prefer to work for large firms as large firms tend to offer steeper career paths to their employees (Mueller et al. (2017); Di Porto et al. (2024)). To test this, we estimate our baseline specification for two subsets of firms: those with above-median average annual salary change and those with below-median average annual salary change. We observe that the coefficient estimates are larger for large firms characterized by steeper career growth as opposed to those with less steep career growth. Individuals with 1 standard deviation higher output are more likely to move to large firms characterized by steeper (less steep) career growth by 7% (2%) relative to the sample mean, following the shock. This potential for increased career growth may come with increased labor risk. To this end, we divide our sample by firm stability, defined as the probability of employees experiencing career downgrades when working for a firm. Our findings indicate that productive employees tend to move to firms with below-median stability following the shock, suggesting that talented employees prefer riskier jobs with increased career advancement opportunities.

On the other hand, large firms tend to invest more in technology to screen workers (Eckel and Yeaple (2017)), and are thus better able to discover talent once the GitHub signal becomes available. To test this, we split firms based on whether one of their Human Resource (HR) personnel uses GitHub or not. We then show that more talented individuals move to large firms with tech-savvy HR departments as opposed to large firms whose HR departments are not tech-savvy. Similarly, large firms already using the GitHub platform will be better versed to identify talent based on their contributions. To this end, we observe that talent flow to large firms is more pronounced when these firms are GitHub users as compared to large firms not on GitHub.

At the firm level, we show that heterogeneity in labor reallocation creates winners and losers. For this analysis, we define the treatment group as firms exhibiting above-median average employee productivity, with productivity determined by the AKM-based individual

component of private contributions during the pre-period. Small firms with above-median individual productivity experienced slower quarter-on-quarter employment growth after the shock, by 13% relative to the sample mean. We also observe a higher likelihood of left-tail outcomes, as the probability of small treated firms exiting the market increased by 36% of the mean after the shock. Small firms with above-median individual productivity also experienced a decline in productivity, proxied by public GitHub contributions (which are not directly affected by GitHub's policy change). Specifically, productivity in small firms employing highly productive individuals pre-treatment, dropped by 5 to 7% post-announcement. Instead, large firms employing high-productivity talent pre-treatment experienced a 5% increase in quarterly employment growth over the mean and a 2 to 5% rise in public GitHub contributions following the policy change.

While there has been considerable discussion on the aggregate trends in local labor market concentration, there is a broad consensus that national industry-level labor concentration has risen in the U.S. since 1980 (Autor et al., 2023).[4] We aggregate our data at the industry level to evaluate the relationship between GitHub private contributions and changes in industry-level labor market concentration in the U.S. from 2011 to 2021. Although these results are based on correlations, they offer interesting insights into the overarching trends. We find that industries with more GitHub contributions before the shock (2013-2016) saw a greater increase in national labor market concentration in the five years after the shock (average concentration over 2016 to 2021) compared to the five years before it (average concentration over 2011 to 2016). This correlation suggests that the availability of labor productivity information through platforms like GitHub may enable larger firms to more effectively identify and recruit top talent from smaller firms, thereby intensifying their dominance in the industry.

Our paper makes significant contributions to three strands of the existing literature. A

---

[4]The market for workers in our sample is national with more than 35% of users in the same firms and roles as our sample reporting work across at least two different locations (states) in the three years prior to the shock.

substantial body of literature has attributed the rise of large firms to technological trends such as the rise of the platform economy (Crouzet and Eberly (2018); Lashkari et al. (2018); Bessen (2020); Firooz et al. (2022); Brynjolfsson et al. (2023); Hsieh and Rossi-Hansberg (2023)) and big data (Tambe and Hitt (2012); Brynjolfsson and McElheran (2016); Begenau et al. (2018); Farboodi et al. (2022); Aghion et al. (2023)), arguing that these technologies provide product market advantages (Prat and Valletti (2022); Eeckhout and Veldkamp (2022)), superior returns on investment (Tambe and Hitt (2012)), and better forecasting abilities (Begenau et al. (2018)) to large firms. Our paper introduces a novel explanation, examining how big data, available via platform technologies, enhance large firms' ability to attract and retain top talent, thus contributing to their growth and dominance.

Second, this paper contributes to the ongoing discussion on the decline in U.S. business dynamism by identifying a new channel influencing this phenomenon. Existing literature has explored various explanations, including a reduction in knowledge diffusion (Akcigit and Ates, 2023; Andrews et al., 2015), decreased responsiveness of firms to productivity shocks (Decker et al., 2020), inefficient allocation of productive resources (Decker et al., 2017; Davis and Haltiwanger, 2014), a rising dependence on intangible assets (De Ridder, 2024), demographic shifts (Karahan et al., 2024), declining interest rates (Liu et al., 2022), strategic hoarding by large firms (Akcigit and Goldschlag, 2023), and deteriorating local housing markets (Davis and Haltiwanger, 2024). Our findings suggest that technological innovations, particularly the enhanced signaling capabilities of web-based platforms, have led to the migration of key talent from smaller to larger firms. This talent shift has contributed to the diminished growth and productivity of smaller firms, offering a novel perspective in the debate on declining business dynamism.

Finally, our paper contributes to the existing literature on labor reallocation between firms of different sizes across firms within the economy. Prior research has extensively documented how regulatory policies influence the distribution of talent between established companies and startups. For instance, Acemoglu et al. (2018) examine R&D taxation,

Pagano and Picariello (2023) study unemployment insurance, Baghai et al. (2024) analyze bankruptcy, Jeffers (2024) explores the impact of non-compete agreements, and Gupta (2023) investigates the effects of immigration policies. However, relatively little attention has been given to the role technological changes play in this trend. Our study fills this gap by exploring how the emergence of big data on employee performance influences the shift of talent towards larger firms.

2 Setting and Data

We use data from GitHub, the world's largest web-based platform in the software development ecosystem. GitHub's user base comprises of over 100 million developers–with more than 20 million in the U.S. alone. Additionally, over 200,000 U.S. firms, including 90% of the Fortune 100 companies use GitHub for software management. In this section, we provide a brief description of GitHub, followed by a summary of how we utilize GitHub's data to create a proxy for employee productivity.

2.1 GitHub: A Web-Platform for Software Management

GitHub is a web-based platform for version control and collaboration in software development projects. GitHub allows users–independent developers and firms alike–to store, manage, and share their data and code streamlined within projects known as *repositories*. By providing centralized online repositories for code storage and management, GitHub enables developers, who might be dispersed geographically, to work collaboratively on projects. It also helps stakeholders track modifications to repositories over time ensuring the integrity and coherence of project versions. Whenever a repository is modified, a snapshot of the repository's contents is automatically created and forever logged with a timestamp. With each change, along with the repository's snapshot, GitHub also records rich metadata: who made the change, what it consists of, and when exactly it was made. Before a proposed

modification is integrated and the new version is released, it must, however, be approved by the repository owners (or designated managers).

An important feature of GitHub repositories is the flexibility of keeping them private or public. Public repositories, used for open-source projects, are accessible to all users, promoting collaboration and knowledge sharing within the entire developer community. Any GitHub user can view the contents, previous versions, contributors or contribution history, and even make changes (subject to approval by repository owners) to a public repository. In contrast, private repositories are restricted to specified users or teams (members invited by the repository owners or designated managers), providing a secure environment for proprietary or sensitive projects. While the content of private contributions is not publicly accessible, the requirement that each contribution must be approved by a repository manager for incorporation into a project makes these contributions difficult to manipulate. While today GitHub users can create an unlimited number of public or private repositories for free, in the past only paid GitHub accounts (developers or firms paying a monthly subscription fee) could maintain private repositories.[5] Subsequently, until recently, private repositories were mostly held by organizations seeking to protect intellectual property and maintain confidentiality while collaborating on development projects.

2.2   GitHub Contributions: A Measure of Productivity

In addition to facilitating version control and collaboration on software projects, GitHub provides users with individual profile pages akin to those found on traditional social media platforms. These user profile pages allow individuals to personalize their online presence by including their name, bio, location, affiliation with firms, and other relevant information. A sample GitHub profile is shown in Figure A.1. GitHub's user profile pages prominently feature the *contribution calendar*, offering a visual representation of a user's GitHub activity

---

[5]GitHub previously charged per repository, with prices ranging from $7 per month for up to 5 private repositories to $200 per month for up to 125 private repositories. In October 2015, they changed this policy to allow unlimited private repositories for $7 per month (Conti et al. (2021)). GitHub made private repositories free in 2019.

over time. This calendar serves as a dynamic record of a user's contributions (in the form of modifications) to GitHub repositories, providing a real-time high-frequency update on a user's work output.

The precise records of users' GitHub contributions provide a way to quantify the output of software developers. GitHub hosts a vast and diverse user base comprising millions of software developers globally, resulting in extensive coverage of coding activity across various domains and projects. Moreover, the granularity and high-frequency of the contributions and the large-scale coverage of GitHub make it an appealing measure of individual output, and thereby of productivity of high-skilled professionals. Although we cannot quantify the importance of contributions (especially the private contributions whose content is not publicly available), all contributions have to be vetted by a repository manager before they are recorded on GitHub, which assures a minimum level of quality.

We validate the usability of GitHub contributions by comparing them to other commonly used measures of employee productivity and quality. Internet appendix Figure A.2 presents binned scatter plots of the logarithm of salary and publications against log contributions. Both plots display strong positive and linear correlations, indicating that total contributions effectively capture employee quality and productivity. Furthermore, GitHub contributions correlate with firm financing for startups (Conti et al., 2021), underscoring their value to sophisticated investors, such as Venture Capitalists (VCs), as a measure of human capital. Other studies similarly use individual GitHub contributions to assess the productivity of software developers, thereby reinforcing our interpretation (Holub and Thies, 2023; El-Komboz and Fackler, 2023; Conti et al., 2023).

Unsurprisingly, recruiters and hiring managers increasingly rely on GitHub user activity to source and evaluate potential tech-sector employees. Figure A.3 presents examples of numerous third-party web platforms that help organizations in recruiting technical talent via GitHub. This shift towards leveraging GitHub contributions as a key metric in the recruitment process underscores the platform's growing significance as a talent marketplace.

Motivated by these facts, we use GitHub as our laboratory to study the impact of web-based labor platforms on labor reallocation and utilize GitHub contributions as a measure of employee productivity.

2.3 GitHub Policy Shift

GitHub implemented a policy change that allowed users to reflect their GitHub contributions more accurately. Before May 19, 2016, users could only display contributions made to public GitHub repositories on their profiles. However, on May 19, 2016, GitHub introduced the option for users to include the number of (anonymized) private contributions to their existing contribution calendar. With a click of a button, users could instantaneously reveal their previously hidden private contributions, with no additional costs. While private contributions were always available internally to the firm, this policy change allowed users to credibly signal their within-firm productivity (verified by GitHub) to external employers for the first time. As an illustration, Figure A.4 shows how the policy changed users' GitHub contribution calendar.

GitHub's new policy added a feature on users' profile pages to "turn on their private contributions" and include the number of (anonymized) private contributions in the calendar. We use users' pre-policy private contributions to construct an exogenously available measure of their output. It is important to note that we observe only the number, not the content, of private contributions, precluding the creation of a quality-adjusted contribution measure. However, the requirement that each contribution be approved by a repository manager for incorporation into a project helps ensure the quality of each contribution. Additionally, we verify in the previous section that the raw number of contributions correlates with other measures of productivity such as salary and publications.

2.4 Data Sources and Sample Construction

We construct a novel dataset using three main data sources: (i) user profiles from GitHub, (ii) GitHub contributions from GitHub API, and (iii) employment records from Revelio Labs. Revelio Labs continuously collects employees' online resumes from various websites (such as LinkedIn) to create one of the world's first universal HR database.

We obtain data on publicly sourced profiles of U.S.-based GitHub users from Humanpredictions, a public data aggregator that developed a proprietary database of over 150 million technical talent worldwide. This data includes users' names, location, coding languages, date of joining GitHub as well as their LinkedIn profile links. We further append this dataset with users' GitHub contributions from GitHub's official API (Application Protocol Interface)[6] The API provides the capability to programmatically access and retrieve all publicly available data on GitHub. We retrieve monthly-level public and private contributions for users from 2011 to 2021. Finally, we merge GitHub users with their employment details from LinkedIn, the largest global online professional networking platform. We obtain individual-level LinkedIn profiles from Revelio Labs. LinkedIn data provide comprehensive information regarding users' educational background (including programs pursued and graduation dates) as well as their employment history (listing firms, positions held, and tenure duration). The dataset also provides us with the direct LinkedIn URLs for each GitHub profile. These merges are made based on publicly available information, where users list both LinkedIn and GitHub URLs together. We are able to successfully merge over 1 million GitHub users with employment data from LinkedIn. Roughly 82% of the matches are obtained by directly merging the LinkedIn URLs in the two datasets. For the remaining matches, we merge on users' names and work history (employer name, start date of job, etc.).

We then construct a matched employer-employee panel detailing each employee's firm

---

[6]We use GraphQL to query GitHub API using the endpoint: `https://api.github.com/graphql`. The documentation can be found at `https://docs.github.com/en/graphql/guides/forming-calls-with-graphql`.

and title each quarter to track employee mobility. In case of an employee listing multiple jobs on LinkedIn at the same time (e.g., volunteering or part-time work along with a primary job), we retain only one primary job. We do so by excluding jobs where job titles include words related to volunteering or part-time employment activity. In case multiple jobs still persist, we retain the one that is listed higher on the user's LinkedIn profile (likely to be more important) and/or that corresponds to a higher seniority level.[7]

As a next step, we append the users' employment records with firm-level identifiers and attributes. For each job, the LinkedIn data provides a URL corresponding to the firm's LinkedIn page that has details on its size, year of incorporation, and industry among others. We use a snapshot of all LinkedIn firm pages as of 2017 to merge these details into our data. We derive the size and age of firms, close to the time of GitHub's policy change, using these appended details. Additionally, we also merge the firm URLs with an auxiliary dataset containing firm identifiers from FactSet (also provided by Revelio Labs). This dataset includes a firm's unique FactSet ID, its parent firm information, as well as its GVKEY in case the firm is publicly listed. For a small fraction of firms that remain unmatched, we also use firm data from Crunchbase to obtain the firm's size, age, and other identifiers.

In our empirical analysis, we utilize GitHub's policy change regarding users' contribution calendar and focus on its impact on their employment outcomes. Subsequently, we filter out a relevant set of GitHub users for our empirical setting from the merged sample in two steps. First, we restrict the sample to users who had joined GitHub and had made at least one GitHub contribution before the policy change on 19th May, 2016. And second, we limit the sample to users who had completed their highest educational degree before the policy change, namely we exclude users who were still studying at some point after May 2016. These filters result in an initial sample size of approximately 380,000 US-based employees across 103,000 firms.

We next filter our data to create a consistent sample of employees where we can estimate

---

[7]Revelio assigns each job a seniority score between 1-4 based on the title and job description among other details.

individual component of productivity using the Abowd et al. (1999) (AKM) methodology. Table 2 Panel A presents the sample estimates. 93% of the initial sample are part of the connected set, representing the largest grouping of firms where employees move between the same organizations. After filtering to retain only firms with at least two employees observed in any given month, we are able to estimate regression coefficients for 89% of the sample. Further refinement involves focusing solely on employees employed at for-profit firms at the time of the policy change. This filtering process yields a final sample size of around 300,000 employees, constituting 80% of employees in the initial sample. Within this final sample, we identify approximately 36,000 firms, representing 35% of the raw sample, as the AKM requires discarding firms where multiple employees are not observed within the same role and month.

2.5   Summary Statistics

Table 1 provides summary statistics for our sample. Panel A presents cross-sectional distributions, showing that 80% of employees are male, 70% are white, and 40% hold software engineering roles. Internet Appendix Table A.1 offers further breakdowns by industry, coding language, and roles. The Software, Internet, and Information Technology sectors constitute over 45% of the sample. Figure A.2 presents the geographic distribution of our sample. The data reveals that most counties in the U.S. with significant technology clusters are represented, with New York, San Francisco, and Seattle comprising a quarter of the sample. On average, employees have eight years of experience and a salary of nearly ninety thousand dollars. Comparatively, the Bureau of Labor Statistics (BLS) data indicates that software engineers are 78% male, 60% white, with over 15% living in tech hubs like New York, Seattle, or San Francisco, earning an average salary close to 98,000 dollars.[8] Thus, our sample is representative of the median American software worker.

Panel B presents time-varying employee outcomes. The average probability of an

---

[8]Data obtained from Bureau of Labor Statistics table: `https://www.bls.gov/oes/tables.htm`

employee moving is 7% per quarter, indicating an average tenure of 3.6 years, which aligns with the median tenure for employees in technical services (3.9 years) as reported by the BLS. Of those who move, 40% transition to large firms (over a thousand employees), 30% to medium firms (fifty to a thousand employees), and 30% to small firms (less than fifty employees). More than 55% of these moves result in a rank or salary increase, while 2% of employees are promoted within the same firm each quarter.

