

Takeovers and Knowledge Worker Productivity

András Danis*

Evgeny Gushchin[†]

Beate Thies[‡]

June 22, 2026

Abstract

We study how takeover announcements affect the productivity of knowledge workers employed by target firms. Using data from GitHub to measure individual work output and a stacked event study specification, we find that work output declines by 14% following takeover announcements. We also find suggestive evidence that code quality deteriorates. The effect on quantity is more pronounced for acquisitions associated with a larger risk of layoffs, specifically within-industry takeovers. Also, the results are weaker in states where takeovers are likely to be motivated by the wish to acquire skilled employees. These patterns are consistent with stress and anxiety induced by takeover uncertainty. The findings highlight a previously under-explored channel through which takeovers can be costly for acquirers. More broadly, the welfare effects of takeovers might extend beyond capital and product markets to include non-trivial productivity and mental health costs for affected workers.

JEL classification: G34, J50.

Keywords: mergers, acquisitions, workers, productivity, GitHub.

We are grateful for comments from Gabor Bekes, Andrey Golubov, Abhinav Gupta, Alice Kuegler, Ernst Maug, Timea Molnar, Elena Simintzi, Adam Zawadowski, and from seminar and conference participants at CEU Department of Economics, CEU Department of Network and Data Science, Mergers and Acquisitions Research Centre (MARC) Conference 2026 at Bayes, FIRS 2026, the Edinburgh Corporate Finance Conference 2026, and the 2025 VGSF anniversary conference at WU Vienna University of Economics and Business. We are grateful to Andrey Golubov for sharing data on non-compete agreement enforcement.

*Central European University, Department of Economics; DanisA@ceu.edu

[†]Central European University, Department of Economics; Gushchin_Evgeny@phd.ceu.edu

[‡]University of Vienna, Department of Economics; beate.thies@univie.ac.at

1. Introduction

What happens to the productivity of a firm’s employees when the firm becomes the target of a takeover? This paper examines the effect of takeover announcements on the productivity of knowledge workers employed by the target firm. We measure productivity by the output of individual knowledge workers. As a proxy for output, we calculate the weekly number of contributions to GitHub, the world’s premier code-sharing platform.

This question is economically important for several reasons. First, the market for mergers and acquisitions (M&As) represents one of the largest components of global capital reallocation. In the United States alone, the total M&A transaction value in 2024 was 1.2 trillion dollars according to S&P Global (Toomey, 2025). This takeover activity directly affects a large number of employees each year. Second, the importance of knowledge work within the modern economy is steadily increasing, as innovation, research, and professional services account for a growing share of value added in advanced economies. Understanding how takeovers affect the productivity of these workers is thus central to assessing both firm-level and aggregate welfare consequences.

While the existing literature has extensively analyzed the effects of takeovers on shareholders, bondholders, consumers, suppliers, and employment, much less is known about how takeovers influence worker productivity (see Golubov (2025) and Eckbo, Malenko, and Thorburn (2025) for recent reviews). From a managerial perspective, productivity changes represent a key channel through which takeovers can create — or destroy — firm value. From a policy perspective, the effects on workers constitute an important component of the aggregate welfare effects of takeovers that have been under-researched in the literature.

Using a novel measure of individual-level output among knowledge workers, we find that takeover announcements lead to a statistically and economically significant decline in worker productivity. The effect is robust across a wide range of empirical specifications and persists beyond the short term, suggesting that there might be a long-term effect. Quantitatively, the decline corresponds to roughly 9–14% relative to the sample average of output, implying

a meaningful economic magnitude.

We collect information on completed acquisitions of firms located in the U.S. or Canada from Capital IQ that were announced between 2016–2023. We use data from GHTorrent to obtain a detailed GitHub user-by-year panel between 2015 and 2021. This allows us to match users’ affiliation to the names of target firms.

GitHub contains public and private projects, or repositories. Public repositories are visible to everyone, while private repos are only visible to the project team. Since we are interested in professional developers, we focus on users who contribute to public projects owned by organizations rather than by individual users.

We augment the merged sample with data from the GitHub API, which allows us to count the number of contributions such as commits, issues, pull requests, etc., for each user in each week. Most existing studies use only GHTorrent which provides information on public repositories. The GitHub API has a key advantage: it counts contributions not only to public repositories but also to private ones, although only for users who choose to report their private activity.

Through the use of GitHub data, we focus on a select group of occupations, namely programmers, software developers, software engineers, data engineers, data scientists, and similar roles. However, this is arguably an important group of occupations in today’s economy. Since we merge the sample of takeover targets with the GitHub data, our final sample is tilted towards technology firms or companies with at least some employees in the occupations listed above. We assume that our measure of GitHub contributions is at least correlated with the kind of productivity that firms care about. This assumption is supported by Gupta, Nishesh, and Simintzi (2024), who show that programmers who have more contributions on GitHub have better job outcomes later.

Our main methodology is an event study regression, which exploits the fact that different takeovers are announced at different points in time. Our control group consists of employees at firms that will be acquired at a future point in time (i.e., not yet treated users). We

examine an event window from 12 weeks before to 12 weeks after the announcement week. The main outcome variable is the number of contributions to GitHub at the user-week level. The reason for the narrow event window is that we want to be sure that a restructuring of the target firm does not mechanically reduce the number of GitHub contributions. Workers could be reassigned or even fired after a takeover, and we do not know the exact career history of each user. With a narrow event window around the announcement date, where many takeovers have not even become effective yet, it is unlikely that target workers are fired or reassigned to different projects.

We find a treatment effect of -1.66 contributions, averaged across all 12 post-treatment weeks. The average user makes 18.4 contributions per week in our sample, so the affected employees contribute 9% less to GitHub. The median is six contributions per user-week, which implies an even larger reduction in output relative to the median. This is a non-negligible reduction in work output. The negative effect remains statistically significant throughout the end of the event window. The estimation does not show any pre-trends, suggesting that treated users behave similarly to control users leading up to the announcement date.

Our results are robust to a large number of alternative specifications. We use a stacked event study specification similar to Gormley and Matsa (2011) or Cengiz, Dube, Lindner, and Zipperer (2019). In simple terms, we create a separate control group for each event, or cohort. Then, we control for cohort-specific event-time fixed effects and cohort-specific user fixed effects. We follow the recommendations for stacked event studies in Wing, Freedman, and Hollingsworth (2024). The treatment effect is slightly larger, at -2.59 ($p < 1\%$) contributions to public and private repositories per user per week, which corresponds to 14% of the unconditional mean.

Are there any long-term effects on worker productivity? Since workers could be reassigned or even fired after a takeover, it is difficult to estimate any long-term effects. With this caveat in mind, we expand the event window to include 26 weeks, or half a year, after the takeover announcement date. We find persistent negative effects on worker output throughout the

end of the extended event window. These findings provide suggestive evidence that there might be a long-term effect on worker productivity. To approximate the dollar amount of the loss in worker productivity, we follow the methodology in Holub and Thies (2023) and estimate a \$6,048 loss per worker over a six-month period.

We also separately estimate the contributions to private and public repos. We find that the combined effect across all 12 post-announcement weeks is -1.7 contributions ($p = 0.006$) for private repos and -0.87 contributions ($p = 0.18$) for public repos. This suggests that most of the effect comes from contributions to private repos.

The deal with the largest number of users in our sample is Red Hat, which was acquired by IBM in 2019. In order to rule out the possibility that this single takeover drives our results, we drop all Red Hat users from our sample and repeat the stacked event study regression. We show that the treatment effect across all 12 post-announcement weeks is -2.73 contributions ($p < 0.01$), which is an even larger reduction than without dropping Red Hat.

In large M&A deals, the completion date is several months later than the announcement date. For most small takeovers in our dataset, the announcement date and the completion date fall on the same day, which is presumably the actual completion date. This could pose a problem if the true announcement date is much earlier than the indicated announcement date. This would imply that some workers at the target firm would be aware of the planned takeover before our “event week”. Also, after the completion date, the target firm might be restructured and workers might be reassigned to different tasks, or even fired, which could mechanically reduce contributions to repos. In another robustness test, we restrict ourselves to deals where the announcement date is strictly earlier than the completion date. Even though our sample is smaller, the results hold.

In addition to measuring the quantity of work, we construct a measure of work quality for a small sample of public repositories. We use the number of software bugs reported for each repository, normalized by repo size, as an inverse proxy for code quality. We find that

following a takeover announcement, the bug rate increases. This effect is less statistically significant than the effect on work quantity, but it is nonetheless suggestive of a deterioration in code quality.

There are several mechanisms through which takeovers could have an effect on productivity, which are not mutually exclusive. Takeover announcements often generate uncertainty regarding job security, workplace structure, and future leadership, potentially inducing stress and anxiety that lower worker performance. They may also divert attention away from core tasks as employees discuss the implications of the transaction. On the other hand, workers might also want to exert high effort in order to keep their job or to improve the chance of obtaining a new job. Finally, employees might just pretend to work hard, also with the intention of keeping their employment.

Our findings of a decrease in work output are consistent with the “stress and anxiety” channel and with the “distraction” hypothesis. Especially the stress and anxiety channel is based on a fear of layoffs, and we design three additional tests that are consistent with this fear.

In our first test, we distinguish between takeovers that are likely motivated by cost-cutting measures (such as layoffs), and deals that are probably motivated by other reasons (acquiring new technology, new products, etc.). We assume that takeovers where the target and the acquirer are in the same industry, it is more likely that layoffs will follow. The same proxy is used in the empirical takeover literature, such as Dessaint, Golubov, and Volpin (2017) and Chen, Gao, and Ma (2021). The negative effects on worker productivity tend to be stronger in within-industry takeovers.

