

The Effects of Generative AI on High Skilled Work: Evidence from Three Field Experiments with Software Developers*

Kevin Zheyuan Cui, Mert Demirer, Sonia Jaffe,
Leon Musolff, Sida Peng, and Tobias Salz

September 2024

Abstract

This study evaluates the impact of generative AI on software developer productivity by analyzing data from three randomized controlled trials conducted at Microsoft, Accenture, and an anonymous Fortune 100 electronics manufacturing company. These field experiments, which were run by the companies as part of their ordinary course of business, provided a randomly selected subset of developers with access to GitHub Copilot, an AI-based coding assistant that suggests intelligent code completions. Though each separate experiment is noisy, combined across all three experiments and 4,867 software developers, our analysis reveals a 26.08% increase (SE: 10.3%) in the number of completed tasks among developers using the AI tool. Notably, less experienced developers showed higher adoption rates and greater productivity gains.

*Cui: Princeton. Demirer: MIT. Jaffe: Microsoft. Musolff: The Wharton School of the University of Pennsylvania. Peng: Microsoft. Salz: MIT. We are grateful to Avi Goldfarb, Shane Greenstein, Anton Korinek, Ethan Mollick, as well as the participants of the IIOC2024 conference and the AI, Cognition, and the Economy Workshop. Thanks also for data help from employees at Microsoft, GitHub, and Accenture: Phillip Coppney, Wen Du, Ya Gao, Lizzie Redford, Ryan J. Salva, Daniel A. Schocke, Amanda Silver, An-Jen Tai, Dan Tetrack, Jeff Wilcox.

1 Introduction

Many economists expect generative AI to profoundly affect the organization of economic activity (Agrawal, Gans, and Goldfarb 2019; Frank et al. 2019; Furman and Seamans 2019). Eloundou et al. 2023 estimate that generative AI can perform tasks associated with over 80% of U.S. jobs and that AI task coverage is notably higher for occupations that require advanced degrees. The ability of generative AI to perform tasks required in such high-skilled occupations — allowing it to assist doctors in diagnosing diseases, lawyers in drafting legal documents, and software engineers with code development — has led to predictions of substantial productivity gains from the adoption of such technologies (Baily, Brynjolfsson, and Korinek 2023). Others, however, are less optimistic about such productivity gains (Acemoglu 2024).

Uncertainty around firms’ willingness to adopt these technologies and their capacity to make necessary complementary investments (Bresnahan 2024; Brynjolfsson, Rock, and Syverson 2021) make it currently difficult to empirically assess whether or not optimism about productivity gains is justified.¹ Nevertheless, some applications of generative AI have already matured and are integrated into existing workflows. An example is software development, where commercial coding assistants based on generative AI have gained widespread adoption.²

In this project, we ask how generative AI affects the productivity of knowledge workers, using software developers as an example. We analyze three large-scale randomized controlled trials in a real-world environment. These experiments randomly assigned access to Copilot, a coding assistant developed by GitHub in collaboration with OpenAI, to just under five thousand software developers at Microsoft, Accenture, and an anonymous Fortune 100 electronics manufacturing company (henceforth Anonymous Company). These experiments were run as part of the ordinary course of business at these companies to decide whether or how extensively to adopt these technologies, and the companies kindly shared the resulting data with us.³

Our preferred estimates suggest that usage of the coding assistant causes a 26.08% (SE: 10.3%) increase in the weekly number of completed tasks. When we look at outcomes of secondary interest, our results support this interpretation, with a 13.55% (SE: 10.0%) increase in the number of code updates (commits) and a 38.38% (SE: 12.55%) increase in the number of times code was compiled. For Microsoft we observe both the developers’ tenure and their seniority as measured by job title. We find that Copilot significantly raises task completion for more recent hires and those in more junior positions but not for developers with longer tenure and in more senior positions. Prior work has shown that

¹In addition, it hard to predict further breakthroughs in the architecture of these models, which may lead to further improvements in quality or decrease the cost of inference and training.

²Prior academic work has shown that generative AI can pass mock interviews for coding jobs at Amazon in the top decile of human performance, performs at human level in a database of coding challenges that measure programming logic and proficiency, and can write entire programs for simple video games from several lines of instructions (Bubeck et al. 2023) Copilot is used by 1.3 million subscribers and more than 50,000 businesses.

³The exact implementation of these experiments was rather ad-hoc as they were driven by business considerations at these companies rather than research goals.

when workers are conducting the same tasks, generative AI helps lower-ability or lower-experience workers more (e.g., Brynjolfsson, Li, and Raymond 2023; Noy and Zhang 2023). Our results extend this finding by showing that even when workers are performing tasks according to their tenure or seniority, generative AI increases productivity more for lower-ability workers.

Our preferred estimate pools estimates across all three experiments and places more weight on periods with larger differences in treatment status. We make these choices because our analysis must confront challenges related to statistical power, despite the large number of software developers in the experiments. These challenges arise due to large variation in measured outcomes and factors that reduce the take-up and duration of the three experiments.⁴ The experiment at Microsoft started before ChatGPT and Copilot were widely known, and initial uptake was low. Shortly after a larger fraction of developers in the treatment group started using it, the control group was also allowed access. At Accenture, only a few hundred people participated in the experiment. Lastly, at Anonymous Company, the treatment consisted of a staggered roll-out with only a short period of time with differences in treatment status.

Most studies of the impact of generative AI tools on worker productivity have been conducted in controlled lab-like experiments (Peng et al. 2023; Vaithilingam, Zhang, and Glassman 2022, Campero et al. 2022, Noy and Zhang 2023). In a lab-in-the-field experiment on consultants employed by Boston Consulting Group, Dell’Acqua et al. 2023 find that productivity on 18 tasks designed to mimic the day-to-day work at a consulting company increased by 12%-25%. Evidence from these experiments generally suggests significant productivity effects of generative AI. The exception is Vaithilingam, Zhang, and Glassman 2022, which did not find a statistically significant difference in completion time. A second common observation from these studies is that generative AI has the largest effect on the least productive and least experienced workers (e.g., Noy and Zhang 2023).

