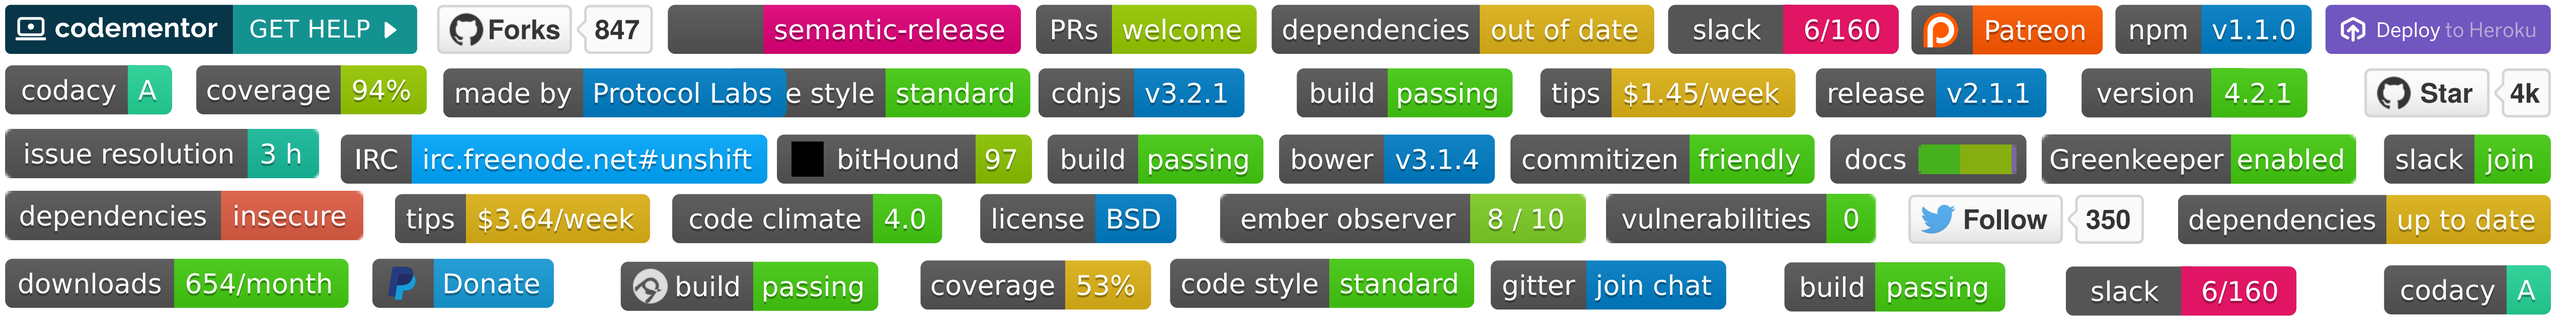

Sixth Dutch National Symposium on Software Engineering

Signals in Social Coding Environments:
What Do They Mean and How Much Can You Trust Them?

Bogdan Vasilescu
@b_vasilescu

w/ Asher Trockman, Shurui Zhou, Christian Kästner

STRUDEL
SOCIO-TECHNICAL RESEARCH
USING DATA EXCAVATION LAB

Carnegie Mellon University

isr institute for SOFTWARE RESEARCH

First: Some bad news

STRUDEL
Carnegie Mellon University

Social Coding: Code is meant to be shared

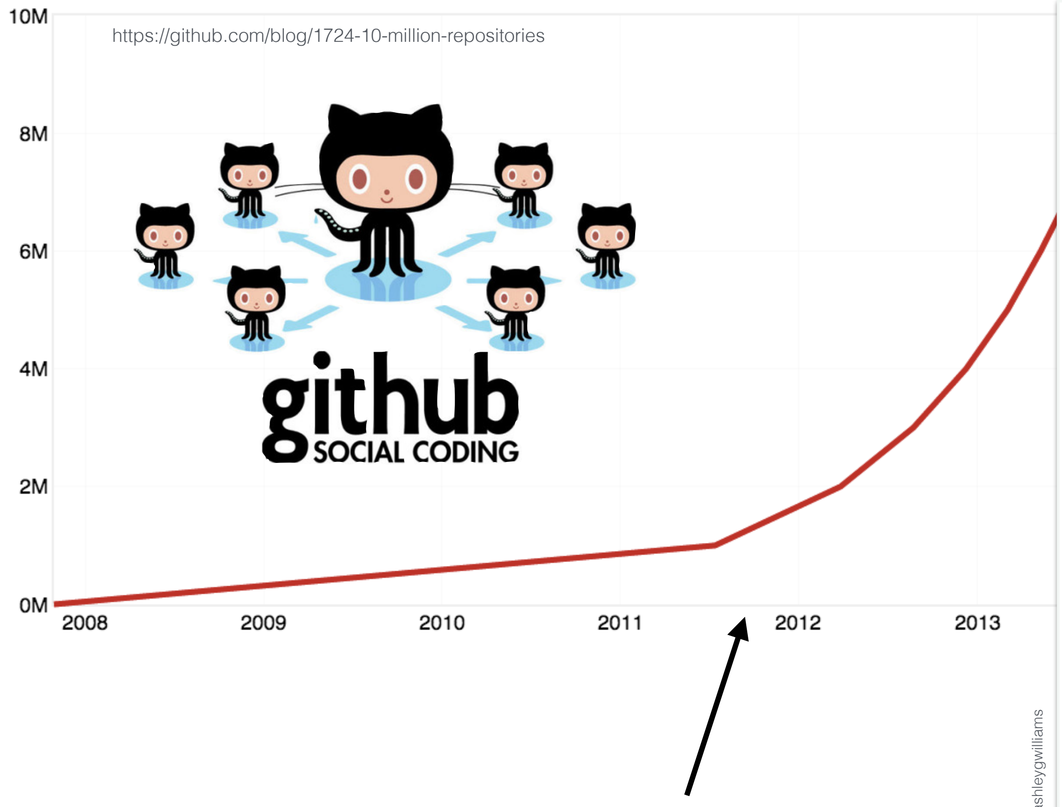

I started my PhD

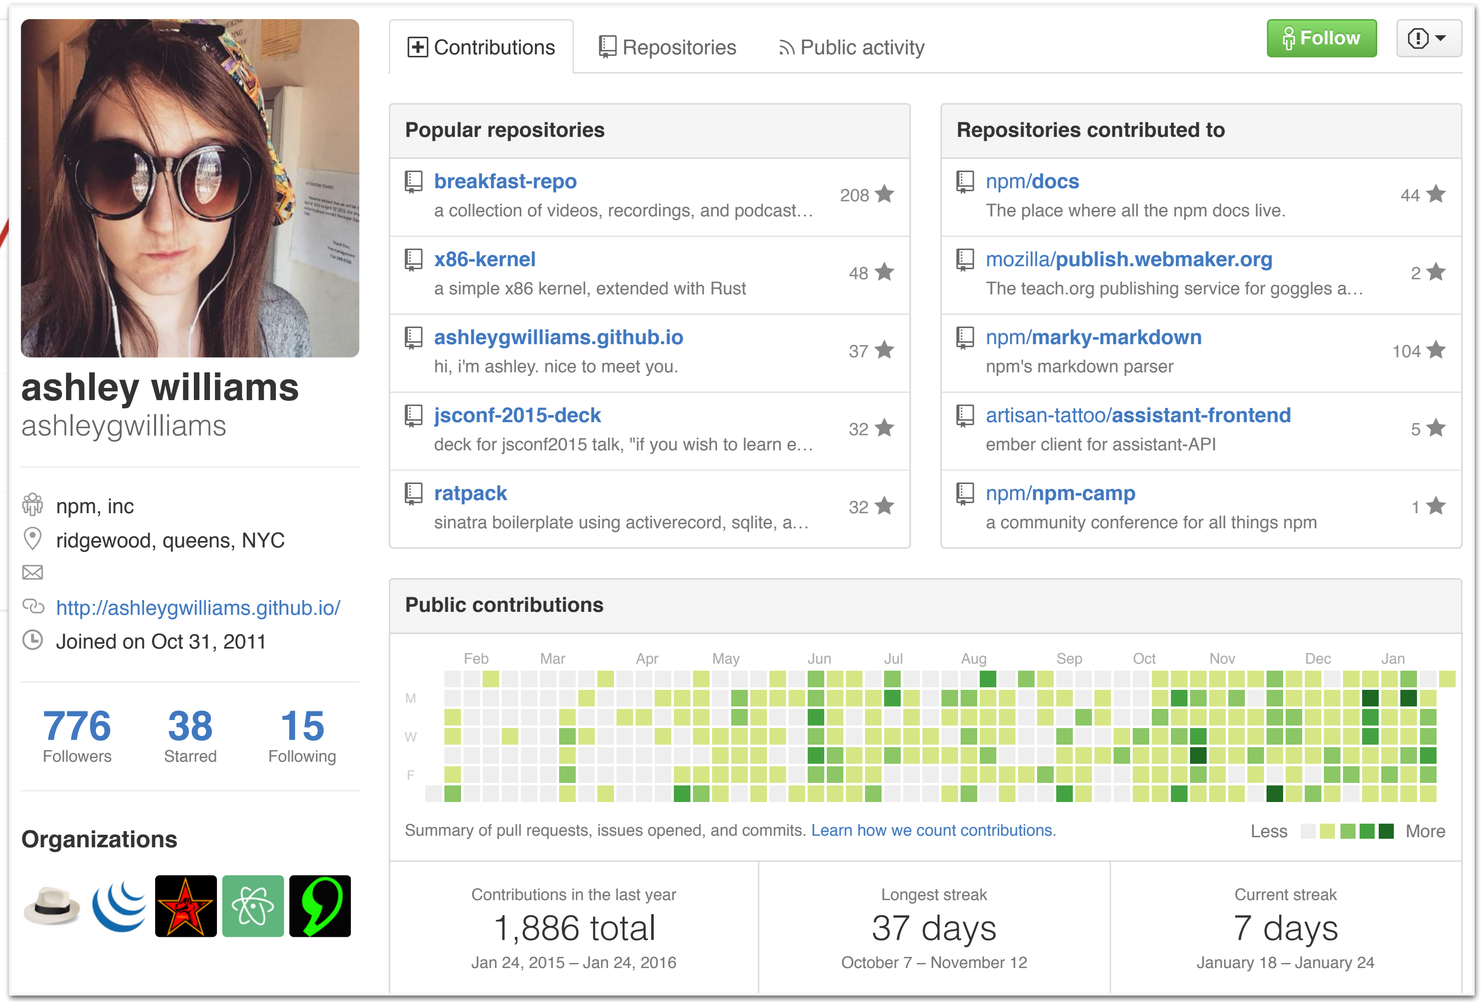

STRUDEL
Carnegie Mellon University

- Programming in a socially networked world: the evolution of the social programmer
  C Treude, F Figueira Filho, B Cleary, MA Storey.
  FutureCSD-CSCW 2012
- Social coding in GitHub: transparency and collaboration in an open software repository
  L Dabbish, C Stuart, J Tsay, J Herbsleb.
  CSCW 2012
- Social networking meets software development: Perspectives from GitHub, MSDN, Stack Exchange, and TopCoder
  A Begel, J Bosch, MA Storey.
  IEEE Software 2013

Knowledge about

Social Coding: Code is meant to be shared

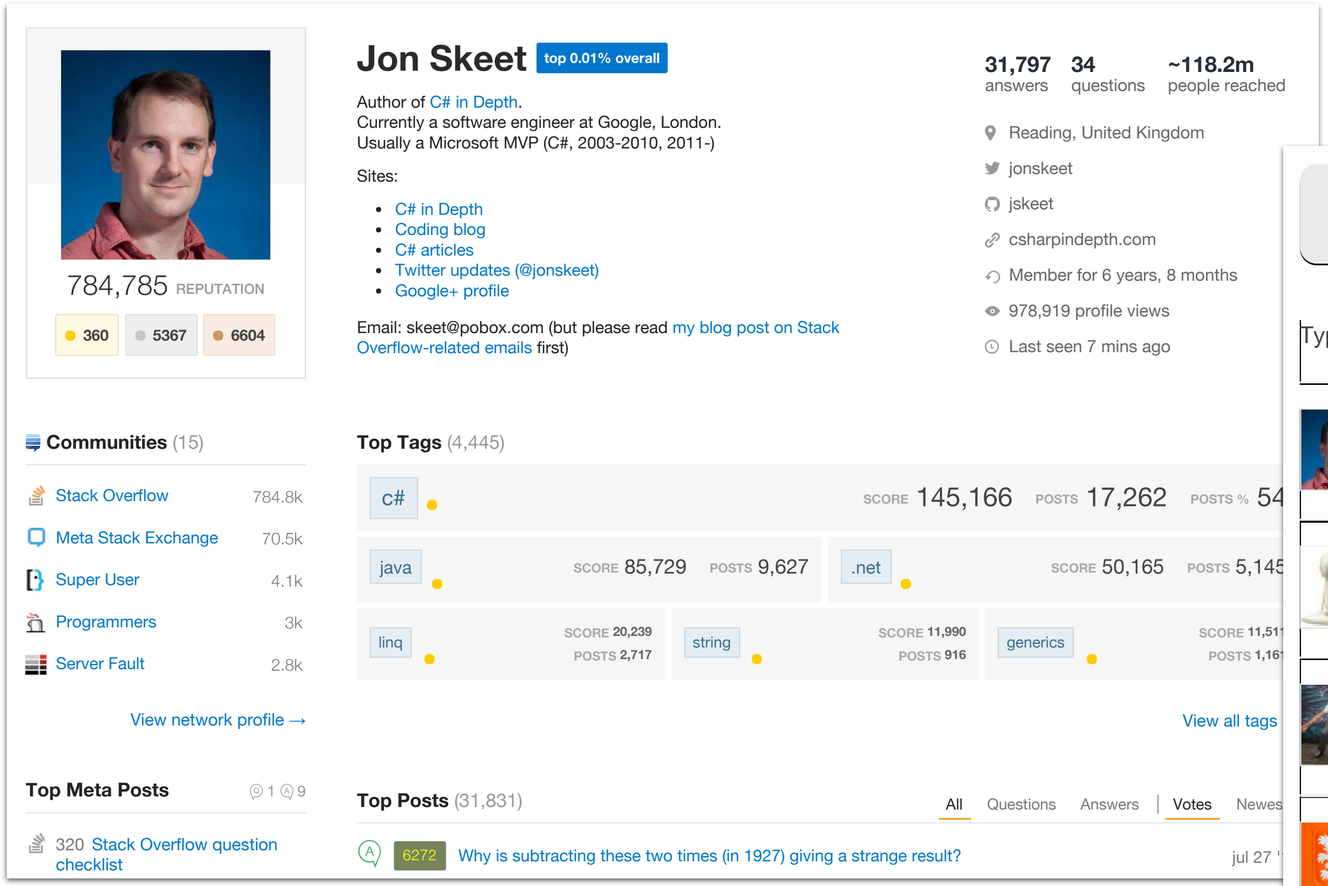

Jon Skeet top 0.01% overall

Author of C# in Depth.
Currently a software engineer at Google, London.
Usually a Microsoft MVP (C#, 2003-2010, 2011-)

