

How does Division of Labor Affect Team Productivity? Evidence from GitHub

Jinci Liu

IIES, Stockholm University

March 12, 2025

Motivation

- ▶ Large differences in organizational structure between/within firms (Bloom et al., 2013, 2014; Giroud et al., 2022)
- ▶ Given a set of workers and tasks, how to allocate them?
- ▶ Seminal theory: **division of labor increases productivity** (Smith, 1776)
 - * To make a pin, divide into 18 distinct operations, each carried out by different people
 - * Gain: human capital accumulation (Rosen, 1983; Becker and Murphy, 1992; Young, 1928)

Routine versus Non-routine Production



- ▶ Explicit rules & repeat procedures Autor et al. (2003)



- ▶ Problem-solving & complex communication

Autor et al. (2003)

Routine versus Non-routine Production



- ▶ Explicit rules & repeat procedures Autor et al. (2003)
- ▶ Empirical evidence of specialization benefits
 - * Gong and Png (2024): **cashiers**
 - * Kohlhepp (2024): **hair salon**



- ▶ Problem-solving & complex communication Autor et al. (2003)

Routine versus Non-routine Production



- ▶ Explicit rules & repeat procedures Autor et al. (2003)
- ▶ Empirical evidence of specialization benefits
 - * Gong and Png (2024): **cashiers**
 - * Kohlhepp (2024): **hair salon**



- ▶ Problem-solving & complex communication Autor et al. (2003)
- ▶ Cross-task feedback and knowledge sharing are important

Research Q: How does Team Specialization Affect Productivity?

- ▶ Effect of specialization on team productivity is ex-ante **ambiguous**
 - * Increase member's task-specific human capital
 - * Suffer from coordination cost Becker and Murphy (1992) and learning myopia Levinthal and March (1993)
- ▶ This project:
 - * **New data:** millions of task assignments for software developers
 - * **New measures:** task allocation specialization and productivity
 - * **New result:** specialization is detrimental for software development team productivity

Roadmap

1. Data
2. Measures
3. Facts
4. Empirical
5. Conclusion

Data Construction

- ▶ Use data from the **largest online coding platform** GitHub
- ▶ **64,400** software development **teams** (public repository) under firms (organization) [◀ Detail](#)
 - * E.g., Teams in Microsoft, Meta, Google
- ▶ **35 million** code files [◀ Detail](#)
 - * Unsupervised learning algorithm to classify into **10 task types** (E.g., frontend, backend) [◀ Detail](#)
- ▶ **292,840** team members (defined by GitHub) [◀ Detail](#)
- ▶ **2017-2023**

Measuring Team Specialization

Example

- ▶ Each member has inelastic 1 unit of labor supply (row sum)
- ▶ Each task requires a different unit of labor supply (column sum: **task share**)

	Task			
	1	2	3	
A	1/2	0	1/2	1
B	1/2	1/2	0	1
C	1/2	0	1/2	1
	3/2	1/2	1	

Actual(A)

Example

- ▶ Each member has inelastic 1 unit of labor supply (row sum)
- ▶ Each task requires a different unit of labor supply (column sum: **task share**)

	Task			
	1	2	3	
A	1/2	0	1/2	1
B	1/2	1/2	0	1
C	1/2	0	1/2	1
	3/2	1/2	1	

Actual(A)

	Task			
	1	2	3	
A	1/2	1/6	1/3	1
B	1/2	1/6	1/3	1
C	1/2	1/6	1/3	1
	3/2	1/2	1	

Generalized(G)

Example

- ▶ Each member has inelastic 1 unit of labor supply (row sum)
- ▶ Each task requires a different unit of labor supply (column sum: **task share**)

	Task			
	1	2	3	
A	1/2	0	1/2	1
B	1/2	1/2	0	1
C	1/2	0	1/2	1
	3/2	1/2	1	

Actual(A)

	Task			
	1	2	3	
A	1/2	1/6	1/3	1
B	1/2	1/6	1/3	1
C	1/2	1/6	1/3	1
	3/2	1/2	1	

Generalized(G)

	Task			
	1	2	3	
A	1	0	0	1
B	1/2	1/2	0	1
C	0	0	1	1
	3/2	1/2	1	

Specialized(S)

Example

- ▶ Each member has inelastic 1 unit of labor supply (row sum)
- ▶ Each task requires a different unit of labor supply (column sum: **task share**)

	Task			
	1	2	3	
A	1/2	0	1/2	1
B	1/2	1/2	0	1
C	1/2	0	1/2	1
	3/2	1/2	1	

Actual(A)

	Task			
	1	2	3	
A	1/2	1/6	1/3	1
B	1/2	1/6	1/3	1
C	1/2	1/6	1/3	1
	3/2	1/2	1	

Generalized(G)

	Task			
	1	2	3	
A	1	0	0	1
B	1/2	1/2	0	1
C	0	0	1	1
	3/2	1/2	1	

Specialized(S)

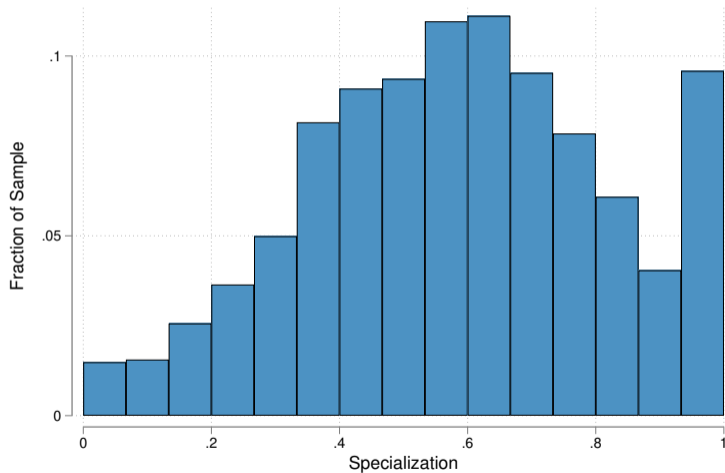
- ▶ Define team specialization index $SPE: \frac{d(A,G)}{d(S,G)}$ (d: Euclidean distance)
- ▶ SPE is **higher (more specialized)** if A is far from G