While lab experiments offer a valuable opportunity to examine the short-term implications of generative AI, challenges and complex interactions arise when these tools are deployed in real-world environments (Jaffe et al. 2024). There are some observational studies of the effects of generative AI in an actual workplace setting (Hoffmann et al. 2024; Yeverechyahu, Mayya, and Oestreicher-Singer 2024) that do not have the benefit of random experimental assignment of these technologies.

Our work complements both the literature on lab experiments as well as these observational studies by studying the impact of generative AI using a field experiment in an actual workplace setting. To date, there is still a dearth of experimental studies examining the effect of generative AI in a field setting. In a notable exception, Brynjolfsson, Li, and Raymond 2023 find that an AI-based conversational assistant increases the productivity of customer chat support agents by 14%. Our study complements theirs by examining a field experiment with high-skilled and highly paid knowledge workers, a group that is particularly relevant given the prediction that high-skilled jobs will be most affected by this technology. Although we examine a different part of the skill distribu-

⁴We observe large variation in the output of software developers due to significant heterogeneity in their seniority, with more senior managers being less likely to engage in coding activities.

tion, we find similar productivity increases. Like Brynjolfsson, Li, and Raymond 2023, we also find that these gains are primarily driven by improved output from recent hires and employees in more junior roles.

2 Setting and Experiments

2.1 What Is AI-Assisted Software Development?

AI Assistants for software development offer intelligent code suggestions and auto-completion within integrated development environments. Prominent examples include GitHub Copilot, Amazon CodeWhisperer, and Replit Ghostwriter. In our study, we examine the effects of one of these tools, GitHub Copilot. GitHub Copilot was developed by GitHub in partnership with OpenAI. The development of Copilot involved combining advanced machine learning techniques and natural language processing. A substantial amount of code from public GitHub repositories was used to train Copilot. This extensive dataset allowed the AI model to learn from real-world coding practices, patterns, and styles across various programming languages and frameworks.

As developers write software code or plain text comments, Copilot analyzes the context and generates relevant code snippets, comments, and documentation. It can auto-complete code that developers might manually type or suggest snippets they would otherwise need to search for online. This capability can save developers time and potentially improve code quality by offering suggestions the developer might not be aware of. However, like all LLM-based tools, Copilot can make mistakes. If developers rely on it without review, it could potentially introduce errors or decrease code quality.

2.2 Experiments

We analyze three randomized experiments conducted with software developers at Microsoft, Accenture, and Anonymous Company. In the Microsoft and Accenture experiments, one group of developers (the treated group) was randomly assigned to be able to access GitHub Copilot, whereas the other group (the control group) did not have access to the tool for a period of seven (Microsoft) or four (Accenture) months. In the Anonymous Company experiment, all users gained access to the tool over a period of two months, but access dates were randomized, with some teams gaining access six weeks before others.

Microsoft The experiment at Microsoft started in the first week of September 2022, involving a sample size of 1,746 developers primarily located in the United States. Of these developers, 50.4% were randomly selected to receive access to Github Copilot. Randomization was implemented at both the individual and the team levels. The developers work on building a wide range of software within Microsoft, with tasks that include engineering, designing, and testing software products and services. They occupy various positions in the company, ranging from entry-level developers to team managers. They may work in a team or individually, depending on their task and team structure.

Participants in the treated group were informed via email about the opportunity to sign up for GitHub Copilot. The email also included information introducing GitHub Copilot as a productivity-enhancing tool and outlining its potential impact on their coding tasks (see Figure 6 in Appendix). Beyond this email, treated participants did not receive any specific instructions regarding their workload or workflow to ensure they use GitHub Copilot in their natural work environment. Control group participants did not receive any communication as part of the study.⁵ The experiment ended on May 3rd, 2023, as growing awareness of AI-assisted coding tools led the control group participants to seek access to Copilot.

Accenture The Accenture experiment started in the last week of July 2023 and included a number of Accenture offices located in Southeast Asia. Randomization occurred at the developer level, with 61.3% of the 320 developers assigned to the treatment group. Treatment group participants were informed over email that they were eligible to sign up for GitHub Copilot. They also participated in a training session, which explained what GitHub Copilot is, how to use it, and the potential benefits. Finally, the participating managers were asked to encourage their reports' adoption of GitHub Copilot. The control group was granted access to Copilot in December 2023, though uptake was lower than in the treatment group.

Anonymous Company The Anonymous Company experiment started in October 2023. It involved 3,054 developers who were all eventually invited to use Copilot. The invitation dates were randomized, with new invites being sent out weekly between September 2023 and October 2023.

2.3 Variables and Outcome Measures

Measuring the productivity of modern knowledge work is notoriously difficult. Our setting has the advantage that almost all professional software development follows a highly structured workflow, where specific tasks are defined and tracked through version control software. Internally defined goals and tasks are, therefore, quantifiable. All three participating organizations use the version control software GitHub. By observing the developers' GitHub activity, we can observe many of the variables that are part of their workflow.

A main outcome of interest is "pull requests", which can be thought of as a unit of work for software developers. Within an organization, the scope of a pull request is likely to remain relatively stable over time, shaped by organizational norms and conventions, even though different organizations may define this scope differently. For instance, a pull request may be asking for a feature to be added to a larger software project. A pull request will lead to a code review, often by a more senior software developer. If this review is passed, the code will be merged and thereby become part of the larger software project.

We use three additional outcome variables related to the developers' workflow. Before submitting a pull request, a developer will work separately on her code, tracking smaller

⁵A small number of developers in the control group nevertheless got access to Copilot because they were working on related tools.

changes through “commits.” Periodically, the developer will “build” the code they are working on, and we can observe whether it compiled successfully. While commits and builds do not directly correspond to a deliverable, we expect them to be nevertheless monotone in the amount of accomplished work. Lastly, we use the build success rate as a measure of code quality.

In addition to these output measures, we observe how developers use Copilot. For each developer who uses Copilot, we observe both the number of suggestions by Copilot and the number of accepted suggestions.

For the Microsoft experiment only, we also see the hire date of developers and their level at the company, allowing us to separate the analysis by tenure and seniority.