In a second test, we exploit U.S. state-level variation in the Inevitable Disclosure Doctrine (IDD). In some states, firms can easily argue in court that employees who take a job at a competing firm will inevitably reveal the trade secrets of the previous employer. Chen et al. (2021) find that firms headquartered in an IDD state are significantly more likely to be acquired than companies in other states. They argue that acquirers want to have access to

the target firm’s employees. Therefore, we hypothesize that target employees in IDD states are not particularly worried about a future layoff. We find that takeover announcements in IDD states on average do not have any negative effects on worker productivity at the target firms.

A third test of our hypothesis that the results are driven by the fear of layoffs, we conduct a placebo test exploiting partial acquisitions. Since such transactions are less likely to imply a change of control in the firm, its announcement should induce less anxiety about job security among employees at target firms. We indeed find weaker results in partial acquisitions.

All of these tests support the stress and anxiety channel, but are also somewhat consistent with the distraction hypothesis. However, the stress and anxiety channel is further backed up by Bach, Bos, and Silva (2025), who show that takeovers in Sweden have a negative effect on workers’ mental health.

While we are not able to clearly distinguish between the stress and distraction channels, our results suggest that workers’ effort signaling — whether honest or fake effort — is not the dominant mechanism.

The results have implications for both firms and policy makers. For acquirers, they suggest that target firm valuations should incorporate the expected productivity costs associated with takeovers, potentially implying an optimal downward adjustment in bid prices. Our data does not allow us to determine to what extent bidders already internalize these costs. More broadly, the findings indicate that the welfare effects of mergers and acquisitions extend beyond capital and product markets to include non-trivial productivity and mental health costs for affected workers. If the main channel is distraction, takeovers lead to temporary losses in output; if stress and anxiety dominate, they generate both output losses and private disutility from reduced well-being. Either way, the cost-benefit analysis of M&As should take into account these additional costs.

Our paper is related to the literature on the effects of M&As on workers. Most studies in this literature find that takeovers have negative implications for employees at target firms.

Dessaint et al. (2017) show that a large fraction of the financial value created in the average takeover is driven by layoffs. Lagaras (2020) uses Brazilian administrative data to show that takeovers have large and persistent negative effects on wages. Arnold (2025) uses U.S. census data to show that acquisitions have negative effects for workers, emphasizing the role of labor market concentration.¹ These studies focus on wages and employment, rather than the productivity of knowledge workers.

A closely related paper is Andres, Bazhutov, Cumming, Köchling, and Limbach (2025), on the effect of takeover rumors on firm productivity. The authors look at speculative rumors that do not materialize, while we look at announcements of actual takeovers. Their main outcome variable is firm sales per employee, which is a crude measure of productivity, as it could be affected by product prices or customer demand. Our contribution is to provide a new and more granular measure of worker output. While they have a robustness test at the employee level, it focuses on the number of patents by inventors and the number of recommendations by stock analysts. GitHub data allows us to examine a different and larger set of occupations. Finally, we use a stacked event study regression, which has fewer weaknesses than difference-in-differences estimators.

This literature includes studies that look at other worker-related outcome variables as well. Using detailed administrative data, Bach et al. (2025) find that firm acquisitions have a negative effect on the mental health of employees in Sweden. For a review of the literature on labor and corporate finance, see Maug (2025). We contribute to this literature by showing that takeovers have a negative effect on knowledge worker productivity.

Finally, this paper contributes methodologically by developing and applying new measures of knowledge-worker productivity, a notoriously difficult construct to quantify. Prior studies have used diverse proxies for individual professions: patent counts for innovators (Trajtenberg, 1990), student outcomes for teachers (Rivkin, Hanushek, and Kain, 2005),

¹There are several other studies showing a negative effect of takeovers on wages and employment: Gugler and Yurtoglu (2004), Siegel and Simons (2010), Maksimovic, Phillips, and Prabhala (2011), Li (2013), Prager and Schmitt (2021), Gehrke, Maug, Obernberger, and Schneider (2021), and Hosken, Larson-Koester, and Taragin (2024).

publication counts for researchers (Cui, Ding, and Zhu, 2022; Yang, Cai, and Li, 2024), and case completions for judges (Chen, He, and Yamashita, 2025; Grajzl and Silwal, 2020). These existing measures, however, often suffer from low frequency and a narrow focus on specific occupations. Our approach provides a higher-frequency measure of output that applies to a broader set of occupations, which allows for more precise identification of productivity responses to takeover events.

2. Setting and data

2.1. *Setting: GitHub*

GitHub, founded in 2008, is the world’s largest coding platform, allowing users to store and collaboratively work on coding projects. It is based on Git, a version control system, that tracks who has changed which lines of a (code) file at what point in time to facilitate especially collaborative work on coding projects (Chacon and Straub, 2014). An individual change to a code file recorded by the version control system is called a *commit*, and thus reflects coding activity by the commit author. On top of this core functionality, GitHub offers several additional tools, e.g. *issues*, which can be used to track bugs, plan tasks or discuss ideas on the project; and *pull requests* which allow to suggest code changes, without directly implementing them, such that the main branch of the project contains only approved work. Users can discuss issues and pull requests using the *comment* function.

GitHub has been growing a lot over time in popularity, and was acquired by Microsoft in 2018. In 2022, the platform had 94 million registered users, up from 30 million in 2019 (GitHub, 2022). The U.S. is the country with the largest number of users.

Importantly, two types of projects (*repositories*) can be created on GitHub: private and public projects. Public projects are visible to everyone, and every user can contribute, e.g., via pull requests (PRs), whereas only team members can make commits. Private projects, by contrast, are only visible to the project team. Users’ activities to public projects are

visualized on their profile in a “contribution graph” and, if users opt to have them included, the number of private activities is also included in the count (but without any details on the activity). For the graph, the following types of activities are counted: making commits, opening of PRs and issues, creating a new repository, and submitting a PR review.

The development of open-source software has a strong tradition on GitHub. Until 2019, public repositories dominated because free accounts could only create public repositories. Since 2019, free accounts can also create unlimited private repositories. Since then, activity in private repos has been growing fast, accounting for almost 80% of all contributions on GitHub in 2022 (GitHub, 2022). Still, the volume of activity in public projects is high. The largest open source projects are mostly commercially-backed, including e.g. Microsoft’s VSCode (20,000 contributors in 2022) or Vercel’s next.js (6,000 contributors). Many tech firms operate organizations on GitHub, hosting their open-source projects. Some of the organizations with the most open source contributors include Microsoft, Meta, Docker, and Google (see GitHub, 2016).

Figure 1 shows the distribution of public actions across hours of the day as reported in Holub and Thies (2023), indicating that they are concentrated during standard business hours. This suggests that contributions to public repositories are likely work-related and not simply side-projects of developers. However, we provide a robustness test for the possibility that a reduction in side-projects explains the drop in contributions after a takeover announcement.

2.2. *Data*

2.2.1. *M&A events*

We collect information on completed acquisitions of firms located in the U.S. or Canada from Capital IQ, which were announced between 2016 and 2023. This database is also used in Golubov and Xiong (2020), among others, and contains some deals that are not included in

LSEG Refinitiv (formerly SDC Platinum).² We consider all deals in which a majority stake in the company was acquired. The data includes information on the name of the target company and the acquirer, their respective industries³, and the dates that the transaction was announced and completed, respectively. While we consider all industries a priori, the matching with the GitHub data (details below) implies that the relevant deals involve mostly tech firms as targets. Some examples for transactions that affected a large number of GitHub users in our data are the acquisition of Red Hat by IBM (announced in 2018), the takeover of VMware by Broadcom (announced in 2022), and the acquisition of GitHub by Microsoft (announced in 2018). Figure 2 shows the full distribution of M&A events considered in our analysis over time.

2.2.2. *GitHub data*

GHTorrent. The GHTorrent database provides information on all GitHub users at multiple points in time, specifically user ids, login names, registration dates, and information provided by the users on their profiles, like location and company affiliation (Gousios and Spinellis, 2012). For each year between 2015 and 2021, we have one snapshot of the user information. It is a longitudinal dataset, in contrast to the GitHub API (see below). Besides, the database includes complete records of users’ activities in public repositories up to June 2019, separately for commits, opening and closing of PRs and issues, as well as comments written on PRs, issues, and commits. It includes information on the project in which the action was conducted and the precise timestamp.

Using the GHTorrent data, we identify all user-by-year observations where users report a company/employer on their profile. We compile a list of company-operated repositories, and further restrict the sample to users who conducted at least one commit in a company-backed repository during the sample period. This complements the self-reported company information to restrict the sample to professional software developers. We drop a handful of

²We compare the deal coverage of the two data vendors Capital IQ and LSEG Refinitiv in Appendix B.

³Buyer industry is only reported if the acquirer is a company.

users who are likely bots based on login names.

We then identify users affected by M&A events by matching employers reported on the GitHub profiles to target company names reported in the Capital IQ database, using a fuzzy string-matching approach, followed by manual elimination of incorrect matches. The matching is based on the company affiliation reported by GitHub users in the year before the takeover. If a user has been affected by multiple takeovers throughout the sample period, we consider only the first event. Figure 3 shows the geographic distribution of the users affected by a takeover. As we consider M&A events involving firms in the U.S. and Canada, unsurprisingly, a large share of the users are located in these two countries, especially in the well-known tech hubs along the east and west coasts. But we also observe users located in other states and even other countries, especially in Europe.

For these users who have ever been affected by a takeover, we then construct a weekly panel of their GitHub activity. With the GHTorrent data, we can measure weekly commits, comments, PR, and issue events, as well as the sum of all activities (total actions) conducted in public repositories, but only between 2014 and June 2019, meaning many takeovers occurring after this period cannot be analyzed with the data. Therefore, we complement GHTorrent with information from the GitHub API.