▶ Math formula

◀ Distance between two matrices

Distribution of Team Specialization Index

Team-month level



Measuring Productivity

Outcome Measures

Team-month level

- ▶ **Output Quality:** stars per month (Users' Appreciation) (Borges and Valente, 2018)
 - * 75% developers check stars before using [▶ Link](#)
 - * Stars have monetary value \approx \$0.88/star [▶ Link](#)
- ▶ **Output Quantity:** lines-of-code (Vasilescu et al., 2015),(Casalnuovo et al., 2015),(Wagner and Ruhe, 2018)
- ▶ **Problem-Solving Speed:** time from users' bug report to solve [◀ Example](#)
- ▶ **Code Acceptance Rate:** success rate of member code submissions [◀ Manage to merge](#) [◀ Fail to merge](#)
- ▶ **Discussion:** comments sent by team members

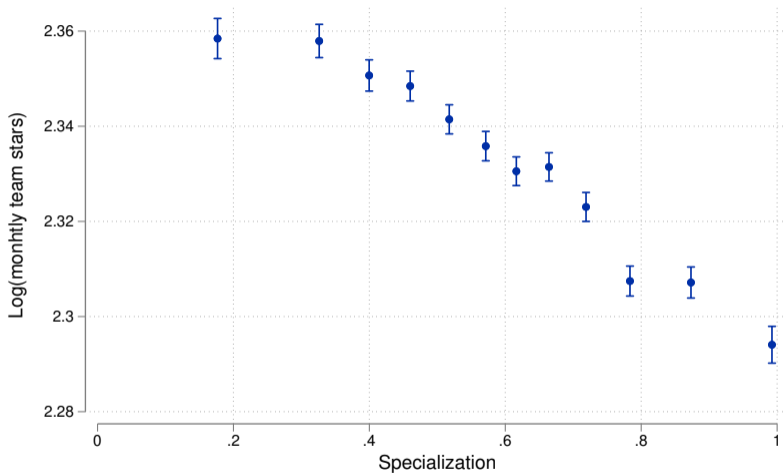
[◀ Summary statistics](#)

Team Specialization and Productivity

$$Y_{mt} = \beta_1 SPE_{mt} + \theta_i + \gamma_m + \phi_{mt} + \beta_2 \mathbf{X}_{mt} + \varepsilon_{mit}$$

- ▶ Y_{mt} : outcome for team m in month t
- ▶ SPE_{mt} : degree of specialization for team m in t
- ▶ β_1 : coefficient of interest
- ▶ θ_i : member fixed effect
- ▶ γ_m : team fixed effect
- ▶ ϕ_{mt} : team age, team size fixed effect
- ▶ X_{mt} : 10 task type distribution
- ▶ Standard error is clustered at team level
- ▶ Weight by 1/team size

Fact 1 Higher Specialization, Lower Quality



Controls: task distribution

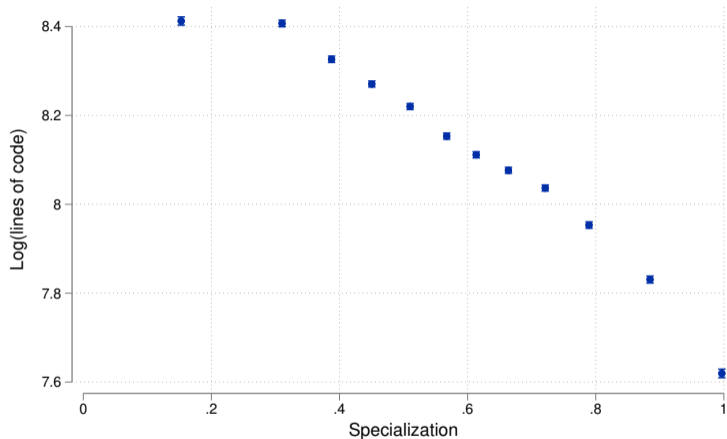
Fixed effect: team age, team size, team, member

Output Quality

Dep Var: Log (stars per month)	(1)	(2)	(3)
SPE_{mt}	-0.320*** (0.021)	-0.078*** (0.007)	-0.086*** (0.007)
Dependent mean	2.13	2.13	2.13
R-squared	0.648	0.903	0.910
Observations	1,823,750	1,823,750	1,770,310
Task type share control	Y	Y	Y
Team age FE	Y	Y	Y
Team size FE	Y	Y	Y
Firm FE	Y		
Team FE		Y	Y
Member FE			Y

NOTE. Weighted by 1/team size.

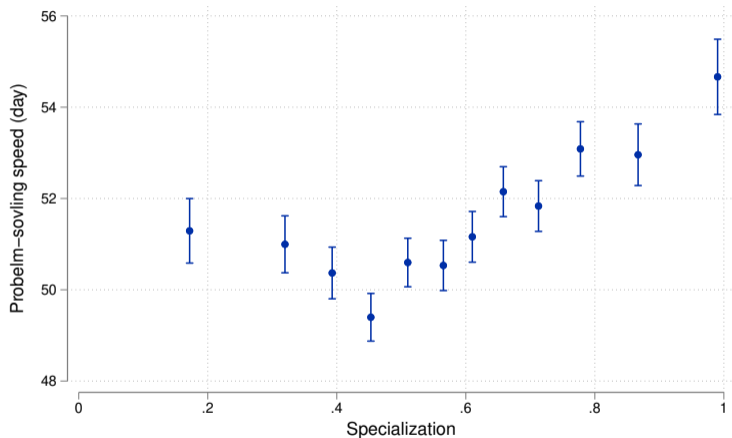
Fact 2 Higher specialization, Lower Quantity