Table 1 shows summaries of the treatment and control groups across all three experiments, as well as balance tests for the outcome variables. With the exception of commits in the Accenture experiment, randomization successfully balanced the average pre-treatment outcomes across the control and treatment groups. However, the table also shows that for all outcomes (with the exception of the Build Success Rate), the standard deviation exceeds the pre-treatment mean, and sometimes by a lot. This high variability will limit our power in our experimental regressions below.

3 Adoption of Copilot

This section reports the adoption of Copilot in the experiments. Understanding adoption patterns is important for assessing the effectiveness of the experiments in generating random variation in Copilot usage. Furthermore, these patterns offer insights into the adoption of AI tools in the workplace. We define the adoption period as the first time a software engineer uses GitHub Copilot and consider an engineer as having adopted the tool even if they later stop using it. This approach captures the initial willingness to try and use the technology.

Figure 1 presents the cumulative adoption rates for the three experiments. In Panel (a), we observe that during the first two weeks of the Microsoft experiment, only 8.5 percent of the treated group signed up for GitHub Copilot. This low adoption rate might have been due to inattentiveness to the initial email notification. Consequently, Microsoft sent two additional email reminders on Feb 15th, 2023, and Feb 28th, 2023. These additional emails increased the take-up rate to 42.5% within two weeks of being sent. The initial compliance in the control group was not perfect, with 0.5 percent of individuals in the control group adopting Copilot. At the conclusion of the experiment in April 2023, the control group was given access to Copilot, and we observed a rapid adoption rate in the control group. However, by Jan 2024, adoption in the control group (75.6%) remained below the adoption in the treated group (64.0%), providing limited long-run variation in adoption generated by the experiment.

Panel (b) reports the adoption rate of Copilot in the Accenture experiment. In this experiment, the treated group’s adoption rate was rapid at the start of the experiment, but after 1-2 months, it plateaus at around 60%. When the control group is allowed to adopt in December 2023, we see a slower but steady increase in adoption rate amongst control users. By April 2024, the adoption rate for the treated group is 69.4%, while the

	Control		Treatment		Difference	p-value
	Mean	SD	Mean	SD		
Panel A: Microsoft						
Pull Requests	0.86	1.49	0.87	1.50	0.01	0.88
Commits	9.43	14.86	9.36	14.80	-0.07	0.94
Builds	7.76	12.99	7.67	12.73	-0.09	0.91
Build Success Rate	0.72	0.30	0.75	0.29	0.02	0.33
Short Tenure	0.48	0.50	0.52	0.50	0.04	0.23
Junior Level	0.55	0.50	0.61	0.49	0.06	0.03**
Panel B: Accenture						
Pull Requests	0.13	0.47	0.14	0.47	0.00	0.85
Commits	2.56	6.00	3.64	7.25	1.08	0.01**
Builds	0.96	2.54	1.10	2.68	0.14	0.38
Build Success Rate	0.51	0.37	0.54	0.38	0.03	0.40
Panel C: Anonymous						
Pull Requests	0.73	1.23	0.73	1.19	-0.00	0.99

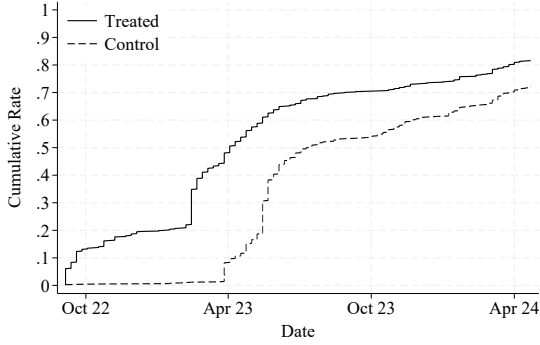
Table 1: Balance Table

Notes: This table presents a comparison of pre-experimental outcomes in control and treatment groups across experiments. For each measure, we present its mean and standard deviation in the control group and in the treatment group. We also show the mean difference across these groups and the p-value associated with an underlying test of a difference in means. We do not present other outcome measures in Panel C because we do not have access to these data. The differences in p-values are calculated using standard errors clustered at the level of treatment assignment, which varies across experiments (Microsoft: mixed team-level and individual assignment; Accenture: individual assignment; Anonymous Company: team-level assignment.).


adoption rate for the control group is 24.4%.

Panel (c) shows the staggered invitation to access Copilot (represented by the solid line) and the cumulative adoption rate among all participants (dashed line). As evident from the plot, all developers gained access to Copilot after six weeks. During the invitation rollout in September and October, we observed a steady increase in adoption as developers got access to Copilot. Following the initial increase in adoption, the adoption rate plateaued, showing only small, steady increases for the remainder of the experiment.

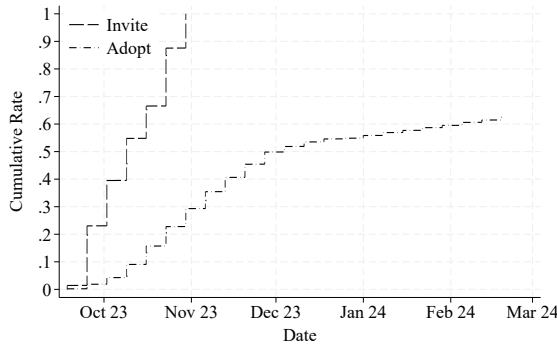
It is worth noting that adopting Copilot is relatively easier and less costly compared to other AI tools in the workplace. Copilot does not require any complementary investment, can be adopted at the individual level, and is already integrated into the software development environment. However, despite these advantages, the adoption rate is significantly below 100% in all three experiments, with around 30-40% of the engineers not even trying the product. Moreover, we observe that the adoption rates are remarkably similar across the experiments. This suggests that factors other than access, such as individual preferences and perceived utility of the tool, play important roles in engineers' decisions to use this tool.



(a) Microsoft Experiment



(b) Accenture Experiment



(c) Anonymous Company Experiment

Figure 1: Cumulative Adoption Rates

Notes: The first two graphs show the cumulative rate of adoption over time for software developers in both the treatment and control groups across various experiments. In the Anonymous Company experiment (Panel c), unlike the other experiments, all developers were granted access to Copilot in a staggered fashion, with the order of access randomized among participants. Hence, we show the cumulative fraction of users invited to participate and who adopted Copilot.