Sites:

- C# in Depth
- Coding blog
- C# articles
- Twitter updates (@jonskeet)
- Google+ profile

Email: skeet@pobox.com (but please read my blog post on Stack Overflow-related emails first)

31,797 answers · 34 questions · ~118.2m people reached

Reading, United Kingdom
jonskeet
jskeet
csharpindepth.com
Member for 6 years, 8 months
978,919 profile views
Last seen 7 mins ago

784,785 REPUTATION

360 · 5367 · 6604

Communities (15)

Stack Overflow	784.8k
Meta Stack Exchange	70.5k
Super User	4.1k
Programmers	3k
Server Fault	2.8k

View network profile →

Top Tags (4,445)

c#	SCORE 145,166	POSTS 17,262	POSTS % 54		
java	SCORE 85,729	POSTS 9,627	.net	SCORE 50,165	POSTS 5,145
linq	SCORE 20,239	POSTS 2,717	string	SCORE 11,990	POSTS 916
generics	SCORE 11,511	POSTS 1,161			

View all tags

Top Meta Posts 1 9

320 Stack Overflow question checklist

Top Posts (31,831) All · Questions · Answers · Votes · Newes

6272 Why is subtracting these two times (in 1927) giving a strange result? jul 27

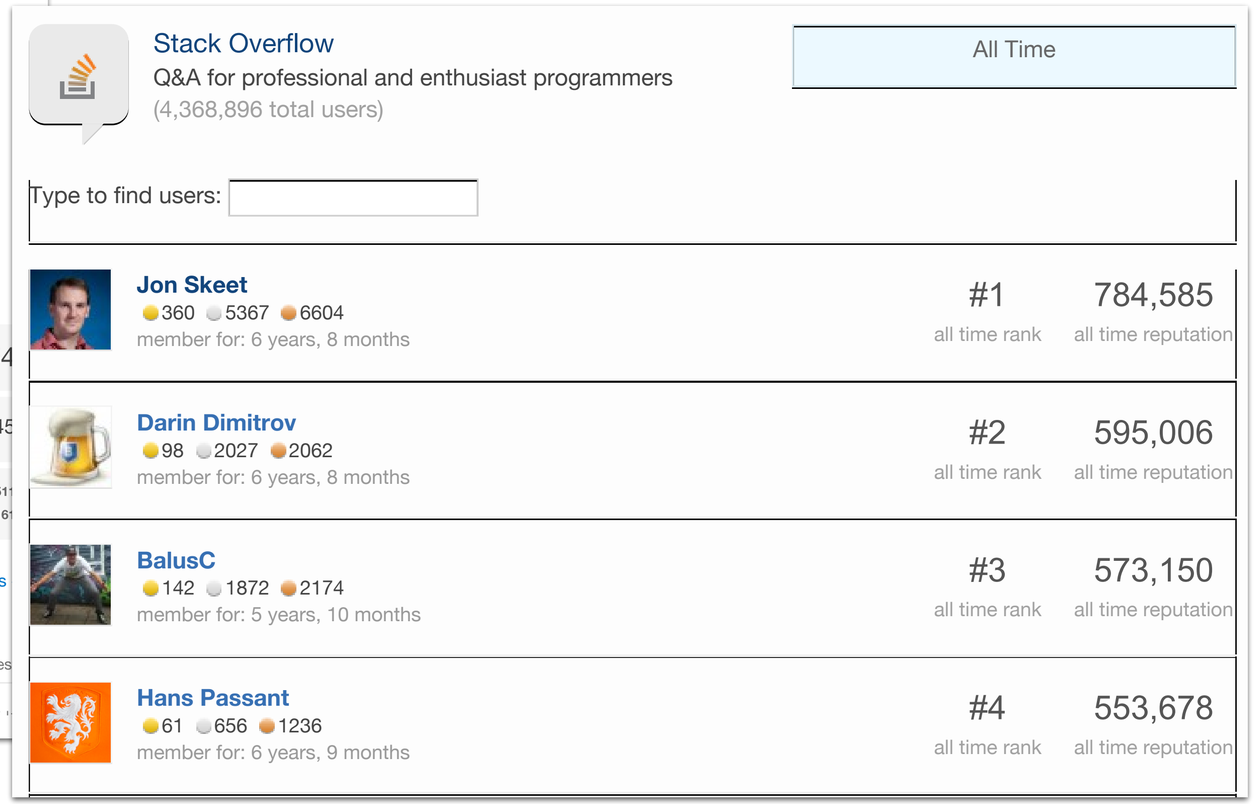

Stack Overflow
Q&A for professional and enthusiast programmers
(4,368,896 total users)

All Time

Type to find users:

Jon Skeet
360 · 5367 · 6604
member for: 6 years, 8 months
#1 all time rank · 784,585 all time reputation

Darin Dimitrov
98 · 2027 · 2062
member for: 6 years, 8 months
#2 all time rank · 595,006 all time reputation

BalusC
142 · 1872 · 2174
member for: 5 years, 10 months
#3 all time rank · 573,150 all time reputation

Hans Passant
61 · 656 · 1236
member for: 6 years, 9 months
#4 all time rank · 553,678 all time reputation

STRUDEL
Carnegie Mellon University

- Programming in a socially networked world: the evolution of the social programmer
  C Treude, F Figueira Filho, B Cleary, MA Storey.
  FutureCSD-CSCW 2012
- Social coding in GitHub: transparency and collaboration in an open software repository
  L Dabbish, C Stuart, J Tsay, J Herbsleb.
  CSCW 2012
- Social networking meets software development: Perspectives from GitHub, MSDN, Stack Exchange, and TopCoder
  A Begel, J Bosch, MA Storey.
  IEEE Software 2013

Knowledge about Social Coding: Code is meant to be shared

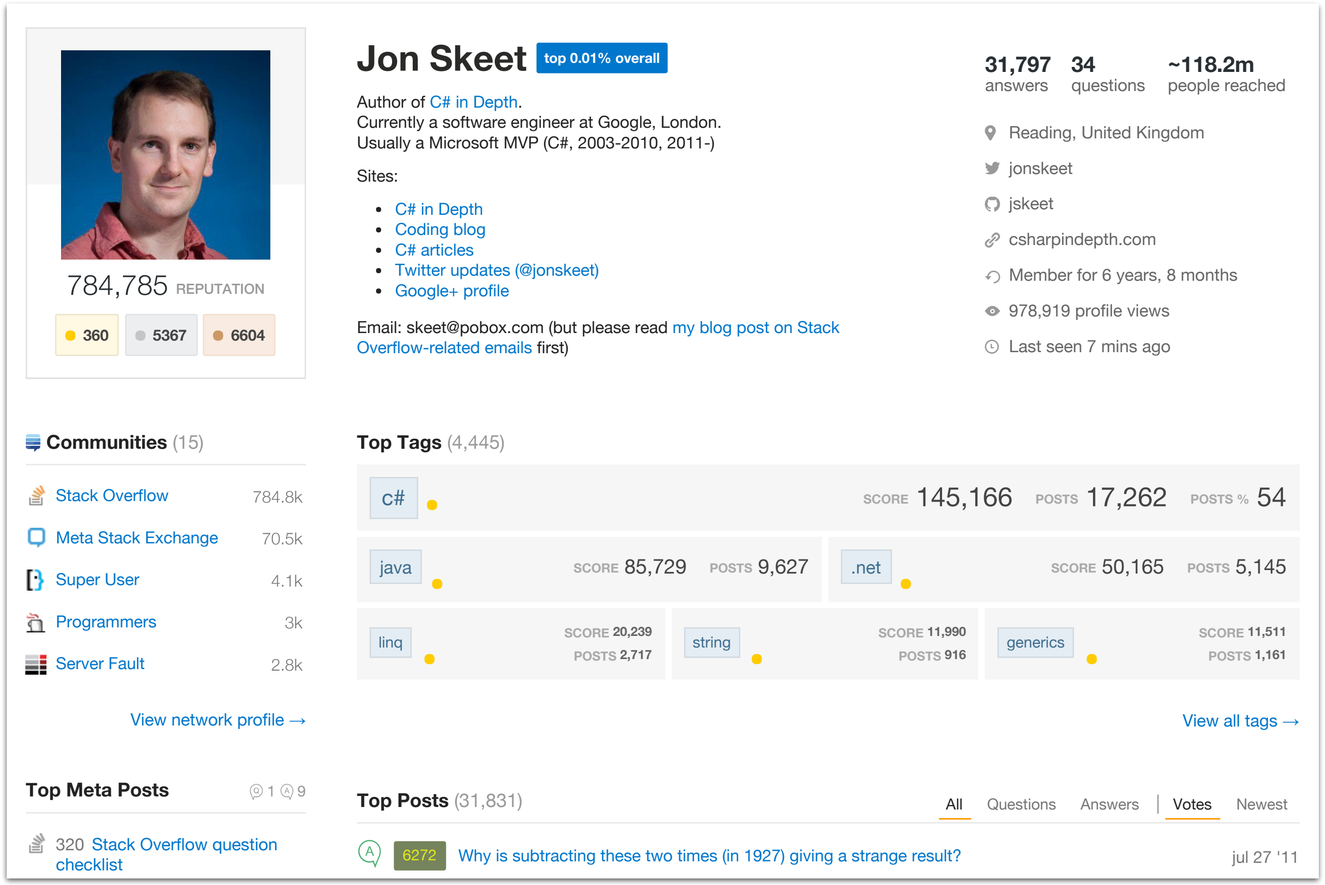

Jon Skeet top 0.01% overall

Author of C# in Depth.
Currently a software engineer at Google, London.
Usually a Microsoft MVP (C#, 2003-2010, 2011-)

Sites:

- C# in Depth
- Coding blog
- C# articles
- Twitter updates (@jonskeet)
- Google+ profile

Email: skeet@pobox.com (but please read my blog post on Stack Overflow-related emails first)

784,785 REPUTATION

360 5367 6604

31,797 answers 34 questions ~118.2m people reached

Reading, United Kingdom
jonskeet
jskeet
csharpindepth.com
Member for 6 years, 8 months
978,919 profile views
Last seen 7 mins ago

Communities (15)

Stack Overflow 784.8k
Meta Stack Exchange 70.5k
Super User 4.1k
Programmers 3k
Server Fault 2.8k

View network profile →

Top Tags (4,445)

c# SCORE 145,166 POSTS 17,262 POSTS % 54
java SCORE 85,729 POSTS 9,627
.net SCORE 50,165 POSTS 5,145
linq SCORE 20,239 POSTS 2,717
string SCORE 11,990 POSTS 916
generics SCORE 11,511 POSTS 1,161

View all tags →

Top Meta Posts 1 9

320 Stack Overflow question checklist

Top Posts (31,831) All Questions Answers | Votes Newest

6272 Why is subtracting these two times (in 1927) giving a strange result? jul 27 '11

- "Jon Skeet **can divide by zero.**"

- "When Jon Skeet's code fails to compile the **compiler apologises.**"

- "Jon Skeet does not use revision control software. **None of his code has ever needed revision.**"

http://meta.stackexchange.com/questions/9134/jon-skeet-facts

STRUDEL
Carnegie Mellon University

- Programming in a socially networked world: the evolution of the social programmer
C Treude, F Figueira Filho, B Cleary, MA Storey.
FutureCSD-CSCW 2012

- Social coding in GitHub: transparency and collaboration in an open software repository
L Dabbish, C Stuart, J Tsay, J Herbsleb.
CSCW 2012

- Social networking meets software development: Perspectives from GitHub, MSDN, Stack Exchange, and TopCoder
A Begel, J Bosch, MA Storey.
IEEE Software 2013

The “social programmer”