GitHub API. The GitHub API offers two advantages. It allows the collection of information on GitHub activities up to the current date, i.e., a longer time period than GHTorrent, and it includes information on activity levels in private repositories for users who have chosen that these actions should be reported on their profiles.⁴ It covers exactly those types of activities that are visualized in users’ “contribution graph” on their profile, namely opening and closing of PRs, issues, commits, and creation of new repositories. An important difference from the GHTorrent data is that comments are not included. We collect information on weekly contributions in public repositories, and if available, in private repositories, for

⁴In contrast to contributions to public repositories, which are visible in detail, e.g., the exact lines of code changed in a commit, for activities in private repositories, no details are reported, but only the number of contributions made during a specific time period of interest.

the sample of users affected by takeovers that we identified using the GHTorrent data. The panel covers the period 2014 to 2022. We use the sum of all contributions to private and public repositories as our main outcome variable to capture as much of a developer’s weekly work activity as possible.

To construct the sample, we restrict the panel to users who have meaningful activity levels before the takeover event. Specifically, we require users to have non-zero activity levels in both public and private repos during at least 5 weeks in a baseline period. This baseline period is 2014–15 for users affected by M&A events between 2016–2019, and 2016–17 for users affected by M&A events between 2020–2023. This results in a sample of 1,235 users at 554 different target firms. The full sample has approximately 571,000 user-week observations.

We report summary statistics on the user-by-week panel in Table 1. On average, users conduct 17.9 contributions in public and private projects per week. Out of these contributions, 6.5 are in public repos, while 11.3 are in private repos. For public repos, we also know that there are 4.9 commits per week on average, which makes commits the most frequent type of contribution to a public repo.

The number of contributions is a noisy approximation of work output. There are more sophisticated measures of software developers’ work output, especially in the computer science literature. Examples are the lines of code changed, cyclomatic complexity, the Levenshtein distance, or the Halstead effort measure. However, Gote, Mavrodiev, Schweitzer, and Scholtes (2022) show that these measures are highly correlated with the number of commits.⁵

Output quality. To assess the effects of takeover announcements on developers’ code quality, we exploit the fact that many project teams use the GitHub issue tracker to organize tasks. Issues are opened to report a bug, to request a new feature, or to discuss new ideas, among others. While every user can open an issue in public repositories, only project

⁵Figure 3 in Gote et al. (2022) shows that the number of commits has a Pearson correlation of 0.81 with the number of lines added/deleted/modified, 0.78 with cyclomatic complexity and with Levenshtein distance, and 0.67 with the Halstead effort measure.

members can attach labels to issues to classify them. GitHub provides nine default labels (bug, documentation, duplicate, enhancement, good first issues, help wanted, invalid, question, wontfix) available in every repo, and project members can create additional customized labels. This feature allows us to construct the monthly number of newly found bugs as a measure of error frequency at the project level. We count the number of newly opened issues that receive a label indicating a bug, i.e., a label containing the string “bug”. We use the frequency of bugs as an inverse proxy for the quality of code in a repository.

As this measure is only available at the project level, rather than the developer level, we have to assign treatment status to projects. To do this, we identify all projects operated by target firms. First, we match target firms to organizations on GitHub. We observe all organizations registered by June 2019 in the GHTorrent data. After cleaning target firm and organization names, we match 818 out of all 1,152 firms that face a takeover announcement between 2016 and 2023 to one or more GitHub organizations. In total, we find 835 relevant organizations, as some firms operate multiple organizations. For example, the target firm Thoughtworks, Inc. is matched to the organizations *thoughtworks*, *ThoughtWorksInc*, and *ThoughtWorksStudios*. Second, we identify all public repositories created by these organizations as of June 2019 using GHTorrent. This yields 46,360 repositories.

For these repos, we collect data on issues, including information on attached labels, and commits, including the number of lines added and deleted in the commit, via the GitHub API. As we want to analyze changes in the number of bugs opened in a repo over time, we restrict this sample to repos that have had at least three labeled issues in the baseline period before the takeover announcement, which reduces the sample to 979 repositories, operated by 188 target firms that are treated during 83 distinct announcement months. In these repos, the average number of new issues per month is 5.2, of which 3.2 get labeled, and 0.89 are classified as bugs. We construct the monthly bug rate for all these repositories, defined as the number of new issues that are labeled as bugs, divided by the number of commits created in the repository during the same and the previous month. We scale the number of

new bugs by recent coding activity as a drop in the amount of code written can be expected to mechanically reduce the number of bugs. The average bug rate is 0.043, or 4.3 bugs per 100 commits.

3. Empirical strategy

3.1. Event study

To identify the causal effect of takeovers on employee productivity, we exploit the temporal variation in takeover announcements across the 554 target firms in our sample in a staggered difference-in-differences (DiD) design. The treatment is the announcement of firm takeover, i.e. it varies at the firm $\times$ week level, while the unit of analysis is user $\times$ week. Every takeover announcement can be considered as a sub-experiment.

We first present results from a simple event study. For this, we include every user i for the period 12 weeks before to 12 weeks after the week a_i when the takeover of i 's employer was announced. Dropping observations outside of this relatively narrow event window leads to a large reduction in sample size, from approximately 571,000 to 29,000 (see Table 1, Panels A and B). The summary statistics for the two samples are very similar. In the event study sample, the average user makes 18.4 contributions per week. Table 1, Panel C, separately shows summary statistics for the pre- and post-treatment subsamples, respectively. The number of contributions changes from a pre-treatment average of 19.6 to a post-treatment level of 17.3, suggesting a drop in activity after the takeover announcement. Using the event study sample, we estimate the following regression specification:

$$Y_{it} = \sum_{\tau=-12, \tau \neq -1}^{12} \beta_{\tau} \mathbb{1}\{t - a_i = \tau\} + \mu_i + \theta_{j,y(t),q(t)} + \gamma \text{Holidays}_t + u_{it} \quad (1)$$

where Y_{it} denotes GitHub activity by user i , who is affected by a takeover announcement in week a_i , during calendar week t . We include user fixed effects μ_i to control time-invariant

productivity differences between users. We also use industry-year-by-quarter fixed effects $\theta_{j,y(t),q(t)}$, for general time trends and business cycle dynamics, which we allow to vary across employers’ industries. Moreover, we control for U.S. public holidays, as the majority of users are based in the U.S. The coefficients of interest are given by β_τ . They measure users’ performance in event week τ relative to the omitted week prior to the announcement.

We cluster standard errors at the firm level. The event study specification allows us to test the parallel trends and the no-anticipation assumptions, and to analyze effect dynamics post treatment. In addition, we also report the DiD estimate that reflects the change from the entire pre-announcement period to the entire post-announcement period.

3.2. *Stacked event study*

There are several issues that can arise in staggered treatment designs (e.g., Goodman-Bacon, 2021; Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; Baker, Larcker, and Wang, 2022). Therefore, we also present results from a stacked event study, following the suggestions in Wing et al. (2024). First, to avoid “forbidden comparisons” between users at firms treated towards the end of the sample period and those treated early — which can lead to bias in case of non-constant treatment effects — we build a stacked data set including only “clean controls” for every takeover event. For all users affected by an acquisition in announcement week a , we use only users at not-yet-treated firms as comparison units. Specifically, we consider firms as controls if they receive treatment, i.e. the takeover announcement, between 12 and 18 months later than announcement week a . This ensures that control firms remain untreated during the full event window (12 weeks), but should be on similar trends as the treated firms since they will also experience a takeover soon. After determining the relevant clean controls for every takeover announcement a , we stack the resulting sub-data sets. Events for which we cannot find clean controls are not included as treatment events.

Second, also following Wing et al. (2024), to avoid bias from compositional changes, we

trim the data set to ensure that all users affected by a takeover are observed during the full event window, i.e., 12 weeks before and after the announcement week. This means, the sample includes the same number of both treated and control users in every event time period e , where event time is measured in weeks relative to the announcement week. For this, we have to drop users affected by takeovers in the final months of 2022 from the set of treated units, as the data does not cover them for 12 weeks post treatment.⁶ Given that the GitHub activity data goes back to 2014, while the first takeover announcement is in 2016, all treated users are observed during the full pre-treatment period.

Using the stacked and trimmed analysis sample, we can estimate the following event study model:

$$Y_{iea} = \sum_{\tau=-12, \tau \neq -1}^{12} \beta_{\tau} D_{ia} \cdot \mathbb{1}\{e = \tau\} + \theta_{ia} + \gamma_{ea} + \epsilon_{iea} \quad (2)$$

where i denotes an individual user at a target firm whose takeover is announced in week a . Y_{iea} measures GitHub activity by user i in event week e relative to the announcement week. Following Cengiz et al. (2019) and others, we control for user $\times$ announcement week fixed effects, θ_{ia} , identifying a user in a specific sub-experiment – as users can be part of several sub-experiments if they are a clean control for at least one other takeover. γ_{ea} controls for sub-experiment specific event time fixed effects (e.g., week 9 after the announcement of the Red Hat takeover), which captures general time trends in GitHub use, holidays, or business cycle dynamics. The dummy D_{ia} identifies users who are treated in announcement week a . Thus, β_{τ} should indicate how the productivity of treated users changed from the last week pre-announcement to event week τ relative to the productivity of their control units.

However, as pointed out by Wing et al. (2024), the coefficients β_{τ} in general do not identify a meaningful aggregate of the sub-experiment specific treatment effects. Thus, we

⁶While we would only need to drop events from the last twelve weeks of 2022 in the trimmed dataset to ensure compositional balance, we instead exclude all events occurring after August 2022. This allows us to find a sufficient number of clean controls with announcement dates 12 to 18 months later for all treated units, given that our M&A announcement data end in 2023.

estimate the model using weighted least squares and corrective weights proposed by Wing et al. (2024). This ensures that the coefficients from the weighted least square regressions identify a weighted average of the sub-experiment specific ATTs (average treatment effects on the treated). The weights are given by the share of the stacked analytic sample used in the respective sub-experiment, i.e., the share of treated and control units in that announcement week a . We also report robustness tests with different weights as well as without weights in Section 4.