Panel C presents firm-level characteristics. In our sample, 15% of firms are large (over a thousand employees), 37% are medium (fifty to a thousand employees), and 48% are small (less than fifty employees). However, as shown in Internet Appendix Table A.2, large firms employ more than half of the workers in our sample, similar to estimates from the Longitudinal Business Database (LBD), where large firms account for 16% of firms but employ 47% of the workforce. The average firm in our sample grows at 3% per quarter, comparable to 2.5% employee growth for all firms in the LBD. Most of these firms are active on GitHub, with the average firm seeing ten public contributions per employee per quarter.

2.6 Extracting Employee Productivity: AKM Model

We first create our measure of individual productivity based on individuals' output-GitHub contributions. The GitHub policy change revealed new information on the total number of employees' private contributions, encapsulating both individual and firm-level productivity. Hiring departments often benchmark these contributions relative to peers, factoring in the influence of the specific firm environment on employee performance. To disentangle the individual component of productivity, independent of the firm influence of the firm they work for, we adopt the methodology proposed by Abowd et al. (1999) (AKM):

$$Log(PvtContributions)_{i,f,t} = \vartheta_i + \vartheta_f + \vartheta_t + \vartheta_k + X_{i,f,t} + \varepsilon_{i,f,t} \tag{1}$$

The dependent variable represents the logarithm of one plus the number of private

contributions made by employee i, while employed at firm f during month t. We partition each employee's private contributions into individual, firm, role, and monthly components using employee (ϑ_i), firm (ϑ_f), role (ϑ_k) and month (ϑ_t) fixed effects. Additionally, $X_{i,f,t}$ includes controls for current job duration, and total work experience, ensuring that we account for variations attributable to different roles or employee seniority. Our regression analysis spans May 2011 to April 2016, precisely five years preceding the GitHub policy change. This time-frame enables us to discern how pre-existing productivity information influenced employee job outcomes. Specifically, we utilize the employee-specific private contributions over this period (ϑ_i) as our proxy for individual-level productivity information.

Table 2 presents summary statistics for the AKM. Panel A shows that our connected set covers 93% of the employees in the initial sample. We are able to calculate AKM estimates for 80% of the initial sample employees, after applying filters as detailed in Section 2.4. Panel B unveils that 67% of the total variation in contributions stems from differences in employee output, with an additional 8.5% attributed to disparities in average firm output. Thus, individual human capital explains 7.8 times the variation in contributions at the firm level. This finding aligns with prior research on innovation, which suggests that individual-level capabilities account for 5-10 times more variation in inventor patenting output compared to firm-level factors (Bhaskarabhatla et al. (2021)).We also test the persistence of individual productivity estimates from the pre-shock period (2011–2016) to the post-shock period (2016–2021). Figure A.5 presents these results, indicating that the majority of employees remain within similar productivity quintiles before and after the shock. This persistence suggests that our estimates capture inherent differences in individual quality, rather than any attempt by employees to manipulate these metrics. The results highlight that the GitHub policy change regarding employee private contributions can be interpreted as a new signal of employee productivity, with inherent and persistent differences among employees accounting for the majority of variation in these contributions.

To address recent concerns about the validity of using the logarithm in case of count-like

16

variables with several zeros and skewed distributions (Cohn et al. (2022); Chen and Roth (2024)), we further perform two analyses. First, we employ alternative AKM specifications, including continuous measures of contributions and categorical variables. Second, we demonstrate that our findings hold even without using the AKM framework by utilizing private contributions and the share of private contributions to define employees productivity. Our results remain consistent across all these approaches, reinforcing the robustness of our findings. Detailed discussions and the outcomes of these robustness checks are provided in Section 3, where we describe the employee-level results.

3 Employee Level Reallocation

3.1 Mobility to large firms

In our baseline analysis, we examine whether improved information on employee quality influences the reallocation of employees between firms. We carry out a differences-in-differences analysis where we compare the change in employee mobility for individuals with higher AKM-implied productivity after the GitHub policy change, as compared to those with lower productivity. Our baseline specification is as follows:

$$Y_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \delta_t + \gamma_{it} + \varepsilon_{i,t} \tag{2}$$

where, the dependent variable Y is a binary indicator that equals 1 if employee i switches jobs in quarter t, and 0 otherwise. $Productivity_i$ represents the normalized version of the individual component of private GitHub contributions (ϑ_i), as obtained from Equation 2 in section 2.6. We normalize $Productivity_i$ by subtracting the sample mean and dividing by the standard deviation to enhance the interpretability of our coefficient estimates. The coefficient of interest, denoted as β, captures the interaction between the AKM-derived productivity estimate for each employee i and an indicator that switches on after GitHub's

policy shift in May 2016. We include individual (α_i) and (δ_t) quarter fixed effects to control for any time-invariant individual differences and time trends. To ensure the robustness of our findings, we sequentially add several controls to our model. First, we account for employee's current salary and seniority (X_{it}). We also include fixed effects that interact together the employee's cohort of graduation, most used coding language, first job role, and quarter fixed effects to account for time-varying job trends. Additionally, we control for employee's first industry (based on LinkedIn's 144 industry categories) interacted with the first geographical area (metropolitan statistical area) and time. This ensures that we are comparing employee transitions within the same market-quarter context. Our sample covers a period of six years: from the third quarter of 2013 (July to September quarter, three years before the shock) to the second quarter of 2019 (April to June quarter, three years after the shock). We cluster our standard errors at the individual employee level, which is the treatment level in our analysis.

Table 3 presents our findings. Individuals with productivity estimates higher by 1 standard deviation augmented their likelihood of transitioning to larger firms (employing over 1000 individuals) by 4 to 5.7% above the sample mean following the shock, depending on the specification. Note our results remain statistically significant and consistent in magnitude across columns 1 to 5, as we progressively introduce controls from solely employee and time fixed effects to incorporating time-varying employee-level controls (salary and seniority), dynamic job-level (cohort-programming language-role-time), and market-level (industry-M.S.A-time) fixed effects. We also test whether the moves to large firms are accompanied by changes in industry and location, as shown in Table A.7. We find that the probability of moving to a larger firm increases by 6% over the sample mean when accompanied by a location change, compared to a 4% increase when the move occurs within the same location. Furthermore, the likelihood of transitioning to a large firm across both similar and different industries rises by 5% over the sample mean. These findings support the idea that enhanced signaling plays a more crucial role in cross-location moves, as it helps

18

compensate for the absence of localized networks.

We also estimate the dynamic version of Equation 2 as follows:

$$Y_{i,t} = \sum_{\tau=-8}^{12} \beta_\tau Productivity_i \times \mathbb{1}(\tau = t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t} \tag{3}$$

The model is similar to the one detailed in Equation 2, except we estimate a separate β_τ for each quarter τ relative to the shock. The omitted quarter is the first quarter of 2016 (January to March 2016), one quarter before the shock. We test the parallel trends identifying assumption of the difference-in-differences methodology, by checking for any pre-trends in β_τ before the shock.

The event study corresponding to Table 3 Column 5, depicted in Figure 2, illustrates that employees with higher private contributions pre-treatment exhibit parallel trends with employees with lower private contributions. In other words, each pre-period coefficient is statistically indistinguishable from zero. Notably, there is a sudden and sustained surge in employee mobility among treated employees in the quarter following the GitHub announcement, indicative of a rapid change in information disclosure driving the observed outcomes.

3.1.1 Identification challenges

Two primary challenges may undermine a causal interpretation of our estimates. First, concurrent changes in factors correlated with employee productivity could potentially yield similar results. However, as shown in Table 3, the coefficients remain robust to granular dynamic controls for both employees, firms, and markets. Additionally, we use a placebo sample of employee public contributions (those unaffected by the shock) to further rule out the impact of unobservable employee characteristics. Unobserved employee traits should equally impact both public and private contributions. Therefore, if our results are driven by concurrent changes in unobserved individual-specific factors, we should observe comparable

results when using public contributions as our independent variable. We present the results in Table 4, Panel A. We decompose pre-shock public contributions using the AKM methodology, similar to our baseline measure, to obtain the individual component of public contributions. We find no significant change, either economically or statistically, in the mobility of employees with 1 standard deviation higher public contributions-based productivity after the shock. This null result supports our conclusion that the observed effects are driven by signaling based on private contributions.

We further validate our results by examining heterogeneity based on the enforcement of non-compete agreements. Previous research has demonstrated that non-competes significantly restrict employee mobility (Jeffers, 2024). Hence, we expect a stronger impact of signaling on mobility in states where non-competes are not enforceable. Table A.8 presents our findings. As anticipated, the data reveals a substantially larger increase in mobility for employees residing in states without enforceable non-compete agreements, consistent with the hypothesis that we capture the impact of signaling on labor mobility. Additionally, we test if our results are influenced by trends correlated with the 2016 presidential elections or other local economic characteristics. If our results were driven by the elections or differences in local economic conditions, we would expect to find different impacts in areas based on their political leanings or economic characteristics. However, Table A.9 shows no differential impact on counties that were more Republican, white, immigrant-heavy, richer, or more unequal.

As a final validation, we analyze changes in firm stock prices following the GitHub announcement. This analysis helps isolate the immediate daily impact of the GitHub announcement, controlling for any longer-term changes.[9] Using CompuStat data from 1,679 firms, we categorize firms two ways: by their average productivity (above or below median), and by size (with indicators for the top and bottom quartiles). Table A.11 shows a 0.65 basis point decline in returns for small firms (bottom quartile) with high GitHub private

---

[9]Bloomberg does not report any major tech-related news announcements from 17 May 2016 to 20 May 2016

contributions (above median), indicating market anticipation of talent loss. No significant impact is observed for larger firms (top quartile) or smaller firms without GitHub exposure (bottom quartile in size and below median average productivity).

Moreover, there's a risk that individuals may endogenously opt to signal their pre-shock productivity when they anticipate moving to larger firms, potentially biasing our estimates upwards. To address this concern, we present results limited to active GitHub users. We test our findings across four alternative samples: users active on GitHub before the shock (with more than 10 pre-period total contributions), users whose employers were active on GitHub at the time of the shock (firms with at least 100 total contributions as of the shock date), users with at least one private contribution before the shock, and users with at least ten private contributions before the shock. Table 4, Panel B, presents the results. Although filtering on users reduces our sample size significantly, our results persist both in magnitude and statistical significance for all samples, diminishing the likelihood of such a bias driving our findings.

We also test the robustness of our results to different specification choices. Table A.3 presents these findings. First, as explained in section 2.6, we might be worried about the impact of using logarithm of one plus contributions as the dependent variable in the AKM calculations. We show that our results are robust to using other measures in the AKM calculation in Panel A, including using the arcsine transformation of contributions, categorical variables, and the raw contribution numbers. Panel B shows that the results persist even if we use the raw number of private contributions instead of AKM-implied productivity. The consistency across specifications reinforces the validity of our findings. Another concern might be that our results are capturing differences in working styles (e.g., making multiple small contributions versus several large ones). We address this by showing that our results are robust when using the share of private contributions over total contributions as dividing by total contributions helps control for these individual-level idiosyncrasies. A concern might stem from the usage of a continuous variable as our

21

treatment in a diff-in-diff setting, as detailed by Callaway et al. (2024). We show our results are robust to using an indicator for non-negative productivity in Panel C.[10] Panel D examines the sensitivity of our findings to the time period over which we estimate the AKM, Panel E explores the impact of alternate clustering methods, Panel F tests the effect of more granular fixed effects, and Panel G shows that are results are robust to estimating our coefficients on a matched treatment and control sample.[11]

3.2 Mobility to medium and small firms

We next examine the impact of the GitHub announcement on the mobility of productive employees to small and medium-sized firms. We use a specification akin to Equation 2, except we now consider transitions to medium-sized firms (50 to 1000 employees) and small firms (less than 50 employees) as the outcome variables. The results, outlined in Table 5, reveal that individuals with productivity estimates higher by 1 standard deviation reduced their likelihood of transitioning to small and medium firms by 6% and 4.8% of the sample mean post-shock, respectively. Event studies depicted in Figure 3 validate the absence of pre-trends in all of those outcomes pre-treatment, supporting the parallel trends assumption.

We further replicate our analysis by non-parametrically estimating our baseline specification across different firm size categories, as illustrated in Figure 4.[12] Our findings indicate that the mobility of employees with 1 standard deviation higher productivity increases almost monotonically with firm size, ranging from -8% of the sample mean for firms with 11-50 employees to +7% for firms with more than 10,000 employees. An exception to

---

[10]Callaway et al. (2024) argue that the parallel trends assumption may not rule out selection bias in difference-in-differences designs with a continuous treatment variable. Subsequently, in Figure A.6, we show that our parallel trends assumption is robust to using a dummy indicator for treatment.

[11]In Panel D, we estimate the AKM using contributions from the pre-sample period, up until the the second quarter (April to June) of 2014. Panel E presents results after double clustering at the industry and location as well as at the employee and time-varying firm levels. Panel F involves interacting all fixed effects, including cohort, role, language, location, industry, and time. We match treated employees (those with positive normalized productivity estimates) with control employees (having negative normalized productivity estimates) on gender, race, industry, location, experience, salary, and seniority in Panel G1. We additionally match on public contributions in Panel G2.

[12]Table A.4 shows the regression counterparts of Figure 4.

this trend is observed for firms with fewer than 10 employees, where the decrease in mobility is economically small, at -2%, and statistically insignificant. This could be attributed to entrepreneurs starting their own firms within this category, making it less directly comparable to other categories consisting of employees joining established firms.

Overall, our findings indicate that the largest gains are captured by the largest employers in the economy, the so-called "superstar" firms. We confirm the robustness of our results with various definitions of superstar firms. Internet Appendix Table A.5 shows that the increase in mobility holds when defining superstars as the top 1% and 10% by market capitalization and revenue, as defined in Tambe et al. (2020). Moreover, internet appendix Table A.6 reveals that even within young firms, mobility increases to only "superstar" startups with over \$500 million in funding.[13] Thus, the GitHub policy change facilitates a redistribution of employees from small and medium-sized firms to the largest firms in the economy (irrespective of age).

3.3 Heterogeneity by initial firm size

The previous specification examines the changes in employee mobility following GitHub's policy change. Next, we aim to additionally identify the origin firms for these employees to establish patterns of labor reallocation. Therefore, we analyze the post-GitHub policy change in the mobility of productive employees by firm size. Specifically, we estimate a triple-difference specification:

$$
\begin{aligned}
\mathbb{1}(MoveToLargeFirm)_{i,t} = {} & \beta_1 Productivity_i \times \mathbb{1}(Post_t) + \\
& \beta_2 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Mid_{i(t-1)}) \\
& + \beta_3 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Small_{i(t-1)}) + \lambda X_{it} + \alpha_i + \delta_t + \gamma_{it} + \varepsilon_{i,t}
\end{aligned}
\tag{4}
$$

The specification is an extension of Equation 2, incorporating a triple interaction involving three indicators: the productivity estimate of employee i derived from the AKM decomposition, an indicator that equals 1 for quarters following GitHub's policy change,

---

[13]Startups with at least \$500 Million have a median size of 1000-5000 employees.

and indicators representing the size of the firm where the individual worked in the previous quarter. Here, *Mid* denotes an indicator variable that equals 1 for firms with 50-1000 employees, and *Small* represents an indicator for firms with fewer than 50 employees. While we also estimate the coefficients for double interaction terms between individual productivity and firm-size indicators, and firm-size indicators interacted with indicators for GitHub policy change, we do not report them in the specification or results for brevity. β_1 captures the impact of the policy change on more productive employees working at large firms, while β_2 and β_3 assess the additional impact on employees currently employed at mid-sized and small firms, respectively. This specification, thus, enables us to examine how the joint interaction of employee productivity, the GitHub policy change, and firm size influences the outcome of interest.

We also estimate the dynamic version of Equation 4, to verify the absence of pre-trends, as follows:

$$
\begin{aligned}
\mathbb{1}(MoveToLargeFirm)_{i,t} = &\sum_{\tau=-8}^{12} \beta_{1,\tau} Productivity_i \times \mathbb{1}(\tau = t) \\
&+ \sum_{\tau=-8}^{12} \beta_{2,\tau} Productivity_i \times \mathbb{1}(\tau = t) \times \mathbb{1}(Mid_{i(t-1)}) \\
&+ \sum_{\tau=-8}^{12} \beta_{3,\tau} Productivity_i \times \mathbb{1}(\tau = t) \times \mathbb{1}(Small_{i(t-1)}) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}
\end{aligned}
$$

Table 6 presents the results. We find that mobility to large firms for individuals with productivity estimates higher by 1 standard deviation, who were currently employed at large firms, decreased by 12% of the sample mean, post-GitHub announcement. Conversely, the probability of individuals with 1 standard deviation higher productivity, currently employed at medium and small firms, increased by 4% and 5% of the sample mean post-shock, respectively. The absence of pre-trends, when plotting the triple-difference coefficient for both mid-sized small firms in Figure 5, underscores that we are capturing the impact of the change in GitHub policy. Our results hence indicate that productive employees

24

currently employed at small and medium firms migrated to larger firms, leading to an overall re-allocation of talent following the GitHub policy change.

3.4   Impact on promotions

Our previous analysis has primarily focused on job changes as the outcome variable. Now, we aim to understand the impact of the GitHub policy change on promotions, encompassing both salary increases and title advancements, within and across firms. The results, presented in Table 7, dissect promotions into three distinct categories: those occurring when employees transition to large firms, those happening when employees transition to small and medium firms, and those occurring within the same firm. We maintain the baseline specification outlined in Equation 2 for odd columns and adopt Equation 4 for even columns. Additionally, we introduce firm interacted employee cohort fixed effects for within-firm promotions to ensure comparisons across employees within the same firm and with equal tenure. Our findings, summarized below, document the impact of the GitHub policy change on promotions.