… signals

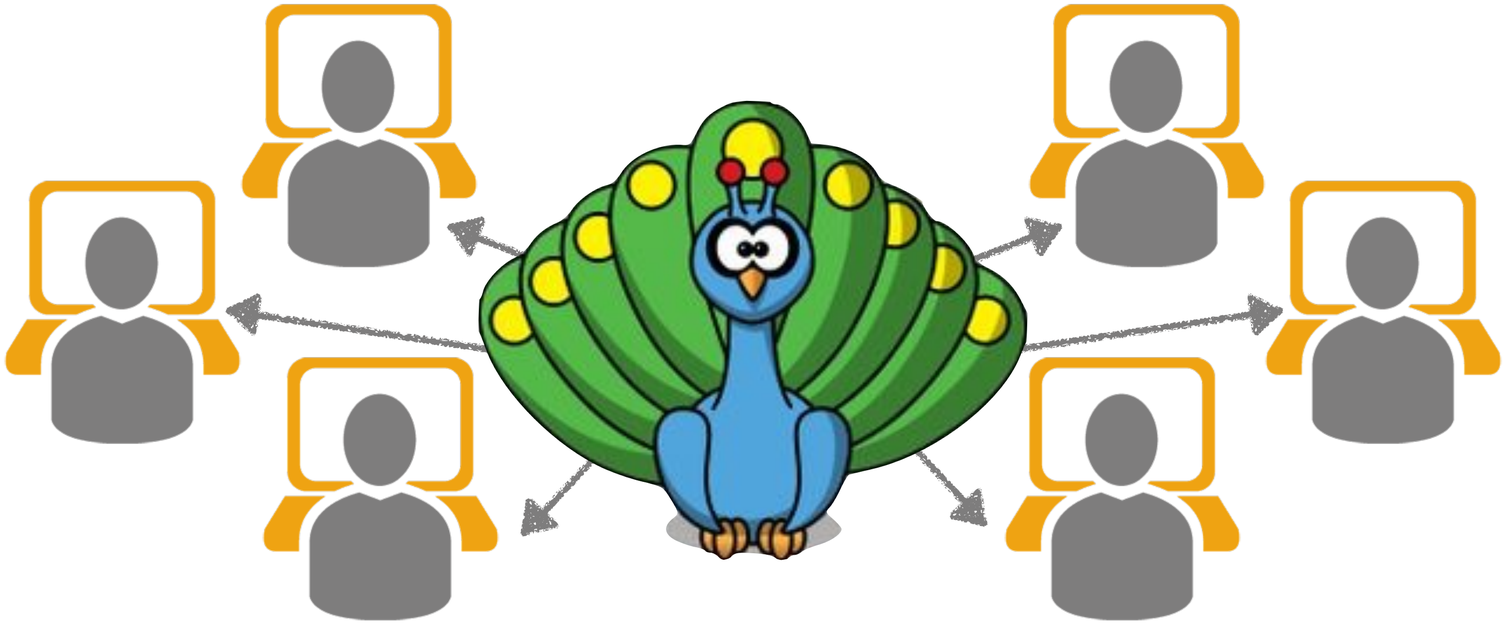

- Assessing technical candidates on the social web
  A Capiluppi, A Serebrenik, L Singer. *IEEE Software 2013*
- Mutual assessment in the social programmer ecosystem: an empirical investigation of developer profile aggregators
  L Singer, F Figueira Filho, B Cleary, C Treude, MA Storey, K Schneider. *CSCW 2013*
- Impression formation in online peer production: activity traces and personal profiles in GitHub
  J Marlow, L Dabbish, J Herbsleb. *CSCW 2013*
- Activity traces and signals in software developer recruitment and hiring
  J Marlow, L Dabbish. *CSCW 2013*

OPEN-SOURCE IS GROWING

BlackDuck 2015 survey
https://goo.gl/Ltaqqs

Companies:
- 78% run OSS
- 66% build on OSS

"Digital dark matter"

Apache httpd server: $7-10b

SOCIAL CODING IS GROWING

27 million people

76 million repositories

22 million software dev's (2017)

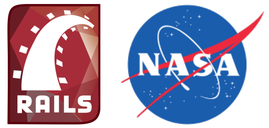

15,000+ people

https://news.ycombinator.com/item?id=2982987

CULTURE CHANGE

"it's just so uncool not sharing the code in the age of social coding"

HIRING

- **$100+** /hour:
  - owns popular OSS products;
  - stack**overflow** score > 20K; …
- **$50+** /hour:
  - active OSS contributor;
  - stack**overflow** score > 5K; …

INDUSTRIAL INVOLVEMENT & ADOPTION

Microsoft
Open source, from Microsoft with love
Redmond, WA http://www.microsoft.com…

Google
https://developers.google.com/

Facebook
We work hard to contribute our work back to the web, mobile, big data, & infrastructure communities.
Menlo Park, California https://code.facebook.com/projects/

STRUDEL
Carnegie Mellon University

- GitHub stats from: https://github.com/about
- Open source-style collaborative development practices in commercial projects using GitHub
  E Kalliamvakou, D Damian, K Blincoe, L Singer, DM German. *ICSE 2015*

- World estimates from: https://evansdata.com/reports/viewRelease.php?reportID=9

- How Much Do You Cost? Yegor Bugayenko http://goo.gl/N0mL3F
- Activity traces and signals in software developer recruitment and hiring
  J Marlow, L Dabbish. *CSCW 2013*

Transparency

Projects

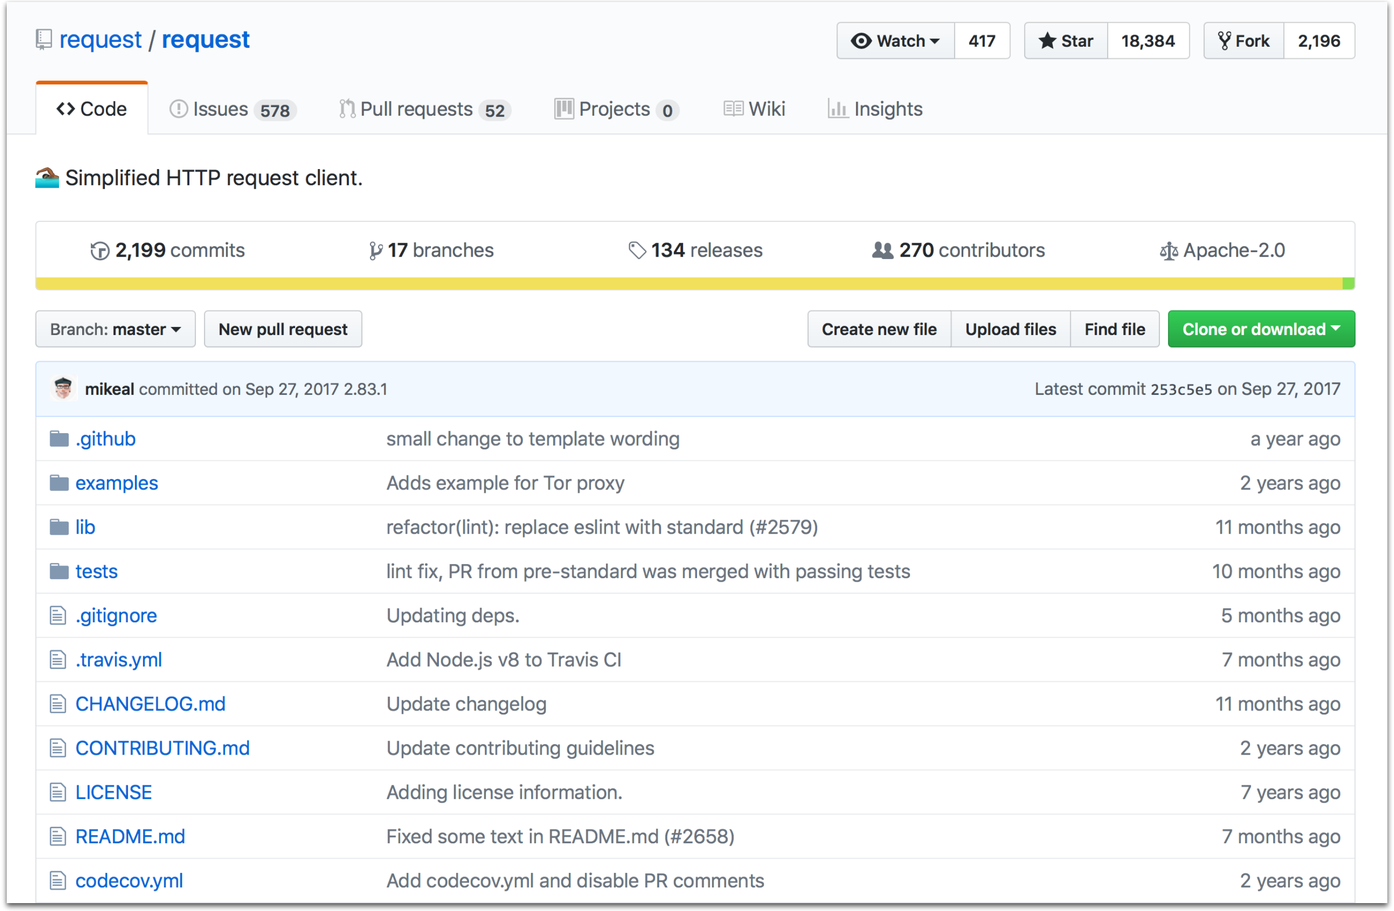

request / request

Watch 417 | Star 18,384 | Fork 2,196

Code | Issues 578 | Pull requests 52 | Projects 0 | Wiki | Insights

Simplified HTTP request client.

2,199 commits | 17 branches | 134 releases | 270 contributors | Apache-2.0

Branch: master | New pull request | Create new file | Upload files | Find file | Clone or download

mikeal committed on Sep 27, 2017 2.83.1 — Latest commit 253c5e5 on Sep 27, 2017

.github	small change to template wording	a year ago
examples	Adds example for Tor proxy	2 years ago
lib	refactor(lint): replace eslint with standard (#2579)	11 months ago
tests	lint fix, PR from pre-standard was merged with passing tests	10 months ago
.gitignore	Updating deps.	5 months ago
.travis.yml	Add Node.js v8 to Travis CI	7 months ago
CHANGELOG.md	Update changelog	11 months ago
CONTRIBUTING.md	Update contributing guidelines	2 years ago
LICENSE	Adding license information.	7 years ago
README.md	Fixed some text in README.md (#2658)	7 months ago
codecov.yml	Add codecov.yml and disable PR comments	2 years ago

People

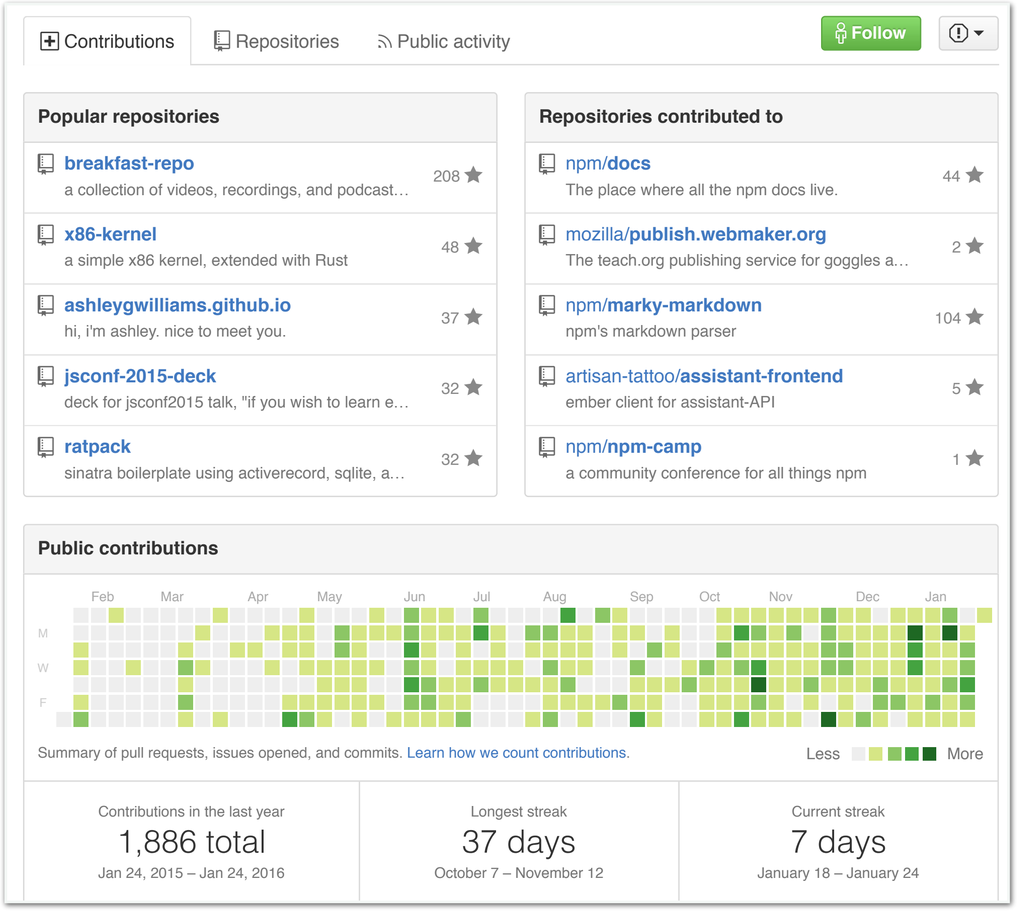

Contributions | Repositories | Public activity — Follow

Popular repositories

breakfast-repo — a collection of videos, recordings, and podcast…	208
x86-kernel — a simple x86 kernel, extended with Rust	48
ashleygwilliams.github.io — hi, i'm ashley. nice to meet you.	37
jsconf-2015-deck — deck for jsconf2015 talk, "if you wish to learn e…	32
ratpack — sinatra boilerplate using activerecord, sqlite, a…	32

Repositories contributed to

npm/docs — The place where all the npm docs live.	44
mozilla/publish.webmaker.org — The teach.org publishing service for goggles a…	2
npm/marky-markdown — npm's markdown parser	104
artisan-tattoo/assistant-frontend — ember client for assistant-API	5
npm/npm-camp — a community conference for all things npm	1

Public contributions

Feb Mar Apr May Jun Jul Aug Sep Oct Nov Dec Jan

M W F

Summary of pull requests, issues opened, and commits. Learn how we count contributions. Less More

Contributions in the last year	Longest streak	Current streak
1,886 total	37 days	7 days
Jan 24, 2015 – Jan 24, 2016	October 7 – November 12	January 18 – January 24

STRUDEL
Carnegie Mellon University

Yet … despite transparency

- Neglected infrastructure
  - “heartbleed” security bug in OpenSSL
  - leftpad incident
- In open source ecosystems, no signal (yet) to help **balance supply and demand for labor**
  - Similar to price in market-based production
  - Could **badges** help?