Since we build individual control groups for every takeover event in the stacked dataset, the number of user $\times$ week observations in the stacked event study sample becomes quite large (543,625), due to the large number of control group observations (515,350). Summary statistics for the stacked sample are presented in Panel D of Table 1. The average number of contributions to GitHub is similar to the simple event study sample.

4. Results

4.1. Main results

We first look at the number of total contributions to GitHub around takeover announcement, but without control variables. Figure 4 shows the weekly average number of contributions to private and public repos in event time. The event is defined as the announcement week of a takeover. We clean the number of contributions by subtracting user-by-Christmas fixed effects and holiday fixed effects. Figure 4 shows that the average number of actions on GitHub falls right after the M&A announcement. Lower activity is still observed 100 weeks after the announcement. For completeness, Figures A1 and A2 in the appendix show separate results for private and public repos, respectively.

This provides suggestive evidence that there might be a drop in contributions after takeover announcements. It also shows that there is no large run-up in activity in the months leading up to a takeover. Next, we estimate an event study regression which controls for a

lot of confounding shocks.

The results of the event study analysis confirm these findings in Figure 5. Firstly, the estimated coefficients for the 12 weeks preceding the announcement week are never statistically significant at the 95% significance level, and the point estimates are close to zero. This suggests that there is no parallel trends violation before the treatment. Since there are no never-treated firms in the sample, the comparison is relative to week -1 and relative to firms that receive treatment at a different point in time. This is important because there might be an anticipation effect leading up to the event. The insignificant pre-treatment coefficient alleviates this concern. However, already during the announcement week and in most weeks after it, the point estimates are negative and statistically significant at the 95% significance level. The estimated coefficients stay at roughly the same level throughout the end of the event window, suggesting no short-term reversal. The estimated treatment effect across the entire post-treatment period is -1.66 ($p < 1\%$), indicating a drop in activity among users affected by the takeover.

The average user makes 18.4 contributions per week in our sample, so the affected employees contribute 9% less to GitHub. The median is six contributions per user-week, which implies an even larger 28% reduction in output relative to the median. This is a non-negligible reduction in work output. A simplified version of this regression specification can be found in Table 2.

Next, we estimate a stacked event study regression that has some advantages relative to a simple event study or to a general two-way fixed effects regression (see Wing et al., 2024, among others). These results can be found in Figure 6. Compared to the simple event study, we get even larger point estimates (in absolute terms). The weekly post-treatment effects are statistically significant at the 95% level for 9 out of 13 treated weeks. There is no violation of the parallel trends assumption before treatment. The estimated treatment effect across the entire post-treatment period is -2.59 ($p < 1\%$), which confirms our simple event study results. Relative to the stacked event study sample mean of 18.98, this implies

a 14% reduction in work output.

We estimate the stacked event study regression separately for contributions to private and public repos in Figure 7. We find that the result is mainly driven by private repos. Although point estimates after the M&A announcement are always negative in both groups, they are never statistically significant for public repos. For private repos, we are observing point estimates that are on average two times larger. The weekly coefficients are statistically significant only for 3 out of 13 treatment weeks, but the standard errors are also wider than in the public repo sample. Nevertheless, the treatment effect across the entire post-treatment period is -1.7 and still significant at the 1% level. A simplified version of the regressions in Figures 6 and 7 can be found in Table 3.

Next, we ask if there are any long-term effects on worker productivity. Since workers could be reassigned or even fired after a takeover, and since we do not know the entire career trajectory of the users in our sample, it is difficult to estimate any long-term effects. The longer the event window, the higher the likelihood that some target workers in our sample will have stopped working for the target, which might mechanically reduce the number of contributions to repositories. We extend the event window to 26 weeks before and 26 weeks after the announcement week in Figure 8. We find persistent negative effects on worker output throughout the end of the extended event window. The average point estimate is -2 in the post treatment period, being statistically significant even on the 25th week after the event. These findings provide suggestive evidence that there might be a long-term effect on worker productivity. However, these results are based on the assumption that firing or internal reassignment decisions do not explain the drop in activity. A simplified version of this regression can be found in Table 4.

4.2. Robustness tests

In Figure 9 we present the results of Poisson Pseudo Maximum Likelihood (PPML) regressions instead of OLS. Poisson regressions might be more appropriate for count data

and might be better able to capture nonlinear effects (Silva and Tenreyro, 2006). The results are even stronger when using PPML. The point estimates are statistically significant for 8 out of 13 treatment weeks, while the pre-trend assumption is still valid. The average estimated coefficient after the event is -0.13 ($p < 1\%$). The point coefficient of a Poisson regression cannot be directly compared to an OLS coefficient. To translate the PPML coefficient into a percentage change, we use the formula $(e^\beta - 1) \times 100$ (see, e.g., Wooldridge, 2010), which implies a 12% reduction in work output. A slightly extended regression specification as well as separate regressions for public and private repos can be found in Table 6.

We perform various additional robustness checks. We confirm that our results are robust to applying alternative weights from Wing et al. (2024) and using no weights as in Cengiz et al. (2019) (see Figure 10 and Table 5). Removing the M&A involving Red Hat, which is an outlier in terms of the number of users in our sample, does not materially alter our results either (see Figure 11).

In large M&A deals, the completion date is several months later than the announcement date. For most small takeovers in our dataset, the announcement date and the completion date fall on the same day, which is presumably the actual completion date. This could be a problem if the true announcement date is much earlier than the indicated announcement date. This would imply that some workers at the target firm would be aware of the planned takeover before our “event week”. Also, after the completion date the target firm might be restructured and workers might be reassigned to different tasks, or even fired, which could mechanically reduce contributions to repos. In Figure 12, we restrict ourselves to deals where the announcement date is strictly earlier than the completion date. Even though our sample is smaller, the results hold.

In the appendix, we replicate our main tests with only GHTorrent data, i.e., without data from the GitHub API. This is possible because GHTorrent has data on public repos, but only for a shorter sample period. We get more statistically significant results if we use only GHTorrent (see Figure A3). GHTorrent data also shows more significant post-treatment

weeks when applying a wider event window (see Figure A4).

4.3. *Code quality*

We have documented a substantial drop in output quantity following the merger announcement. In high-skill jobs like software development, however, the *quality* of output is also of first order importance. We consider the monthly bug rate, i.e., the number of newly detected bugs divided by the number of commits in the same and the previous month, as an inverse measure of code quality.

We conduct the analysis at the monthly level because labeled issues are created much less frequently than commits. We consider an event window from three months before to three months after the month of the takeover announcement, in line with the window used in the weekly analysis. When constructing the sample for the stacked DiD analysis, we only include repositories that have non-missing observations for the bug rate during every month in the observation window, to avoid changes in sample composition across periods. This, together, with the narrow event window, reduces the number of repos from 979 to 343, by 94 target firms. The average bug rate per commit is 0.046 in the control group.

Stacked difference-in-differences estimates are reported in Table 7, indicating that the bug rate increases significantly post announcement. The number of bugs per commit increases by 0.024, which amounts to 52% of the mean level in the control group, indicating a substantial increase in error frequency (column 1). The coefficients are less statistically significant than in the tests for code quantity. We interpret the increase in the bug rate as suggestive evidence for a deterioration in code quality.

We find no evidence for diverging pre-trends as indicated by the insignificant, small coefficient on the interaction term between the treatment group dummy and an indicator for the third month before the takeover announcement, i.e., the first month in the event window (column 4).

We also use two alternative definitions of the bug rate, where we normalize the number

of bugs by the total size rather than the number of commits. Specifically, in columns 2 and 3, we replace commits in the denominator with the number of new lines of codes written, and the number of changed lines of code, respectively.⁷ We find the same pattern of results with these alternative definitions of the outcome, i.e., a significant increase in error frequency post announcement, and no evidence for diverging pre-trends. In terms of magnitude, the estimates indicate .008 additional bugs per line of code changed, which is more than a twofold increase relative to the control group mean.

In Appendix Table A1 we show that the stacked Diff-in-Diff results are robust to using no regression weights or alternative regression weights.

Since the analysis of the bug rate is based on a different sample than the analysis of output quantity, we show estimates of the effect of the takeover announcement on the number of commits in this sample in Appendix Table A2. We find that commits, i.e., coding output, also drops at the repo level in this sample. In summary, the results indicate that on top of the drop in output quantity, the takeover announcement is also associated with a deterioration of work quality. Statistical significance is lower than in the tests for code quantity, which could be driven by the smaller sample in the quality regressions.

4.4. *Mechanism*

There are several a priori reasons why takeovers could have an effect on productivity, which are not mutually exclusive. Takeover announcements often generate uncertainty regarding job security, workplace structure, and future leadership, potentially inducing stress and anxiety that lower worker performance. They may also divert attention away from core tasks as employees discuss the implications of the transaction. On the other hand, workers might also want to exert high effort in order to keep their job or to improve the chance of

⁷The size of public commits in terms of lines of code added and deleted is available only at the repository level, not for commits by individual developers as used in the quantity analysis. In column two, we use the total number of lines of code added in the denominator, while in column 3 we use the maximum of the number of lines of code added and the number of lines of code deleted, i.e., the total number of lines worked on.

obtaining a new job, as in Abou El-Komboz and Goldbeck (2025). Finally, employees might just pretend to work hard, also with the intention of keeping employment.

Our findings are consistent with the “stress and anxiety” channel and with the “distraction” hypothesis. The stress and anxiety channel is further backed up by Bach et al. (2025), who show that takeovers in Sweden have a negative effect on workers’ mental health. Unfortunately, we are not able to clearly distinguish between the stress and distraction channels. However, our results suggest that workers’ effort signaling — whether honest or fake effort — is not the dominant mechanism.

In order to support the notion that the main mechanism behind these results has to do with the possibility of layoffs, we distinguish between takeovers that are likely motivated by cost-cutting measures (such as layoffs), and deals that are probably motivated by other reasons (acquiring new technology, new products, etc.).