Controls: task distribution
Fixed effect: team age, team size, team, member

Fact 3 Slower to Solve Users' Problems

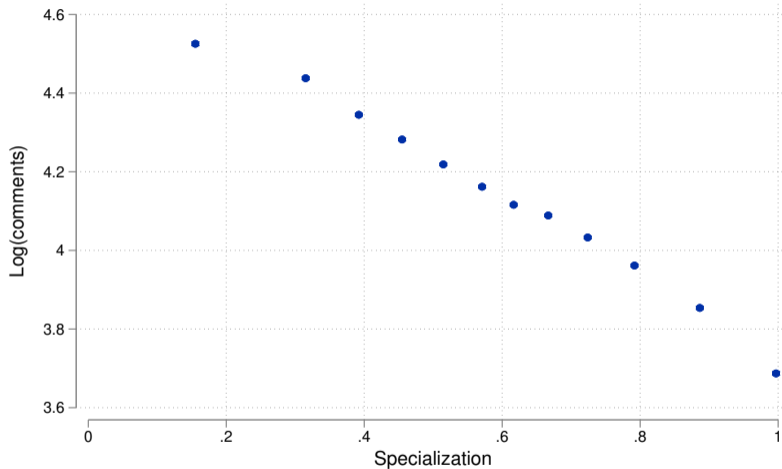
Cross-task knowledge



Controls: task distribution

Fixed effect: team age, team size, #issues, team, member

Fewer Discussions between Team Members

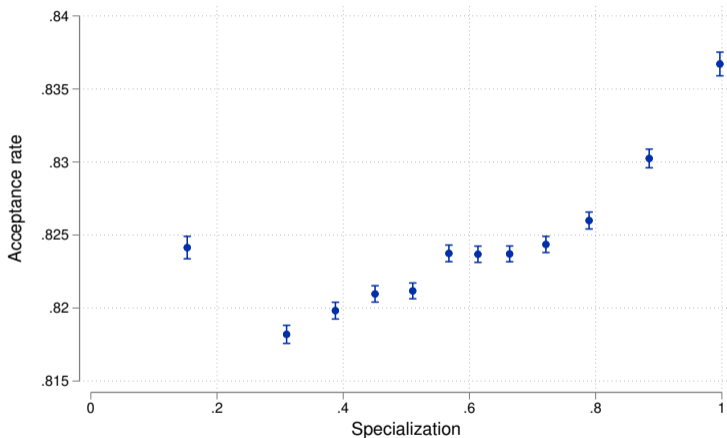


Controls: task distribution

Fixed effect: team age, team size, team, member

Fact 4 Higher Code Acceptance Rates

Task-specific knowledge



Controls: task distribution, team age
Fixed effect: team size, #submissions, team, worker

Summary so far

- ▶ Negative correlation between specialization and output quality and quantity
- ▶ Specialized teams take longer to solve users' problems
- ▶ Specialized teams have higher code acceptance rate

Summary so far

- ▶ Negative correlation between specialization and output quality and quantity
- ▶ Specialized teams take longer to solve users' problems
- ▶ Specialized teams have higher code acceptance rate
- ▶ But task allocation is endogenous
- ▶ Ideal experiment:
 - * Randomly change a team's task allocation to increase or decrease specialization
 - * Impact on team productivity

Automatic Task Assignment

Research Design

- ▶ Automatic task assignment **decreases specialization**
 - * **Evenly distributes** some **tasks** among team members
 - * Teams enable it via configuration files

Research Design

▶ Automatic task assignment **decreases specialization**

- * **Evenly distributes** some **tasks** among team members
- * Teams enable it via configuration files

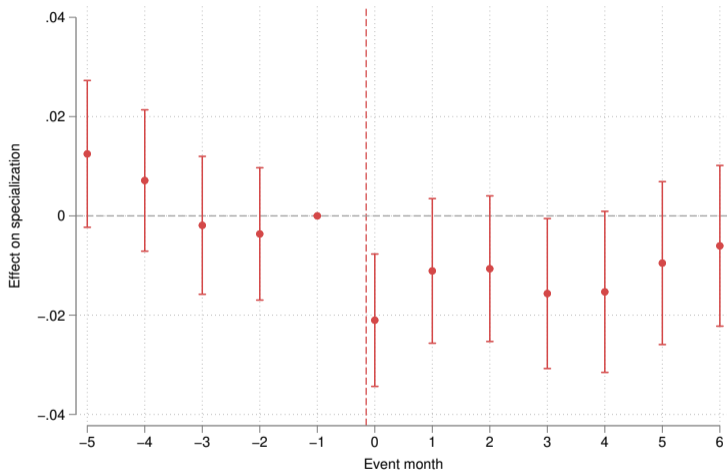
▶ **Concern:** Adoption is not exogenous

- * **Solution:** Create a control group for teams that adopted

1. 1:1 match on team size, task types, and activities in $t-1$ to $t-5$
2. Use matched groups to construct treatment and control groups (98.2% matching rate)

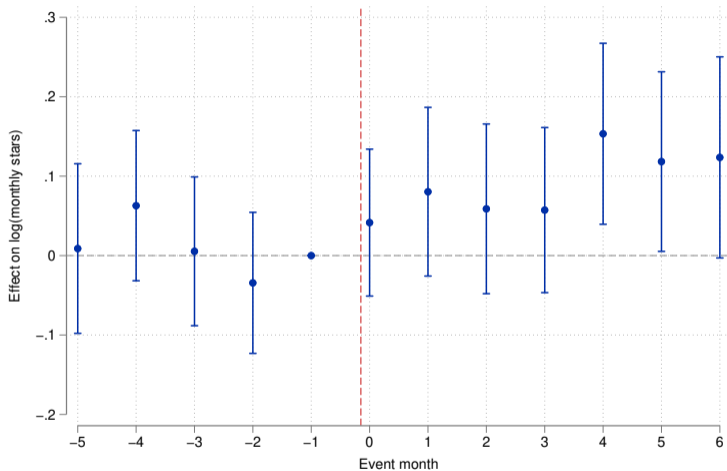
- * **Validate** with empirical test of parallel trends assumption and compare outcomes using **diff-in-diff**