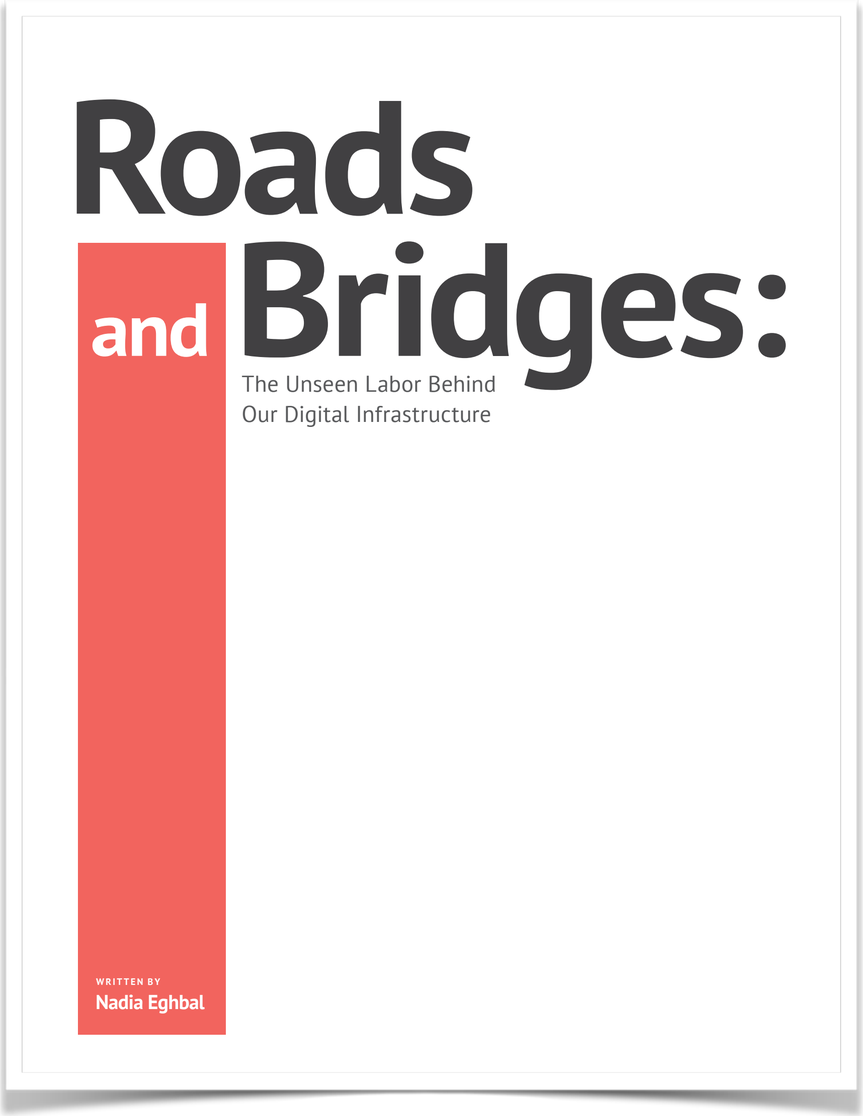

Today: Signaling with badges

codementor GET HELP ▶ | npm v1.1.0 | Donate

build passing | code style standard | tips $3.64/week

ember observer 8 / 10 | dependencies out of date

build passing | license BSD | vulnerabilities 0

bitHound 97 | Star 4k | cdnjs v3.2.1

bower v3.1.4 | downloads 654/month | codacy A

issue resolution 3 h | coverage 53% | Follow 350

slack join | build passing | dependencies insecure

dependencies up to date | IRC irc.freenode.net#unshift

coverage 94% | gitter join chat | code climate 4.0

build passing | commitizen friendly | version 4.2.1

release v2.1.1 | Greenkeeper enabled | docs

Patreon | slack 6/160 | made by Protocol Labs

code style standard | tips $1.45/week | Deploy to Heroku

Forks 847 | semantic-release | PRs welcome

Adding Sparkle to Social Coding: An Empirical Study of Repository Badges in the npm Ecosystem.
Trockman, A., Zhou, S., Kästner, C., and Vasilescu, B.
International Conference on Software Engineering, ICSE, ACM (2018).

Built-in (GitHub)

This repository Search Pull requests Issues Marketplace Explore

request / request

Watch 395 Star 16,836 Fork 2,023

Code Issues 523 Pull requests 40 Projects 0 Wiki Insights

Simplified HTTP request client.

2,190 commits 17 branches 130 releases 273 contributors Apache-2.0

Custom

README.md

Request - Simplified HTTP client

npm install request
22 dependencies version 2.81.0
22,431 dependents updated 6 months ago
1,275 ★

build passing coverage 92% coverage 93% dependencies up to date vulnerabilities 0 gitter join chat

Types of badges

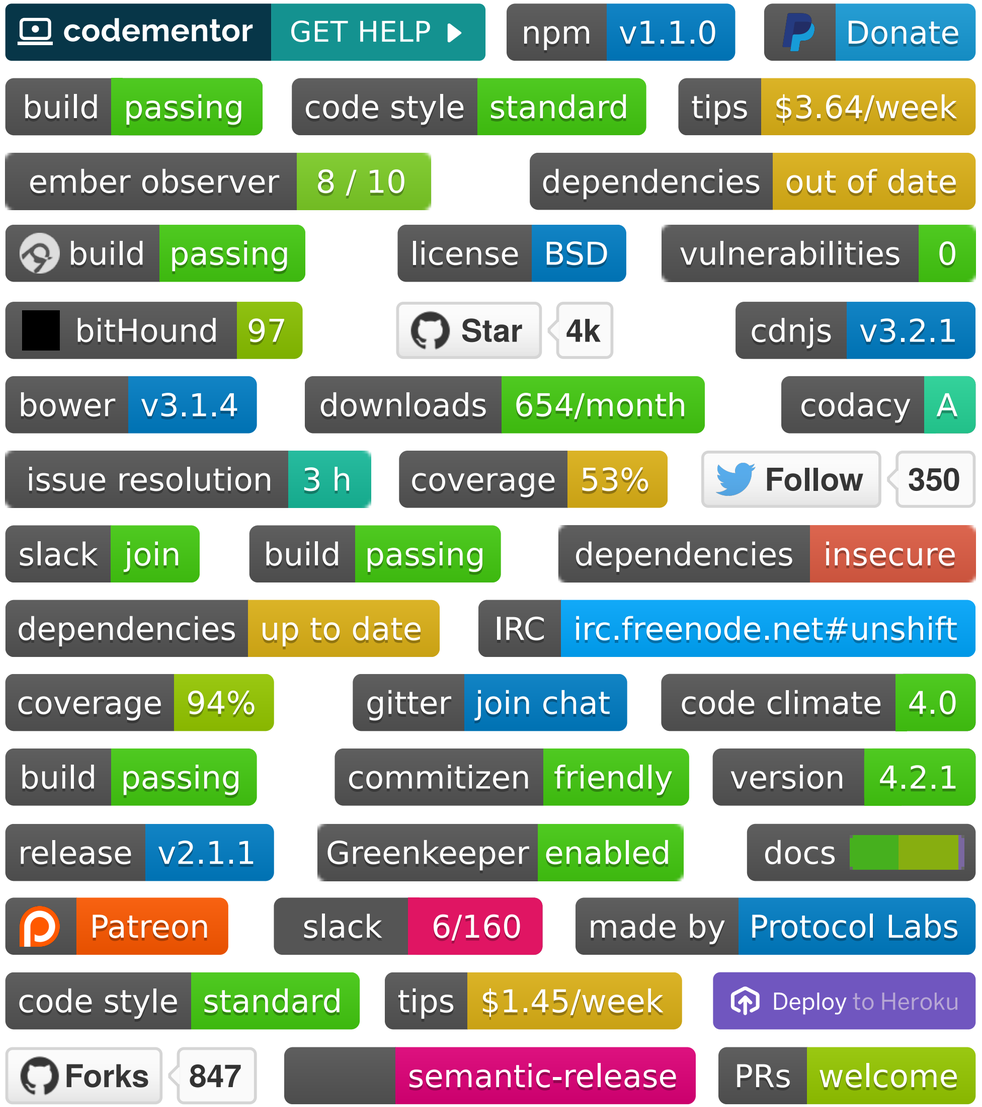

STRUDEL
Carnegie Mellon University

Types of badges

- Quality assurance
  - Build status, test coverage, static analysis, …

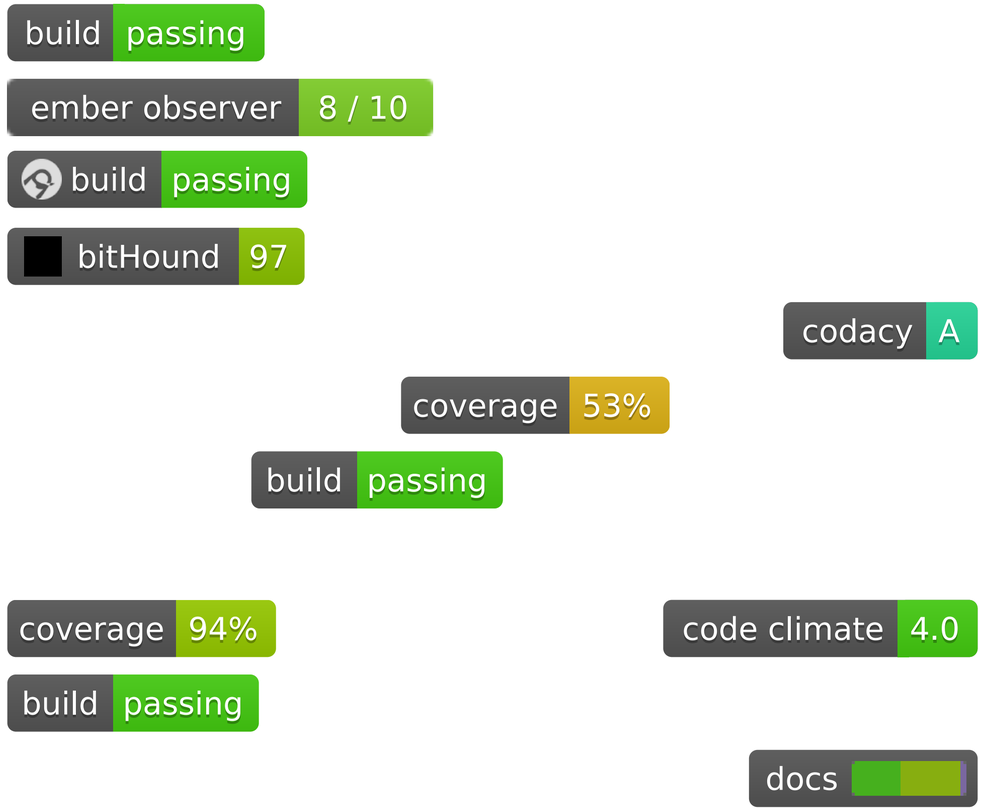

Types of badges

- Quality assurance
  - Build status, test coverage, static analysis, …
- Dependency management
  - Version tracking, vulnerability tracking, …

STRUDEL
Carnegie Mellon University

Types of badges

- Quality assurance
  - Build status, test coverage, static analysis, …
- Dependency management
  - Version tracking, vulnerability tracking, …
- Information
  - *npm* version, license, coding style, release strategy, commit message conventions, …

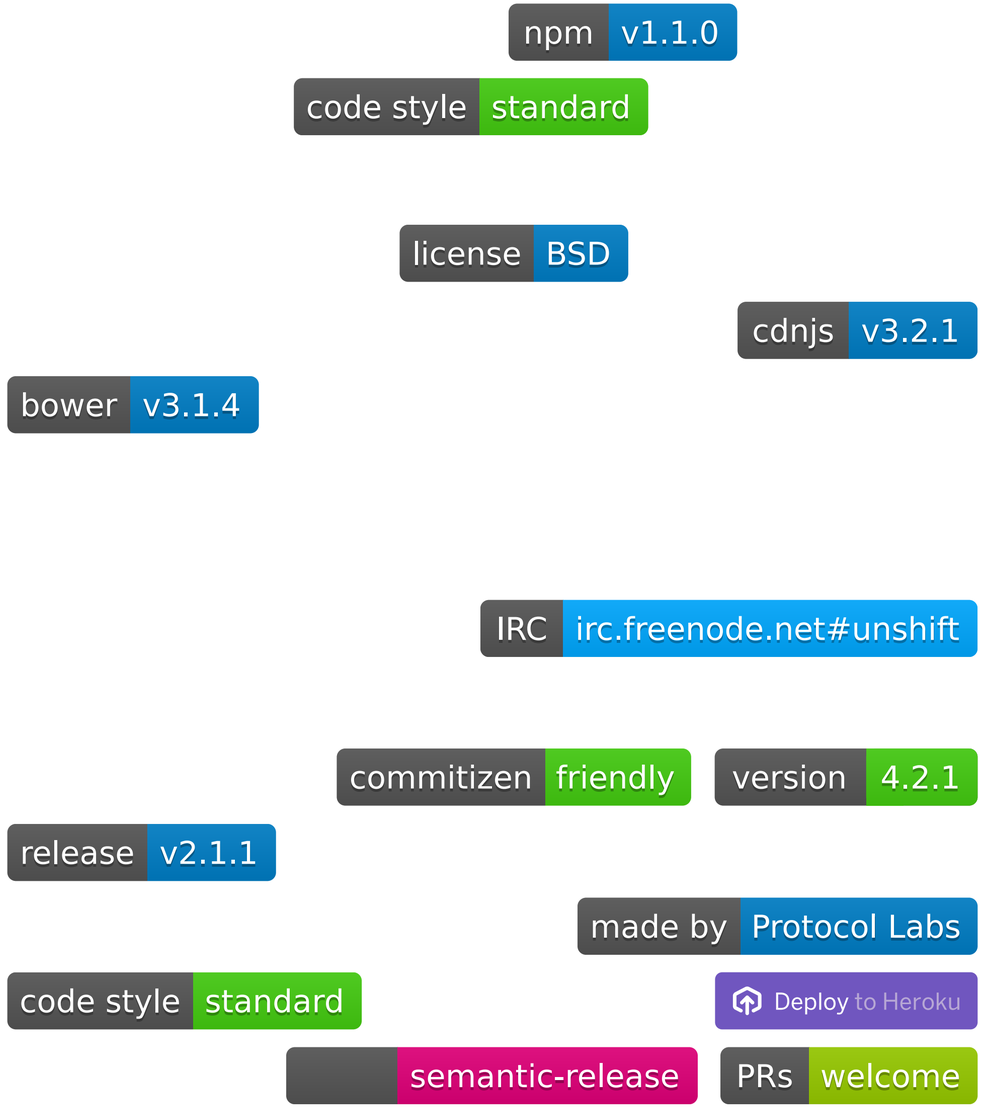

Types of badges

- Quality assurance
  - Build status, test coverage, static analysis, …
- Dependency management
  - Version tracking, vulnerability tracking, …
- Information
  - *npm* version, license, coding style, release strategy, commit message conventions, …
- Popularity
  - *npm* downloads, GitHub stats, Twitter, …