To this end, we assume that takeovers where the target and the acquirer are in the same industry, it is more likely that layoffs will follow. Analogously, we assume that between-industry deals have a different motivation and are less likely to lead to layoffs. The same proxy is used in the empirical takeover literature, such as Dessaint et al. (2017) and Chen et al. (2021).⁸

We create a dummy variable *Within* that takes a value of one for within-industry takeovers. We estimate both OLS and Poisson regressions where we interact the *Within* dummy with a *Post announcement* dummy as well as with a dummy for weeks that are more than three weeks prior to the announcement week. We also interact *Within* with quarter-by-year-by-industry fixed effects.

The results are reported in Table 8. In the case of the OLS regression (column 1), the point estimate for the interaction term $Within \times Post\ announcement$ has the expected sign, but it is not statistically significant. However, when estimating a Poisson regression (column

⁸Another study is Ma, Ouimet, and Simintzi (2025). They examine a larger set of U.S. industries than our sample, but they only use horizontal mergers. They find that it is routine workers that are most likely to be laid off in the average horizontal takeover.

2), the coefficient becomes statistically significant at the 10% confidence level.

As a second mechanism test, we exploit variation in the IDD, which allows firms in certain U.S. states to block employees from joining competitors by arguing inevitable trade secret disclosure. There is a large literature on the effect of the IDD on firms and on workers, going back at least to Klasa, Ortiz-Molina, Serfling, and Srinivasan (2018). In particular, Chen et al. (2021) find that firms headquartered in an IDD state are significantly more likely to be acquired than companies in other states, because acquirers want to have access to the target firm’s employees. Without an acquisition, they could not hire critical workers from other firms due to the IDD. Based on this result, we hypothesize that target employees in IDD states are not particularly worried about a future layoff, whereas workers in non-IDD states are either stressed or distracted by the risk of a layoff. Therefore, we expect the negative effect on worker productivity to be stronger in non-IDD states than in IDD states.

Our GitHub data is particularly appropriate for the IDD analysis. Klasa et al. (2018) write that “the relevant jurisdiction for the application of the IDD is the state where the employee works.” Most studies in the IDD literature use headquarter location as a proxy for employee location. Our data, however, allow us to measure the location of each individual GitHub user in our sample, which is more precise than just headquarter location. We drop non-U.S. users and those without location data.

Using state IDD definitions from Driver, Kolasinski, and Stanfield (2025), Table 9 reports event-study estimates interacting the post-announcement dummy with an IDD indicator. Column (1) confirms the baseline effect on the U.S. subsample. The coefficient of the post takeover announcement dummy is still significant at the 5% level, even though the sample size shrinks from approximately 29,000 to 16,000. In column (2), we interact the post announcement dummy with the IDD dummy. The post announcement dummy is even more negative at -3.6 ($p < 1\%$), while the interaction with the IDD dummy is 4.7 ($p < 1\%$). In column (3), we add another term that interacts the post announcement dummy with the state’s non-compete agreement (NCA) enforcement dummy. This dummy takes the value

one if the state’s NCA score was above 3 in 2013, which is the median in our sample of users. We take state-level NCA scores from Golubov and Zhong (2025). After including the interaction term with NCA enforcement, the post announcement dummy becomes -4.4 ($p < 5\%$), while the interaction with the IDD dummy is 4.2 ($p < 1\%$). Consistent with our prior hypothesis, the negative effect of takeovers on work productivity is even larger in non-IDD states than in the average state. While in IDD states, this negative effect is offset by the large positive interaction term.

Results are robust to alternative IDD definitions from Klasa et al. (2018), Li and Tao (2024), and Na (2020) (Table A3) and to dropping top states (California, New York, and North Carolina) one at a time despite $\sim 35\%$ sample shrinkage (Table A4).

These results collectively suggest that our main findings have something to do with the possibility of future layoffs. This rules out a possible alternative explanation where our result is driven by startups that *want* to be acquired, as is anecdotally often the case in Silicon Valley. It also rules out the alternative explanation where software developers cut back on side projects and focus on work tasks following the announcement of a takeover.

Placebo test. As an additional test of our hypothesis that the results are driven by possible layoffs, rather than other changes induced by a takeover announcement, we conduct a placebo test exploiting partial acquisitions. The Capital IQ database allows us to identify partial acquisitions, defined as acquisitions of a minority stake, or an asset/branch. Since such a transaction is less likely to imply a change of control in the firm, its announcement should induce less anxiety about job security among employees at target firms.

Analogously to our main analysis, we identify GitHub users who report to work for firms that were targets in partial acquisitions in the year prior to the announcement of the transaction. When imposing the same activity restrictions for admission into the sample as for the main analysis, we obtain a sample of 283 users at 52 target firms. The average (median) number of weekly contributions during the last 12 weeks pre-announcement is

slightly lower than in the main sample with 13.6 (5) contributions.

We conduct standard and stacked event studies as well as DiD regressions. The results are presented in Figures 13 (non-stacked event study) and 14 (stacked event study). The aggregate DiD estimates are reported in Table 10.

In the simple event study, we see a temporary drop in public activity levels right after the takeover announcement. It is, however, much shorter lived than the effects we found for acquisitions of majority stakes, completely vanishing by week seven. Private contributions seem to be unaffected. The DiD estimates (Panel A of Table 10) are insignificant for both public and private, as well as for total contributions. The point estimates are also small, indicating that, on average, users contribute 0.17 fewer actions per week during the 12 weeks post-treatment relative to the last 3 weeks pre-announcement, which amounts to only 1.25% of the pre-announcement mean.

The absence of a significant effect is also confirmed in the stacked event study (Figure 14) and the stacked DiD estimates are again insignificant and close to zero (Panel B of Table 10).

Majority acquisitions are more likely to create a change in control at the target firm and, therefore, make it easier to impose large layoffs. Our results are stronger for majority acquisitions than for partial deals, which is consistent with the hypothesis that the reduction in productivity has to do with the fear and/or distraction associated with future layoffs.

4.5. *Effect magnitude*

Is this reduction in output economically meaningful? To assess this, we approximate the total loss in output across our sample, and the implied monetary loss for the users' employers. Our main specification indicates a reduction in weekly output of 1.66 contributions over a thirteen-week period, starting in the week of the takeover announcement. Based on the monetary value of a commit, which Holub and Thies (2023) quantify as \$112 in 2021, this output reduction translates into a monetary loss of \$2,417 per worker.⁹ Based on the DiD

⁹While Holub and Thies (2023) derive the value of a commit, whereas our estimates refer to *contributions*, the correlation between (public) contributions and (public) commits in our analysis sample is very high, at

estimate of -2 contributions per week for the longer 27 week period, the implied loss would amount to as much as \$6,048 per worker.

Across all 1,181 users we observe in our event study sample, the takeover announcements impose productivity losses valued at \$2.8m (short run) to \$7.1m (medium-run). However, this underestimates the aggregate effect, as we only observe a subset of the workers at each firm. Also, we only observe the subset of target firms that use GitHub for collaboration.

5. Conclusion

We find that after the announcement of an M&A deal, the work output of knowledge workers at the target firm declines. The data source is GitHub, which allows us to measure the output of software developers, data scientists, and similar occupations. Our results suggest that the reduction in productivity is a causal effect. Compared to the unconditional mean of output, we find a reduction of 9% in a conservative regression specification, and even larger effects in alternative specifications.


Although we cannot pin down the exact mechanism, the evidence suggests that it is related to the threat of future layoffs at the target firm. Our findings are consistent with the “stress and anxiety” hypothesis, but we cannot rule out other mechanisms. The stress and anxiety channel is further backed up by Bach et al. (2025), who show that takeovers in Sweden have a negative effect on workers’ mental health.

The results have implications for both firms and policy makers. For acquirers, they suggest that target firm valuations should incorporate the expected productivity costs associated with takeovers, potentially implying an optimal downward adjustment in bid prices. Our data does not allow us to determine to what extent bidders already internalize these costs.

More broadly, the findings indicate that the welfare effects of mergers and acquisitions extend beyond capital and product markets to include non-trivial productivity and mental health costs for affected workers. If the main channel is distraction, takeovers lead to tempo-

0.966.

rary losses in output; if stress and anxiety dominate, they generate both output losses and private disutility from reduced well-being. Either way, the cost-benefit analysis of M&As should take into account these additional costs.



Holub and Thies (2023 WP)

Fig. 1. Work on public repos mostly happens during work hours.




Fig. 2. Distribution of takeover announcements over time.




Fig. 3. Geographic distribution of users in the year prior to takeover.




Fig. 4. Number of contributions around takeover announcements. The figure shows the weekly average number of contributions to private and public repos in event time. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. We clean the number of contributions by subtracting user-by-Christmas fixed effects and holiday fixed effects. There is no control group in this analysis.




Fig. 5. Effect of takeover announcement on the number of total contributions. The figure shows an event study regression of the number of contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level.




Fig. 6. Effect of takeovers on the number of total contributions, stacked event study. The figure shows a stacked event study regression of the number of contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.




Fig. 7. Effect of takeover on the number of contributions to public vs. private repos. The figure shows a stacked event study regression of the number of contributions to private vs. public repos, respectively. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.




Fig. 8. Long-term effect of takeovers on the number of total contributions. The figure shows an event study regression of the number of contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. The event window is defined over 26 weeks before to 26 weeks after the announcement date.




Fig. 9. Effect of takeovers on the number of total contributions. This figure contains a robustness with a stacked event study Poisson regression. The outcome variable is the number of contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.




Fig. 10. Effect of takeovers on the number of total contributions. Robustness test with stacked event study regressions with different model weights. The left panel, labeled “Alternative Weights”, implements Wing et al. (2024) with corrective weights to obtain a consistent estimate of the trimmed aggregate ATT. The right panel, labeled “Cengiz et al. 2019”, uses no weights.