Specialization Decreased by 1.7%



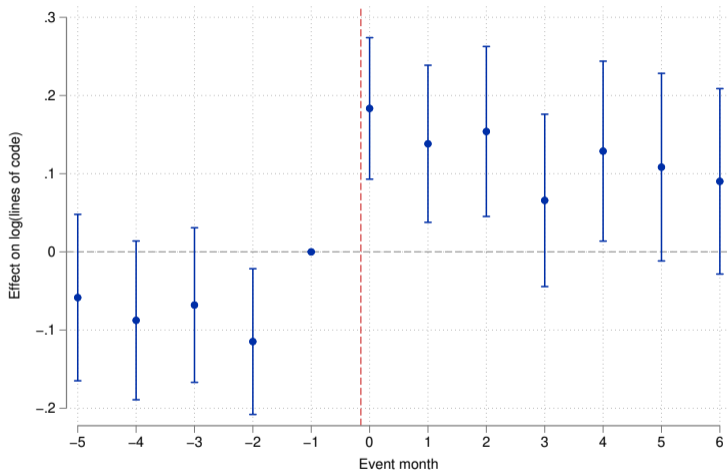
Implied constant treatment effect: -0.017^{***} ($se=0.004$)

Output Quality (Star) Increased after 3 months



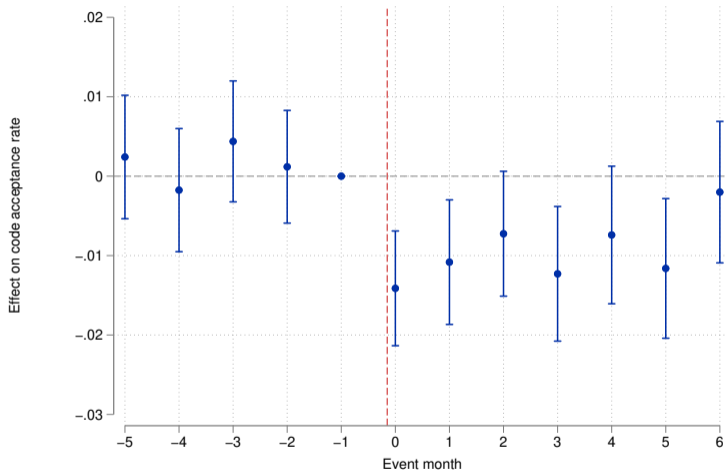
Implied constant treatment effect: 0.021 (se=0.022)

Output Quantity (Code) Increased by 17.7 %



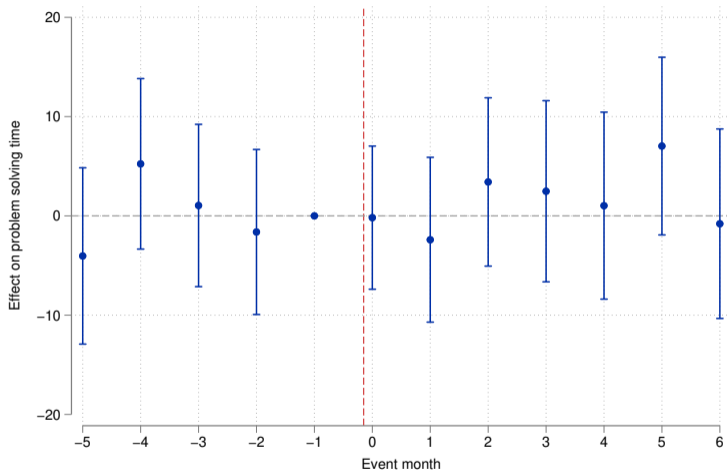
Implied constant treatment effect: 0.177*** (se=0.027)

Code Acceptance Rate Decreased by 2.4%



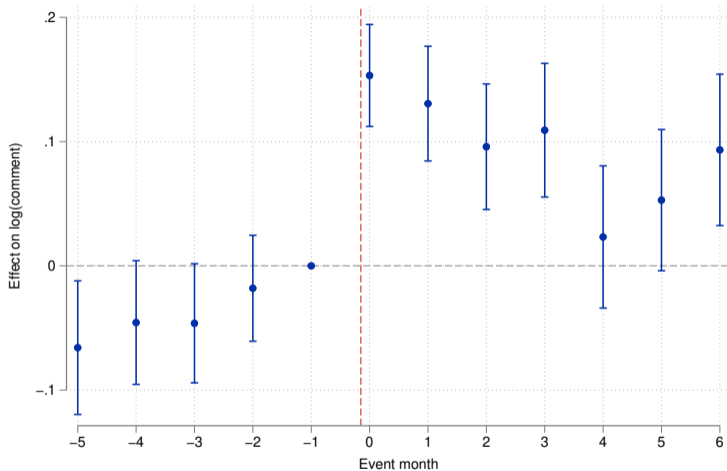
Implied constant treatment effect: -0.01^{***} (se=0.0.02)

No Effect for Problem Solving Speed



Implied constant treatment effect: 0.384 (se=1.970)

Potential Mechanism: Discussion Increased by 14.2%



Implied constant treatment effect after : 0.142^{***} (se=0.0.017)

Conclusion

- ▶ Measure task allocation specialization and team productivity
- ▶ Specialization is **negatively** associated with output **quality** and **quantity**
- ▶ Automatic assignment ↓ specialization, ↑ productivity
- ▶ Potential Mechanism: **loss of cross-task knowledge**
 - * Reduced specialization increases team discussions
 - * **Future:** text analysis to capture team communication patterns and knowledge spillover
- ▶ Specialization restricts knowledge exchange for innovation

Thank you

jinci.liu@iies.su.se

Outcome: Code acceptance rate

	(1)	(2)	(3)	(4)
SPE	-0.010*** (0.002)	0.020*** (0.001)	0.015*** (0.001)	0.015*** (0.001)
Task type share control	Y	Y	Y	Y
Team age fixed effect	Y	Y	Y	Y
Team size fixed effect	Y	Y	Y	Y
Firm fixed effect	Y	Y		
Project type fixed effect		Y		
# code submission		Y	Y	Y
Team fixed effect			Y	Y
Member fixed effect				Y
R-squared	0.217	0.214	0.417	0.442
Dependent mean	0.824	0.825	0.824	0.824
Observations	3,213,677	2,982,846	3,213,677	3,125,570