4 Empirical Strategy & Main Results

4.1 Empirical Strategy

Our empirical strategy exploits the experimental variation while accommodating imperfect compliance by utilizing the experimental assignment as an instrument for the adoption of GitHub Copilot. For each experiment, we observe data at the developer-week level. To gain precision, we control for both developer and week fixed effects (to account for, e.g., differences in developer skills and holidays). This leaves us with the following regression as our main specification:

$$y_{it} = \beta D_{it} + \mu_i + \gamma_t + \epsilon_{it}. \quad (1)$$

Here, β is the coefficient of interest, D_{it} is an adoption dummy that turns on after a developer first uses GitHub Copilot, μ_i is a developer fixed effect and γ_t is a week fixed effect.⁶ We estimate Equation 1 by two-stage least squares (2SLS), instrumenting D_{it} with a dummy Z_{it} that turns on for all developers randomized into the treatment group after the start of the experiment.

Before we move on to the results, we have to discuss a key complication: our data comes from experiments where the control group was eventually also allowed to access GitHub Copilot. This is not a challenge for identification, but it reduces the power of our instrument if we employ the naive strategy detailed above. In particular, consider a hypothetical experiment that lasts just one month: at $t = 0$, developers are randomized into treatment and control groups, where control is not allowed access to Copilot until $t = 4$. Suppose further that, starting at $t = 4$, the differences in uptake between the two groups start declining over time, asymptoting a zero uptake difference. If we naively estimate (1) by 2SLS in this setting, the power of our instrument (and hence the precision of our estimates) will be strictly declining in the number of periods we observe – in the limit, with infinite periods, our instrument violates the relevance condition because initial treatment assignment eventually fails to predict uptake. One potential solution to this dilemma involves cutting out some data and only using $t \in \{1, \dots, 4\}$ to estimate the model. During this period, the instrument has maximal relevance. However, to the extent that there is still an adoption difference between treatment and control groups at $t = 5$, this strategy is wasteful in that it does not exploit all possible identifying variation.

To avoid ad-hoc decisions about how many periods after the experiment ends should be included in the analyzed data, we weight the 2SLS estimates by the (period-by-period) difference in adoption across treatment and control groups. The resulting weighted IV regression gracefully handles the issue of declining instrument relevance, and it has been previously proposed and analyzed in the context of uptake differences across individuals by (Coussens and Spiess 2021; Huntington-Klein 2020); in the context of uptake differences over time, a similar strategy was employed by Bloom et al. (2012) to improve precision.

The impact of our weighting on the interpretation of our results is straightforward: the weighted regression weights periods based on the difference in Copilot adoption between control and treatment groups. As we can see in Appendix Figure 5, the weighted IV estimates place greater emphasis on treatment effects during periods like March 2023 in the Microsoft experiment, where there was a significant difference in adoption between the treatment and control groups. Thus, to the extent that one would expect treatment effects that are heterogeneous over calendar time (e.g. because Copilot improved over time), the weighting will affect which estimand our estimator is targeting. If, on the other hand, treatment effects are unchanged over the course of the experiment, the weighting will purely improve the precision of our estimates without affecting their interpretation.

⁶For the initial phase of the Microsoft experiment, we do not observe intensive usage data. Hence, we say a developer at Microsoft has adopted Copilot after they either register to use it (relevant in the initial phase) or we see any usage of Copilot (relevant in the later phase.)

4.2 Results

We present our main (weighted) results in Table 2, split by experiment.⁷ At Microsoft, we find positive effects of Copilot on the number of completed pull requests, the number of commits made, and the number of times that code was build (i.e., compiled.) However, only the effect on the number of pull requests is statistically significant at conventional significance levels. We find no negative effect on the build success rate, which would be negatively affected if Copilot was writing code that does not compile and these mistakes are not caught by developers. At Accenture and Anonymous Company we find comparable effect sizes, but these effects lack statistical significance.

Outcome	Microsoft	Accenture	Anon. Comp.	Pooled
Pull Requests	27.38** (12.88)	17.94 (18.72)	54.03 (42.63)	26.08** (10.3)
Commits	18.32 (11.25)	-4.48 (21.88)	- -	13.55 (10.0)
Builds	23.19 (14.20)	92.40*** (26.78)	- -	38.38*** (12.55)
Build Success Rate	-1.34 (4.23)	-17.40** (7.12)	- -	-5.53 (3.64)
N Developers	1,521	316	3,030	4,867
N Clusters	690	316	432	1,438

Table 2: Experiment-by-Experiment Results (Weighted IV).

Notes: This table provides estimates of the effect of GitHub Copilot adoption on the number of Pull Requests, Commits and Successful Builds across three experiments at Microsoft, Accenture, and Anonymous Company. Each entry corresponds to an estimate of β in (1) expressed as a percentage of the control mean. Standard errors are clustered at the level of treatment assignment, which varies across experiments (Microsoft: mixed team-level and individual assignment; Accenture: individual assignment; Anonymous Company: team-level assignment.)

In the final column, we combine estimates across experiments (taking the precision-weighted average across our three experiment-by-experiment estimates) to get the most precise estimate of the effect of the coding assistance tool.⁸ While standard errors are consistently large, we find evidence of the productivity-enhancing effects of GitHub Copilot: on average, the number of weekly pull requests made by developers increases by 26.08% (SE: 10.3%), the number of weekly commits increases by 13.55% (SE: 10.0%), and the number of weekly builds increases by 38.38% (SE: 12.55%).

Before moving on, we discuss an additional experiment run at Accenture that was abandoned due to a large layoff affecting 42% of participants, resulting in a lack of data on Copilot usage (and hence adoption status). Because of these data quality issues, we relegate this experiment to Appendix D. However, if we conservatively impute Copi-

⁷We relegate unweighted results (which are less precise) to Appendix A.