Fig. 11. Effect of takeovers on the number of total contributions without Red Hat. We drop users at Red Hat, the target firm with the largest number of users in our sample. The figure shows a stacked event study regression of the number of contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.




Fig. 12. Effect of takeovers on the number of total contributions. Subsample of takeovers where the announcement date is strictly before the completion date. The figure shows a stacked event study regression of the number of contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.




Fig. 13. Placebo Test: Effect of partial acquisitions on the number of weekly contributions. Partial acquisitions are defined as acquisitions of a minority stake or an asset or branch. The figure shows estimates from an event study regression of the number of contributions to private repos, contributions to public repos, and total contributions (to public and private repos). The event is defined as the week of announcement of a partial acquisition of the target company. The unit of observation is a user-week pair. Standard errors are clustered at the firm level.




Fig. 14. Placebo Test: Effect of partial acquisitions on the number of weekly contributions. Partial acquisitions are defined as acquisitions of a minority stake or an asset or branch. The figure shows estimates from a stacked event study regression of the number of contributions to private and public repos. The event is defined as the week of announcement of a partial acquisition of the target company. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.

<i>Panel A: Full Sample</i>				
	1st Quartile	Median	Mean (SD)	3rd Quartile
Total Contributions	0.0	5.0	17.90 (38.48)	21.0
Any Contribution	0.0	1.0	0.69 (0.46)	1.0
Public Contributions	0.0	0.0	6.54 (23.22)	4.0
Private Contributions	0.0	1.0	11.26 (26.17)	12.0
Public Commits	0.0	0.0	4.87 (24.90)	2.0
Developer × Week Observations		571,570		
Unique Developers		1235		
Target Firms		554		
<i>Panel B: Event Study Sample</i>				
	1st Quartile	Median	Mean (SD)	3rd Quartile
Total Contributions	0.0	6.0	18.40 (38.22)	22.0
Any Contribution	0.0	1.0	0.70 (0.46)	1.0
Public Contributions	0.0	0.0	6.62 (23.48)	4.0
Private Contributions	0.0	1.0	11.71 (26.30)	13.0
Public Commits	0.0	0.0	4.81 (25.41)	2.0
Developer × Week Observations		29,035		
Unique Developers		1181		
Target Firms		521		
<i>Panel C: Pre- and Post-Treatment Period</i>				
	Pre-Treatment	Post-Treatment	Total	
Total Contributions	19.63 (41.6)	17.26 (34.7)	18.40	
Any Contribution	0.71 (0.46)	0.68 (.46)	0.70	
Public Contributions	7.2 (26.3)	6.1 (20.4)	6.6	
Private Contributions	12.4 (27.8)	11.1 (24.8)	11.7	
Public Commits	5.4 (30.9)	4.3 (18.8)	4.8	
Developer × Week Observations		14,056	14,979	29,035
<i>Panel D: Stacked Event Study Sample</i>				
	Treatment Group	Control Group	Total	
Total Contributions	18.5 (38.6)	19.0 (37.4)	18.98	
Any Contribution	0.70 (0.45)	0.72 (0.46)	0.72	
Public Contributions	6.7 (23.7)	6.7 (19.5)	6.73	
Private Contributions	11.7 (26.5)	12.1 (27.0)	12.09	
Public Commits	4.7 (20.6)	4.7 (16.3)	4.71	
Developer × Week Observations		28,275	515,350	543,625

Table 1: Summary statistics of user activity. Panel A shows the full sample after merging the takeover data from Capital IQ, the GitHub user-by-year panel from GHTorrent, and the user activity data from the GitHub API. Panel B shows the sample used in the event study regressions. The reduction in sample size is due to the relatively short event window of 12 weeks before to 12 weeks after the takeover announcement week. Panel C splits the event study sample into pre- and post-treatment periods. Panel D shows the sample used in the stacked event study regressions. The reason for the increase in sample size is that the same user can serve as a control user for multiple treated users. Values in Panels C and D are means, with standard deviations in parentheses.

Dependent Variables: Model:	Total Contributions (1)	Public Contributions (2)	Private Contributions (3)
<i>Variables</i>			
Post Announcement	-1.7*** (0.62)	-0.89** (0.35)	-0.72* (0.39)
<i>Fixed-effects</i>			
xmas-developer	Yes	Yes	Yes
year-quarter(week)-industry	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	29,035	29,035	29,035
R ²	0.59574	0.60917	0.63035
Within R ²	0.00088	0.00035	0.00111

Clustered (target) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: Effect of takeover announcement on the number of contributions. The table shows event study regressions with three outcome variables: the number of total contributions to private and public repos, the number of contributions to public repos, and the number of contributions to private repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level.

Dependent Variables: Model:	Total Contr. (1)	Public Contr. (2)	Private Contr. (3)
<i>Variables</i>			
Post Announc. $\times$ Treatment Group	-2.59*** (0.957)	-0.868 (0.653)	-1.70*** (0.615)
<i>Fixed-effects</i>			
developer-stack	Yes	Yes	Yes
stack-event time	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	543,625	543,625	543,625
R ²	0.66415	0.59729	0.69000
Within R ²	0.00018	6.03×10^{-5}	0.00016
<i>Clustered (target) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Table 3: Effect of takeovers on the number of contributions, stacked event study. The table shows stacked event study regressions with three outcome variables: the number of total contributions to private and public repos, the number of contributions to public repos, and the number of contributions to private repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. We follow Wing et al. (2024) and use corrective weights to obtain a consistent estimate of the sample-share-weighted ATT.

Dependent Variable: Model:	Total Contributions	
	(1)	(2)
<i>Variables</i>		
Post Takeover Announcement	−2.0*** (0.59)	
> 10 Weeks Post Announcement		−1.9** (0.87)
0 to 10 Weeks Post Announcement		−2.0*** (0.59)
> 10 Weeks Pre Announcement		0.92 (0.56)
<i>Fixed-effects</i>		
xmas-developer	Yes	Yes
year-quarter(week)-industry	Yes	Yes
<i>Fit statistics</i>		
Observations	61,474	61,474
R ²	0.56788	0.56792
Within R ²	0.00090	0.00099
<i>Clustered (target) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Table 4: Long-term effect of takeovers on the number of total contributions. The table shows an event study regression of the number of total contributions to private and public repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level. The event window is defined over 26 weeks before to 26 weeks after the announcement date.

Dependent Variable: Model:	Total Contributions	
	(1) Alternative weights	(2) No weights
<i>Variables</i>		
Post Announcement $\times$ Treatment Group	-1.65^{***} (0.609)	-2.01^{***} (0.553)
<i>Fixed-effects</i>		
developer-stack	Yes	Yes
stack-event time	Yes	Yes
<i>Fit statistics</i>		
Observations	543,625	543,625
R ²	0.72192	0.65962
Within R ²	9.31×10^{-5}	7.84×10^{-5}
<i>Clustered (target) standard-errors in parentheses</i>		
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>		

Table 5: Stacked event study regressions with various regression weights. Column (1) presents the estimation with alternative regression weights from Wing et al. (2024) to obtain a consistent estimate of the trimmed aggregate ATT. Column (2) shows the results without weights, as in Cengiz et al. (2019). The outcome variable is the number of total contributions to private and public repos.

Dependent Variables: Model:	Total Contr. (1) OLS	Public Contr. (2) OLS	Private Contr. (3) OLS	Total Contr. (4) Poisson	Public Contr. (5) Poisson	Private Contr. (6) Poisson
<i>Variables</i>						
Post $\times$ Treatment Group	-3.1*** (1.0)	-0.77 (0.71)	-2.3*** (0.67)	-0.15*** (0.04)	-0.09 (0.08)	-0.16*** (0.04)
> 4 Weeks Pre $\times$ Treatment Group	-0.68 (1.1)	0.13 (0.57)	-0.78 (0.84)	-0.04 (0.05)	0.05 (0.09)	-0.06 (0.04)
<i>Fixed-effects</i>						
developer-stack	Yes	Yes	Yes	Yes	Yes	Yes
stack-event time	Yes	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>						
Observations	543,625	543,625	543,625	518,950	473,475	457,450
Squared Correlation	0.65925	0.59674	0.68518	0.67823	0.63465	0.69172
Pseudo R ²	0.10581	0.10065	0.12239	0.66585	0.68451	0.67059
BIC	5,262,580.1	4,660,196.7	4,862,692.2	7,548,957.8	4,165,575.0	5,378,445.2

Clustered (target) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 6: Stacked difference-in-differences regresssion, OLS & Poisson. Columns (1)–(3) contain estimates from OLS regressions, while columns (4)–(6) present the results of Poisson regressions. The outcome variables are the total number of contributions to public and private repos, the number of contributions to public repos, and the number of contributions to private repos, respectively.

	$\frac{\text{bugs}}{\text{commits}}$ (1)	$\frac{\text{bugs}}{\text{lines changed}}$ (2)	$\frac{\text{bugs}}{\text{lines added}}$ (3)	$\frac{\text{bugs}}{\text{commits}}$ (4)	$\frac{\text{bugs}}{\text{lines changed}}$ (5)	$\frac{\text{bugs}}{\text{lines added}}$ (6)
Post $\times$ Treated	0.024* (0.013)	0.008* (0.004)	0.008* (0.004)	0.021** (0.009)	0.008* (0.004)	0.008* (0.004)
3 Months Pre $\times$ Treated				-0.008 (0.018)	-.00001 (0.003)	0.0005 (0.004)
<i>Fixed-effects</i>						
repo-stack	Yes	Yes	Yes	Yes	Yes	Yes
event time-stack	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,517	6,517	6,517	6,517	6,517	6,517
Dep. Var Mean	0.046	0.005	0.006	0.046	0.005	0.006
<i>Clustered (target) standard-errors in parentheses</i>						
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>						

Table 7: Effect of takeovers on code quality. The table shows estimates from stacked difference-in-differences regressions for the bug rate at the repository-by-month level. The outcome variables are defined as the number of new bugs divided by a measure of code quantity, namely commits created in the current and previous months (Columns 1 and 4), new lines of code added in commits in the current and previous months (Columns 2 and 5), or total lines changed in commits in the current and previous months (Columns 3 and 6). All variables refer to public repositories only.