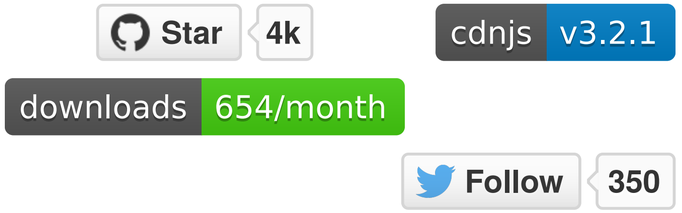

Types of badges

- Quality assurance
  - Build status, test coverage, static analysis, …
- Dependency management
  - Version tracking, vulnerability tracking, …
- Information
  - *npm* version, license, coding style, release strategy, commit message conventions, …
- Popularity
  - *npm* downloads, GitHub stats, Twitter, …
- Support
  - chat & collaboration, issue stats, …

STRUDEL
Carnegie Mellon University

Types of badges

- Quality assurance
  - Build status, test coverage, static analysis, …
- Dependency management
  - Version tracking, vulnerability tracking, …
- Information
  - *npm* version, license, coding style, release strategy, commit message conventions, …
- Popularity
  - *npm* downloads, GitHub stats, Twitter, …
- Support
  - chat & collaboration, issue stats, …
- Misc:
  - Paypal, donations, Gittip, …

STRUDEL
Carnegie Mellon University

Signaling theory (Spence, 1973)

- Badges are signals:
  - reduce **information asymmetry**
- **Conventional signals** vs **assessment signals**
  - **assessment signals**: more costly to produce —> more reliable

- Badges vary widely in production cost
  - Expensive:

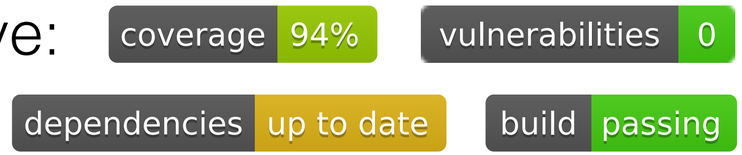

  - Cheap:

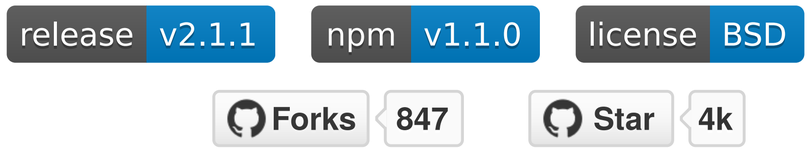

  - No cost:

Can you trust them?

seanmonstar / intel

Watch 10 | Star 191 | Fork 29

Code | Issues 11 | Pull requests 0 | Projects 0 | Wiki | Insights

I need more intel! http://seanmonstar.github.io/intel/

224 commits | 3 branches

Branch: master | New pull request

README.md

intel

build passing | coverage 100% | npm package 1.2.0

An abbreviation of intelligence. In this case, the acquirement of information.

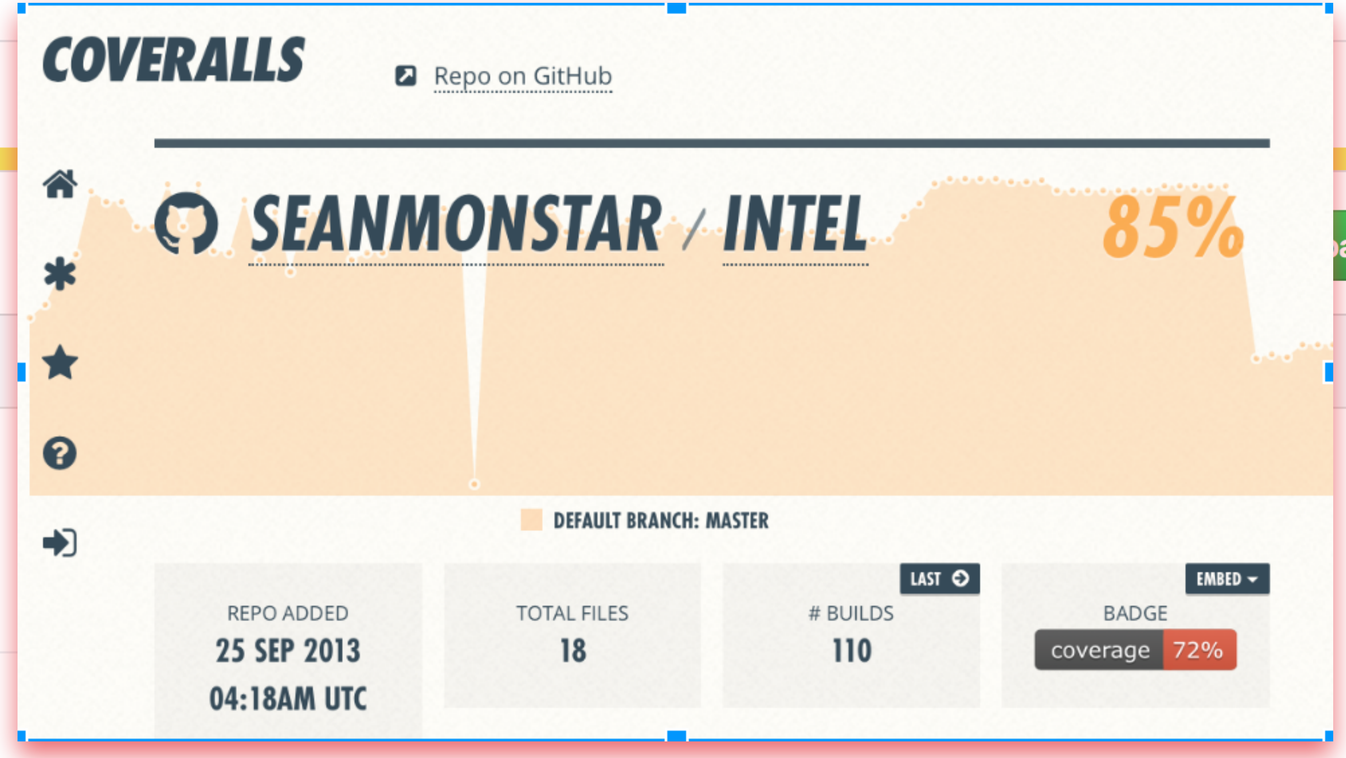

Research questions

- How are badges used?
- What do they tell about a project?
- How much can you trust them?

Adding Sparkle to Social Coding: An Empirical Study of Repository Badges in the npm Ecosystem

Asher Trockman,[†‡] Shurui Zhou,[‡] Christian Kästner,[‡] Bogdan Vasilescu[‡]
[†]University of Evansville, USA [‡]Carnegie Mellon University, USA
atrockma@andrew.cmu.edu, shuruiz@cs.cmu.edu, {kaestner, vasilescu}@cmu.edu

ABSTRACT

In fast-paced, reuse-heavy software development, the transparency provided by social coding platforms like GitHub is essential to decision making. Developers infer the quality of projects using visible cues, known as signals, collected from personal profile and repository pages. We report on a large-scale, mixed-methods empirical study of *npm* packages that explores the emerging phenomenon of repository badges, with which maintainers signal underlying qualities about the project to contributors and users. We investigate which qualities maintainers intend to signal and how well badges correlate with those qualities. After surveying developers, mining 294,941 repositories, and applying statistical modeling and time series analysis techniques, we find that non-trivial badges, which display the build status, test coverage, and up-to-dateness of dependencies, are mostly reliable signals, correlating with more tests, better pull requests, and fresher dependencies. Displaying such badges correlates with best practices, but the effects do not always persist. In short, signals mostly reliable.

1 INTRODUCTION

Contemporary software development is characterized by increased reuse and speed. Open-source software forges like GitHub host millions of repositories of libraries and tools, which developers reuse liberally [25], creating complex, often fragile networks of interdependencies [11]. This has earned GitHub a reputation as a one-stop shop for software development [38], an influencer of practices in both open-source and industry [32]. Furthermore, the widely-adopted DevOps culture [30, 45] also contributes to this acceleration, with its emphasis on automation and rapid deployment. As a result, developers are expected to make more decisions at higher speed, finding which libraries to depend and which projects to contribute to.

A key enabler of this decision making process is the *transparency* provided by social coding platforms like GitHub [20, 21]. The development history of open-source GitHub projects is archived and publicly accessible in a standardized format; and GitHub user profile pages display aggregate information about one's contributions and social standing in the community (*e.g.*, through repository *stars* and *watchers*). This unprecedented level of transparency in social coding can enhance collaboration and coordination [21]. Using visible cues—known in the literature as *signals*—collected from personal profile and repository pages, developers can better manage their projects and their dependencies, communicate more efficiently, become informed about action items requiring their attention, learn, socialize, and form impressions about each other's coding ability, personal characteristics, and interpersonal skills [21, 37, 39, 56].

However, open-source ecosystems are also competitive. In order to survive and thrive, projects must successfully attract and retain contributors, and fend off competitors [16, 35, 40, 44]. In a social coding environment, the visible signals enabled by transparency can, therefore, be seen as a survival mechanism, with high profile signalers benefiting the most. For example, more popular and famous projects attract more contributors [61], coding "rock stars" collect thousands of followers [20], and visible traces of developer actions and interactions are used in recruitment and hiring [13, 36].

Here we focus on *repository badges*, images such as build passing, embedded into a project's README, often generated on-demand, reflecting the current status of online services the project is using, *e.g.*, continuous integration and dependency management. From a *signaling theory* [51] perspective (Section 2), badges can be seen as easily observable signals used by maintainers to convey underlying qualities of their projects to users and contributors, *e.g.*, code quality, adherence to best practices; this increases transparency (hard to observe qualities become salient), hence may impact users' and contributors' decision making process and the project's chances of survival. Badges can also be seen as a *gamification mechanism* [23], *i.e.*, a game-like incentive designed to engage participants (Section 2); *e.g.*, a badge with real-time code coverage information may act as an incentive for contributors to improve the quality of the project's test suite. In summary, badges are a potentially impactful feature in transparent, social coding environments. However, the value and effects of badges are not well understood.

In this paper, we explore two main research questions regarding badges. First, we explore the phenomenon of badges quantitatively and qualitatively, and ask *What are the most common badges and what does displaying them intend to signal?* (**RQ**$_1$). Second, we analyze whether badges indeed signal what developers expect, and ask *To what degree do badges correlate with qualities that developers expect?* (**RQ**$_2$). To answer these questions, we perform a large-scale mixed-methods empirical study of the badges in the *npm* ecosystem, a large and vibrant open-source ecosystem for JavaScript with documented interdependency-related coordination challenges [11], wherein many badges originated. We observe the frequency and historical adoption of badges in practice among 294,941 *npm* packages, we survey maintainers and contributors about their intentions and perceptions, and we build regression models to check hypotheses regarding developer perceptions (collected when exploring **RQ**$_1$), such as, "coverage badges signal the importance of tests and therefore attract more pull requests with tests."

Our investigation reveals that badges are popular in *npm*, adopted in 46 % of all packages. The most frequent show the build status or version of the latest release, but dependency managers, code

Conference'17, July 2017, Washington, DC, USA
2017. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

Adding Sparkle to Social Coding: An Empirical Study of Repository Badges in the npm Ecosystem.
Trockman, A., Zhou, S., Kästner, C., and Vasilescu, B.
International Conference on Software Engineering, ICSE, ACM (2018).

STRUDEL
Carnegie Mellon University

Mixed methods study

+

- 32 maintainers, 57 contributors (15% resp. rate)
- Maintainers:
  - What do you intend to signal?
  - What effects do you expect?
- Contributors:
  - What do badges tell you?

- 294,941 *npm* packages
- Mined badge adoptions/removals from README files
- Measured proxies for code quality, test suite quality, popularity, dependency freshness, …

STRUDEL
Carnegie Mellon University

Popular badges

- Only few badges are broadly adopted
- Badges tend to be adopted in groups and hardly ever change