Column 1 demonstrates that the likelihood of transitioning to a large firm with an increase in salary and title rose by 5.4% above the mean for individuals with 1 standard deviation higher productivity, following the GitHub announcement. Column 2 shows that this increase is mainly driven by employees from mid and small-sized firms. Specifically, the probability of moving from one large firm to another with promotion decreased by 9% of the sample mean for those individuals with 1 standard deviation higher productivity after the GitHub shock. Conversely, the probability of transitioning from small and medium firms to large firms with promotion increased by 5% of the sample mean for those with 1 standard deviation higher productivity post the GitHub announcement. These findings suggest that large firms attracted productive employees from small and medium firms by offering improved salary and title, after the GitHub change.

In contrast, Column 3 reveals that the probability of moving to a small or medium firm

with an improvement in salary and title decreased by 6.1% of the sample mean for individuals with 1 standard deviation higher productivity, following the GitHub announcement. Column 4 reveals that this decrease is mainly attributed to employees from large firms. Specifically, the probability of moving from one large firm to small or medium firm with promotion decreased by 12.7% of the sample mean for those with 1 standard deviation higher productivity after the GitHub shock. Conversely, there was less than 1% change in probability of moving from one small and medium firm to another with promotion for those with 1 standard deviation higher productivity after the GitHub shock. These results suggest that small and medium firms were unable to attract highly productive employees from large firms with better title or pay.

Column 5 notably illustrates that current employers also react to the disclosure of new productivity information, showing a significant increase of 4.9% over the mean in the probability of within-firm promotions for employees with 1 standard deviation higher productivity post-shock. We also plot the dynamic event study for these results in Figure 6 and observe no pre-trends. Column 6 further demonstrates that this effect is particularly pronounced for large firms, which experience a post-shock increase of 6.2% over the mean in within-firm promotion probability for employees with 1 standard deviation higher productivity. In contrast, the same effect is 4.2% of the mean for small firms, diminished by the impact of productive employees transitioning to more senior positions at large firms following the shock. This increase in internal promotion probability aligns with existing theoretical models (e.g., Waldman (1984); Bernhardt and Scoones (1993)). These models suggest that firms are hesitant to promote talented employees without credible signals because such promotions could alert other firms to their quality, increasing the risk of losing them to competitors.

In summary, our findings indicate that the sudden disclosure of productivity information on GitHub was advantageous for productive employees, facilitating their career progression. While large firms demonstrated an ability to retain employees through internal promotions,

26

small firms appeared to lose valuable talent to larger firms, which offered these employees better career opportunities.

4 Potential mechanisms

4.1 Heterogeneity by Individual Characteristics

Previous literature has highlighted significant information asymmetry regarding employee quality, suggesting that access to productivity information could help employers better screen candidates mitigating this asymmetry (Waldman (1984); Greenwald (2018)). We examine this mechanism by exploring the heterogeneity of our results across three types of alternate individual signals: work experience, degree from a high-ranked university, and education-based employment networks. We present these results in Table 8.

Columns 1 and 2 examine the heterogeneity of our results based on the total number of years of employee experience since graduation at the time of the shock. We find that more productive employees with less than median experience (i.e., 7 years) at the time of shock saw a 1.5 times larger increase in post-shock mobility to large firms compared to those with more than median experience. This suggests that productivity information from GitHub served as a substitute for work experience, enabling less experienced but productive employees to be employed by larger firms.

Columns 3 to 6 assess the heterogeneous impact of having a degree from a highly-ranked university on our results. We define elite universities as those within the top 30 ranked universities according to U.S. News and World Report (Columns 3 and 4) and QS Global University Rankings (Columns 5 and 6). Our findings indicate that productive employees without an elite university affiliation experienced a mobility increase to large firms that was 2.5 to 4 times greater than the change observed for employees with an elite university degree. Our results, so far, suggest that the increased availability of information on employee productivity may reduce reliance on hard information, such as work experience or university

rankings, as a criterion for mobility to large firms.

Finally, columns 7 and 8 examine the differential impact of mobility based on the employment networks. We use the employees' schools at large firms as a proxy for such networks. Specifically, we define high-network schools as the top five recruitment destinations for all large firms in our sample. Our analysis shows that more productive employees from low-network schools experienced a 1.5 times greater increase in post-shock mobility to large firms compared to employees from high-network schools. These results suggest that the provision of productivity information reduces reliance on relationship-based information, enabling firms to recruit talented employees from outside their traditional recruitment networks.

Overall, our results support the hypothesis that web-based platforms have reduced reliance on both hard signals, such as experience and education, and network-based screening, leading to a "democratization" of labor markets. This has facilitated a less constrained flow of talent across firms, allowing employees to advance based on productivity rather than traditional metrics.

4.2   Heterogeneity by Firm Characteristics

The flow of talent towards large firms can result from both supply and demand-side factors. On the supply side, previous literature has documented that large firms offer steeper career paths than smaller firms (Di Porto et al. (2024)). These steeper career paths, and the resulting higher wage inequality, have also been demonstrated to correlate with enhanced employee and firm productivity (Lemieux et al. (2009); Mueller et al. (2017); Wallskog et al. (2024)). Hence, a preference for such steeper career paths may drive more productive employees toward larger firms. On the demand side, larger firms may have better IT resources (Eckel and Yeaple (2017)), enabling them to better utilize and recruit based on GitHub signals. We test these hypotheses using heterogeneity tests.

First, we classify firms based on their probability of offering salary increases. We

calculate the average annual salary change for all employees in a firm using LinkedIn data for individuals with the same job roles as those in our GitHub sample during the pre-treatment sample period (quarter 3 2013 to quarter 2 of 2016). Firms are then divided into two subsets: those with above-median average annual salary change and those with below-median average annual salary change. Table 9, columns 1-2, present the results. We find that individuals with productivity 1 standard deviation above the mean were 6.8% more likely (compared to the sample average) to move to large firms with steeper salary growth after the GitHub policy change. In contrast, the increase was only 1.5% of the sample mean, and economically insignificant, for mobility to large firms with below-median salary growth. T-tests confirm that the difference between these two coefficients is significant at the 1% level. These results support the notion that employees prefer to join large firms due to the potential for faster career growth.

Second, we assess whether the increase in career growth potential increases job risk. Using the entire LinkedIn dataset for roles matching those in our GitHub sample during the pre-treatment sample period (quarter 3 2013 to quarter 2 of 2016), we determine the probability of forced turnover for each firm. Forced turnover is defined as job separation with a gap of at least three months, accompanied by a decrease in salary or seniority. We categorize firms into two groups based on their forced turnover probability relative to the median for all firms in our sample. Columns 3-4, present the results. Individuals with 1 standard deviation higher productivity increased their mobility to larger firms with above-median forced turnover by 6.2% above the sample mean following the GitHub policy change. In contrast, there was a statistically insignificant decrease of 1.2% of the sample mean for mobility to large firms with below-median forced turnover. T-tests confirm that the difference between these two coefficients is significant at the 1% level. These findings show that the increase in steep career growth is accompanied by a shift to "riskier" jobs for more productive employees following the GitHub shock, suggesting that productive human capital prefers jobs with high upside.

Third, we examine the role for tech-savvy recruiters and firms' own GitHub usage. Columns 5-6 in Table 9 present the results, splitting firms based on whether their Human Resource (HR) personnel use GitHub. We find that individuals with productivity estimates 1 standard deviation above the mean increased their probability of moving to large firms with tech-savvy HR by 8.4% of the sample mean after the GitHub policy change. In contrast, the increase was only 1.3% of the sample mean for moves to large firms without any recruiters on GitHub. T-tests show these coefficients to be significantly different from each other at the 1% confidence level.

Columns 7-8 in Table 9 further split the sample based on whether firms themselves have created GitHub profiles. Large firms with a presence on GitHub may be more adept at utilizing these signals for recruitment. Consistent with this idea, we find that the probability of an employee with 1 standard deviation higher productivity moving to a large firm with its own GitHub account post-shock increased by 5.4% of the sample mean. In comparison, the increase was 3.6% of the sample mean for large firms not on GitHub. The difference between these two coefficients is statistically significant at the 5% confidence level. Together, these tests support the hypothesis that firms who invest more in, or possess better, screening technologies leverage new information from web-based platforms for recruiting talent.

5　Firm Level Impact

Section 3 demonstrates that high-productivity employees reallocated from small and medium-sized firms to large firms as a result of the GitHub productivity disclosure. It is crucial to assess whether this labor reallocation had an aggregate impact on firm outcomes. If small and medium firms are capable of replacing workers or internally redistributing projects, we would expect the disclosure to have no impact on firm growth or productivity. However, if these firms are constrained in their ability to hire and train talented employees, we would anticipate observing an aggregate impact.

We, thus, examine next the effect of the GitHub policy change on firm growth and productivity. To this end, we aggregate our data to the firm level and use a triple-differences specification. The basic difference-in-differences setup compares the changes in firms with a higher proportion of high-productivity (treated) employees with firms with a low proportion of high-productivity employees, before and after the GitHub announcement. Subsequently, the third difference dissects this coefficient based on firm size, allowing us to discern the divergent impact of the GitHub policy change on treated small and medium-sized firms compared to treated large firms. We estimate the following regression equation:

$$
\begin{aligned}
FirmOutcome_{j,t} = {} & \beta_1 \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(Post_t) \\
& + \beta_2 \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(Post_t) \times \mathbb{1}(Mid_j) \\
& + \beta_3 \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(Post_t) \times \mathbb{1}(Small_j) + \lambda X_{jt} + \alpha_j + \gamma_{j,t} + \varepsilon_{j,t}
\end{aligned}
\tag{5}
$$

In our regression model, the dependent variable is an outcome variable for firm g at quarter t. The coefficient β_1 captures the interaction between $\mathbb{1}(Productivity_j > Median)$, an indicator that takes the value 1 if the average AKM-derived productivity estimate is greater than the median for firm j, and $\mathbb{1}(Post_t)$, an indicator that switches on after GitHub's policy shift in May 2016. Coefficients β_2 and β_3 further interact these variables with indicators for firm size. Our sample covers a period of six years: from the third quarter of 2013 (three years before the shock) to the second quarter of 2019 (three years after the shock). To ensure robustness of our analysis, we control for firm and industry-state-time fixed effects, controlling for constant firm-level differences and market level time trends on our results. Additionally, we incorporate other relevant controls. In the case of employment-related outcomes, we control for lagged employment level, while for contribution-based productivity outcomes, we control for lagged total contribution stock and the median pre-period contribution interacted with $\mathbb{1}(Post_t)$. We cluster our standard errors at the firm level.

We also estimate a dynamic version of the same regression as follows:

$$
\begin{aligned}
FirmOutcome_{j,t} = & \sum_{\tau=-8}^{12} \gamma_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \\
& + \sum_{\tau=-8}^{12} \theta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \times \mathbb{1}(Mid-SizedFirm_j) \quad (6) \\
& + \sum_{\tau=-8}^{12} \beta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \times \mathbb{1}(SmallFirm_j) + \lambda X_{jt} + \alpha_j + \gamma_{jt} + \varepsilon_{j,t}
\end{aligned}
$$

The model is similar to the one detailed in Equation 5, except we estimate a separate β_τ for each quarter τ relative to the shock. The omitted quarter is the first quarter of 2016, one quarter before the shock. We test for the parallel trends assumption by checking for any pre-trends in β_τ before the shock.

We present the results for employee growth and productivity in Table 10. Columns 1 and 2 present the results for firm employment growth. We find that large firms with more than median individual productivity grew 3.5 to 4.8% faster than the mean after the GitHub announcement, depending on specification. Above-median productivity med-sized firms grew less by 2.5% than large firms following the policy shift, although the difference is not statistically significant. On the other hand, employment growth in above-median productivity small firms decreased by 12.9% of the mean after the announcement. Column 2 creates a proxy indicator for firm exit, defined as a firm experiencing a headcount reduction of more than half compared to the previous quarter. Our findings show no significant change in the probability of closures for large or medium firms with above-median productivity levels post-GitHub shock. However, small firms with above-median productivity witnessed a substantial 35% increase over the mean in the probability of closures, after the GitHub announcement.[14] Figure 7 demonstrates that our results are not driven by differential pre-trends across firms. Overall, these findings support the notion that large firms benefited from the information revealed as they could discover new talent, while small firms experienced

---

[14]The magnitude appears large because the pre-period mean of closures is very low (0.6%), particularly due to large firms that are much less likely to exit. The pre-period mean of exits among small firms is 1%, suggesting that the shock leads to about 21% increase in their probability to exit.

losses and were unable to replenish talent effectively.[15]

Columns 3 to 6 use instead employee productivity as the outcome variable. Here, we utilize the public contributions for each employee, as these were not directly influenced by the shock. We exclude any contributions made by an employee to their own repository, as these may represent personal activity. The results show the number of public contributions increased by 5.3% for large firm with productivity estimates exceeding the median post the GitHub announcement. There was a 1.7% (though statistically insignificant) increase in public contributions per employee for large firms with above-median productivity levels post-GitHub shock. We find statistically indistinguishable changes for med-sized firms. Notably, the entire positive impact is negated for small firms with above-median productivity, which witnessed a 4.5% reduction in the number of contributions and a 6.5% reduction in the number of contributions per employee after the GitHub shock. We find that these results are driven by changes in the probability of experiencing very high (top quartile) contributions, reducing concerns that we are capturing bias in logarithm due to count like variables. Figure 8 displays the corresponding event studies. The absence of any discernible pre-trends before the shock is consistent with capturing the impact of the change in GitHub policy. In summary, these results highlight that large firms were able to enhance productivity, whereas small firms struggled to offset talent losses, resulting in an overall reduction in productivity. We also test our results using different types of public contributions and various cut-offs for high GitHub activity. Table A.10 presents these findings. Panel A shows that our results hold when using all public contributions, as well as contributions to organization-owned repositories. Panel B confirms robustness across alternate cut-offs, including the top tercile, top decile, non-zero contributions, and contributions exceeding 10.

---

[15]In un-reported analysis we find that the results in Column 1 are robust to dropping firms which exited from the sample. We also find that a maximum of 1.5% of our sample is aqui-hires (where more than 80%) of the team moves to a new firm. Our results are also robust to excluding acqui-hire deals from the regression analysis.

6 Market Concentration

Sections 3 and 5 document the effects of GitHub productivity disclosure on individual reallocation to larger firms and the subsequent impact on the growth and productivity of most impacted firms by size. We now proceed to evaluate whether this reallocation has influenced broader industry dynamics, potentially leading to increased concentration in more affected industries. To address this, we aggregate our data to the industry level and investigate the correlation between changes in concentration, as measured by the Herfindahl-Hirschman Index (HHI), and the average number of private contributions in each industry. We estimate the following regression equation:

$$\Delta Concentration_h = \beta Productivity_h + \varepsilon_h \tag{7}$$

The dependent variable, $\Delta Concentration_h$, denotes the change in concentration for industry h over the period spanning five years before the shock (2011-2016) to five years after the shock (2016-2021). We investigate changes in market concentration across various specifications. The coefficient β captures the correlation between the outcome and a measure of private contributions, $Productivity_h$. We present our findings for both the average productivity and the proportion of employees with productivity greater than zero during the pre-shock period, within industry h. We conduct the analysis for both 144-LinkedIn industries and 3-digit SIC industry codes, clustering standard errors at the industry level. We abstain from asserting causality for these regressions, recognizing the difficulty of inferring causal relationships from aggregate correlations. Nevertheless, these correlations offer interesting insights into overarching trends, particularly concerning the observed increase in firm concentration over the past decade (Autor et al. (2023)).

Table A.12 presents our results. Columns 1-2 and 3-4 test whether there was a greater increase in labor market HHI in industries with a higher proportion of GitHub private contributions, based on LinkedIn and SIC industry definitions, respectively. Industries with 1

standard deviation higher share of GitHub employees with private contributions experienced a 6.4% (14.0%) larger increase in labor concentration from 2016 to 2021, compared to 2011 to 2016, at the LinkedIn (SIC 3-digit) industry level. Similarly, industries with a 1 standard deviation higher average number of private contributions saw a 6.7% (13.7%) higher increase in labor concentration over a five-year period at the LinkedIn (SIC 3-digit) industry level. These findings demonstrate that there were greater concentration changes in labor markets within industries where more information was unveiled following the GitHub announcement.

7   Conclusion

This paper represents one of the first examinations of the impact of individual-productivity-related big data on labor distribution within the economy. Our investigation centers on GitHub, the world's largest repository of individual-level productivity information. Leveraging a quasi-natural experiment, we exploit a scenario where employees were exogenously enabled to disclose existing private contributions on their GitHub profiles.

Our findings unveil a significant paradigm shift: higher signaling ability by individuals triggered the reallocation of talented individuals toward larger firms. Smaller and medium-sized firms encountered challenges in both recruiting and retaining talented individuals post-shock. In contrast, large firms effectively retained existing talent through internal promotions and attracted new talent. Consistent with a reduction of information asymmetry on the labor market, the surge in productivity information mitigated reliance on traditional signals of hard information and existing networks, thereby democratizing opportunities for talent from lower-ranked, non-networked universities.