► [Graph version](#)

NOTE. Weighted by 1/team size

Outcome: Problem Solving Speed

	(1)	(2)	(3)	(4)
SPE	17.655*** (1.496)	11.247*** (1.540)	4.411** (1.385)	4.155** (1.332)
Task type share control	Y	Y	Y	Y
Team age fixed effect	Y	Y	Y	Y
Team size fixed effect	Y	Y	Y	Y
Firm fixed effect	Y	Y		
Project type fixed effect		Y		
# question fixed effect		Y	Y	Y
Team fixed effect			Y	Y
Member fixed effect				Y
R-squared	0.166	0.166	0.355	0.398
Dependent Mean	51.182	51.372	51.182	51.579
Observations	2,364,271	2,267,892	2,364,271	2,294,737

► [Graph version](#)

NOTE. Weighted by 1/team size

Outcome:Log (Lines of Code)

	(1)	(2)	(3)	(4)
SPE	-1.224*** (0.021)	-1.260*** (0.022)	-0.954*** (0.015)	-0.934*** (0.015)
Dependent mean	8.124	8.128	8.124	8.118
R-squared	0.398	0.403	0.597	0.616
Observations	3,212,871	2,982,107	3,212,871	3,124,751
Task type share control	Y	Y	Y	Y
Team age fixed effect	Y	Y	Y	Y
Team size fixed effect	Y	Y	Y	Y
Firm fixed effect	Y	Y		
Project type fixed effect		Y		
Team fixed effect			Y	Y
Member fixed effect				Y

► Graph version

NOTE. Weighted by 1/team size

Team specialization index

- ▶ \mathbb{K} the set of all tasks
- ▶ Each team m consists of $N_m \in \mathbb{N}$ members, handles a set of task $K_m \subseteq \mathbb{K}$
- ▶ Define A_m to be task allocation matrix of team m
 - * $A_m(i, j)$ represents member i 's labor input on task j
- ▶

$$A_m = \begin{pmatrix} A_m(1, 1) & A_m(1, 2) & \cdots & A_m(1, |K_m|) \\ A_m(2, 1) & A_m(2, 2) & \cdots & A_m(2, |K_m|) \\ \vdots & \vdots & \ddots & \vdots \\ A_m(N_m, 1) & A_m(N_m, 2) & \cdots & A_m(N_m, |K_m|) \end{pmatrix}$$

- ▶ Member i 's labor share: $l_i := \sum_j A_m(i, j)$
- ▶ Task share: $\alpha_j := \sum_i A_m(i, j)$

Team specialization index

- ▶ Euclidean distance

$$d(A_m - G_m) = \sqrt{(A_m(i, j) - G_m(i, j))^2}$$

$$d(S_m - G_m) =$$

$$\sqrt{\sum_j \left(\underbrace{[\alpha_j]}_{\# \text{ of 1s}} \cdot \left(1 - \frac{\alpha_j}{N_m}\right)^2 + \underbrace{(N_m - [\alpha_j])}_{\# \text{ of 0s}} \cdot \left(\frac{\alpha_j}{N_m}\right)^2 + ([\alpha_j] - [\alpha_j]) \cdot \left(\underbrace{\alpha_j - [\alpha_j]}_{\text{residual workload}} - \frac{\alpha_j}{N_m} \right)^2 \right)}$$

- ▶ Kullback-Leibler divergence

$$d(A_m - G_m) = \sum_i \sum_j A_m(i, j) \log \left(\frac{A_m(i, j)}{G_m(i, j)} \right)$$

$$d(S_m - G_m) = \sum_j \left([\alpha_j] \cdot \log \left(\frac{1}{\frac{\alpha_j}{N_m}} \right) + ([\alpha_j] - [\alpha_j]) \cdot (\alpha_j - [\alpha_j]) \cdot \log \left(\frac{\alpha_j - [\alpha_j]}{\frac{\alpha_j}{N_m}} \right) \right)$$

Project Type

Topic 1



Topic 2



Topic 3



Topic 4



Topic 5



Topic 6



▶ Back

Specialization by Team Size

Final sample only includes teams with size >3



▶ Back

Summary Statistics

	N	Mean	St. Dev.	Min	Median	Pctl(75)	Pctl(95)
SPE	439079	0.59	0.24	0.00	0.59	0.75	1.00
Team size	439079	7.39	8.77	4	5	7.0	17
Task type	439079	4.69	2.24	2	4	6	9
Lines of code	439079	43558.76	372784.61	0.00	3325	12926	119168
Activities	439079	750.20	2322.29	0.00	381	784	2379
Monthly Stars	439079	44.55	254.21	0.00	3	18	197
Comments	439079	144.98	362.40	0.00	54	145	533
Solving time	300034	51.18	122.35	0.00	13.86	41.36	224.82
Edited files	439079	570.47	2985.87	0.00	123	363	1917
Code acceptance rate	433833	0.82	0.17	0.00	0.86	0.94	1.00
Create year	439079	2018.14	2.82	2011	2018	2020	2022

Notes: This table provides the summary statistics for the main variables of interest at the team-month level. Data is from 2017-01 to 2023-12.

▶ Back

What is star



Version: Free, Pro, & Team ▾

← Home

Get started

Start your journey ▾

Onboarding ▾

Using GitHub ▾

Learning about GitHub ▾

Accessibility ▾

Writing on GitHub ▾

Explore projects ▲

About stars [↗](#)

Starring makes it easy to find a repository or topic again later. You can see all the repositories and topics you have starred by going to your [stars page](#).

You can star repositories and topics to discover similar projects on GitHub. When you star repositories or topics, GitHub may recommend related content on your personal dashboard. For more information, see "[Finding ways to contribute to open source on GitHub](#)" and "[About your personal dashboard](#)."

Starring a repository also shows appreciation to the repository maintainer for their work. Many of GitHub's repository rankings depend on the number of stars a repository has. In addition, [Explore GitHub](#) shows popular repositories based on the number of stars they have.