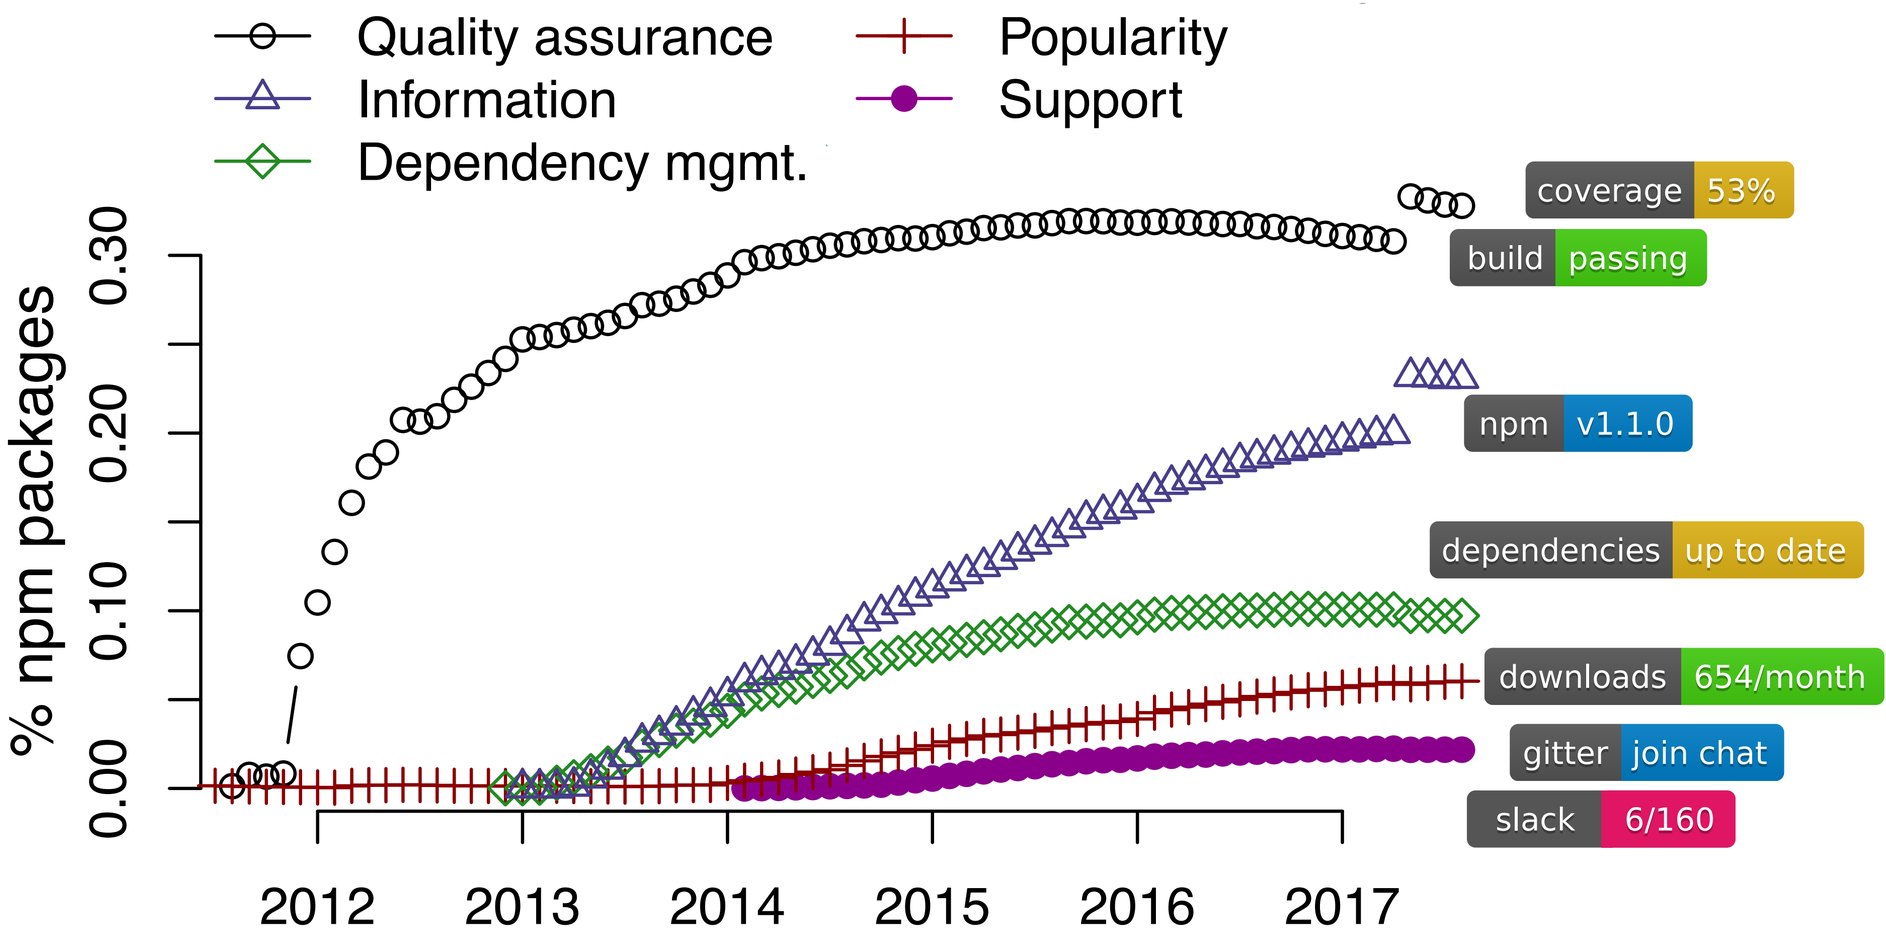

STRUDEL
Carnegie Mellon University

Analysis

If all you saw was the badge, how much would that tell you?

How much more does the badge tell you, relative to existing signals?

How do things change after adding the badge?

STRUDEL
Carnegie Mellon University

Signals of fresh dependencies

- Hyp: The adoption of dependency-management badges correlates with fresher dependencies
- Hyp: Information badges have no effect

- In aggregate: both badges correlate with having the quality

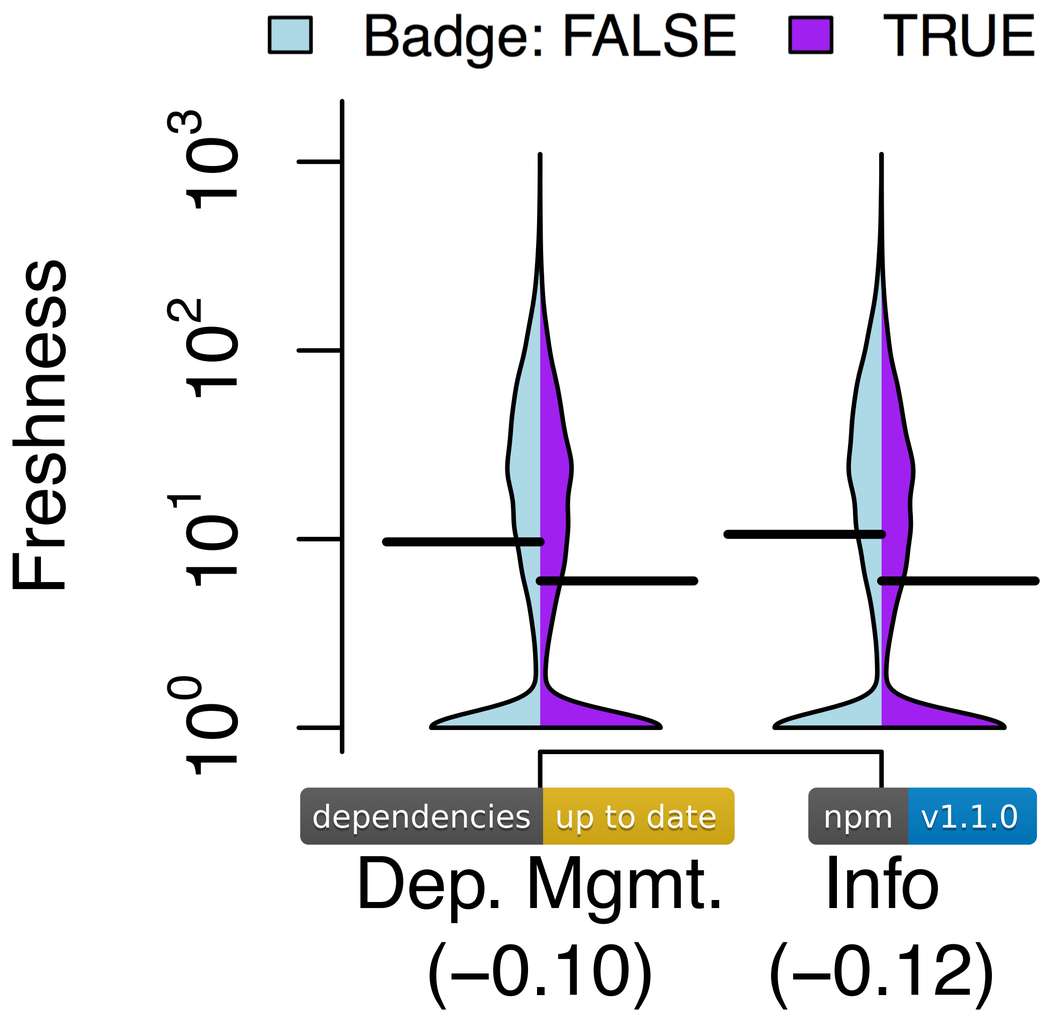

Signals of fresh dependencies

- Hyp: The adoption of dependency-management badges correlates with fresher dependencies
- Hyp: Information badges have no effect

- **Both badges add information beyond other readily observable signals**

	Basic Model response: *freshness = 0* 17.3% deviance explained		**Full Model** response: *freshness = 0* 17.4% deviance explained	
	Coeffs (Err.)	LR Chisq	Coeffs (Err.)	LR Chisq
(Interc.)	$3.54\ (0.03)^{***}$		$3.50\ (0.03)^{***}$	
Dep.	$-1.78\ (0.01)^{***}$	32077.8^{***}	$-1.79\ (0.01)^{***}$	32292.8^{***}
RDep.	$0.22\ (0.01)^{***}$	610.3^{***}	$0.21\ (0.01)^{***}$	560.6^{***}
Stars	$-0.08\ (0.00)^{***}$	301.4^{***}	$-0.09\ (0.00)^{***}$	311.2^{***}
Contr.	$-0.24\ (0.01)^{***}$	500.5^{***}	$-0.25\ (0.01)^{***}$	548.7^{***}
lastU	$-0.65\ (0.01)^{***}$	12080.9^{***}	$-0.64\ (0.01)^{***}$	11537.9^{***}
dependencies: up to date			$0.24\ (0.03)^{***}$	116.1^{***}
npm: v1.1.0			$0.11\ (0.02)^{***}$	48.3^{***}
dependencies: up to date : npm: v1.1.0			$-0.05\ (0.04)$	1.9
hasOther			$0.01\ (0.01)$	

$^{***}p < 0.001, {}^{**}p < 0.01, {}^{*}p < 0.05$

STRUDEL
Carnegie Mellon University

Signals of fresh dependencies

- Hyp: The adoption of dependency-management badges correlates with fresher dependencies
- Hyp: Information badges have no effect