However, this talent redistribution had adverse effects on the growth and productivity of smaller firms. Further correlational evidence suggests that this talent reallocation led to an increase in market concentration in industries with greater informational disclosure

facilitated by GitHub. Our findings unearth a novel labor channel for big data, wherein increased employee productivity information reallocates labor from small to large firms, amplifying the significance of large firms in the economy.

References

ABOWD, J. M., F. KRAMARZ, AND D. N. MARGOLIS (1999): "High wage workers and high wage firms," *Econometrica*, 67, 251–333.

ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): "Innovation, Reallocation, and Growth," *American Economic Review*, 108, 3450–91.

AGHION, P., A. BERGEAUD, T. BOPPART, P. J. KLENOW, AND H. LI (2023): "A theory of falling growth and rising rents," *Review of Economic Studies*, 90, 2675–2702.

AKCIGIT, U. AND S. T. ATES (2023): "What Happened to US Business Dynamism?" *Journal of Political Economy*, 131, 2059–2124.

AKCIGIT, U. AND N. GOLDSCHLAG (2023): "Where Have All the "Creative Talents" Gone? Employment Dynamics of US Inventors," Tech. rep., National Bureau of Economic Research.

ANDREWS, D., C. CRISCUOLO, AND P. N. GAL (2015): "Frontier Firms, Technology Diffusion and Public Policy," .

AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 135, 645–709.

AUTOR, D., C. PATTERSON, AND J. VAN REENEN (2023): "Local and national concentration trends in jobs and sales: The role of structural transformation," Tech. rep., National Bureau of Economic Research.

BAGHAI, R. P., R. C. SILVA, AND L. YE (2024): "Teams and Bankruptcy," *The Review of Financial Studies*, 37, 2855–2902.

BEGENAU, J., M. FARBOODI, AND L. VELDKAMP (2018): "Big data in finance and the growth of large firms," *Journal of Monetary Economics*, 97, 71–87.

BERNHARDT, D. AND D. SCOONES (1993): "Promotion, turnover, and preemptive wage offers," *The American Economic Review*, 771–791.

BESSEN, J. (2020): "Industry concentration and information technology," *The Journal of*

Law and Economics, 63, 531–555.

BHASKARABHATLA, A., L. CABRAL, D. HEGDE, AND T. PEETERS (2021): "Are Inventors or Firms the Engines of Innovation?" *Management Science*, 67, 3899–3920.

BRADSHAW, R. (2024): "Apollo Technical Research, Surprising Social Media Recruiting Statistics," `https://www.apollotechnical.com/social-media-recruiting-statistics/`.

BROWN, C. AND J. MEDOFF (1989): "The employer size-wage effect," *Journal of political Economy*, 97, 1027–1059.

BRYNJOLFSSON, E., W. JIN, AND X. WANG (2023): "Information technology, firm size, and industrial concentration," Tech. rep., National Bureau of Economic Research.

BRYNJOLFSSON, E. AND K. MCELHERAN (2016): "The rapid adoption of data-driven decision-making," *American Economic Review*, 106, 133–139.

CALLAWAY, B., A. GOODMAN-BACON, AND P. H. SANT'ANNA (2024): "Difference-in-differences with a continuous treatment," Tech. rep., National Bureau of Economic Research.

CHEN, J. AND J. ROTH (2024): "Logs with zeros? Some problems and solutions," *The Quarterly Journal of Economics*, 139, 891–936.

COHN, J. B., Z. LIU, AND M. I. WARDLAW (2022): "Count (and count-like) data in finance," *Journal of Financial Economics*, 146, 529–551.

CONTI, A., V. GUPTA, J. GUZMAN, AND M. P. ROCHE (2023): "Incentivizing Innovation in Open Source: Evidence from the GitHub Sponsors Program," Tech. rep., National Bureau of Economic Research.

CONTI, A., C. PEUKERT, AND M. P. ROCHE (2021): "Beefing it up for your investor? open sourcing and startup funding: Evidence from github," *Open Sourcing and Startup Funding: Evidence from GitHub (August 25, 2021)*.

CROUZET, N. AND J. EBERLY (2018): "Intangibles, investment, and efficiency," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305,

Nashville, TN 37203, vol. 108, 426–431.

DAVIS, S. J. AND J. HALTIWANGER (2014): "Labor Market Fluidity and Economic Performance," Working Paper 20479, National Bureau of Economic Research.

——— (2024): "Dynamism Diminished: The Role of Housing Markets and Credit Conditions," *American Economic Journal: Macroeconomics*, 16.

DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 135, 561–644.

DE RIDDER, M. (2024): "Market Power and Innovation in the Intangible Economy," *American Economic Review*, 114.

DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2017): "Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown," *American Economic Review*, 107.

——— (2020): "Changing Business Dynamism and Productivity: Shocks versus Responsiveness," *American Economic Review*, 110.

DI PORTO, E., M. PAGANO, V. PEZONE, R. SAGGIO, AND F. SCHIVARDI (2024): "Careers and Wages in Family Firms: Evidence from Matched Employer-Employee Data," *Unpublished Manuscript*.

ECKEL, C. AND S. R. YEAPLE (2017): "Too much of a good thing? Labor market imperfections as a source of exceptional exporter performance," Tech. rep., National Bureau of Economic Research.

EECKHOUT, J. AND L. VELDKAMP (2022): "Data and market power," *NBER Working Paper*.

EL-KOMBOZ, L. A. AND T. FACKLER (2023): "Productivity Spillovers among Knowledge Workers in Agglomerations: Evidence from GitHub," *Available at SSRN 4600283*.

FARBOODI, M., A. MATRAY, L. VELDKAMP, AND V. VENKATESWARAN (2022): "Where has all the data gone?" *The Review of Financial Studies*, 35, 3101–3138.

FIROOZ, H., Z. LIU, AND Y. WANG (2022): "Automation and the rise of superstar firms,"

Available at SSRN.

GIBBONS, R. AND L. F. KATZ (1991): "Layoffs and lemons," *Journal of labor Economics*, 9, 351–380.

GOLDSTEIN, I., C. S. SPATT, AND M. YE (2021): "Big data in finance," *The Review of Financial Studies*, 34, 3213–3225.

GREENWALD, B. C. (2018): *Adverse selection in the labor market*, Routledge.

GUPTA, A. (2023): "Labor Mobility, Entrepreneurship, and Firm Monopsony: Evidence from Immigration Wait-Lines," *Entrepreneurship, and Firm Monopsony: Evidence from Immigration Wait-Lines (May 16, 2023)*.

HOLUB, F. AND B. THIES (2023): "Air quality, high-skilled worker productivity and adaptation: Evidence from github," *High-Skilled Worker Productivity and Adaptation: Evidence from Github (March 6, 2023)*.

HSIEH, C.-T. AND E. ROSSI-HANSBERG (2023): "The industrial revolution in services," *Journal of Political Economy Macroeconomics*, 1, 3–42.

JEFFERS, J. S. (2024): "The impact of restricting labor mobility on corporate investment and entrepreneurship," *The Review of Financial Studies*, 37, 1–44.

KARAHAN, F., B. PUGSLEY, AND A. ŞAHIN (2024): "Demographic Origins of the Start-up Deficit," *American Economic Review*, 114.

LASHKARI, D., A. BAUER, AND J. BOUSSARD (2018): "Information technology and returns to scale," *Available at SSRN 3458604*.

LEMIEUX, T., W. B. MACLEOD, AND D. PARENT (2009): "Performance pay and wage inequality," *The Quarterly Journal of Economics*, 124, 1–49.

LIU, E., A. MIAN, AND A. SUFI (2022): "Low Interest Rates, Market Power, and Productivity Growth," *Econometrica*, 90, 193–221.

MUELLER, H. M., P. P. OUIMET, AND E. SIMINTZI (2017): "Within-firm pay inequality," *The Review of Financial Studies*, 30, 3605–3635.

OI, W. Y. AND T. L. IDSON (1999): "Firm size and wages," *Handbook of labor economics*,

3, 2165–2214.

PAGANO, M. AND L. PICARIELLO (2023): "Talent discovery, layoff risk and unemployment insurance," *European Economic Review.*

PRAT, A. AND T. VALLETTI (2022): "Attention oligopoly," *American Economic Journal: Microeconomics*, 14, 530–557.

SMITH, A. (2015): "Pew Research, Searching for Work in the Digital Era," `https://www.pewresearch.org/internet/2015/11/19/searching-for-work-in-the-digital-era`.

TAMBE, P., L. HITT, D. ROCK, AND E. BRYNJOLFSSON (2020): "Digital capital and superstar firms," Tech. rep., National Bureau of Economic Research.

TAMBE, P. AND L. M. HITT (2012): "The productivity of information technology investments: New evidence from IT labor data," *Information systems research*, 23, 599–617.

TROSKE, K. R. (1999): "Evidence on the employer size-wage premium from worker-establishment matched data," *Review of Economics and Statistics*, 81, 15–26.

VELDKAMP, L. AND C. CHUNG (2024): "Data and the aggregate economy," *Journal of Economic Literature*, 62, 458–484.

WALDMAN, M. (1984): "Job assignments, signalling, and efficiency," *The RAND Journal of Economics*, 15, 255–267.

——— (1990): "Up-or-out contracts: A signaling perspective," *Journal of Labor Economics*, 8, 230–250.

WALLSKOG, M., N. BLOOM, S. W. OHLMACHER, AND C. TELLO-TRILLO (2024): "Within-Firm Pay Inequality and Productivity," Tech. rep., National Bureau of Economic Research.

Figure 1: **Our Sample GitHub Users by MSA**

This figure presents the distribution of our sample by employee's location (MSA) based on their first job.

Share of GitHub Users (%)

| 0.00 | 0.05 | 0.10 | 0.50 | 1.00 | 5.00 | 15.00 |
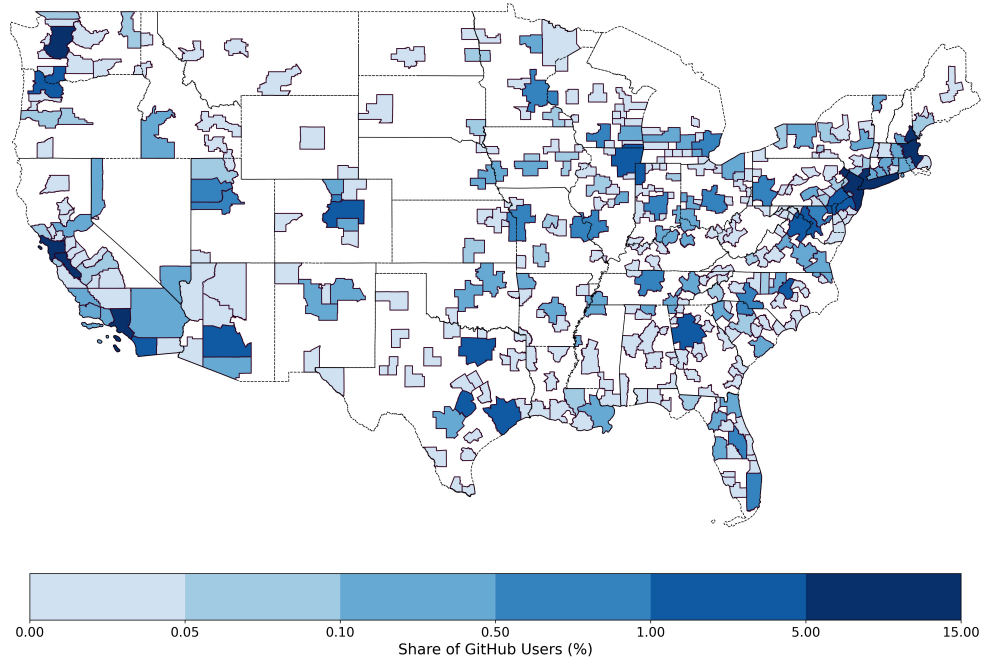
Figure 2: **GitHub Shock & Employee Mobility to Large Firms**

This figure presents an event study on employee mobility to large firms around the disclosure of employee productivity information on GitHub. The x-axis plots the quarter around the shock. The dashed line indicates the quarter 2 (April-June) of 2016, during which GitHub introduced the policy that allowed users to display their productivity information more accurately. The quarter 1 of 2016 (Jan-March) is the omitted quarter set equal to zero. On the y-axis, we plot the differential impact of the shock on more productive employees estimated by the β coefficients obtained from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \sum_{\tau=-8}^{12} \beta_\tau Productivity_i \times \mathbb{1}(\tau = t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with the time indicator for each quarter. The dependent variable is an indicator that equals 1 when employee i switches job to a large firm (>1000 employees) in quarter t and is 0 otherwise. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Light blued lines show 95% confidence intervals for estimates.

Figure 3: **GitHub Shock & Employee Mobility**

This figure presents event studies on employee mobility to small-sized and medium-sized firms around the disclosure of employee productivity information on GitHub. The x-axis plots the quarter around GitHub's policy change. The dashed line indicates the quarter 2 (April-June) of 2016, during which GitHub introduced the policy that allowed users to display their productivity information more accurately. The quarter 1 of 2016 (Jan-March) is the omitted quarter set equal to zero. On the y-axis, we plot the differential impact of the shock on more productive employees estimated by the β coefficients obtained from regression estimate:

$$\mathbb{1}(Move)_{i,t} = \sum_{\tau=-8}^{12} \beta_\tau Productivity_i \times \mathbb{1}(\tau = t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with the time indicator for each quarter. The dependent variable is an indicator that equals 1 when employee i switches job to a mid-sized firm ($50 <$ employees $<= 1000$) for Panel (a), to a small firms ($<= 50$ employee) for Panel (b), in quarter t and is 0 otherwise. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Light blued lines show 95% confidence intervals for estimates.


(a) Move to Mid-Sized Firms

(b) Move to Small Firms

Figure 4: **GitHub Shock & Employee Mobility to Different Size Buckets**

This figure presents the economic magnitude of employee mobility to firms in different size buckets around the disclosure of employee productivity information on GitHub. On the y-axis, we plot the economic magnitude of the differential impact of the shock on more productive employees using the β coefficients obtained from regression estimate:

$$\mathbb{1}(Move)_{i,t} = \sum_{\tau=-8}^{12} \beta_\tau Productivity_i \times \mathbb{1}(\tau = t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with the time indicator for each quarter. The dependent variable is an indicator that equals 1 when employee i switches job to a firm in a given size bucket. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Light blued lines show 95% confidence intervals for estimates.

Figure 5: **GitHub Shock & Employee Mobility to Large Firms: Heterogeneity by Initial Firm Size**

This figure presents event studies on employee mobility to large firms from small and mid-sized firms around the disclosure of employee productivity information on GitHub. The x-axis plots the quarter around the shock. The dashed line indicates the quarter 2 (April-June) of 2016, during which GitHub introduced the policy that allowed users to display their productivity information more accurately. The quarter 1 of 2016 (Jan-March) is the omitted quarter set equal to zero. On the y-axis, we plot the differential impact of the shock on more productive employees, particularly at small and mid-sized firms, estimated by the β coefficients obtained from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \sum_{\tau=-8}^{12} \delta_\tau Productivity_i \times \mathbb{1}(\tau = t) + \sum_{\tau=-8}^{12} \beta_\tau Productivity_i \times \mathbb{1}(\tau = t) \times \mathbb{1}(Mid-Sized_{i(t-1)})$$

$$+ \sum_{\tau=-8}^{12} \beta_\tau^{'} Productivity_i \times \mathbb{1}(\tau = t) \times \mathbb{1}(Small_{i(t-1)}) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is a triple interaction of three indicators: normalized productivity estimate of the employee i derived from the AKM decomposition, indicators for the size of the firm where the employee worked in the previous quarter, and the time indicator for each quarter. We also include indicator variables *Mid* and *Small* as well as their double interactions with *Productivity* and time, but do not show them in the above equation for brevity. Panel A plots interaction with *Mid*, an indicator variable that takes the value of 1 for firms with 50-1000 employees. Panel B plots the interaction with *Small*, an indicator for firms with less than 50 employees. The dependent variable is an indicator that equals 1 if the employee switches job to a large firm (> 1000 employees) in that quarter. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Light blued lines show 95% confidence intervals for estimates.
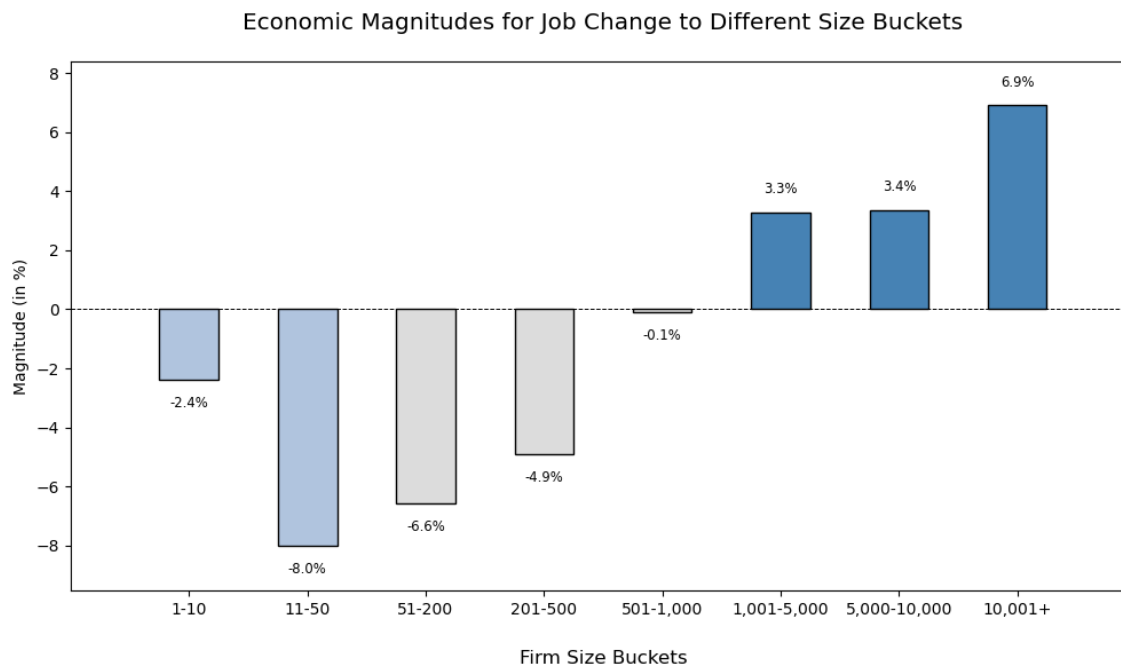