⁸Thus, $\hat{\beta}_{Pooled} = (\sum_{e \in E} 1/\hat{\sigma}_e^2)^{-1} \sum_{e \in E} \hat{\beta}_e/\hat{\sigma}_e^2$ where $E = \{\text{Microsoft, Accenture, Anon. Comp}\}$, $\hat{\beta}_e$ refers to the estimate in experiment e , and $\hat{\sigma}_e$ refers to the standard error.

lot adoption status, we find a negative point estimate of -39.18% (SE: 36.78%) on the number of completed pull requests. The estimates on the number of commits 43.04% (SE: 38.8%) and builds 12.33% (SE: 53.6%) are both positive, though also not statistically significant. The combined results do not change much when we include this second Accenture experiment: the number of pull requests made by developers increases by 21.34% (SE: 9.92%), the number of commits increases by 15.39% (SE: 9.69%), and the number of builds increases by 37.03% (SE: 12.22%).

5 Heterogeneity

Previous literature has noted that productivity enhancements driven by large language models are heterogeneous across skill level and education (Brynjolfsson, Li, and Raymond 2023; Noy and Zhang 2023). In particular, large language models have been found to help the least educated, least skilled workers the most. Because we have access to developer characteristics for the Microsoft experiment, we can separate out our estimates of the productivity effects of Copilot by tenure and seniority. In doing so, we are (to our knowledge) the first to confirm these prior findings for highly-paid experts.

In particular, we now break out results by (i) the tenure and (ii) the level of employees at Microsoft⁹. We split developers into short and long tenure based on the median observed tenure in our data.¹⁰ Any developers that have been with Microsoft for less than the median time at the start of the experiment are considered “short tenure,” and all other developers “long tenure.” Similarly, we split developers into “junior” and “senior” based on the level at which they are hired at the company.

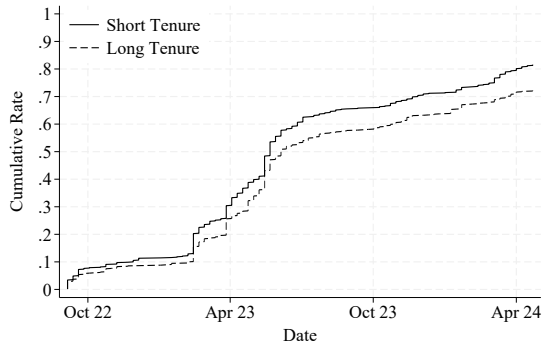
We begin our analysis by considering the heterogeneity in adoption patterns in Figure 2(a-b). We see that short-tenure developers are 9.5 percentage points (SE: 2.2pp) more likely to adopt Copilot (84.3% vs 74.8%), consistent with prior research suggesting that younger workers (who naturally have lower tenure on average) are more likely to adopt new technologies (Meyer 2011). The same effect is at play for the junior developers, which are 5.3 percentage points (SE: 2.1pp) more likely to adopt (82.1% vs 76.8%), though the adoption difference is slightly smaller in this dimension.

The next panel, Figure 2(c), reveals that employees of shorter tenure are more likely to continue using Copilot more than one month after initial adoption, suggesting that they perhaps expect larger benefits from the technology than their more experienced counterparts. Judging by Figure 2(d), this effect does not seem present when comparing junior to senior developers. Finally, Figure 2(e) reveals that higher-tenure developers are approximately 4.3% (or 1.0 percentage points) less likely to accept code suggested by Copilot. When comparing junior to senior developers, this difference in acceptance rates is much smaller at 1.8% (or 0.5 percentage points), though it goes in the same direction of senior developers being less likely to accept AI suggestions.


Moving on towards output measures in Figure 4, we see that the productivity-enhancing effects of Copilot are stronger for lower tenure and more junior developers. While our

⁹We measure level as of March 1st, 2023, which is the earliest date available in our data.

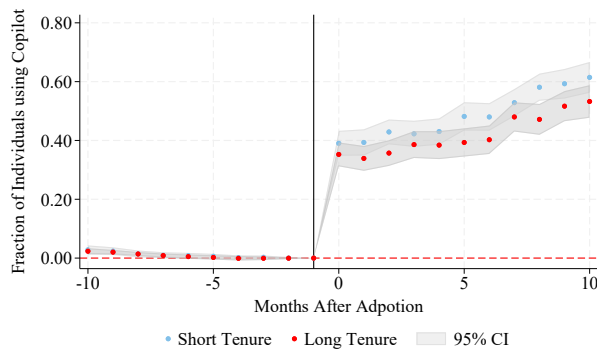
¹⁰As it is considered sensitive information, we cannot reveal the exact median tenure, but it is between 2 and 4 years.