	Total Contributions	
	(1)	(2)
	OLS	Poisson
> 3 weeks pre-announcement	0.2554 (0.8614)	0.0009 (0.0463)
Post announcement	-0.9342 (0.6018)	-0.0564* (0.0304)
> 3 weeks pre-announcement $\times$ Within	-1.149 (1.329)	-0.0546 (0.0703)
Within $\times$ Post announcement	-2.047 (1.259)	-0.1119* (0.0652)
Observations	26,824	25,505
Developer fixed effects	✓	✓
Type-quarter-year-industry fixed effects	✓	✓

Table 8: Effect heterogeneity: within- vs. between-industry takeovers. Difference-in-differences regressions, estimated by OLS and Poisson.

Dependent Variable: Model:	Total Contributions		
	(1)	(2)	(3)
<i>Variables</i>			
> 3 weeks pre announcement	-0.37 (0.82)	-0.62 (1.1)	-1.3 (1.4)
Post Takeover announcement	-2.1** (1.0)	-3.6*** (1.4)	-4.4** (1.7)
IDD $\times$ > 3 weeks pre announcement		0.83 (1.4)	0.45 (1.4)
IDD $\times$ Post Takeover announcement		4.7*** (1.5)	4.2*** (1.5)
NCA $\times$ > 3 weeks pre announcement			1.6 (1.6)
NCA $\times$ Post Takeover announcement			2.0 (1.5)
<i>Fixed-effects</i>			
xmas-developer	Yes	Yes	Yes
year-quarter-industry	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	16,356	16,356	16,356
R ²	0.59663	0.59718	0.59724
Within R ²	0.00228	0.00363	0.00380
<i>Clustered (target) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Table 9: Takeovers in states applying the Inevitable Disclosure Doctrine (IDD) vs. other states according to the definition from Driver et al. (2025). Difference-in-differences regressions estimated by OLS. Column (1) is estimated on the subsample of users located in the U.S. and reporting a state level location. Column (2) adds an interaction term of an indicator taking the value one if the user’s state of residence applies the IDD and the post takeover announcement dummy. Column (3) adds an interaction term of an indicator taking the value one if the user’s state of residence enforces non-compete agreements (NCAs) strictly and the post takeover announcement dummy. The NCA dummy takes the value one for states where the enforcement score was above 3 in 2013, which is the median. The enforcement score is from Golubov and Zhong (2025).

Dependent Variables:	Total Contributions	Public Contributions	Private Contributions
Panel A. Simple Diff-in-Diff			
Post Announcement	-0.17 (1.1)	-0.28 (0.72)	-0.21 (0.49)
> 3 Weeks Pre Announcement	0.96 (1.5)	0.25 (1.1)	0.71 (1.0)
<i>Fixed-effects</i>			
xmas-developer	Yes	Yes	Yes
year-quarter-industry	Yes	Yes	Yes
Observations	18,318	18,318	18,318
R ²	0.55	0.507	0.63
Panel B. Stacked Diff-in-Diff			
Post $\times$ Treatment Group	-0.19 (1.8)	0.66 (0.74)	-0.91 (1.8)
> 3 Weeks Pre $\times$ Treatment Group	-0.10 (2.0)	-0.22 (0.50)	0.10 (2.1)
<i>Fixed-effects</i>			
developer-stack	Yes	Yes	Yes
stack-event time	Yes	Yes	Yes
Observations	26,525	26,525	26,525
R ²	0.62	0.596	0.61
Dep. Var. Mean (Pre-Announc.)	13.6	3.5	10.05
<i>Clustered (target) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Table 10: Placebo Test: Effect of partial acquisitions on the number of weekly contributions. Panel A depicts estimates from non-stacked DiD estimations, Panel B shows results from stacked DiD estimations following Wing et al. (2024) with corrective weights to obtain a consistent estimate of the sample-share-weighted ATT. Partial acquisitions are defined as acquisitions of a minority stake or an asset or branch. The table shows estimation results for three outcome variables: the number of total contributions to private and public repos, the number of contributions to public repos, and the number of contributions to private repos. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. Standard errors are clustered at the firm level.

Appendix A. Additional results




Fig. A1. Number of contributions to private repos around takeover announcements. The figure shows the weekly average number of contributions to private repos in event time. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. We clean the number of contributions by subtracting user-by-Christmas fixed effects and holiday fixed effects. There is no control group in this analysis.




Fig. A2. Number of contributions to public repos around takeover announcements. The figure shows the weekly average number of contributions to public repos in event time. The event is defined as the announcement week of a takeover. The unit of observation is a user-week pair. We clean the number of contributions by subtracting user-by-Christmas fixed effects and holiday fixed effects. There is no control group in this analysis.




Fig. A3. Effect of takeovers on the number of public contributions. This figure contains a robustness with an event study based on data on GitHub activities in public repos from GHTorrent.




Fig. A4. Long-run Effect of takeovers on the number of public contributions. This figure contains a robustness with an event study based on data on GitHub activities in public repos from GHTorrent.

	$\frac{\text{bugs}}{\text{commits}}$ (1)	$\frac{\text{bugs}}{\text{lines changed}}$ (2)	$\frac{\text{bugs}}{\text{lines added}}$ (3)	$\frac{\text{bugs}}{\text{commits}}$ (4)	$\frac{\text{bugs}}{\text{lines changed}}$ (5)	$\frac{\text{bugs}}{\text{lines added}}$ (6)
Post $\times$ Treated	0.007 (0.009)	0.004** (0.002)	0.004** (0.002)	0.012* (0.007)	0.005** (0.002)	0.005** (0.002)
<i>Fixed-effects</i>						
repo	Yes	Yes	Yes	Yes	Yes	Yes
event time-stack	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6,517	6,517	6,517	6,517	6,517	6,517
<i>Clustered (target) standard-errors in parentheses</i>						
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>						

Table A1: Effect of takeovers on code quality. The table shows estimates from stacked difference-in-differences regressions for the bug rate at the repository-by-month level. The outcome variables are defined as the number of new bugs divided by a measure of code quantity, namely commits created in the current and previous months (Columns 1 and 4), new lines of code added in commits in the current and previous months (Columns 2 and 5), or total lines changed in commits in the current and previous months (Columns 3 and 6). All variables refer to public repositories only. Columns (1)–(3) show the results with alternative regression weights from Wing et al. (2024) to obtain a consistent estimate of the trimmed aggregate ATT, while Columns (4)–(6) show results from regressions without weights, as in Cengiz et al. (2019).

Dependent Variable:	Commits (Repo Level)		
	(1)	(2)	(3)
<i>Variables</i>			
3 Months Pre Announcement $\times$ Treated	-9.73 (6.45)	1.41 (4.60)	1.08 (3.72)
Treated $\times$ Post Announcement	-15.7** (6.93)	-5.77* (3.41)	-7.17** (3.34)
<i>Fixed-effects</i>			
repo-stack	Yes	Yes	Yes
event time-stack	Yes	Yes	Yes
Observations	6,517	6,517	6,517
Mean Dep. Var.	46.5	46.5	46.5
<i>Clustered (target) standard-errors in parentheses</i>			
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>			

Table A2: Effect of takeovers on the number of commits in the sample used in the code quality estimations. The table shows estimates from stacked difference-in-differences regressions for the number of commits at the repository-by-month level. The estimation is done on the sample used in the regressions for the bug rate. All variables refer to public repositories only. Column (1) shows the results with the corrective regression weights from Wing et al. (2024) to obtain a consistent estimate of the sample-share-weighted ATT, Column (2) shows the results with alternative regression weights from Wing et al. (2024) to obtain a consistent estimate of the trimmed aggregate ATT, while Columns (3) shows the estimate from a regression without weights, as in Cengiz et al. (2019).

Dependent Variable: Model:	(1)	(2)	Total Contributions		
			(3)	(4)	(5)
<i>Variables</i>					
> 3 weeks pre announcement	0.46 (0.39)	0.09 (0.45)	0.0004 (0.48)	0.37 (0.43)	0.43 (0.41)
Post Takeover announcement	-1.3*** (0.37)	-1.7*** (0.44)	-1.9*** (0.46)	-1.6*** (0.41)	-1.4*** (0.39)
> 3 weeks pre announcement $\times$ IDD		1.1 (0.74)	1.2* (0.69)	0.62 (0.82)	0.33 (0.94)
IDD $\times$ Post Takeover announcement		1.3** (0.64)	1.4** (0.63)	1.9** (0.75)	1.2 (0.88)
IDD:	x	Driver et al. (2025)	Klasa et al. (2018)	Fung et al. (2025)	Na (2020)
<i>Fixed-effects</i>					
xmas-developer	Yes	Yes	Yes	Yes	Yes
year-quarter-industry	Yes	Yes	Yes	Yes	Yes
<i>Fit statistics</i>					
Observations	135,687	135,687	135,687	135,687	135,687
R ²	0.55948	0.55952	0.55952	0.55951	0.55950
Within R ²	0.00196	0.00204	0.00204	0.00204	0.00199

Clustered (target) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table A3: Effect heterogeneity by IDD adoption: Robustness tests. Difference-in-differences regressions estimated by OLS. Column (2) repeats the results from column (3) in Table 9. Columns (3)–(5) use alternative IDD state classifications, as indicated in the bottom row.