(a) Mid-Sized Firms

(b) Small Firms
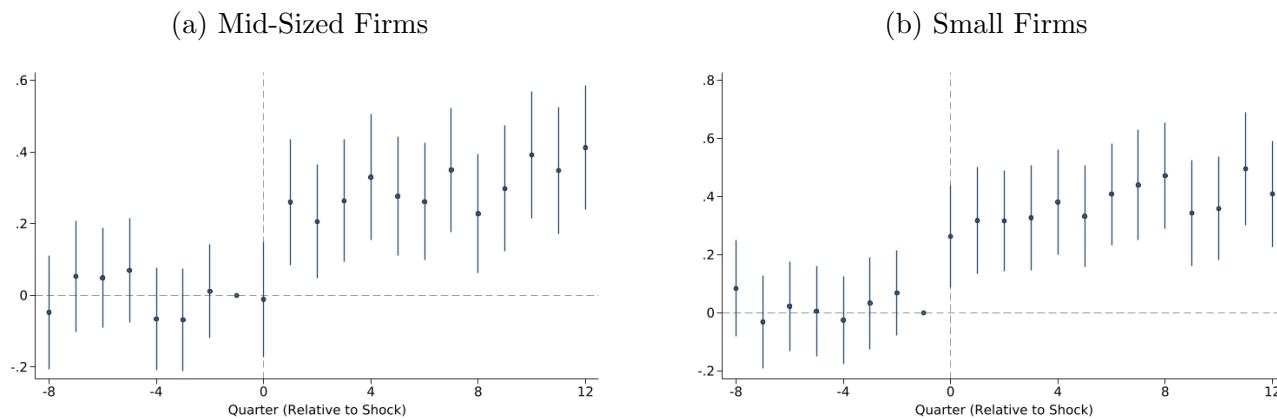
Figure 6: **GitHub Shock & Employee Promotions at Large Firms**

This figure presents an event study on employees' within-firm promotions at large firms around the disclosure of employee productivity information on GitHub. The x-axis plots the quarter around the shock. The dashed line indicates the quarter 2 (April-June) of 2016, during which GitHub introduced the policy that allowed users to display their productivity information more accurately. The quarter 1 of 2016 (Jan-March) is the omitted quarter set equal to zero. On the y-axis, we plot the differential impact of the shock on more productive employees, particularly at large firms, estimated by the β coefficients obtained from the equation:

$$\mathbb{1}(Within-FirmPromotion)_{i,f,t} = \sum_{\tau=-8}^{12} \beta_\tau Productivity_i \times \mathbb{1}(\tau = t)$$

$$+ \sum_{\tau=-8}^{12} \lambda_\tau Productivity_i \times \mathbb{1}(\tau = t) \times \mathbb{1}(Mid-Sized_{i(t-1)})$$

$$+ \sum_{\tau=-8}^{12} \lambda_\tau^{'} Productivity_i \times \mathbb{1}(\tau = t) \times \mathbb{1}(Small_{i(t-1)}) + \alpha_i + \gamma_{it} + \delta_{ft} + \varepsilon_{i,f,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with the time indicator for each quarter. We also include indicator variables *Mid* and *Small* as well as their double interactions with *Productivity* and time, but do not show them in the above equation for brevity. The dependent variable is an indicator that equals 1 when employee i does not switch jobs and her salary or seniority improves in quarter t compared to quarter *t-1* and is 0 otherwise. α_i represents employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. δ_{ft} is firm cohort fixed effects representing an interaction of the employee's current firm and the year of joining the firm. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Light blued lines show 95% confidence intervals for estimates.

Figure 7: **GitHub Shock & Employment Growth at Small Firms**

This figure presents an event study on firm-level employment growth for small firms ($<= 50$ employees) around the disclosure of employee productivity information on GitHub. The x-axis plots the quarter around GitHub's policy change. The dashed line indicates the quarter 2 (April-June) of 2016, during which GitHub introduced the policy that allowed users to display their productivity information more accurately. The quarter 1 of 2016 (Jan-March) is the omitted quarter set equal to zero. On the y-axis, we plot the differential impact of the shock on employment growth of small firms with above-median employee productivity estimated by the β coefficients obtained from the following equation:

$$EmploymentGrowth_{j,t} = \sum_{\tau=-8}^{12} \theta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t)$$

$$+ \sum_{\tau=-8}^{12} \delta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \times \mathbb{1}(MidSized_j)$$

$$+ \sum_{\tau=-8}^{12} \beta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \times \mathbb{1}(Small_j) + \lambda X_{jt} + \alpha_j + \gamma_{jt} + \varepsilon_{j,t}$$

The independent variable is a triple interaction of three binary indicators: $\mathbb{1}(Productivity_j > Median)$ that takes the value of 1 if firm j's average AKM-based employee productivity is above median, $\mathbb{1}(Small_j)$ that takes the value of 1 if firm j is small ($<= 50$ employees), and $\mathbb{1}(\tau = t)$ that is the time indicator for each quarter. We also include the double interaction of indicator variables *MidSized* and *Small* with time, but do not show them in the above equation for brevity. In Panel (a), the dependent variable is the employment change (in %) for firm j in quarter t with respect to quarter *t-1*. In Panel (b), the dependent variable is a dummy that takes the value of 1 if the employment at the firm frops by more than 50% in that quarter and 0 otherwise. X_{jt} represents firm's time-varying lagged employment. α_j is firm fixed effects. γ_{jt} represents the firm's industry × state interacted with time fixed effects. We cluster the standard errors at the firm level. Light blued lines show 95% confidence intervals for estimates.

(a) Employment Growth

(b) $\mathbb{1}$(Employment Reduction > 50%)

Figure 8: **GitHub Shock & Productivity at Small Firms**

This figure presents event studies on firm-level productivity of small firms ($<= 50$ employees) proxied by their public GitHub contributions around the disclosure of employee productivity information on GitHub. The x-axis plots the quarter around GitHub's policy change. The dashed line indicates the quarter 2 (April-June) of 2016, during which GitHub introduced the policy that allowed users to display their productivity information more accurately. The quarter 1 of 2016 (Jan-March) is the omitted quarter set equal to zero. On the y-axis, we plot the differential impact of the shock on public contributions of small firms with above-median employee productivity estimated by the β coefficients obtained from the following equation:

$$\mathbb{1}(Contributions_{j,t}) = \sum_{\tau=-8}^{12} \theta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) + \sum_{\tau=-8}^{12} \delta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \times \mathbb{1}(MidSized_j)$$

$$+ \sum_{\tau=-8}^{12} \beta_\tau \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(\tau = t) \times \mathbb{1}(Small_j) + \lambda X_{jt} + \alpha_j + \gamma_{jt} + \varepsilon_{j,t}$$

The independent variable is a triple interaction of three binary indicators: $\mathbb{1}(Productivity_j > Median)$ that takes the value of 1 if firm j's average AKM-based employee productivity is above median, $\mathbb{1}(Small_j)$ that takes the value of 1 if firm j is small ($<= 50$ employees), and $\mathbb{1}(\tau = t)$ that is the time indicator for each quarter. We also include the double interaction of indicator variables *MidSized* and *Small* with time, but do not show them in the above equation for brevity. The dependent variable in Panel (a) (Panel (b)) is a binary indicator that takes the value of 1 if the firm j's total public contributions (total public contributions per employee) in the quarter is in the top quartile and is 0 otherwise. Firm contributions only include public contributions made by incumbent employees towards repositories not owned by the employees themselves. X_{jt} represents firm's time-varying lagged stock of total contributions as well as the median pre-period public contribution interacted with an indicator for quarters after GitHub's policy change. α_j is firm fixed effects. γ_{jt} represents the firm's industry $\times$ state interacted with time fixed effects. We cluster the standard errors at the firm level. Light blued lines show 95% confidence intervals for estimates.

(a) Public Contributions

(b) Public Contributions per Employee

Table 1: **Descriptive Statistics**

This table presents descriptive statistics for employee-level variables in Panel A, and firm-level variables in Panel B.

Panel A: Employee Level Summary						
	Mean	**Std. Dev.**	**P10**	**Median**	**P90**	**Obs.**
$\mathbb{1}$(Male)	0.79	0.41	0	1	1	302,903
$\mathbb{1}$(White)	0.69	0.46	0	1	1	302,903
$\mathbb{1}$(Software Engg. Role)	0.37	0.48	0	0	1	302,903
$\mathbb{1}$(SF/NY/Seattle)	0.25	0.43	0	0	1	302,581
Experience (Years)	8.82	8.13	1	7	20	302,903
Initial Seniority	2.20	0.93	1	2	4	302,903
Initial Salary (in '000 $s)	87.34	26.85	56.73	84.88	116.66	302,903
Pre-Period Total Contributions (Log)	2.60	2.16	0	2.08	5.87	302,903
Pre-Period Public Contributions (Log)	2.35	1.91	0	1.95	5.17	302,903
Pre-Period Private Contributions (Log)	0.66	1.86	0	0	3.22	302,903
$\mathbb{1}$(Job Change to Large Firms)	0.03	0.04	0	0	0.08	302,902
$\mathbb{1}$(Job Change to Mid-Size Firms)	0.02	0.04	0	0	0.08	302,902
$\mathbb{1}$(Job Change to Small Firms)	0.02	0.03	0	0	0.06	302,902
$\mathbb{1}$(Job Change)	0.07	0.07	0	0.05	0.17	302,903
$\mathbb{1}$(Promotion)	0.06	0.05	0	0.04	0.13	302,903
$\mathbb{1}$(Across-Firm Promotion)	0.04	0.05	0	0.04	0.10	302,903
$\mathbb{1}$(Within-Firm Promotion)	0.02	0.03	0	0	0.07	302,903
Panel B: Firm Level Summary						
	Mean	**Std. Dev.**	**P10**	**Median**	**P90**	**Obs.**
All Sample Firms (LinkedIn)						
$\mathbb{1}$(Large Firm) (Employees > 1000)	0.15	0.36	0	0	1	57,982
$\mathbb{1}$(Mid-size Firm) (50 < Employees <= 1000)	0.37	0.48	0	0	1	57,982
$\mathbb{1}$(Small Firm) (Employees <= 50)	0.48	0.50	0	0	1	57,982
Employment Change	0.03	0.09	-0.03	0.02	0.13	55,335
Public Contributions	46.49	739.07	0.17	4	65	50,307
Public Contributions/Employee	10.26	107.47	0.13	1.75	17.26	50,307
Compustat Firms						
Firm Size (in '000 Employees)	24.05	74.32	0.36	5.10	60.08	1,675
Daily Excess Return (in %)	-0.03	0.14	-0.18	-0.02	0.09	1,675

Table 2: **Sample & Variation Breakdown for AKM Tests**

This table presents the details of sample construction (in Panel A) and variance decomposition (in Panel B) of the AKM (Abowd et al., 1999) tests done to estimate individual employee productivity from private GitHub contributions. We isolate employee productivity from historical private GitHub contributions by running the following regression specification:

$$Log(PvtContributions)_{i,f,t} = \vartheta_i + \vartheta_f + \vartheta_t + \vartheta_k + X_{i,f,t} + \varepsilon_{i,f,t}$$

ϑ_i, the employee fixed effect, represents the estimated employee i's productivity. The dependent variable is the logarithm of the number of private contributions made by employee i in month t. The regression is estimated on data spanning May 2011 to April 2016, exactly five years before GitHub policy change. We control for employee's current job duration, and total work experience. We also include firm fixed effects (ϑ_f), time fixed effects (ϑ_t), and employee's role fixed effects (ϑ_k). Panel A shows the number of observations in our initial sample, the connected set (the largest set of firms linked by employee mobility), the filtered set (connected set with at least two employees in each firm-time observation), the sample for which productivity estimates were obtained, and the final sample remaining after excluding users working in non-profit industries at the time of GitHub's policy change. Panel B summarizes the breakdown of the variance in the dependent variable explained by the employee and firm components.

	(1)	(2)	(3)	(4)
Panel A: Summary Stats				
	Employees		**Parent Firms**	
Sample	N	% Initial Sample	N	% Initial Sample
Initial Sample	381,114		103,308	
Connected Set	356,106	93.4%	77,899	75.4%
Filtered Set	342,769	89.9%	36,358	35.2%
Productivity Estimates	338,523	88.8%	36,358	35.2%
Excluding Non-Profits	302,903	79.5%		

Panel B: Variation Breakdown	
	Var. Explained
Person F.E.	66.7%
Firm F.E.	8.5%
Residual	24.7%

Table 3: **GitHub Shock & Employee Mobility to Large Firms**

This table presents estimates on the impact of GitHub's disclosure of employee productivity information on employee mobility to large firms. We report the differential impact of GitHub's policy change on productive employees estimated by the β coefficient from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \gamma X_{it} + \alpha_i + \theta_t + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with *Post*, an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2 of 2016). The dependent variable is an indicator that equals 1 when employee i switches job to a large firm (>1000 employees) in quarter t and is 0 otherwise. Columns 1 to 5 present estimates with varying controls and fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)
Outcome:	$\mathbb{1}$**(Job Change to Large Firms)**				
Productivity $\times$ $\mathbb{1}$(Post)	0.0013***	0.0012***	0.0017***	0.0016***	0.0015***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0002)
Employee F.E.	Y	Y	Y	Y	Y
Time F.E.	Y	Y			
Salary & Seniority Controls		Y	Y	Y	Y
Cohort $\times$ Time F.E.			Y		
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.				Y	Y
Industry $\times$ Location $\times$ Time F.E.					Y
Observations	6,804,386	6,804,386	6,804,166	3,811,484	3,199,173
# Unique Employees	302,781	302,781	302,772	170,880	143,446
R-squared	0.054	0.054	0.06	0.168	0.212
Y-Mean	0.031	0.031	0.031	0.032	0.033
Magnitude (%)	4.04	3.84	5.32	5.69	5.27

Table 4: **GitHub Shock & Employee Mobility to Large Firms: Placebo & Alternate Samples**

This table presents placebo and alternate sample tests for estimates reported in Table 3. Panel A presents a placebo test with the independent variable being an alternative productivity estimate computed by using employee's *public* contributions as the dependent variable in the AKM decomposition. Panel B estimates results on a subsample of active GitHub Users. Panel B1 only includes users with more than 10 total GitHub contributions during the pre-period (in the 12 quarters before May 2016). Panel B2, only includes users employed at firms (at the time of shock) with 100 or more total GitHub contributions during the pre-period. Panel B3 restricts the sample to users with non-zero private GitHub contributions during the pre-period. Panel B4 restricts the sample to users with at least 10 private GitHub contributions during the pre-period. All tests correspond to our baseline specification, Column 5 in Table 3. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

Outcome:	(1) Coefficient	(2) Std. Error	(3) Observations	(4) R-squared	(5) Y-Mean	(6) Magnitude (%)
Panel A: Placebo Test						
A1: Productivity based on Public Contributions						
Productivity × $\mathbb{1}$(Post)	-0.0004	(0.0003)	3,199,024	0.211	0.033	-1.4
Panel B: Within Active User Comparison						
B1: Pre Period Individual Total Contributions > 10						
Productivity × $\mathbb{1}$(Post)	0.0018***	(0.0002)	1,708,418	0.224	0.033	7.8
B2: Pre-period Firm Contributions >= 100						
Productivity × $\mathbb{1}$(Post)	0.0042***	(0.0002)	1,748,870	0.259	0.0025	12.56
B3: Pre-period Individual Private Contributions > 0						
Productivity × $\mathbb{1}$(Post)	0.0011***	(0.0003)	506,974	0.271	0.043	7.65
B4: Pre-period Individual Private Contributions >= 10						
Productivity × $\mathbb{1}$(Post)	0.0011***	(0.0003)	401,146	0.274	0.024	8.35

Table 5: **GitHub Shock & Employee Mobility to Medium & Small Firms**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information on employee mobility based on destination firm size. We report the differential impact of GitHub's policy change on productive employees estimated by the β coefficient from the equation:

$$\mathbb{1}(Move)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with *Post*, an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). The dependent variable is an indicator that equals 1 when employee i switches job in quarter t to a large firm (>1000 employees) for Column 1, mid-sized firm ($50 < $ employees $<= 1000$) for Column 2, and small firm ($<= 50$ employees) for Column 3. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)
Outcome:		$\mathbb{1}$**(Job Change)**	
	Large Firm	**Medium Firm**	**Small Firm**
Productivity $\times$ $\mathbb{1}$(Post)	0.0015***	-0.0013***	-0.0012***
	(0.0002)	(0.0002)	(0.0002)
Salary & Seniority Controls	Y	Y	Y
Employee F.E.	Y	Y	Y
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Y	Y	Y
Industry $\times$ Location $\times$ Time F.E.	Y	Y	Y
Observations	3,199,173	3,199,173	3,199,173
# Unique Employees	143,446	143,446	143,446
R-squared	0.212	0.211	0.222
Y-Mean	0.033	0.03	0.023
Magnitude (%)	5.27	-4.83	-5.99

Table 6: **GitHub Shock & Employee Mobility to Large Firms: Heterogeneity by Origin Firm Size**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information on employee mobility to large firms based on the *origin* firm size. We report the differential impact of GitHub's policy change on productive employees by origin firm type estimated by the β coefficient from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta_1 Productivity_i \times \mathbb{1}(Post_t) + \beta_2 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Mid_{i(t-1)})$$
$$+ \beta_3 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Small)_{i(t-1)} + \lambda X_{it} + \alpha_i + \theta_t + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is a triple interaction of three indicators: normalized productivity estimate of employee i derived from the AKM decomposition, indicators for the size of the origin firm, i.e. where employee i worked in the previous quarter, and an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). *Mid* is an indicator variable that takes the value of 1 for firms with 50-1000 employees and *Small* is an indicator for firms with less than 50 employees. The dependent variable is an indicator that equals 1 when employee i switches job to a large firm (>1000 employees) in quarter t and is 0 otherwise. We also estimate coefficients for indicator variables *Mid* and *Small* as well as their double interactions with *Productivity* and *Post*, but do not report them for brevity. Columns 1 to 5 present estimates with varying controls and fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: **($p < 0.05$), ***($p < 0.01$).