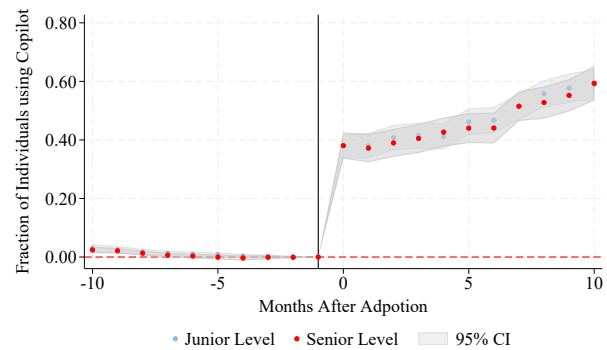
(a) Adoption by Tenure



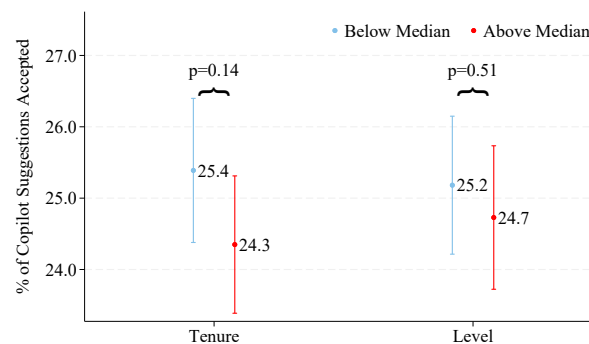
(b) Adoption by Level



(c) Usage (Since Adoption) by Tenure



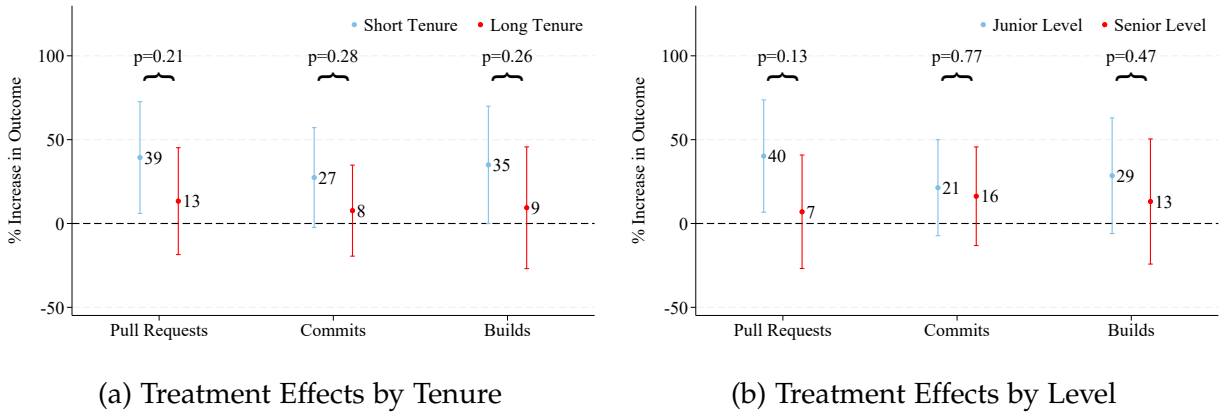
(d) Usage (Since Adoption) by Level



(e) Fraction of Suggestions Accepted

Figure 2: Heterogeneity of Adoption and Usage of Copilot

Notes: This figure explores heterogeneous adoption and usage patterns of Copilot across developer tenure and level. Panel (a) provides the adoption of Copilot over time broken out by whether a developer's tenure with Microsoft at the beginning of the experiment was below or above median; panel (b) does the same for level. Developers who have only been with the company for a short period and junior developers are slightly more likely to adopt. Panels (c) and (d) provide event studies that detail the extensive margin, i.e., how likely a developer is to have used Copilot at all a given number of months after her first usage of Copilot; short-tenure developers are more likely to stick with Copilot. Panel (e) shows the fraction of Copilot suggestions that are accepted by developers; short-tenure developers are slightly more likely to accept suggestions.



(a) Treatment Effects by Tenure

(b) Treatment Effects by Level

Figure 3: Heterogeneity of Copilot Effect by Tenure and Level (Weighted IV).

Notes: This figure provides weighted IV estimates of the effect of adopting Copilot on total number of pull requests, commits, and builds broken out by (a) whether a developer’s tenure with Microsoft at the beginning of the experiment was below median (short tenure) or above median (long tenure) and (b) which level a developer was hired into. The dots are estimates derived from estimating (1) separately for short- and long-tenure developers, and the bars provide 95% confidence intervals based on standard errors clustered at the level of treatment assignment. For all three outcome measures, the effects on productivity are stronger for short-tenure developers.

estimates are noisy and not statistically significant at conventional levels, the pattern is the same across all three outcome measures: short-tenure developers increase their output by 27% to 39% while long-tenure developers have more marginal gains of 8% to 13%. However, we note that because of our emphasis on the average effect of *initially* adopting Copilot and the patterns in Figure 2(c), the estimates for longer-tenure developers may be attenuated by a larger extent of developers abandoning the technology after an initial trial phase. Still, there is no difference in usage conditional on adoption between junior and senior developers in Figure 2(d) and still we see in 4(b) that junior developers increase their output by 21% to 40% while senior developers have more marginal gains of 7% to 16%.

6 Conclusion

To summarize, we find that usage of a generative AI code suggestion tool increases software developer productivity by 26.08% (SE: 10.3%). This estimate is based on observing, partly over years, the output of almost five thousand software developers at three different companies as part of their regular job, which strongly supports its external validity.

References

- Acemoglu, Daron (Apr. 2024). *The Simple Macroeconomics of AI*. Tech. rep. Prepared for Economic Policy. Massachusetts Institute of Technology.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb (2019). "Economic policy for artificial intelligence". In: *Innovation policy and the economy* 19.1, pp. 139–159.
- Baily, Martin, Erik Brynjolfsson, and Anton Korinek (2023). "Machines of mind: The case for an AI-powered productivity boom". In.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (Nov. 2012). "Does Management Matter? Evidence from India *". In: *The Quarterly Journal of Economics* 128.1, pp. 1–51. eprint: <https://academic.oup.com/qje/article-pdf/128/1/1/30624754/qjs044.pdf>.
- Bresnahan, Timothy (2024). "What innovation paths for AI to become a GPT?" In: *Journal of Economics & Management Strategy* 33.2, pp. 305–316. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jems.12524>.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2023). "Generative AI at Work". In: *NBER Working Paper* w31161.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2021). "The productivity J-curve: How intangibles complement general purpose technologies". In: *American Economic Journal: Macroeconomics* 13.1, pp. 333–372.
- Bubeck, Sébastien, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrke, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, et al. (2023). "Sparks of artificial general intelligence: Early experiments with gpt-4". In: *arXiv preprint arXiv:2303.12712*.
- Campero, Andres, Michelle Vaccaro, Jaeyoon Song, Haoran Wen, Abdullah Almaatouq, and Thomas W. Malone (2022). *A Test for Evaluating Performance in Human-Computer Systems*. arXiv: 2206.12390 [cs.HC].
- Coussens, Stephen and Jann Spiess (2021). *Improving Inference from Simple Instruments through Compliance Estimation*. arXiv: 2108.03726 [econ.EM].
- Dell'Acqua, Fabrizio, Edward McFowland, Ethan R Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Kraymer, François Candelon, and Karim R Lakhani (2023). "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality". In: *Harvard Business School Technology & Operations Mgt. Unit Working Paper* 24-013.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2023). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*. arXiv: 2303.10130 [econ.GN].
- Frank, Morgan R, David Autor, James E Bessen, Erik Brynjolfsson, Manuel Cebrian, David J Deming, Maryann Feldman, Matthew Groh, José Lobo, Esteban Moro, et al. (2019). "Toward understanding the impact of artificial intelligence on labor". In: *Proceedings of the National Academy of Sciences* 116.14, pp. 6531–6539.
- Furman, Jason and Robert Seamans (2019). "AI and the Economy". In: *Innovation policy and the economy* 19.1, pp. 161–191.