Dependent Variable: Model:	Total Contributions			
	(1)	(2)	(3)	(4)
<i>Variables</i>				
> 3 weeks pre announcement	0.09 (0.45)	-0.06 (0.51)	0.007 (0.46)	0.07 (0.47)
Post Takeover announcement	-1.7*** (0.44)	-1.2** (0.50)	-1.6*** (0.43)	-1.8*** (0.46)
> 3 weeks pre announcement $\times$ IDD	1.1 (0.74)	1.2 (0.78)	1.0 (0.72)	1.2* (0.74)
IDD $\times$ Post Takeover announcement	1.3** (0.64)	0.83 (0.70)	1.9*** (0.68)	1.3** (0.67)
State dropped:	none	CA	NY	MA
<i>Fixed-effects</i>				
xmas-developer	Yes	Yes	Yes	Yes
year-quarter-industry	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	135,687	91,011	119,499	127,709
R ²	0.55952	0.55009	0.58258	0.56379
Within R ²	0.00204	0.00225	0.00177	0.00208
<i>Clustered (target) standard-errors in parentheses</i>				
<i>Signif. Codes: ***: 0.01, **: 0.05, *: 0.1</i>				

Table A4: Effect heterogeneity by IDD adoption: Robustness tests. Difference-in-differences regressions estimated by OLS. Column (1) repeats the results from column (3) in Table 9. In columns (2)–(4), top states are dropped one at a time from the sample, as indicated in the bottom row.

Appendix B. Comparison of M&A deal coverage in Capital IQ and Refinitiv

In Capital IQ, the dataset contains approximately 24,300 M&A deals announced in 2016. To assess the relative coverage of Capital IQ against the LSEG Refinitiv (formerly SDC Platinum) database, we apply the following filters in Refinitiv:

- Universe: Deals
- Asset Class: M&A
- Announcement Date: January 1 – December 31, 2016
- Nation of Headquarters: United States and Canada

Applying these filters yields a sample of 14,330 deals. When we further restrict the sample to transactions with a Deal Status of Withdrawn, Completed, Tentative, or Pending, only two additional observations are excluded.

For comparability, the filters used in Capital IQ are as follows:

- Transaction Type: M&A
- Private Equity Deals: Included (*no direct analogue identified in Refinitiv*)
- Geographic Scope: United States and Canada
- Announcement Date Range: January 1 – December 31, 2016
- Transaction Status: Announced, Completed, or Terminated/Withdrawn

Repeating the same comparison for 2018 yields 22,900 deals in Capital IQ and 16,523 deals in Refinitiv. Overall, Capital IQ demonstrates broader coverage of M&A transactions. Therefore, for the purposes of this analysis, we rely on Capital IQ as the primary data source for M&A deal information.

References

- About El-Komboz, L., Goldbeck, M., 2025. Career concerns as a public good: The role of signaling for open source software development. *Labour Economics* 97, 102800.
- Andres, C., Bazhutov, D., Cumming, D., Köchling, G., Limbach, P., 2025. Does speculative news hurt productivity? evidence from takeover rumors. *Journal of Financial and Quantitative Analysis* 60, 2952–2996.
- Arnold, D., 2025. Mergers and acquisitions, local labor market concentration, and worker outcomes. Tech. rep., SSRN Working Paper.
- Bach, L., Bos, M., Silva, R., 2025. How do mergers affect the mental health of employees? *Management Science* Forthcoming.
- Baker, A. C., Larcker, D. F., Wang, C. C., 2022. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics* 144, 370–395.
- Callaway, B., Sant’Anna, P. H. C., 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics* 225, 200–230.
- Cengiz, D., Dube, A., Lindner, A., Zipperer, B., 2019. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics* 134, 1405–1454.
- Chacon, S., Straub, B., 2014. *Pro Git*. Apress, second ed.
- Chen, D., Gao, H., Ma, Y., 2021. Human Capital-Driven Acquisition: Evidence from the Inevitable Disclosure Doctrine. *Management Science* 67, 4643–4664.
- Chen, D., He, Y., Yamashita, T., 2025. Unraveling and judge productivity in the market for federal judicial law clerks: Evidence and proposal. Working paper, National Bureau of Economic Research.

- Cui, R., Ding, H., Zhu, F., 2022. Gender Inequality in Research Productivity During the COVID-19 Pandemic. *Manufacturing & Service Operations Management* 24, 707–726.
- Dessaint, O., Golubov, A., Volpin, P., 2017. Employment protection and takeovers. *Journal of Financial Economics* 125, 369–388.
- Driver, J., Kolasinski, A. C., Stanfield, J. R., 2025. Trade secrets, labor mobility, and innovation. Working paper, SSRN.
- Eckbo, B. E., Malenko, A., Thorburn, K., 2025. Corporate takeovers: Theory and evidence. Working paper, SSRN.
- Gehrke, B., Maug, E., Obernberger, S., Schneider, C., 2021. Post-merger Restructuring of the Labor Force. <https://ssrn.com/abstract=3838865>, SSRN Working Paper.
- GitHub, 2016. The state of the octoverse: 2016 report. <https://octoverse.github.com/2016>, gitHub, Inc.
- GitHub, 2022. The state of the octoverse: 2022 report. <https://octoverse.github.com/2022>, gitHub, Inc.
- Golubov, A., 2025. Beyond the shareholders: The impact of M&A on other stakeholders. Working paper, SSRN.
- Golubov, A., Xiong, N., 2020. Post-acquisition performance of private acquirers. *Journal of Corporate Finance* 60, 101545.
- Golubov, A., Zhong, Y., 2025. Non-Compete Agreements and the Market for Corporate Control. Working Paper 978/2024, European Corporate Governance Institute (ECGI) – Finance Working Paper No. 978/2024; HKU Jockey Club Enterprise Sustainability Global Research Institute, available at SSRN.
- Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225, 254–277.

- Gormley, T. A., Matsa, D. A., 2011. Growing Out of Trouble? Corporate Responses to Liability Risk. *The Review of Financial Studies* 24, 2781–2821.
- Gote, C., Mavrodiev, P., Schweitzer, F., Scholtes, I., 2022. Big data= big insights? operationalising brooks’ law in a massive github data set. In: *Proceedings of the 44th International Conference on Software Engineering*, pp. 262–273.
- Gousios, G., Spinellis, D., 2012. Ghtorrent: Github’s data from a firehose. In: *2012 9th IEEE Working Conference on Mining Software Repositories (MSR)*, IEEE, pp. 12–21.
- Grajzl, P., Silwal, S., 2020. Multi-court judging and judicial productivity in a career judiciary: Evidence from nepal. *International Review of Law and Economics* 61, 105888.
- Gugler, K., Yurtoglu, B., 2004. The effects of mergers on company employment in the USA and Europe. *International Journal of Industrial Organization* 22, 481–502.
- Gupta, A., Nishesh, N., Simintzi, E., 2024. Big data and bigger firms: A labor market channel. SSRN Working Paper, available at SSRN: <https://ssrn.com/abstract=4994837> or <http://dx.doi.org/10.2139/ssrn.4994837>.
- Holub, F., Thies, B., 2023. Air Quality, High-Skilled Worker Productivity And Adaptation: Evidence From Github. Discussion paper series, University of Bonn and University of Mannheim, Germany.
- Hosken, D., Larson-Koester, M., Taragin, C., 2024. Labor and product market effects on mergers. Research report, Federal Trade Commission.
- Klasa, S., Ortiz-Molina, H., Serfling, M., Srinivasan, S., 2018. Protection of trade secrets and capital structure decisions. *Journal of financial economics* 128, 266–286.
- Lagaras, S., 2020. M&As, Employee Costs and Labor Reallocation. Tech. rep., SSRN Working Paper.

- Li, J., Tao, J., 2024. Does Knowledge Protection Spur Common Ownership? Evidence from the Inevitable Disclosure Doctrine. Working paper, SSRN.
- Li, X., 2013. Productivity, restructuring, and the gains from takeovers. *Journal of Financial Economics* 109, 250–271.
- Ma, W., Ouimet, P., Simintzi, E., 2025. Mergers and acquisitions, technological change and inequality. *Journal of Financial Economics* Forthcoming.
- Maksimovic, V., Phillips, G., Prabhala, N., 2011. Post-merger restructuring and the boundaries of the firm. *Journal of Financial Economics* 102, 317–343.
- Maug, E. G., 2025. Financing human capital: A survey and synthesis. Tech. rep., ECGI Working Paper.
- Na, K., 2020. CEOs’ outside opportunities and relative performance evaluation: Evidence from a natural experiment. *Journal of Financial Economics* 137, 679–700.
- Prager, E., Schmitt, M., 2021. Employer consolidation and wages: Evidence from hospitals. *American Economic Review* 111, 397–427.
- Rivkin, S. G., Hanushek, E. A., Kain, J. F., 2005. Teachers, schools, and academic achievement. *Econometrica* 73, 417–458.
- Siegel, D. S., Simons, K. L., 2010. Assessing the effects of mergers and acquisitions on firm performance, plant productivity, and workers: new evidence from matched employer-employee data. *Strategic Management Journal* 31, 903–916.
- Silva, J. M. C. S., Tenreyro, S., 2006. The log of gravity. *The Review of Economics and Statistics* 88, 641–658.
- Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* 225, 175–199.

- Toomey, J., 2025. Global M&A by the Numbers: 2024 in Review. Research report, S&P Global Market Intelligence.
- Trajtenberg, M., 1990. A penny for your quotes: Patent citations and the value of innovations. *The RAND Journal of Economics* 21, 172–187.
- Wing, C., Freedman, S. M., Hollingsworth, A., 2024. Stacked difference-in-differences. Tech. rep., National Bureau of Economic Research.
- Wooldridge, J. M., 2010. *Econometric Analysis of Cross Section and Panel Data*. MIT Press, second ed.
- Yang, X., Cai, X. L., Li, T. S., 2024. Does the tenure track influence academic research? An empirical study of faculty members in China. *Studies in Higher Education* 49, 476–492.