Viewing who has starred a repository [↗](#)

You can view everyone who has starred a public repository or a private repository you have access to.

- **Show appreciation**
- **GitHub will recommend related content on your dashboard**
- **Many Github repo rankings depend on stars**
- **Not anonymized**

▶ Back

Buying GitHub Stars

How Much Are GitHub Stars Worth to You?

Wednesday, May 31st 2023



Yassin Eldeeb



Aleksandra Sikora

open-source

OPEN SOURCE

How Much Are GitHub Stars Worth to You?



Announcing GraphQL Hive, the complete GraphQL API manager

The best and most obvious way to judge an open-source project is to look at the code but this can be kind of tedious and sometimes you don't like what you see there, so an alternative that we have all naturally developed on our own or have been advised to, is to see how many people have starred a project, and then pick the one with the most stars.

"For example, React.js has 207K stars compared to Angular's measly 88K stars, so we can conclude that React.js is a better framework" — Ben Awad

▶ Back

github-stars

Confirmation #MSW38UNLL
Thank you bogey!

Your order is confirmed
You'll receive a confirmation email with your order number shortly.

Customer information

Contact information [REDACTED]

Payment method
Express - €19.90

Billing address
bogey man
unemployed
somewhere on a really big tree
C
5th settlement
55001
Egypt

Need help? Contact us

Continue shopping

Buy GitHub Stars Premium
25 Premium GitHub Stars €19.90

Total
including €3.00 in taxes €19.90

19.9 Euros for 25 stars ~ \$0.88/star

And after a month, they are all gone. GitHub detected and banned them.

75% developers check stars metric before using it

What's in a GitHub Star? Understanding Repository Starring Practices in a Social Coding Platform

Hudson Borges , Marco Tulio Valente 

Show more 

[+](#) Add to Mendeley [🔗](#) Share [📄](#) Cite

<https://doi.org/10.1016/j.jss.2018.09.016>

[Get rights and content](#) 

Highlights

- Developers star repositories mainly to show appreciation or to bookmark projects
- 3 out of 4 developers check the stars metric before using or contributing to projects
- But developers also evaluate other factors, such as code quality and documentation
- Fast growth in the number of stars is often a result of promotion in social sites
- When ranking projects, we should check whether stars are result of active promotion

Manage to merge code

Fix tool call params order #226502

Merged roblourens merged 1 commit into main from robLou/pleased-tuna last week

Conversation 0 Commits 1 Checks 5 Files changed 3

The screenshot shows the activity of a GitHub pull request. At the top, a comment by user 'roblourens' is highlighted with a red box. The comment text is 'And convert tool_use part correctly'. Below the comment, a commit by 'roblourens' is shown with a green checkmark and the hash 'ff1778d'. Further down, another commit by 'roblourens' is highlighted with a red box, showing it was merged into the 'main' branch. The commit message is partially visible as 'roblourens merged commit 6afce14 into main last week'. Other activity includes a self-assignment, enabling auto-merge, and an approval by 'rebornix'.

Submit code change

Manage to merge

Back

Fail to merge code

Update main.ts to Refactor Startup and Service Initialization #2265

Closed imsharukh1994 wants to merge 1 commit into microsoft:main from imsharukh1994:imsharukh1994-patch-2

Conversation 0 Commits 1 Checks 2 Files changed 1

imsharukh1994 commented last week

Refactored the main.ts file to improve clarity and maintainability. Key changes include:

- Improved error handling and logging throughout the startup process.
- Organized service initialization into a dedicated method for better readability.
- Added more detailed comments and TypeScript typings for enhanced code understanding.
- Implemented more robust environment patching and IPC server handling.

These updates aim to simplify future maintenance and make the code more resilient to errors.

Submit code change

Update main.ts to Refactor Startup and Service Initialization 9d72a5a

Fail to merge

imsharukh1994 force-pushed the imsharukh1994-patch-2 branch from 274246b to 9d72a5a last week [Compare](#)

vs-code-engineering bot added the triage-needed label last week

vs-code-engineering bot assigned ulugbekna last week

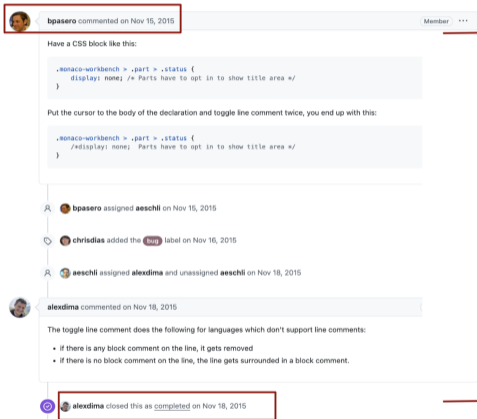
bpasero closed this last week

[← Back](#)

Time to solve problem

CSS: toggle line comment is not preserving #15

 Closed bpasero opened this issue on Nov 15, 2015 · 1 comment



The screenshot shows the following events in chronological order:

- bpasero commented on Nov 15, 2015**: A comment box containing:
 - Text: "Have a CSS block like this:"
 - Code block:

```
.sonaco-workbench > .part > .status {
  display: none; /* Parts have to opt in to show title area */
}
```
 - Text: "Put the cursor to the body of the declaration and toggle line comment twice, you end up with this:"
 - Code block:

```
.sonaco-workbench > .part > .status {
  /*display: none; Parts have to opt in to show title area */
}
```
- bpasero assigned aeschli on Nov 15, 2015**
- chrisdias added the bug label on Nov 16, 2015**
- aeschli assigned alexdima and unassigned aeschli on Nov 18, 2015**
- alexdima commented on Nov 18, 2015**: A comment box containing:
 - Text: "The toggle line comment does the following for languages which don't support line comments:"
 - List:
 - if there is any block comment on the line, it gets removed
 - if there is no block comment on the line, the line gets surrounded in a block comment.
- alexdima closed this as completed on Nov 18, 2015**