- Hoffmann, Manuel, Sam Boysel, Frank Nagle, Sida Peng, and Kevin Xu (July 2024). “Generative AI and Distributed Work: Evidence from Open Source Software”. In: *Unpublished*. Version: July 8, 2024.
- Huntington-Klein, Nick (2020). In: *Journal of Causal Inference* 8.1, pp. 182–208.
- Jaffe, Sonia, Neha Parikh Shah, Jenna Butler, Alex Farach, Alexia Cambon, Brent Hecht, Michael Schwarz, and Jaime Teevan (July 2024). *Generative AI in Real-World Workplaces*. Tech. rep. MSR-TR-2024-29. Microsoft.
- Meyer, Jeremy (2011). “Workforce age and technology adoption in small and medium-sized service firms”. In: *Small Business Economics* 37.3, pp. 305–324.
- Noy, Shakked and Whitney Zhang (2023). “Experimental evidence on the productivity effects of generative artificial intelligence”. In: *Science* 381.6654, pp. 187–192.
- Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirel (2023). “The impact of ai on developer productivity: Evidence from github copilot”. In: *arXiv preprint arXiv:2302.06590*.
- Vaithilingam, Priyan, Tianyi Zhang, and Elena L Glassman (2022). “Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models”. In: *Chi conference on human factors in computing systems extended abstracts*, pp. 1–7.
- Yeverechyahu, Doron, Raveesh Mayya, and Gal Oestreicher-Singer (July 2024). “The Impact of Large Language Models on Open-source Innovation: Evidence from GitHub Copilot”. In: *Unpublished*. Available at SSRN: <https://ssrn.com/abstract=4684662> or <http://dx.doi.org/10.2139/ssrn.4684662>.

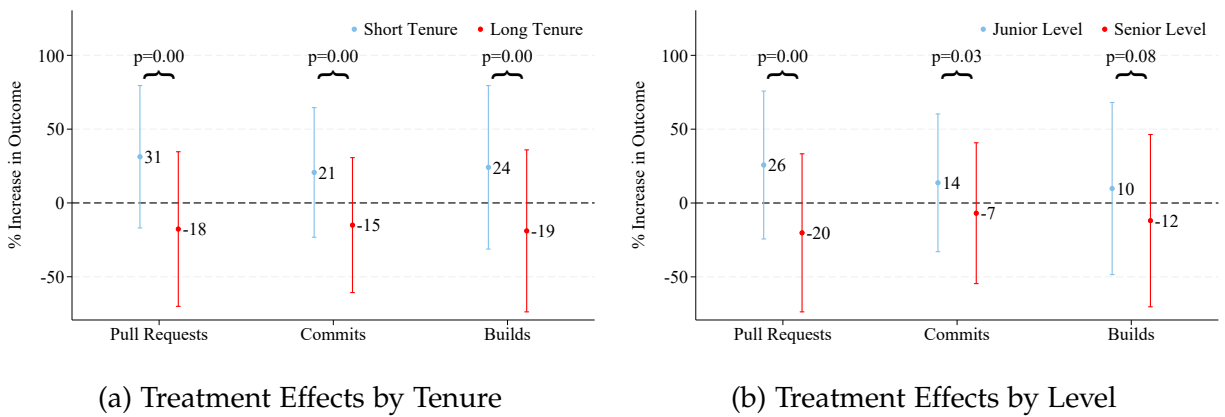
A Weighted vs Unweighted IV

This appendix provides versions of Table 2 and Figure 2 that do not use weights when estimating 1. The resulting results have wider confidence intervals, but point in the same direction as our main estimates. We also exhibit the weights that underlie Table 2 and Figure 2.

Outcome	Microsoft	Accenture	Anon. Comp.	Pooled
Pull Requests	10.53 (24.82)	15.97 (21.26)	54.03 (42.63)	18.73 (15.1)
Commits	5.54 (22.20)	-3.60 (22.19)	- -	0.97 (15.69)
Builds	5.87 (27.25)	96.05*** (28.05)	- -	49.66** (19.55)
Build Success Rate	3.92 (8.17)	-18.10* (9.55)	- -	-5.39 (6.21)
N Developers	1,521	316	3,030	4,867
N Clusters	690	316	432	1,438

Table 3: Experiment-by-Experiment Results (Unweighted IV).


Notes: This table provides estimates of the effect of GitHub Copilot adoption on the number of Pull Requests, Commits and Successful Builds across three experiments at Microsoft, Accenture, and Anonymous Company. Each estimate is expressed as a percentage of the control mean. Standard errors are clustered at the level of treatment assignment, which varies across experiments (Microsoft: mixed team-level and individual assignment; Accenture: individual assignment; Anonymous Company: team-level assignment.)




(a) Treatment Effects by Tenure

(b) Treatment Effects by Level


Figure 4: Heterogeneity of Effect of Copilot By Tenure and Level (Unweighted IV)
Notes: This figure provides unweighted IV estimates of the effect of adopting Copilot on total number of pull requests (top left), commits (top right), and builds (bottom center) broken out by whether a developer's tenure with Microsoft at the beginning of the experiment was below median (short tenure) or above median (long tenure.) The estimates are derived from estimating (1) separately for short- and long-tenure developers. For all three outcome measures, the effects on productivity are stronger for short-tenure developers.



(a) Microsoft Experiment



(b) Accenture Experiment



(c) Accenture Experiment

Figure 5: Regression Weight

Notes: This figure provides the weights used in the IV estimates underlying Table 2. Recall that we are weighting the IV estimates to exploit information from periods where the instrument predicts uptake. Hence, we use the difference in adoption across the control and treatment group by a given date as our weight. This matters most for the Microsoft experiment, in which the control group adopted at an elevated rate when its access was granted in March 2023. For this experiment, we put extra weight on the period just before the control group was allowed to adopt Copilot.

B Data Cleaning

We provide details on which individuals we had to exclude from each raw dataset, and the reasons for their exclusion.