- **The adoption of DM badges is associated with a longer lasting effect**

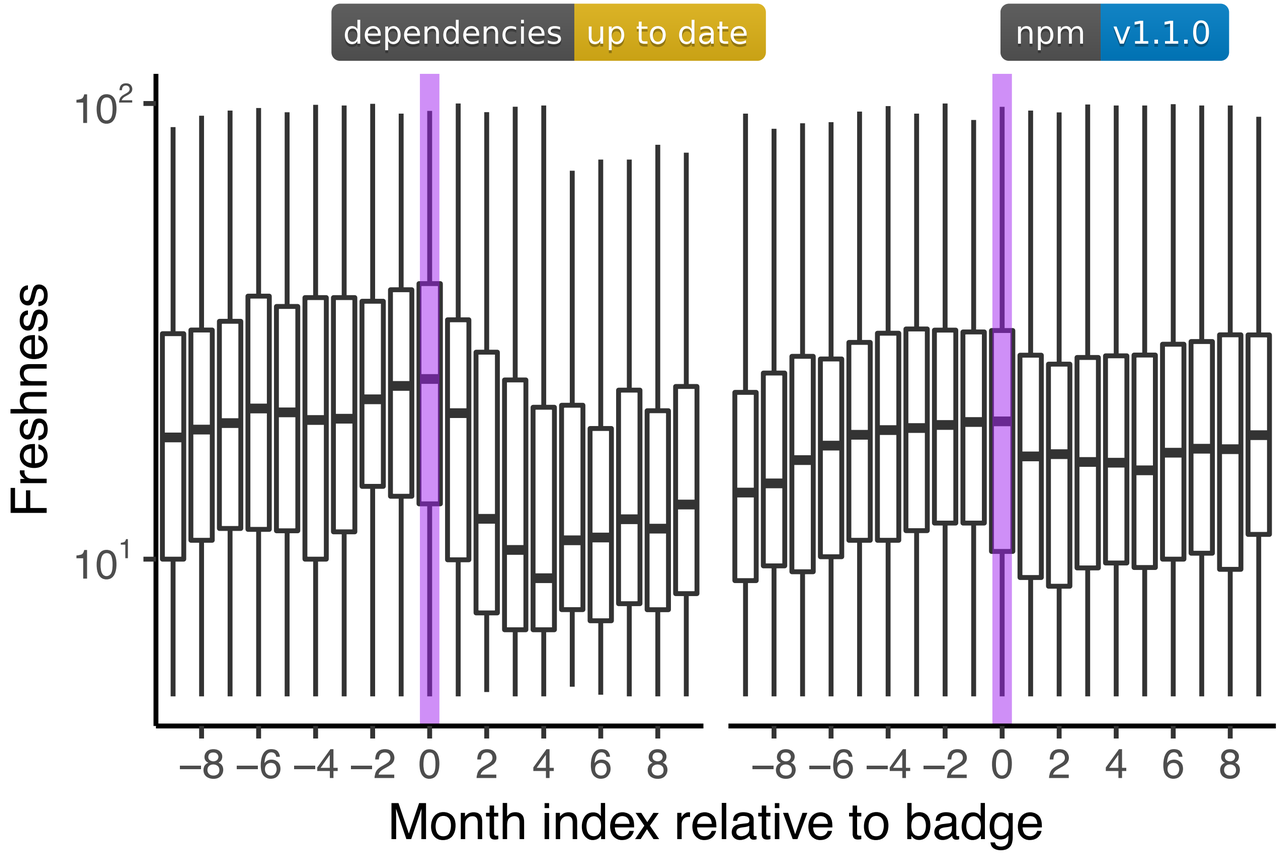

STRUDEL
Carnegie Mellon University

Interrupted time series

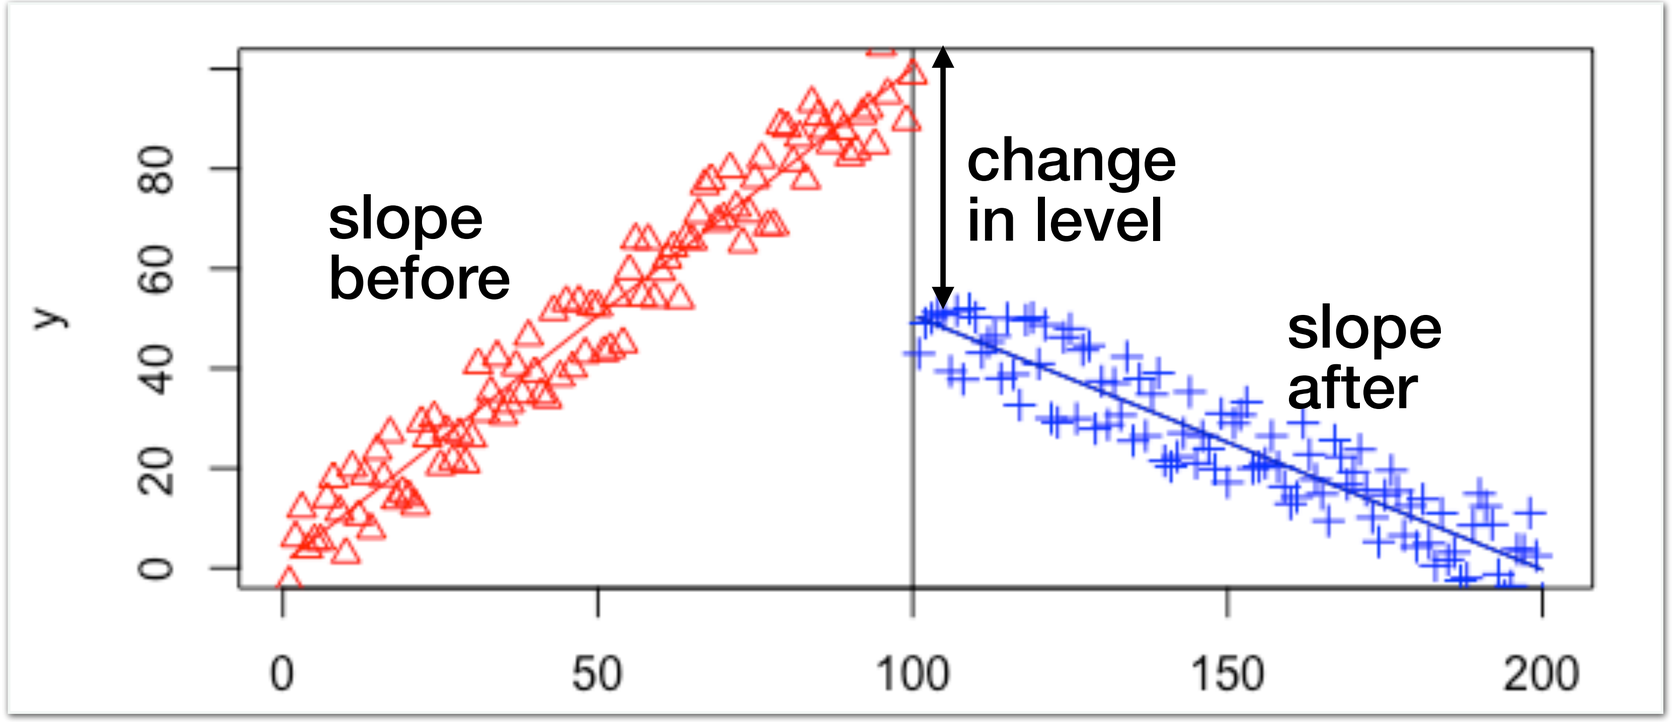

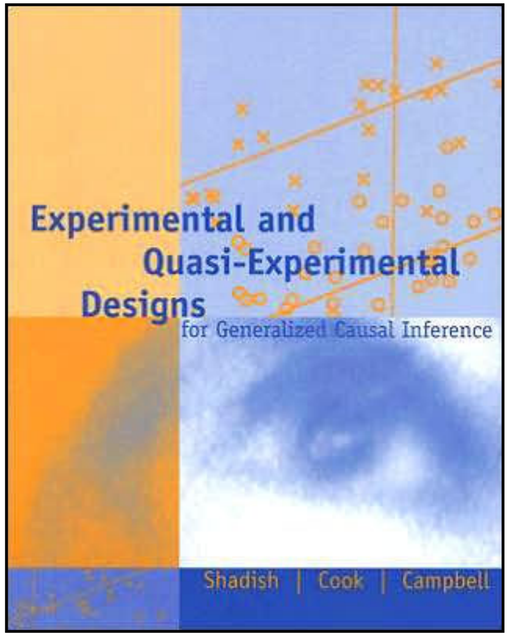

The Impact of Continuous Integration on Other Software Development Practices: A Large-Scale Empirical Study
Zhao, Y., Zhou, Y., Serebrenik, A., Filkov, V., and Vasilescu, B. ASE 2017

Interrupted time series

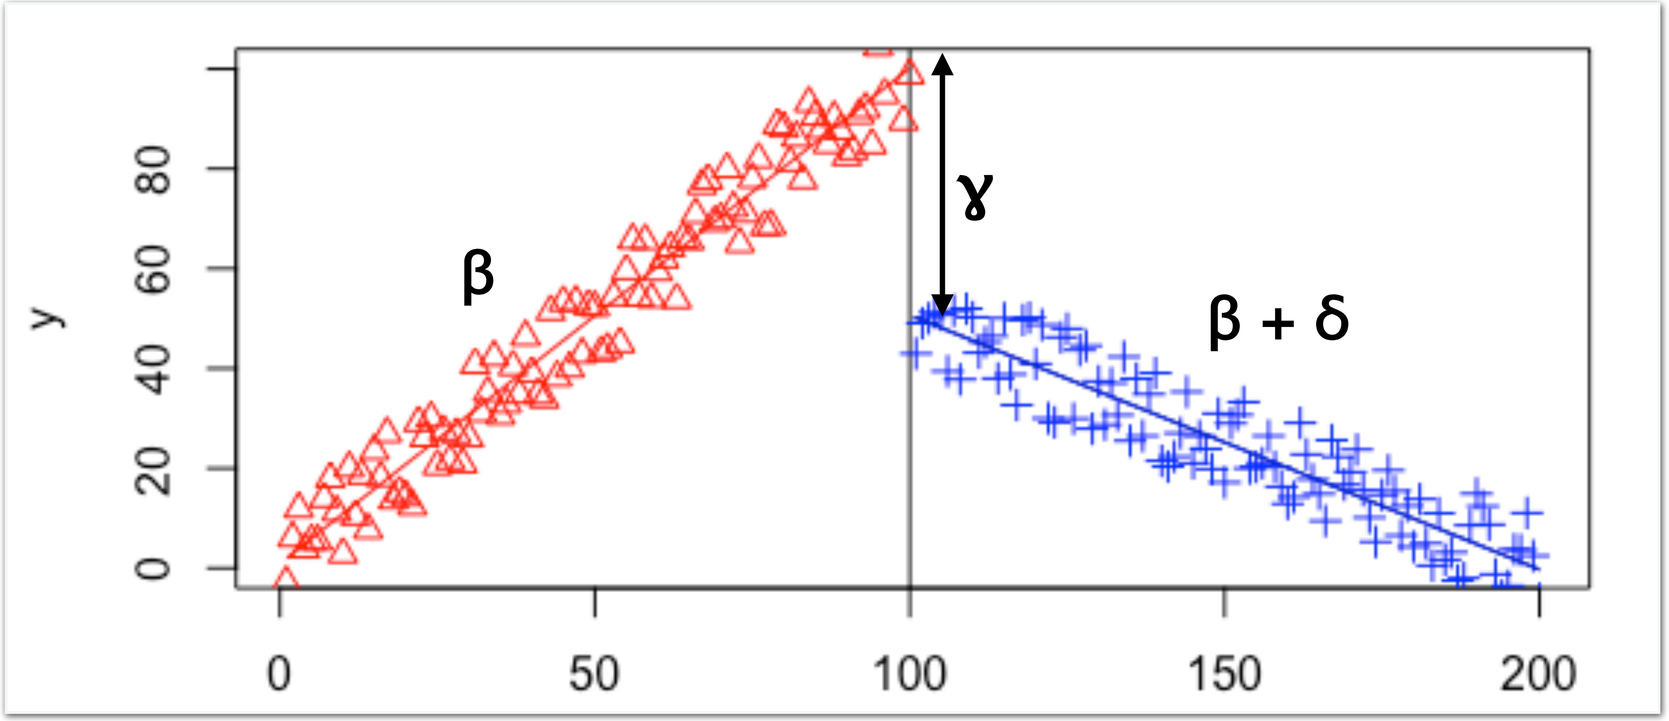

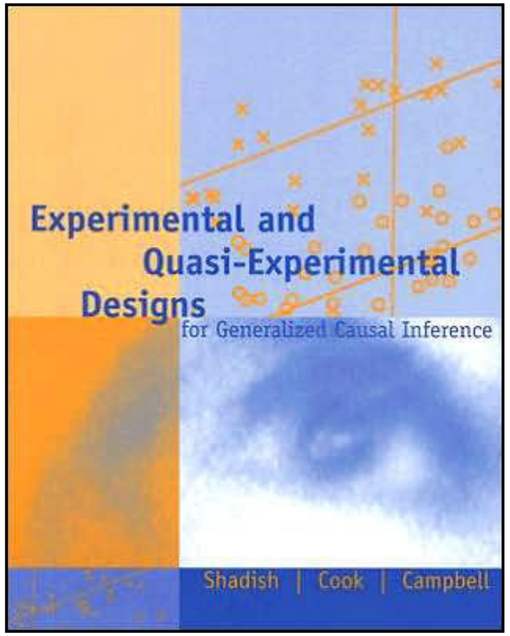

$$y_i = \alpha + \beta \cdot \text{time}_i + \gamma \cdot \text{intervention}_i + \delta \cdot \text{time\_after\_intervention}_i + \varepsilon_i$$

time:	1	2	3	…	…	…	100	101	102	…	…	…	200
time after intervention:	0	0	0	…	…	…	0	1	2	…	…	…	100
intervention:	F	F	F	…	…	…	T	T	T	…	…	…	T

STRUDEL
Carnegie Mellon University

The Impact of Continuous Integration on Other Software Development Practices: A Large-Scale Empirical Study
Zhao, Y., Zhou, Y., Serebrenik, A., Filkov, V., and Vasilescu, B. ASE 2017

Signals of popularity

- Hyp: The adoption of quality-assurance badges makes users more confident in a package and attracts more users
- Hyp: The adoption of popularity-related badges in popular packages correlates with more future downloads

- **Packages with a badge tend to skew toward more downloads than packages without.**

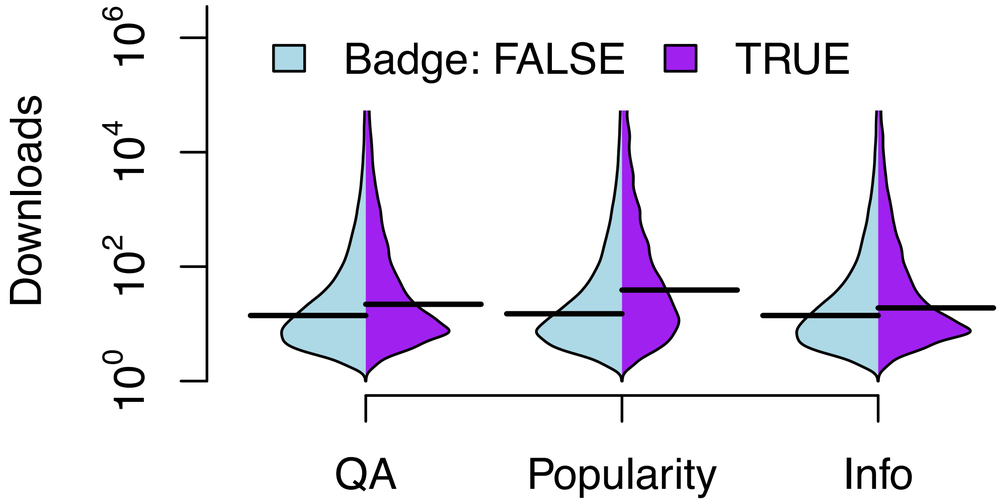

Code was "built with love" or "well written" by an "experienced developer" who pays "attention to quality"

STRUDEL
Carnegie Mellon University

Signals of popularity

- Hyp: The adoption of quality-assurance badges makes users more confident in a package and attracts more users
- Hyp: The adoption of popularity-related badges in popular packages correlates with more future downloads

- **Attractiveness wears off beyond five badges**

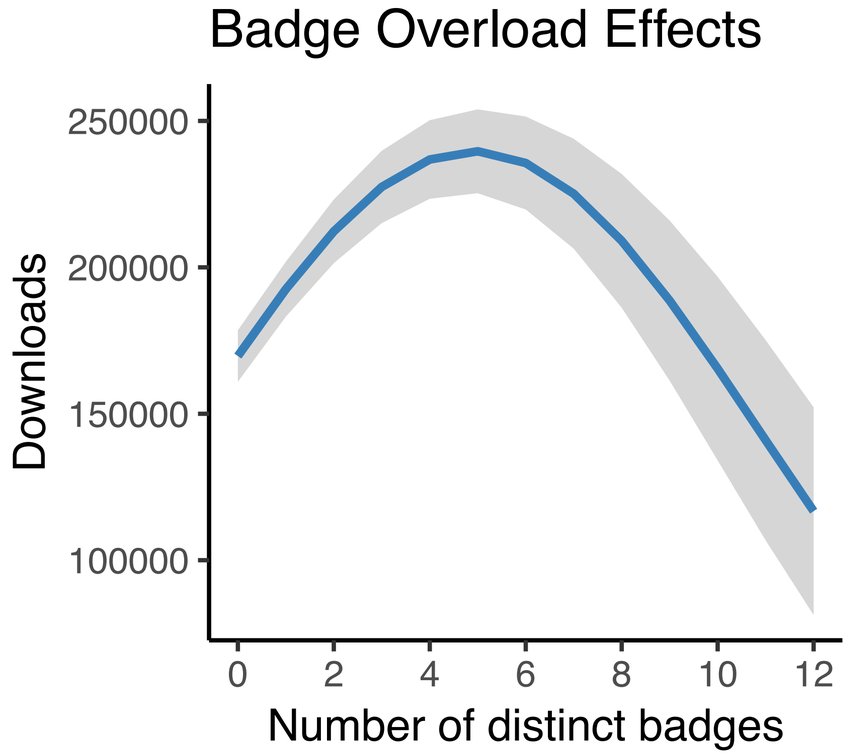

"People tend to overwhelm visitors with too many (useless) badges, thus creating a contra effect and loosing the initial purpose of having useful information."