Open time

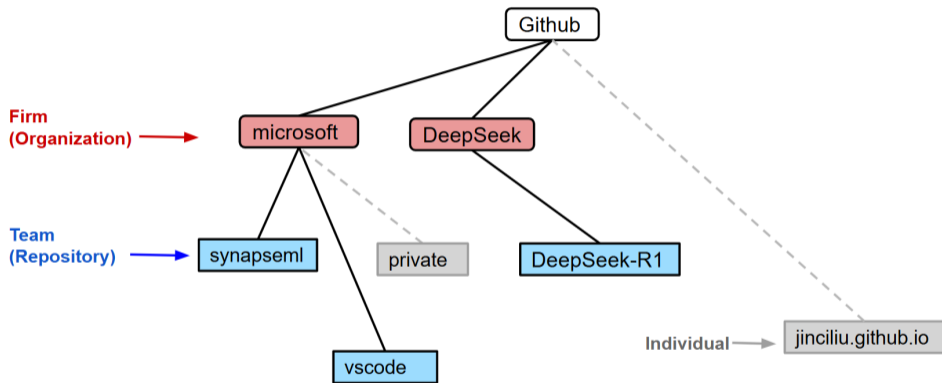
Solving time

Close time

◀ Back

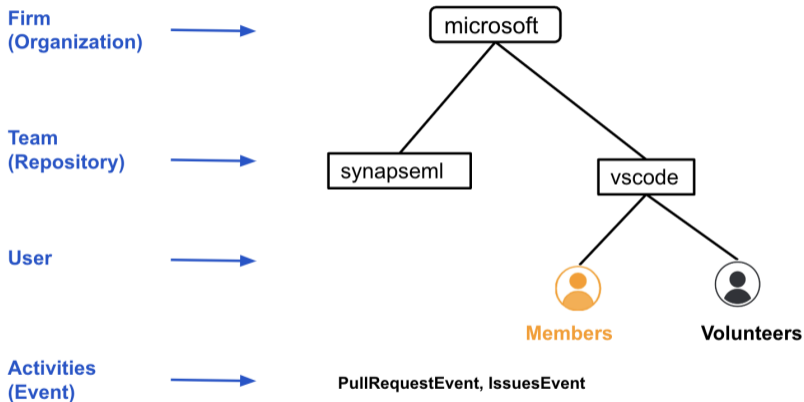
What is GitHub?

- ▶ World's largest open source platform for software development



- ▶ Focus on public teams under firms [◀ Back](#)

Team Member



▶ Task allocation within team members [◀ Back](#)

Code File

The screenshot displays a code repository interface with three main panels:

- File Explorer (Left):** Shows a directory structure with folders like `.configurations`, `.devcontainer`, `.github`, `.vscode`, and `build`. The `build` folder is expanded to show `azure-pipelines`, which contains the `alpine` folder. The file `cli-build-alpine.yml` is highlighted with a red box.
- Commit History (Middle):** Shows a commit by `joaomoreno` committed on Apr 2. Below it, a list of changed files is shown, with `build/azure-pipelines/alpine/cli-build-alpine.yml` highlighted in a red box.
- Code Diff (Right):** Shows the diff for the selected file. The diff includes a commit hash `00-1+1 00` and a list of parameters and steps. The `steps` section is highlighted in a red box.

Annotations:

- A red arrow points from the `cli-build-alpine.yml` file in the file explorer to the commit history view.
- A red arrow points from the commit history view to the `author` label.
- A red arrow points from the `steps` section of the code diff to the `code` label.

Navigation:

- A blue button labeled `Back` is located on the right side of the interface.

Task Type

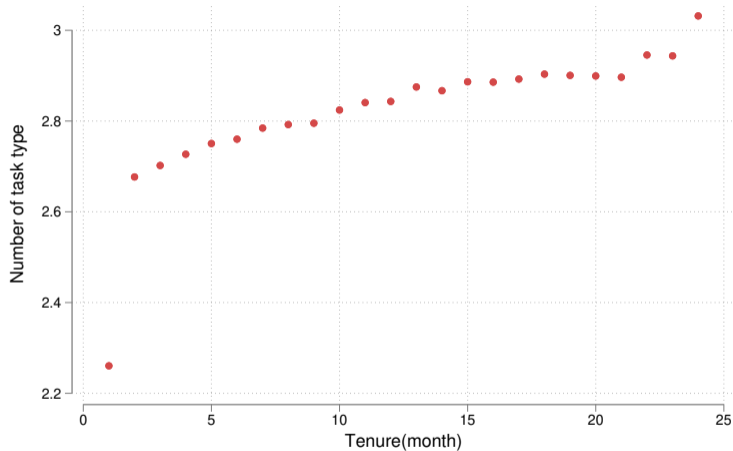
Use Latent Dirichlet allocation (LDA) to classify code files into 10 task types

	Task type	Key words (LDA result)	Occupational role
1	Frontend development	Frontend, UI	Front-end engineer
2	Server management, Platform migration	Client, server	DevOps engineer
3	Android mobile development	Kotlin, runtime	Mobile engineer
4	Cloud feature implementation	Feature, sdk	Cloud engineer
5	Data management	Data, web	Data engineer
6	Internal system management	Internal, apache	System Administrator
7	CLI Development and Framework	User, cli	Technical Writer
8	API and Backend Services	API, controller	Back-end engineer
9	Integration system	Integration, function	System Integrator
10	App Development and UI Design	App, style	App Developer