	(1)	(2)	(3)	(4)	(5)
Outcome:			$\mathbb{1}$(Job Change to Large Firms)		
Productivity $\times \mathbb{1}$(Post)	-0.0037***	-0.0039***	-0.0029***	-0.0025***	-0.0025***
	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0004)
Productivity $\times \mathbb{1}$(Post) $\times \mathbb{1}$(Mid)	0.0045***	0.0047***	0.0042***	0.0031***	0.0030***
	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0005)
Productivity $\times \mathbb{1}$(Post) $\times \mathbb{1}$(Small)	0.0050***	0.0053***	0.0046***	0.0041***	0.0040***
	(0.0003)	(0.0003)	(0.0003)	(0.0004)	(0.0005)
Employee F.E.	Y	Y	Y	Y	Y
Time F.E.	Y	Y			
Salary & Seniority Controls		Y	Y	Y	Y
Cohort $\times$ Time F.E.			Y		
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.				Y	Y
Industry $\times$ Location $\times$ Time F.E.					Y
Observations	6,455,765	6,455,765	6,455,545	3,596,761	3,022,682
# Unique Employees	300,779	300,779	300,770	169,629	142,754
R-squared	0.069	0.069	0.075	0.184	0.227
Y-Mean	0.03	0.03	0.03	0.031	0.032
Magnitude (Mid) (%)	2.8	2.66	4.29	2.27	1.78
Magnitude (Small) (%)	4.56	4.48	5.62	5.99	5.42

Table 7: **GitHub Shock & Employee Promotion**

This table presents estimates on the impact of GitHub's disclosure of employee productivity information on employee career advances. We report the differential impact of GitHub's policy change on productive employees by initial firm type estimated by the β coefficient from the equation:

$$\mathbb{1}(Promotion)_{i,f,t} = \beta_1 Productivity_i \times \mathbb{1}(Post_t) + \beta_2 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Mid_{i(t-1)})$$
$$+ \beta_3 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Small)_{i(t-1)}) + \alpha_i + \gamma_{it} + \delta_{ft} + \varepsilon_{i,f,t}$$

The independent variable is an interaction of two indicators: normalized productivity estimate of the employee i derived from the AKM decomposition, and an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). Even-numbered columns additionally interact this term with indicators for the size of the firm where the employee worked in the previous quarter. *Mid* is an indicator variable that takes the value of 1 for firms with 50-1000 employees, and *Small* is an indicator for firms with less than 50 employees. The dependent variable is an indicator variable that is 1 if there is a promotion (salary or job-seniority increase compared to the previous quarter) and the employee switches job to a large firm (> 1000 employees) in Columns 1 and 2, switches job to a small or mid-sized firm ($<= 1000$ employees) in Columns 3 and 4, remains at the same firm in columns 5 and 6. We also estimate coefficients for indicator variables *Mid* and *Small* as well as their double interactions with *Productivity* and *Post*, but do not report them for brevity. α_i represents employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. In columns 5 and 6, we also control for δ_{ft}, firm cohort representing an interaction of the employee's current firm and the year of joining the firm. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:			$\mathbb{1}$**(Promotion)**			
	Move to Large		*Move to Small & Med.*		*Within-Firm*	
Productivity $\times$ $\mathbb{1}$(Post)	0.0011***	-0.0018***	-0.0020***	-0.0039***	0.0008***	0.0011**
	(0.0001)	(0.0003)	(0.0002)	(0.0006)	(0.0002)	(0.0005)
Productivity $\times$ $\mathbb{1}$(Post) $\times$ $\mathbb{1}$(Mid)		0.0023***		0.0032***		0.0002
		(0.0004)		(0.0007)		(0.0006)
Productivity $\times$ $\mathbb{1}$(Post) $\times$ $\mathbb{1}$(Small)		0.0028***		0.0042***		-0.0004
		(0.0004)		(0.0007)		(0.0006)
Employee F.E.	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y	Y	Y
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Y	Y	Y	Y	Y	Y
Industry $\times$ Location $\times$ Time F.E.	Y	Y	Y	Y	Y	Y
Firm Cohort					Y	Y
Observations	3,199,173	3,022,682	3,199,173	3,022,682	2,860,181	2,828,288
# Unique Employees	143,446	142,754	143,446	142,754	142,827	142,521
R-squared	0.204	0.216	0.212	0.22	0.269	0.271
Y-Mean	0.024	0.023	0.036	0.035	0.019	0.019
Magnitude (%) (Main)	5.42	-9	-6.16	-12.7	4.96	6.25
Magnitude (%) (Small)		5.11		0.84		4.19

Table 8: **GitHub Shock & Employee Mobility to Large Firms: Heterogeneity by Employee Type**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information based on employee-type. We report the differential impact of GitHub's policy change on productive employees, particularly those susceptible to higher information asymmetry, estimated by the β coefficients from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta_1 Productivity_i \times \mathbb{1}(Post_t)$$
$$+ \beta_2 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Group_i) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is a triple interaction of three indicators: the employee's normalized productivity estimate derived from the AKM decomposition, indicators for employee group, and indicator that takes the value of 1 for quarters after GitHub's policy change quarter 2, i.e., April-June 2016). Columns 1-2 include triple interaction for $\mathbb{1}$(Experience < Median), which takes the value of 1 for employees that have less than median work experience at the time of shock. Columns 3-4 (5-6) include triple interactions with $\mathbb{1}$(Non-Elite School) that takes the value 1 if the employee did not attend a school ranked among the top 30 in the US (world) by US News & Ranking (QS University Rankings) in 2016. Columns 7-8 include triple interaction with $\mathbb{1}$(Low-Network School) that takes the value of 1 if the employee did not attend a school that is among the top 5 universities any large firm hired from in the pre-period. We also estimate coefficients for double interaction of *Group* and *Post*, but do not report them for brevity. We control for employees' time-varying salary and seniority. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:				$\mathbb{1}$(Job Change to Large Firms)				
Productivity × $\mathbb{1}$(Post)	0.0013***	0.0012***	0.0007	0.0006	0.0005	0.0004	0.0015***	0.0014***
	(0.0002)	(0.0002)	(0.0005)	(0.0006)	(0.0006)	(0.0007)	(0.0001)	(0.0002)
Productivity × $\mathbb{1}$(Post) × $\mathbb{1}$(Experience < Median)	0.0005*	0.0006**						
	(0.0003)	(0.0003)						
Productivity × $\mathbb{1}$(Post) × $\mathbb{1}$(Non-Elite School - US)			0.0009*	0.0010*				
			(0.0005)	(0.0006)				
Productivity × $\mathbb{1}$(Post) × $\mathbb{1}$(Non-Elite School - Global)					0.0011*	0.0012*		
					(0.0006)	(0.0007)		
Productivity × $\mathbb{1}$(Post) × $\mathbb{1}$(Low-Network School)							0.0006*	0.0007*
							(0.0003)	(0.0004)
Employee F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y	Y	Y	Y	Y
Cohort × Language × Role × Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Industry × Location × Time F.E.		Y		Y		Y		Y
Observations	4,397,956	3,112,896	4,397,956	3,112,896	4,397,956	3,112,896	4,397,956	3,112,896
# Unique Employees	196,369	139,680	196,369	139,680	196,369	139,680	196,369	139,680
R-squared	0.092	0.213	0.092	0.213	0.092	0.213	0.091	0.212

Table 9: **GitHub Shock & Employee Mobility to Large Firms: Heterogeneity by Destination Firm Characteristics**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information based on destination firm characteristics. We report the differential impact of GitHub's policy change on productive employees estimated by the β coefficients from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is an interaction of the employee's normalized productivity estimate derived from the AKM decomposition and an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). The dependent variable is mobility to a subset of large firms ($> 1,000$ employees) and differs in each column based on the subset. In columns 1-2, the subset is based on steepness in career growth defined by whether the firm had an above-median average annual salary change in the pre-period. In columns 3-4, the subset is based on job stability defined by whether the firm had below-median average forced turnovers in the pre-shock period. In columns 5-6, the subset is based on whether any Human Resource professional at the firm used GitHub during the pre-period, and in columns 7-8, the subset is based on whether the firm and its employees use GitHub (derived from firm-level private contributions). X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:				$\mathbb{1}$(**Job Change to Large Firms with**)				
	Steep Career Growth		**High Stability**		**Tech-savvy Recruiters**		**GitHub Usage**	
	Yes	No	Yes	No	Yes	No	Yes	No
Productivity $\times$ $\mathbb{1}$(Post)	0.0014***	0.0001	0.0000	0.0016***	0.0014***	0.0002*	0.0014***	0.0001***
	(0.0001)	(0.0001)	(0.0000)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0000)
Employee F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y	Y	Y	Y	Y
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Industry $\times$ Location $\times$ Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Observations	3,192,168	3,192,168	3,194,104	3,194,104	3,199,173	3,199,173	3,199,173	3,199,173
# Unique Employees	143,445	143,445	143,445	143,445	143,446	143,446	143,446	143,446
R-Squared	0.2	0.221	0.224	0.207	0.189	0.23	0.204	0.266
Y-Mean	0.024	0.007	0.003	0.029	0.018	0.015	0.03	0.003
Magnitude (%)	6.8	1.59	-1.25	6.23	8.43	1.31	5.44	3.58
T-test ($H_0 : \beta_{Yes} = \beta_{No}$)	(0.0000)***		(0.0000)***		(0.0000)***		(0.0189)**	

Table 10: **GitHub Shock & Firm Outcomes**

This table presents estimates on the impact of GitHub's disclosure of employee productivity information on firm employment and productivity. We report the differential impact of GitHub's policy change on small and mid-sized firms with more productive employees pre-treatment estimated by the β coefficient from the equation:

$$FirmOutcome_{j,t} = \beta_1 \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(Post_t) + \beta_2 \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(Post_t) \times \mathbb{1}(Mid_j)$$
$$+ \beta_3 \mathbb{1}(Productivity_j > Median) \times \mathbb{1}(Post_t) \times \mathbb{1}(Small_j) + \lambda X_{jt} + \alpha_j + \gamma_{jt} + \varepsilon_{j,t}$$

The independent variable is a triple interaction of three indicators: $\mathbb{1}(Productivity_j > Median)$ that takes the value of 1 if firm j's average AKM-based employee productivity is above median, an indicator for firm size, and an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). *Mid* is an indicator variable that takes the value of 1 for firms with 50-1000 employees, and *Small* is an indicator for firms with less than 50 employees. In column 1, the dependent variable is the percentage change in employment compared to the previous quarter. In column 2, the dependent variable is a binary indicator which takes the value of 1 if the employment at the firm falls by more than 50% compared to the previous quarter and 0 otherwise. In columns 3-4 and 5-6, the dependent variable is the firm's public contributions and per-capita public contributions, respectively. Firm public contributions only include contributions made by incumbent employees towards repositories not owned by the employees themselves. In columns 3 and 5, the dependent variable is the logarithm of one added to the firm public contributions variable. In columns 4 and 6, the dependent variable takes the value of 1 if the firm contribution variable is in the top quartile and 0 otherwise. We also estimate coefficients for double interaction of indicator variables *Mid* and *Small* with *Post*, but do not report them for brevity. X_{jt} represents firm's time-varying controls (lagged employment in columns 1-2 and lagged stock of total contributions as well as pre-period median public contributions interacted with *Post* in columns 3-6). α_j is firm fixed effects. γ_{jt} represents the firm's industry × state interacted with time fixed effects. Standard errors are clustered at the firm level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	Employment		Public Contributions		Per-capita Public Contributions	
	QoQ Emp. Growth	Emp Reduction > 50%	Log(1+Y)	𝟙(Top Quartile)	Log(1+Y)	𝟙(Top Quartile)
$\mathbb{1}(Productivity > Median) \times \mathbb{1}(Post)$	0.0034**	-0.0001	0.0528**	0.0549**	0.0167	0.0147
	(0.0016)	(0.0004)	(0.0253)	(0.0243)	(0.0162)	(0.0253)
$\mathbb{1}(Productivity > Median) \times \mathbb{1}(Post) \times \mathbb{1}(Mid)$	-0.0018	-0.0003	-0.0068	0.0012	-0.0103	0.0109
	(0.0025)	(0.0007)	(0.0306)	(0.0295)	(0.0201)	(0.0305)
$\mathbb{1}(Productivity > Median) \times \mathbb{1}(Post) \times \mathbb{1}(Small)$	-0.0125***	0.0022**	-0.1079***	-0.1050***	-0.0651***	-0.0719**
	(0.0030)	(0.0011)	(0.0299)	(0.0291)	(0.0207)	(0.0305)
Controls	Y	Y	Y	Y	Y	Y
Firm F.E.	Y	Y	Y	Y	Y	Y
Industry × Location × Time F.E.	Y	Y	Y	Y	Y	Y
Observations	1,110,207	1,110,207	860,044	860,044	860,044	860,044
# Unique Firms	51,144	51,144	44,600	44,600	44,600	44,600
R-squared	0.222	0.166	0.722	0.584	0.642	0.515
Y-Mean	0.071	0.006	1.237	1	0.754	0.998
Magnitude- Base (%)	4.76	-1.39	4.26	5.49	2.21	1.47
Magnitude - Small (%)	-12.88	35.94	-4.45	-5.01	-6.43	-5.73

A.1   Appendix

Figure A.1: **Sample GitHub Profile and Contribution Calendar**

This figure presents a sample GitHub profile and the user's contribution calendar.
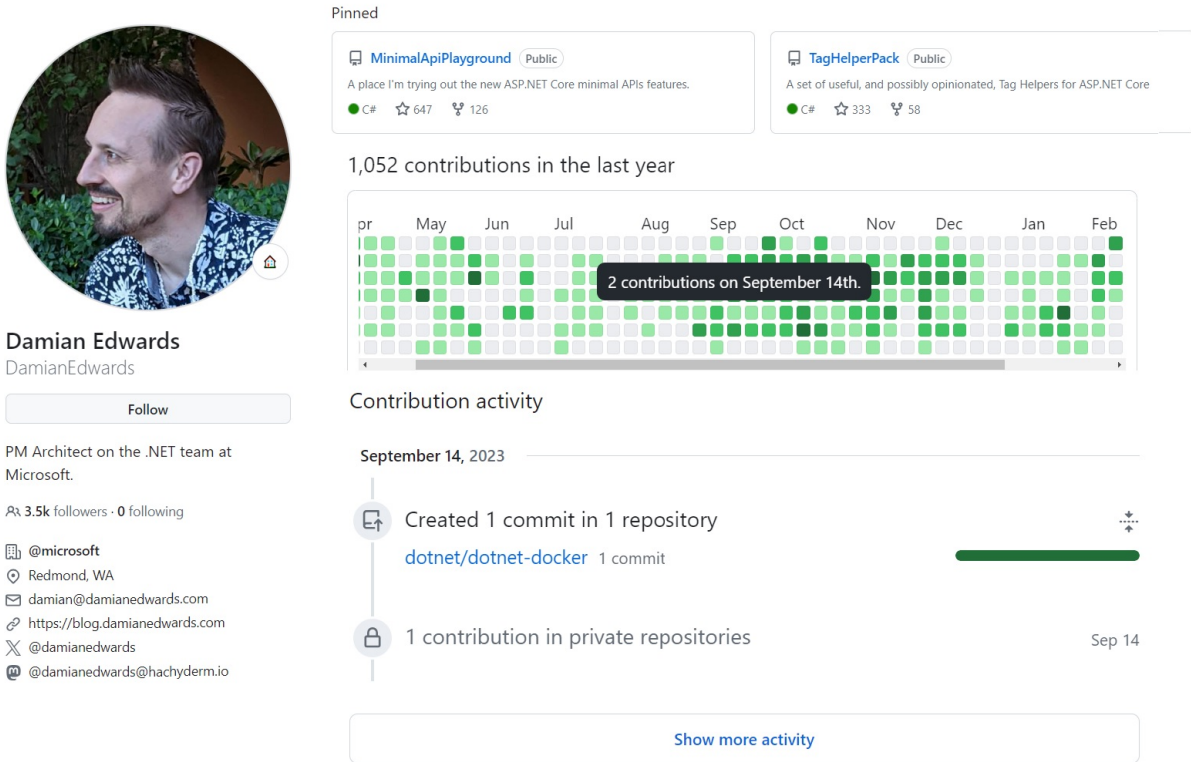
Figure A.2: **Association between Employees' GitHub Contributions, Salaries, and Publications**

This figure presents the scatter plot of employees' GitHub contributions and salary (in panel A) and publications (in Panel B). On the y-axis, we plot the logarithm of employees' average salary in Panel A and the logarithm of the number of articles authored by the employee on Google Scholar, DevTo, and ORCiD in Panel B. On the x-axis, we plot the logarithm of the number of total GitHub contributions by the employees. We collapse the data into equal-sized bins and plot the mean value of the variables in each bin.