B.1 Microsoft

In the original sample of the dataset, we have 1,746 individuals. We kept only software engineers, which leaves us 1,538, and we dropped people who switched organizations, leaving us 1,522. Finally, we dropped one individual who adopted before the experiment started, with a final sample of 1,521.

We also drop the data for the last week of the dataset, since the dataset does not record the full week of activity for the last week.

Finally, note that while the restriction for the control group was lifted in April 2023, ten individuals in the control group adopted before that date. We include these individuals in our regressions, and they naturally weaken the strength of the instrument.

B.2 Accenture

We drop individuals who have no record of data, and people who left the company. We start with the original dataset with 369 individuals, and after dropping individuals with no outcome measure found we are left with 320, and after dropping individuals who left the company we are left with a final sample of 316.

Finally, we note that while individuals in the control group were allowed to adopt starting December 2023, there was one individual in the control group who adopted in October 2023. We include this individual in our regressions.

B.3 Anonymous Company

The original sample has 3,054 individuals. We drop individuals who have shown/adopted before they were given access, and are left with a final sample of 3,030 individuals.

C Experiment Details

TO TREATED GROUP

Intended recipients: Engineers and PM under [REDACTED]
Proposed subject line: Copilot dogfood experiment

Hi there,

We would like to invite you to use [Github Copilot](#) for your day to day work as part of our Copilot research project with Office of Chief Economist. Dogfooding is an important step to ensure we are using our own product and providing feedback to the Copilot team. This is a key step for us to make Copilot better for all developers.

Please visit: <link to onboarding experience> to learn more. If you agree to take part in this study you must first consent to participation, fill out the onboarding form and review the usage guideline. After submitting your onboard form, your account will be manually activated within 3 days and you will get a welcome email with installation instructions.

If you have any question regarding this project, please contact [REDACTED]

If you'd rather not be contacted regarding this in the future or have any questions about this project, please let us know.

Contact: [REDACTED]
[REDACTED] | [Microsoft Data Privacy Notice](#)

CONTROL GROUP: Separate message for the control group who would want access to Copilot

Hello,

Thanks for your interest in Copilot. Your division is participating in an internal controlled research study. Your team was randomly selected to not yet receive Copilot access. As the study concludes, you will be notified that Copilot is now available for your use.

If you have any questions about this study, please reach out to [REDACTED] from the Chief Economist's office.

If you'd rather not be contacted regarding this in the future or have any questions about this project, please let us know.

Contact: [REDACTED]
[REDACTED] | [Microsoft Data Privacy Notice](#)

Figure 6: E-mail Sent To Participants in the Microsoft Experiment

Notes: This figure exhibits the copy that was sent to participants in the Microsoft experiment.

	Control		Treatment		Difference	p-value
	Mean	Std. Dev	Mean	Std. Dev		
Pull Requests	0.08	0.26	0.09	0.29	0.02	0.38
Commits	6.28	11.24	5.28	10.09	-1.00	0.30
Builds	5.32	10.32	5.23	10.52	-0.09	0.93
Build Success Rate	0.49	0.33	0.50	0.33	0.01	0.60

Table 4: Balance Table for First Accenture Experiment

Notes: This table presents a comparison of pre-experimental outcomes in control and treatment groups across experiments. For each measure, we present its mean and standard deviation in the control group, and in the treatment group. We also show the mean difference across these groups, and the p-value associated with an underlying test of a difference in means. The differences in p-values are calculated using standard errors clustered at the level of treatment assignment.

D The First Accenture Experiment

We do not discuss at length in the main text another experiment which was run by Accenture in April 2023 and included a number of Accenture offices located in Southeast Asia. This experiment was abandoned by the company after Accenture laid off 19,000 employees that some month (cnn.com), including 42% of the developers participating in this experiment. Still, this attrition was balanced across treatment and control, and we can thus subset to the 204 developers who were not let go for our analysis; indeed Table 4 confirms that after this subsetting, treatment and control are still balanced. The problem emerges because Microsoft did not log all Copilot usage data for this experiment, as the company considered it abandoned. In particular, we lack adoption data for the control group until October '23. Without this adoption data, any analysis is potentially biased.

Still, because our initial analysis revealed that this experiment was the only experiment across the three in which we have a negative (though statistically insignificant) point estimate for Copilot’s effect on productivity, we proceed to analyze this experiment in this appendix by imputing that nobody in the Control group adopts Copilot until October '23, yielding the adoption path in Figure 7. Thus, in the worst-case scenario, it could be that all the adoptions that we attribute to October 2023 already happened right at the beginning of the experiment. This data quality concern means our treatment effect estimates will be conservative (as we may mistakenly count up to 10% of the control group as non-adopters during half of the sampling period.)

Keeping in mind this caveat that our treatment effect estimates are potentially conservative, we report the results from this first Accenture experiment in Table 5. We find a negative point estimate of -39.18% (SE: 36.78%) on the number of tasks completed. Still, this estimate has a high degree of statistical uncertainty, and we note that the estimates on the number of commits 43.04% (SE: 38.8%) and builds 12.33% (SE: 53.6%) are both positive, though also not statistically significant.




Figure 7: Cumulative Adoption Rates for First Accenture Experiment.

Notes: The first three graphs show the cumulative rate of adoption over time for software developers in both the treatment and control groups across various experiments. In the Anonymous Company experiment (Panel d), unlike the other experiments, all developers were granted access to Copilot in a staggered fashion, with the order of access randomized among participants. Hence, we show the cumulative fraction of users invited to participate, and who adopted Copilot.

Outcome	Accenture #1
Pull Requests	-39.18 (36.78)
Commits	43.04 (38.80)
Builds	12.33 (53.60)
Build Success Rate	-0.99 (16.51)
N Developers	204
N Clusters	204

Table 5: Weighted IV Results for First Accenture Experiment.

Notes: This table provides estimates of the effect of GitHub Copilot adoption on the number of Pull Requests, Commits and builds across three experiments at Microsoft, Accenture, and Anonymous Company. Each estimate is expressed as a percentage of the control mean. Standard errors are clustered at the level of treatment assignment, which varies across experiments (Microsoft: mixed team-level and individual assignment; Accenture: individual assignment; Anonymous Company: team-level assignment.)