▶ Word Cloud10

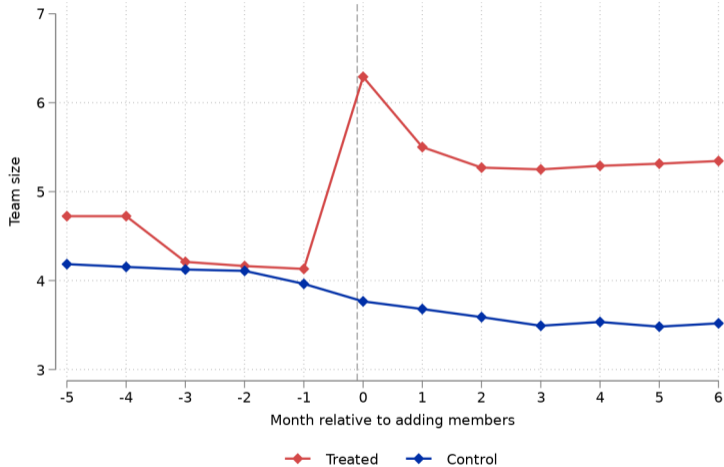
▶ Data Summary

New members start with fewer tasks

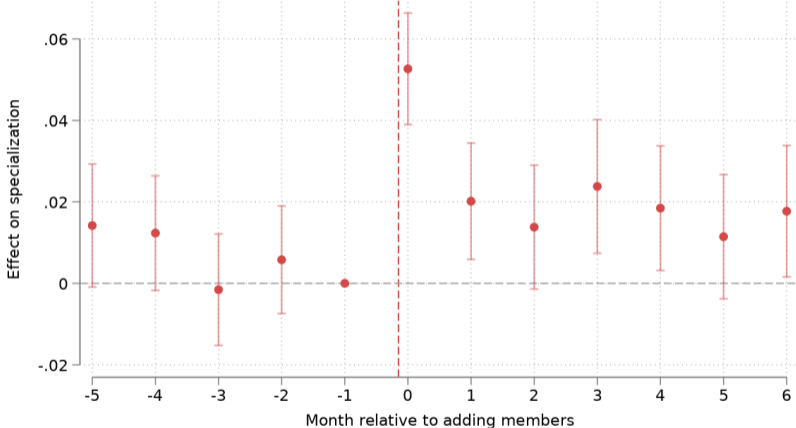


Fixed effect: team, team size

Adding New Members

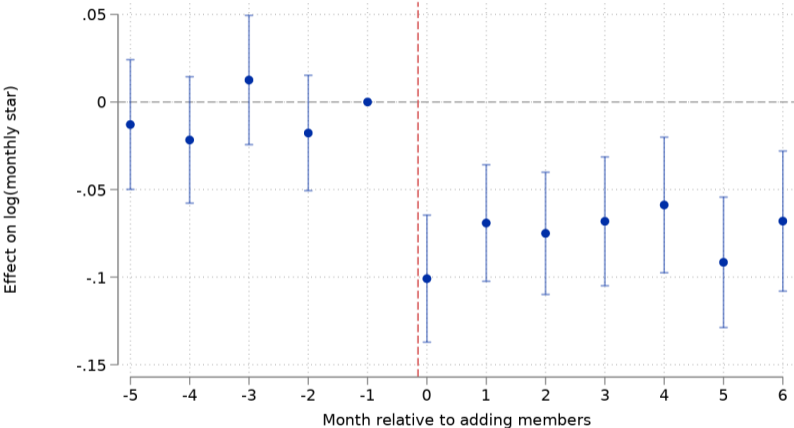


Specialization Increased by 5% in the first month



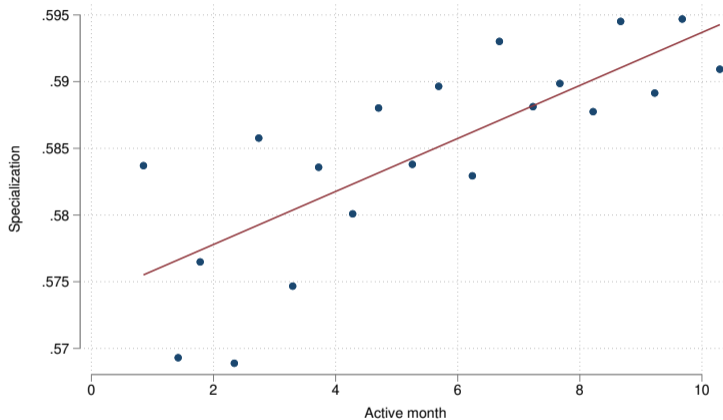
▶ Back

Output Quality Decreased by 10%



▶ Back

Teams are slightly more specialized over time



Teams active at least 10 months
Controls: task distribution
Fixed effect: Team size

References I

- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.
- Becker, G. S. and Murphy, K. M. (1992). The division of labor, coordination costs, and knowledge. *The Quarterly journal of economics*, 107(4):1137–1160.
- Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *The Quarterly journal of economics*, 128(1):1–51.
- Bloom, N., Garicano, L., Sadun, R., and Van Reenen, J. (2014). The distinct effects of information technology and communication technology on firm organization. *Management Science*, 60(12):2859–2885.
- Borges, H. and Valente, M. T. (2018). What's in a github star? understanding repository starring practices in a social coding platform. *Journal of Systems and Software*, 146:112–129.
- Casalnuovo, C., Vasilescu, B., Devanbu, P., and Filkov, V. (2015). Developer onboarding in github: the role of prior social links and language experience. In *Proceedings of the 2015 10th joint meeting on foundations of software engineering*, pages 817–828.
- Giroud, X., Matvos, G., Seru, A., and Silva, R. C. (2022). Managing resource (mis) allocation.

References II

- Gong, J. and Png, I. P. (2024). Automation enables specialization: Field evidence. *Management Science*, 70(3):1580–1595.
- Kohlhepp, J. (2024). The inner beauty of firms.
- Levinthal, D. A. and March, J. G. (1993). The myopia of learning. *Strategic management journal*, 14(S2):95–112.
- Rosen, S. (1983). Specialization and human capital. *Journal of Labor Economics*, 1(1):43–49.
- Smith, A. (1776). *The wealth of nations [1776]*, volume 11937. na.
- Vasilescu, B., Yu, Y., Wang, H., Devanbu, P., and Filkov, V. (2015). Quality and productivity outcomes relating to continuous integration in github. In *Proceedings of the 2015 10th joint meeting on foundations of software engineering*, pages 805–816.
- Wagner, S. and Ruhe, M. (2018). A systematic review of productivity factors in software development. *arXiv preprint arXiv:1801.06475*.
- Young, A. (1928). Increasing returns and economic progress. *The economic journal*, 38(152):527–542.