Panel A

Panel B

Figure A.3: **GitHub as a Source of Technical Talent**

This figure presents snapshots of third-party web platforms that assist firms in sourcing potential employees using GitHub.

Figure A.4: **GitHub Calendar Before and After Policy Change**

This figure provides an illustration of the change in a user's GitHub contribution calendar before and after GitHub's policy change on May 19, 2016. Prior to May 19, 2016, users' contribution calendar only showcased the number of contributions made by users towards public repositories on GitHub. Following the policy change, users could also include the (anonymized) number of private contributions in their calendar, subsequently increasing their total number of contributions and making their contribution calendar more active.
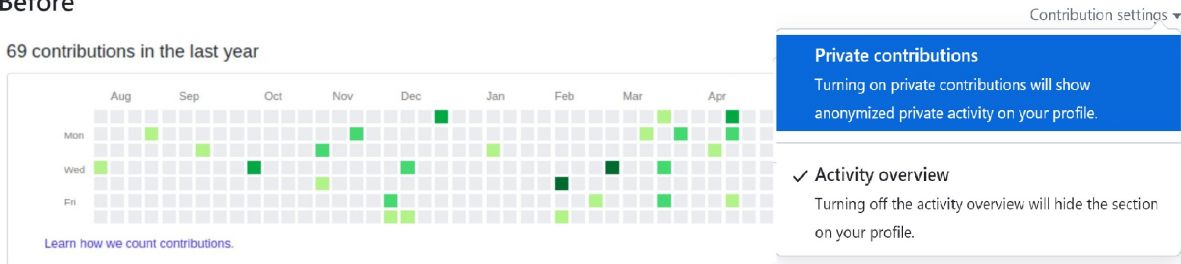
Figure A.5: **Persistence of Employee Productivity: Evidence from AKM Decomposition Before and After Shock**

This figure presents a heat map on the probability of transition of employees across productivity quintiles derived from two different AKMs - before and after GitHub's policy change. The pre-shock productivity quintile sorts employees into 5 buckets based on the productivity estimates from the pre-shock AKM decomposition (AKM sample spanning May 2011 - April 2016). The post-shock productivity quintile sorts employees similarly using AKM decomposition, but in the post-shock period (AKM sample spanning May 2016 - April 2021). A cell (row, column) represents the percentage of people in the pre-shock productivity quintile "row" who transition to post-shock quintile "column". For instance, cell (1, 2) indicates that 21.9% employees, who were initially in quintile 1 based on their productivity estimates from pre-shock AKM move to quintile 2 of productivity when estimates are derived from AKM in the post-period. The analysis is restricted to employees for whom productivity estimates could be derived from the AKM decomposition in both pre- and post-shock periods.
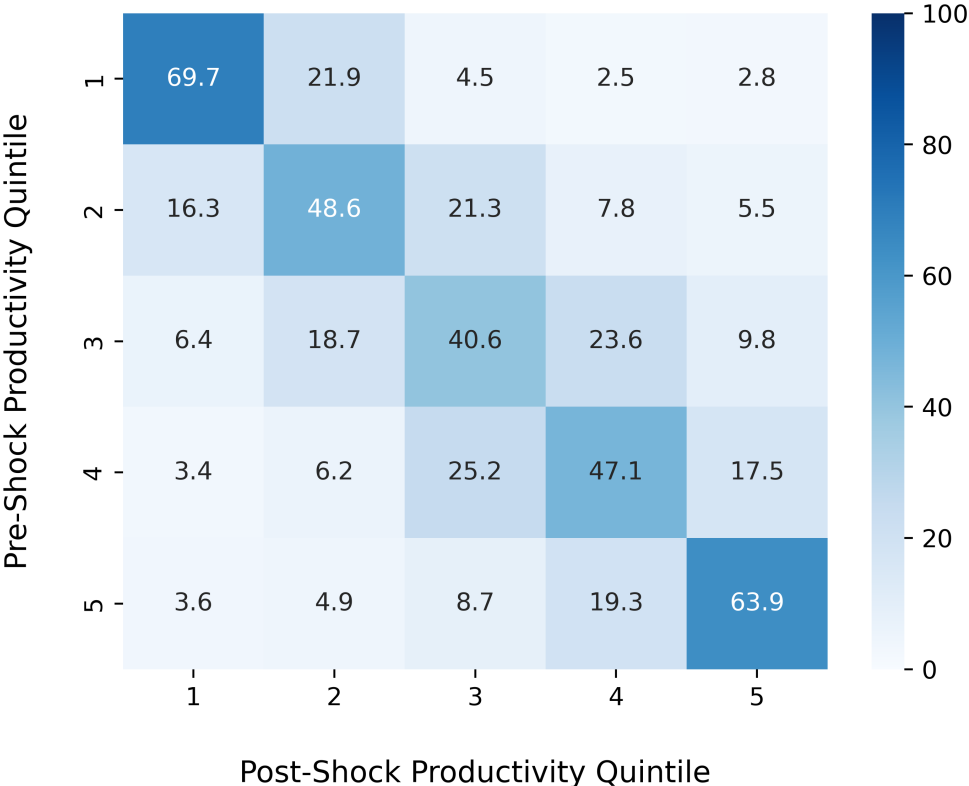
Figure A.6: GitHub Shock & Employee Mobility to Large Firms: Robustness with Dummy Treatment Indicator

This figure presents an event study on employee mobility to large firms around the disclosure of employee productivity information on GitHub. The event study replicates the setup in Figure 2 but uses a dummy measure of *Productivity* as the dependent variable. The dummy indicator takes the value of 1 if the employee's normalized AKM-based productivity is positive and takes the value of 0 otherwise. The rest of the variables and empirical setup remain unchanged.
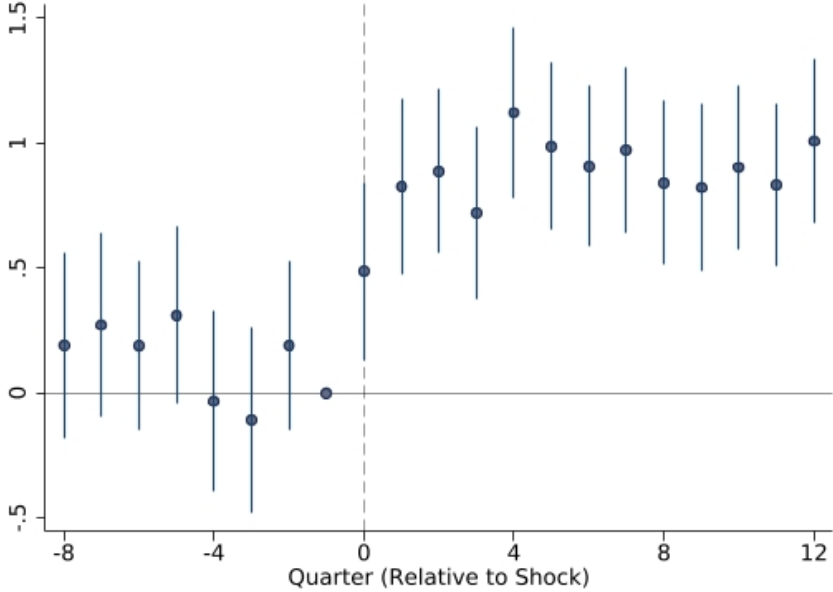
Table A.1: **Sample Characteristics**

This table presents the distribution of employees across five major industries, locations, coding languages, and roles. Language is the employee's most preferred coding language. Role, industry, and location correspond to an employee's initial job role, MSA, and industry respectively.

Panel A: Top 5 Industries	
Industry	**Share of Users (%)**
Computer Software	16.77
Internet	15.17
Information Technology and Services	14.04
Financial Services	4.88
Marketing and Advertising	3.57
Total	54.43

Panel B: Top 5 Locations	
MSA	**Share of Users (%)**
New York-Northern New Jersey-Long Island	11.94
Seattle-Tacoma-Bellevue	7.65
San Francisco-Oakland-Fremont	7.41
San Jose-Sunnyvale-Santa Clara	5.85
Boston-Cambridge-Quincy	5.29
Total	38.14

Panel C: Top 5 Coding Languages	
Language	**Share of Users (%)**
Javascript	26.09
Python	14.03
Java	12.35
Ruby	10.18
C#	5
Total	67.65

Panel D: Top 5 Roles	
Role	**Share of Users (%)**
Software Engineer - Applications	16.66
Software Engineer - Systems	15.7
Full Stack Developer	5.15
Software Developer	3.71
Web Developer	3.26
Total	44.48

Table A.2: **Distribution by Firm Size**

This table presents the distribution of firms by their size buckets and the subsequent classification of firms into small, medium, or large category.

	Firm Size Buckets from LinkedIn							
	[1 - 10]	[11 - 50]	[51 - 200]	[201 - 500]	[501 - 1,000]	[1,001 - 5,000]	[5,001 - 10,000]	[10,001 -)
Share (% All Firms)	19.82	28.25	21.92	9.34	5.39	8.15	2.4	4.73
Percentile	10th	25th	50th	75th	85th	90th	95th	99th
Classification	**Small (Size <= 50)**		**Mid (50 < Size <= 1,000)**			**Large (Size > 1,000)**		
Share (% All Firms)	48.07		36.65			15.28		
Share (% All Employees)	19.35		30.23			50.42		

Table A.3: **GitHub Shock & Employee Mobility to Large Firms: Robustness**

This table presents robustness tests for estimates reported in column 5 of Table 3. In Panel A, we use alternative measures of productivity from alternative AKMs estimated by treating the dependent variable differently. In part A1 (A3), we use the inverse hyperbolic sine (raw value) of user's private contributions as the dependent variable. In part A2, we use five buckets of private contributions (0, [1, 9], [10,99], [100,999], and [1000+)) as input in the AKM. In Panel B, we use different measures of employee productivity computed without estimating the AKM. We define productivity as the logarithm of private contributions (private contributions as a share of total contributions) in part B1 (B2). Panel C presents results from estimates using a binary dummy indicator for employees' productivity. $\mathbb{1}(Productivity > 0)$ is a dummy variable that takes the value of 1 for employees whose normalized productivity is greater than zero and 0 for the remaining employees. In Panel E, we use the original AKM specification but restrict the AKM sample until quarter 2, i.e., April-June quarter of 2014. We then estimate employee mobility in the quarters after quarter 2 (April-June) of 2014. In Panel E, we report our main results by clustering our standard errors alternatively. In part E1 (E2), we double cluster standard errors at the level of employee's industry and location (firm and employee). In Panel F, we control for more granular fixed effects by interacting employee's cohort, role, language, location, and industry with time. Finally, in Panel G, we run the test on a matched sample. Employees with positive normalized productivity estimates are matched with employees with non-positive estimates based on demographic characteristics (gender, white or non-white, whether in software engineering, whether located in SF/NY/Seattle, experience, salary, and seniority) in G1 and demographic characteristics along with public contributions in G2. All panels run tests corresponding to our baseline specification, Column 5 in Table 3. In panels A, B, C, D, F, and G, standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

Outcome:	(1) Coefficient	(2) Std. Error	(3) Observations	(4) R-squared	(5) Y-Mean	(6) Magnitude (%)
Panel A: Alternate AKM Specifications						
A1: AKM with Asine Function						
Productivity × $\mathbb{1}$(Post)	0.0015***	(0.0002)	3,199,173	0.212	0.033	4.64
A2: AKM with Categorial Variables						
Productivity × $\mathbb{1}$(Post)	0.0016***	(0.0002)	3,199,173	0.212	0.033	4.67
A3: AKM with Raw Contribution Numbers						
Productivity × $\mathbb{1}$(Post)	0.0008***	(0.0001)	3,199,173	0.211	0.033	2.5
Panel B: Alternate Productivity Definition						
B1: Log(Private Contributions)						
Productivity × $\mathbb{1}$(Post)	0.0011***	(0.0001)	3,199,173	0.212	0.033	3.63
B2: Share Private Contributions						
Productivity × $\mathbb{1}$(Post)	0.0040***	(0.0008)	3,199,173	0.212	0.033	13.66
Panel C: Dummy Productivity Definition						
C1: Productivity defined as Dummy Indicator Variable						
$\mathbb{1}$(Productivity > 0) × $\mathbb{1}$(Post)	0.0066***	(0.0005)	3,199,173	0.212	0.033	22.92
Panel D: Alternate AKM Sample Period						
D1: Pre Sample AKM (2011-14)						
Productivity × $\mathbb{1}$(Post)	0.0011***	(0.0002)	1,680,047	0.223	0.03	4.2
Panel E: Alternate Clustering						
E1: Double Cluster at Industry & Location						
Productivity × $\mathbb{1}$(Post)	0.0015***	(0.0003)	3,199,173	0.212	0.033	5.27
E2: Double Cluster at Employee & Firm						
Productivity × $\mathbb{1}$(Post)	0.0016***	(0.0003)	3,092,556	0.215	0.034	5.45
Panel F: More Granular Controls						
F1: Control for Cohort × Role × Language × Location × Industry × Time						
Productivity × $\mathbb{1}$(Post)	0.0017***	(0.0004)	745,659	0.37	0.036	5.71
Panel G: Matched Sample						
G1: Match on Individual Demographics						
Productivity × $\mathbb{1}$(Post)	0.0016***	(0.0002)	1,444,081	0.247	0.032	7.38
G2: Match on Individual Demographics and Public Contributions						
Productivity × $\mathbb{1}$(Post)	0.0015***	(0.0002)	1,405,593	0.246	0.032	6.61

Table A.4: **GitHub Shock & Employee Mobility to Different Firm-Size Buckets**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information on employee mobility based on destination firm size. We report the differential impact of GitHub's policy change on productive employees estimated by the β coefficient from the equation:

$$\mathbb{1}(Move)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with *Post*, an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). The dependent variable is a dummy variable that equals 1 when employee i switches job in quarter t to a firm size as indicated in various columns. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:					$\mathbb{1}$**(Job Change)**			
Firm Size:	[1 - 10]	[11 - 50]	[51 - 200]	[201 - 500]	[501 - 1,000]	[1,001 - 5,000]	[5,001 - 10,000]	[10,001 -)
Productivity $\times \mathbb{1}$(Post)	-0.0002	-0.0010***	-0.0009***	-0.0004***	0.0000	0.0003***	0.0001**	0.0011***
	(0.0001)	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Employee F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y	Y	Y	Y	Y
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Industry $\times$ Location $\times$ Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Observations	3,199,173	3,199,173	3,199,173	3,199,173	3,199,173	3,199,173	3,199,173	3,199,173
R-Squared	0.221	0.21	0.206	0.206	0.205	0.205	0.2	0.205
Y-Mean	0.008	0.015	0.016	0.008	0.006	0.01	0.005	0.018
Magnitude (%)	-2.4	-8.03	-6.56	-4.9	-0.09	3.26	3.37	6.9

Table A.5: **GitHub Shock & Employee Mobility to Superstar Firms**

This table presents estimates on the impact of GitHub's disclosure of employee productivity information on employee mobility to superstar firms. We report the differential impact of GitHub's policy change on productive employees estimated by the β coefficient from the equation:

$$\mathbb{1}(MoveToSuperstar)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with *Post*, an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). The dependent variable is an indicator that equals 1 when employee i switches job to a *superstar firm* in quarter t and is 0 otherwise. In each column, we define superstar firms differently based on the average market capitalization or average revenue of the firm in the three years before GitHub's policy change. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:	$\mathbb{1}$(Mobility to Superstar Firms)							
	Within-Industry Top 10%		*Within-Industry Top 1%*		*Across Industry Top 10%*		*Across Industry Top 1%*	
	Market Cap	Revenue	Market Cap	Revenue	Market Cap	Revenue	Market Cap	Revenue
Productivity × $\mathbb{1}$(Post)	0.0008***	0.0008***	0.0005***	0.0004***	0.0008***	0.0008***	0.0005***	0.0003***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Employee F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y	Y	Y	Y	Y
Cohort × Language × Role × Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Industry × Location × Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Observations	3,219,087	3,219,087	3,219,087	3,219,087	3,219,087	3,219,087	3,219,087	3,219,087
# Unique Employees	143,447	143,447	143,447	143,447	143,447	143,447	143,447	143,447
R-Squared	0.191	0.194	0.175	0.18	0.193	0.196	0.18	0.188
Y-Mean	0.017	0.018	0.007	0.006	0.017	0.017	0.009	0.006
Magnitude (%)	5.35	4.84	8.61	7.58	5.38	5.3	6.93	4.9

Table A.6: **GitHub Shock & Employee Mobility to Star Startups**

This table presents estimates on the impact of GitHub's disclosure of employee productivity information on employee mobility to star startups. We report the differential impact of GitHub's policy change on productive employees estimated by the β coefficient from the equation:

$$\mathbb{1}(MoveToStarStartup)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is the normalized productivity estimate of the employee i derived from the AKM decomposition interacted with *Post*, an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). The dependent variable is an indicator that equals 1 when employee i switches job to a *star startup* in quarter t and is 0 otherwise. We define a firm as a startup if the firm's age at the time of GitHub's policy change was less than ten years. Subsequently, as indicated across different columns, we define star startups differently based on the final venture funding round or the total funding raised before GitHub's policy change. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)
Outcome:	\multicolumn{5}{c}{$\mathbb{1}$(Mobility to Star Startups)}				
	Funding > Series B	*Funding > Series D*	*Funding > \$ 100 Mn*	*Funding > \$ 500 Mn*	*Funding > \$ 1 Bn*
Productivity $\times$ $\mathbb{1}$(Post)	-0.0002**	0.0001	0.0001	0.0001***	0.0001***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Employee F.E.	Y	Y	Y	Y	Y
Time F.E.	Y	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y	Y
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Y	Y	Y	Y	Y
Industry $\times$ Location $\times$ Time F.E.	Y	Y	Y	Y	Y
Observations	3,145,396	3,145,396	3,145,396	3,145,396	3,145,396
# Unique Employees	143,437	143,437	143,437	143,437	143,437
R-Squared	0.173	0.168	0.166	0.155	0.151
Y-Mean	0.005	0.003	0.003	0.001	0.001
Magnitude (%)	-5.03	3.07	2.75	13.08	14.77
10th percentile Firm Size	[11 - 50]	[11-50]	[51 - 200]	[201 - 500]	[501 - 1,000]
Median Firm Size	[51 - 200]	[201 - 500]	[201 - 500]	[1,001 - 5,000]	[1,001 - 5,000]
90th percentile Firm Size	[501 - 1,000]	[1,001 - 5,000]	[1,001 - 5,000]	[10,001 -)	[10,001 -)

Table A.7: **GitHub Shock & Employee Mobility to Different Industry and Location**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information based on destination firm characteristics. We report the differential impact of GitHub's policy change on productive employees, estimated by the β coefficients from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta Productivity_i \times \mathbb{1}(Post_t) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is an interaction of the employee's normalized productivity estimate derived from the AKM decomposition and an indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). The dependent variable is mobility to a subset of large firms ($> 1,000$ employees) and differs in each column based on the subset. In columns 1-2, the subset is based on whether the move is associated with a change in the employee's location. In columns 3-4, the subset is based on whether the industry categorization changes with the new job. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)
Outcome:	$\mathbb{1}$(Job Change to Large Firms with)			
	Location Change		Industry Change	
	Yes	No	Yes	No
Productivity $\times$ $\mathbb{1}$(Post)	0.0010***	0.0005***	0.0013***	0.0002***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Employee F.E.	Y	Y	Y	Y
Time F.E.	Y	Y	Y	Y
Salary & Seniority Controls	Y	Y	Y	Y
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Y	Y	Y	Y
Industry $\times$ Location $\times$ Time F.E.	Y	Y	Y	Y
Observations	3,199,173	3,199,173	3,198,865	3,198,865
# Unique Employees	143,446	143,446	143,446	143,446
R-Squared	0.213	0.204	0.214	0.19
Y-Mean	0.018	0.015	0.029	0.004
Magnitude (%)	6.36	4	5.27	5.46
T-test ($H_0 : \beta_{Yes} = \beta_{No}$)	(0.0145)**		(0.9836)	

Table A.8: **GitHub Shock & Employee Mobility: Heterogeneity by Non-Compete Enforceability**

This table presents estimates on the heterogeneous impact of GitHub's disclosure of employee productivity information based on state-level variation in the enforcement of Non-Comepete Agreements (NCAs). We report the differential impact of GitHub's policy change on productive employees, estimated by the β coefficients from the equation:

$$
\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta_1 Productivity_i \times \mathbb{1}(Post_t)
$$
$$
+ \beta_2 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(Non-CompeteNot/LessEnforced_i)
$$
$$
+ \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}
$$

The independent variable is the triple interaction of three variables: employee's normalized productivity estimate derived from the AKM decomposition, an indicator variable that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016), and an indicator that takes the value of 1 for employees in states where NCAs are not/less enforced and 0 otherwise. In columns 1-2 (3-4), the NCA enforceability variable takes the value of 1 for states that almost entirely prohibit NCAs (states that do not have any statute governing NCAs) and 0 otherwise. We use US government's official archives to get state-level variation in the enforcement of NCAs as of 2016.[a] The dependent variable is mobility to large firms ($> 1,000$ employees). X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)
Outcome:	Job Change to Large Firms			
Productivity × $\mathbb{1}$(Post)	0.0013***	0.0013***	0.0010***	0.0011***
	(0.0002)	(0.0002)	(0.0002)	(0.0003)
Productivity × $\mathbb{1}$(Post) x$\mathbb{1}$(NCAs Not Enforced)	0.0006**	0.0007**		
	(0.0003)	(0.0003)		
Productivity × $\mathbb{1}$(Post) x$\mathbb{1}$(No Statute for NCs)			0.0007***	0.0006**
			(0.0003)	(0.0003)
Employee F.E.	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Cohort × Language × Role × Time F.E.	Yes	Yes	Yes	Yes
State × Time	Yes	Yes	Yes	Yes
Industry × Location × Time F.E.	No	Yes	No	Yes
Observations	3,425,232	3,199,173	3,425,232	3,199,173
# Unique Employees	153,374	143,446	153,374	143,446
R-Squared	0.173	0.212	0.173	0.212
Y-Mean	0.033	0.033	0.033	0.033
Magnitude (%)	4.67	4.59	3.48	3.7

[a]Link: `https://obamawhitehouse.archives.gov/sites/default/files/competition/state-by-statenoncompetesexplainer_unembargoedfinal.pdf`

Table A.9: **GitHub Shock & Employee Mobility: Heterogeneity by Local Economic and Demographic Characteristics**

This table presents estimates on any heterogeneous impact of GitHub's disclosure of employee productivity information based on local economic and demographic characteristics. We report the differential impact of GitHub's policy change on productive employees, estimated by the β coefficients from the equation:

$$\mathbb{1}(MoveToLargeFirm)_{i,t} = \beta_1 Productivity_i \times \mathbb{1}(Post_t)$$
$$+ \beta_2 Productivity_i \times \mathbb{1}(Post_t) \times \mathbb{1}(MSAIndicator_i) + \lambda X_{it} + \alpha_i + \gamma_{it} + \varepsilon_{i,t}$$

The independent variable is a triple interaction of three indicators: the employee's normalized productivity estimate derived from the AKM decomposition, indicators for employee's initial location (MSA), and indicator that takes the value of 1 for quarters after GitHub's policy change (quarter 2, i.e., April-June 2016). Each column includes a triple interaction with a dummy indicator that takes the value of 1 if the employee's initial MSA has an above-median value of the economic/demographic characteristic described in the column at the time of GitHub's policy change. X_{it} represents employee's time-varying salary and seniority controls. α_i is employee fixed effects. γ_{it} represents employee's cohort, role, and language interacted with time fixed effects as well as employee's location and industry interacted with time fixed effects. Cohort represents an employee's graduation year corresponding to the last degree. Language is the employee's most preferred coding language. Role, industry, and location represent the employee's first job role, MSA, and industry in our sample period respectively. Standard errors are clustered at the employee level. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:				$\mathbb{1}$(Job Change to Large Firms)		
Productivity $\times$ $\mathbb{1}$(Post)	0.0016***	0.0014**	0.0016***	0.0018***	0.0016***	0.0016***
	(0.0002)	(0.0006)	(0.0002)	(0.0005)	(0.0002)	(0.0003)
Productivity $\times$ $\mathbb{1}$(Post) $\times$ $\mathbb{1}$(MSA Indicator)	-0.0009	0.0002	-0.0001	-0.0002	-0.0001	-0.0001
	(0.0007)	(0.0006)	(0.0004)	(0.0006)	(0.0003)	(0.0004)
MSA Indicator - Above Median	Republican Vote Share	Share of Immigrants	Share of White Population	Median House Value	Change in House Value	Gini Index
Employee F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cohort $\times$ Language $\times$ Role $\times$ Time F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ Location $\times$ Time F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Observations	3,132,014	3,142,580	3,143,808	3,143,808	3,143,808	3,143,808
# Unique Employees	140,433	140,911	140,966	140,966	140,966	140,966
R-Squared	0.209	0.21	0.21	0.21	0.21	0.21
Y-Mean	0.033	0.033	0.033	0.033	0.033	0.033

Table A.10: **GitHub Shock & Firm Productivity: Robustness**

This table presents robustness tests for estimates reported in columns 3-6 of Table 10. In Panel A, we present robustness results using the same specification as in Table 10 but using different measures of public contributions. In columns 1-2 and 5-6, firm contributions include contributions made by incumbent employees towards repositories owned by an organization. In columns 3-4 and 7-8, firm contributions include all employee contributions. In Panel B, we present robustness using the same measure of public contributions as in Table 10 but using different specifications. In columns 1 and 5 (2 and 6), the dependent variable takes the value of 1 if the firm contribution is in the top tercile (top decile) and 0 otherwise. In columns 3 and 7 (4 and 8), the dependent variable takes the value of 1 if the firm contribution is greater than zero (greater than ten) and 0 otherwise. We also estimate coefficients for double interaction of indicator variables *Mid* and *Small* with *Post*, but do not report them for brevity. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

16

	Panel A							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:	Total Contributions				Contributions/Employees			
	Organization Repos		All Contributions		Organization Repos		All Contributions	
	Log(1+Y)	$\mathbb{1}$(Top Quartile)	Log(1+Y)	$\mathbb{1}$(Top Quartile)	Log(1+Y)	$\mathbb{1}$(Top Quartile)	Log(1+Y)	$\mathbb{1}$(Top Quartile)
$\mathbb{1}$(Productivity > Median) × $\mathbb{1}$(Post)	0.0349	0.031	0.0841***	0.0721***	0.009	0.0048	0.0384**	0.0522**
	(0.0243)	(0.0237)	(0.0246)	(0.0239)	(0.0148)	(0.0248)	(0.0163)	(0.0255)
$\mathbb{1}$(Productivity > Median) × $\mathbb{1}$(Post) × $\mathbb{1}$(Mid)	-0.0014	0.0233	0.0379	0.0428	-0.0091	0.0169	0.0278	0.0294
	(0.0291)	(0.0289)	(0.0294)	(0.0286)	(0.0185)	(0.0299)	(0.0198)	(0.0301)
$\mathbb{1}$(Productivity > Median) × $\mathbb{1}$(Post) × $\mathbb{1}$(Small)	-0.0812***	-0.0701**	-0.0970***	-0.0713**	-0.0465**	-0.0364	-0.0507**	-0.0651**
	(0.0285)	(0.0286)	(0.0286)	(0.0280)	(0.0191)	(0.0299)	(0.0202)	(0.0298)
Observations	860,044	860,044	859,337	860,044	860,044	860,044	859,398	860,044
# Unique Firms	44,600	44,600	44,597	44,600	44,600	44,600	44,597	44,600
R-squared	0.716	0.577	0.678	0.54	0.639	0.518	0.571	0.447
Y-Mean	0.946	0.974	1.342	0.998	0.55	0.998	0.824	0.998
Magnitude- Base (%)	3.68	3.18	6.27	7.23	1.63	0.48	4.66	5.23
Magnitude - Small (%)	-4.89	-4.02	-0.96	0.09	-6.82	-3.17	-1.5	-1.29

	Panel B							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Outcome:	Total Contributions				Contributions/Employees			
	$\mathbb{1}$(Top Tercile)	$\mathbb{1}$(Top Decile)	$\mathbb{1}$(Y>0)	$\mathbb{1}$(Y>10)	$\mathbb{1}$(Top Tercile)	$\mathbb{1}$(Top Decile)	$\mathbb{1}$(Y>0)	$\mathbb{1}$(Y>10)
$\mathbb{1}$(Productivity > Median) × $\mathbb{1}$(Post)	0.0371*	0.0306	0.0124*	0.0104*	0.0273	-0.0111	0.0124*	0.0002
	(0.0193)	(0.0498)	(0.0068)	(0.0059)	(0.0205)	(0.0410)	(0.0068)	(0.0043)
$\mathbb{1}$(Productivity > Median) × $\mathbb{1}$(Post) × $\mathbb{1}$(Mid)	0.0163	-0.0242	0.0051	0.0024	0.0023	-0.0291	0.0051	-0.0018
	(0.0235)	(0.0589)	(0.0083)	(0.0072)	(0.0246)	(0.0516)	(0.0083)	(0.0054)
$\mathbb{1}$(Productivity > Median) × $\mathbb{1}$(Post) × $\mathbb{1}$(Small)	-0.0618***	-0.1362**	-0.0143*	-0.0246***	-0.0465*	-0.1358**	-0.0143*	-0.0158***
	(0.0233)	(0.0570)	(0.0082)	(0.0071)	(0.0245)	(0.0529)	(0.0082)	(0.0055)
Controls	Y	Y	Y	Y	Y	Y	Y	Y
Firm F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Industry × Location × Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Observations	860,044	860,044	860,044	860,044	860,044	860,044	860,044	860,044
# Unique Firms	44,600	44,600	44,600	44,600	44,600	44,600	44,600	44,600
R-squared	0.578	0.578	0.552	0.585	0.525	0.492	0.552	0.496
Y-Mean	0.993	1.024	0.457	0.222	0.996	1	0.457	0.107
Magnitude- Base (%)	3.74	2.99	2.71	4.67	2.74	-1.11	2.71	0.16
Magnitude - Small (%)	-2.48	-10.3	-0.43	-6.4	-1.93	-14.7	-0.43	-14.56

Table A.11: **GitHub Shock & Firm Market Reaction**

This table presents estimates on the impact of GitHub's disclosure of employee productivity information on firm value for Compustat firms. The dependent variable is the firm's abnormal return using the Fama-French 3-factor model. The event window starts 252 days (one calendar year) before GitHub's policy change and ends on the day of the change. The independent variable Event Date takes the value 1 on the day of the policy change and 0 otherwise. High-productive (Low-productive) firms have above-median (below-median) AKM-based average employee productivity. Small (Large) firms comprise the bottom (top) 25th percentile firms in terms of employee size.

	(1)	(2)	(3)	(4)	(5)
	All Firms	High Productive		Low Productive	
		Small	Large	Small	Large
Event Date	-0.0017***	-0.0065**	0.0016	-0.0033	0.0004
	(0.0006)	(0.0028)	(0.0021)	(0.0029)	(0.0022)
#Firms	1679	112	106	136	57

Table A.12: **Industry Labor Market Concentration**

This table presents correlations between the change in industry-level concentration around GitHub's policy change and the average productivity of GitHub employees within the industry. The dependent variable is the change in the Herfindahl-Hirschman Index (HHI) for employment five years before (2011-2015) and five years after GitHub's policy change (2016-2020). In columns 1-2, the industry and employment figures are sourced from LinkedIn while in columns 3-4, the data is based on Compustat firms (SIC 3-digit industries). In columns 1 and 3 (2 and 4), the independent variable is the normalized share of employees with positive private contributions (normalized average number of private contributions) at the time of GitHub's policy change within the industry. Standard errors are clustered at the LinkedIn industry level in columns 1-2 and SIC 3-digit industry level in columns 3-4. Significance levels: *(p<0.10), **(p < 0.05), ***(p < 0.01).

	(1)	(2)	(3)	(4)
Outcome:	Change in Industry Labor Market Concentration (ΔHHI)			
	All Firms (LinkedIn)		Compustat Firms (SIC 3-Digit)	
Share Emp with Private Contr > 0	0.064**		0.140***	
	(0.029)		(0.002)	
Average Private Contr		0.067**		0.137***
		(0.027)		(0.013)
Observations	143	143	221	221
R-squared	0.098	0.107	0.133	0.128