STRUDEL
Carnegie Mellon University

Signals of popularity

- Hyp: The adoption of quality-assurance badges makes users more confident in a package and attracts more users
- Hyp: The adoption of popularity-related badges in popular packages correlates with more future downloads

- **Badge adoption correlates with a sudden popularity boost, but the acceleration is not sustained over time.**

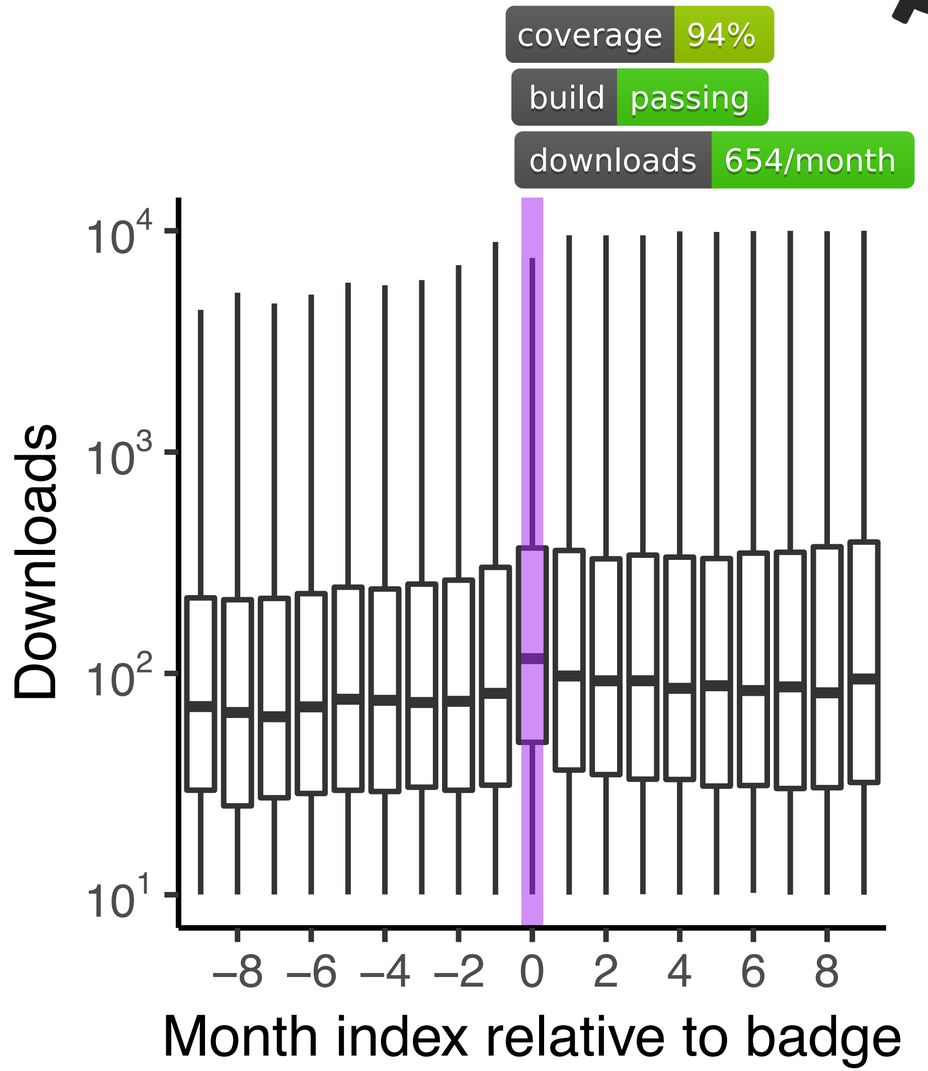

STRUDEL
Carnegie Mellon University

Signals of test suite quality

- Hyp: The adoption of quality-assurance badges correlates with other indicators of code quality (metric: test suite size).

  Code was "built with love" or "well written" by an "experienced developer" who pays "attention to quality"

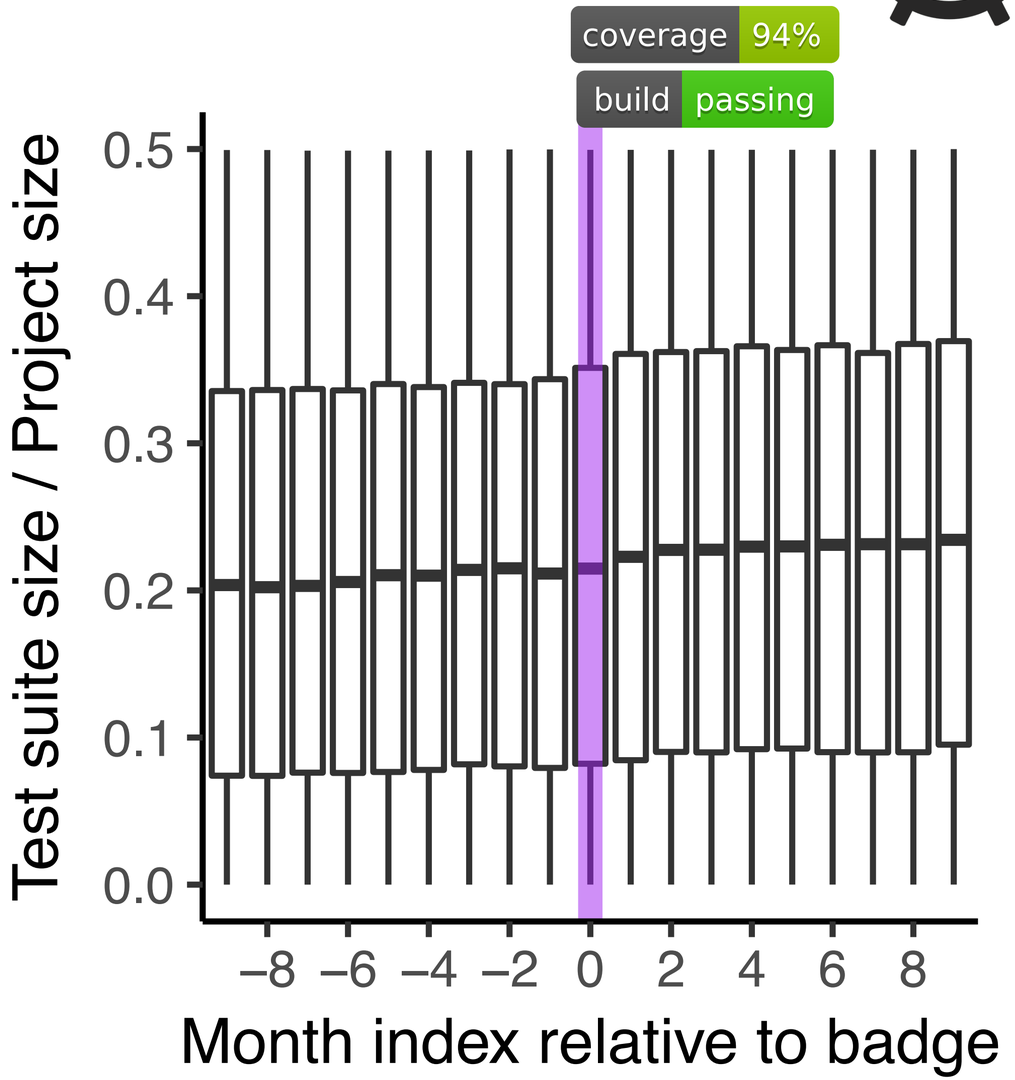

- 18x higher odds of having any tests with QA badge; among those with tests, 18.3% larger test suite with QA badge.
- But, no change in trend

Signals of PR quality

- Hyp: The adoption of a quality-assurance badge, and even more so of a coverage badge, encourages more external contributors to include tests.

  "PRs with new functionality tend to include new tests, as not to decrease coverage."

- Increase in the monthly fraction of PRs containing tests after adopting QA badge

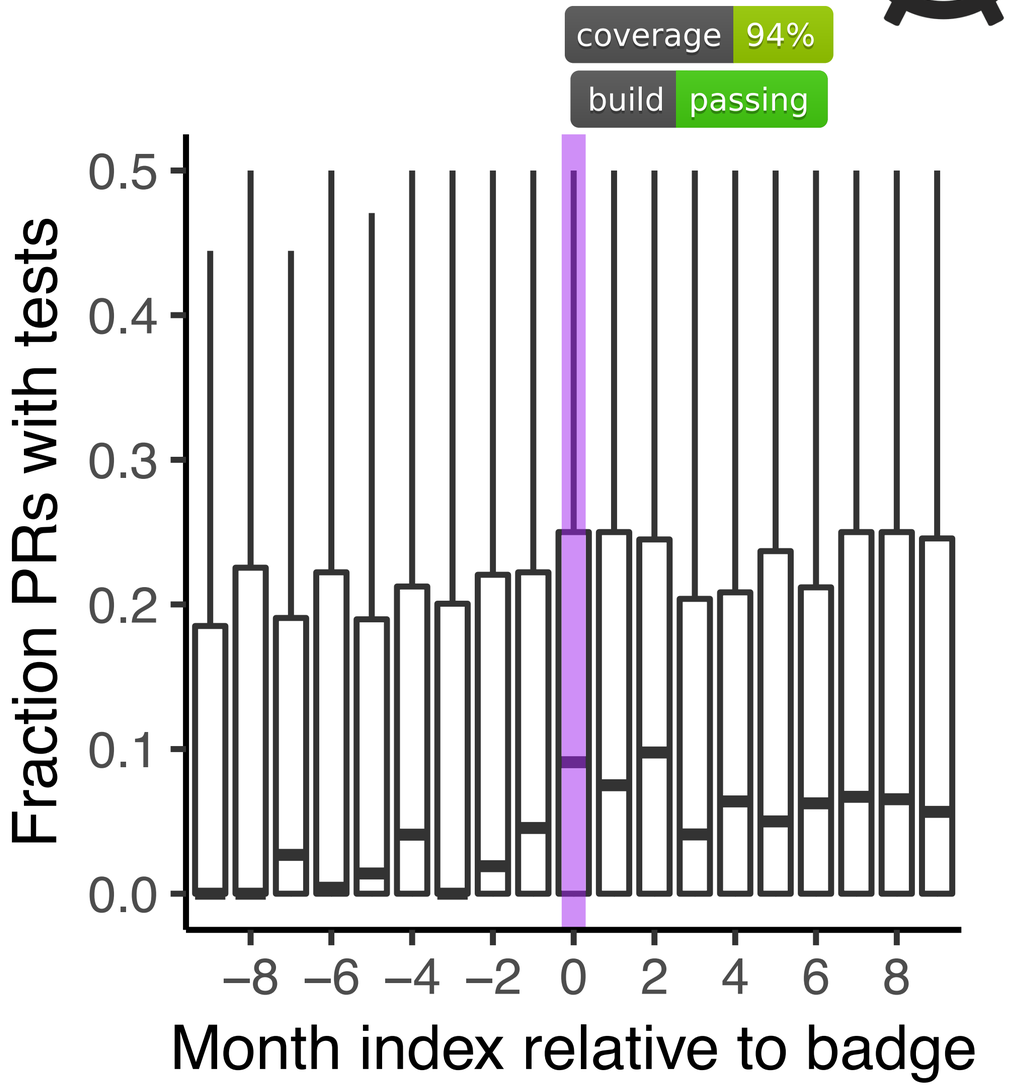

STRUDEL
Carnegie Mellon University

Signals of PR quality

- Hyp: The adoption of a quality-assurance badge, and even more so of a coverage badge, encourages more external contributors to include tests.

  "PRs with new functionality tend to include new tests, as not to decrease coverage."

- Coverage and CI badges interact, amplifying each other's effects.

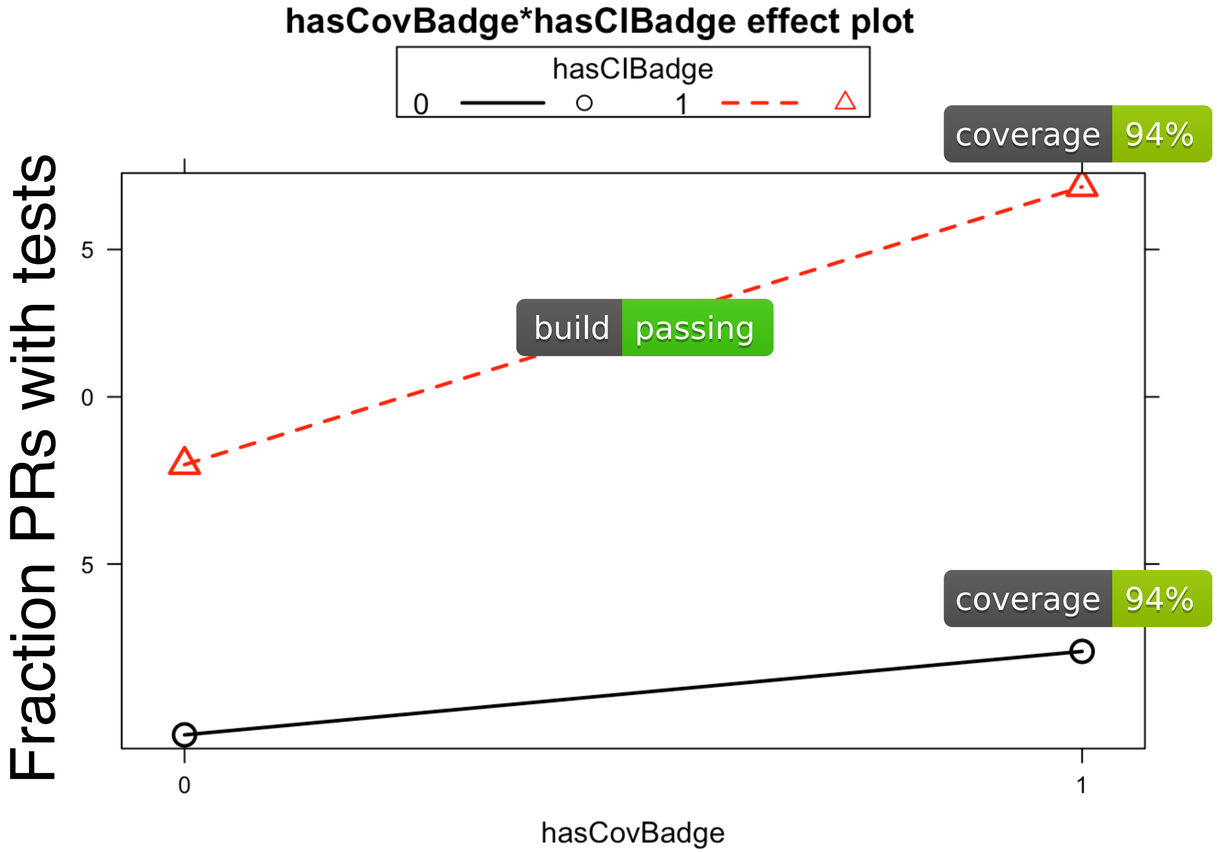

STRUDEL
Carnegie Mellon University

seanmonstar / intel

Watch 10 Star 191 Fork 29

Code Issues 11 Pull requests 0 Projects 0 Wiki Insights

I need more intel! http://seanmonstar.github.io/intel/

224 commits 3 branches

Branch: master New pull request

README.md

intel

build passing coverage 100% npm package 1.2.0

An abbreviation of intelligence. In this case, the acquirement of information.

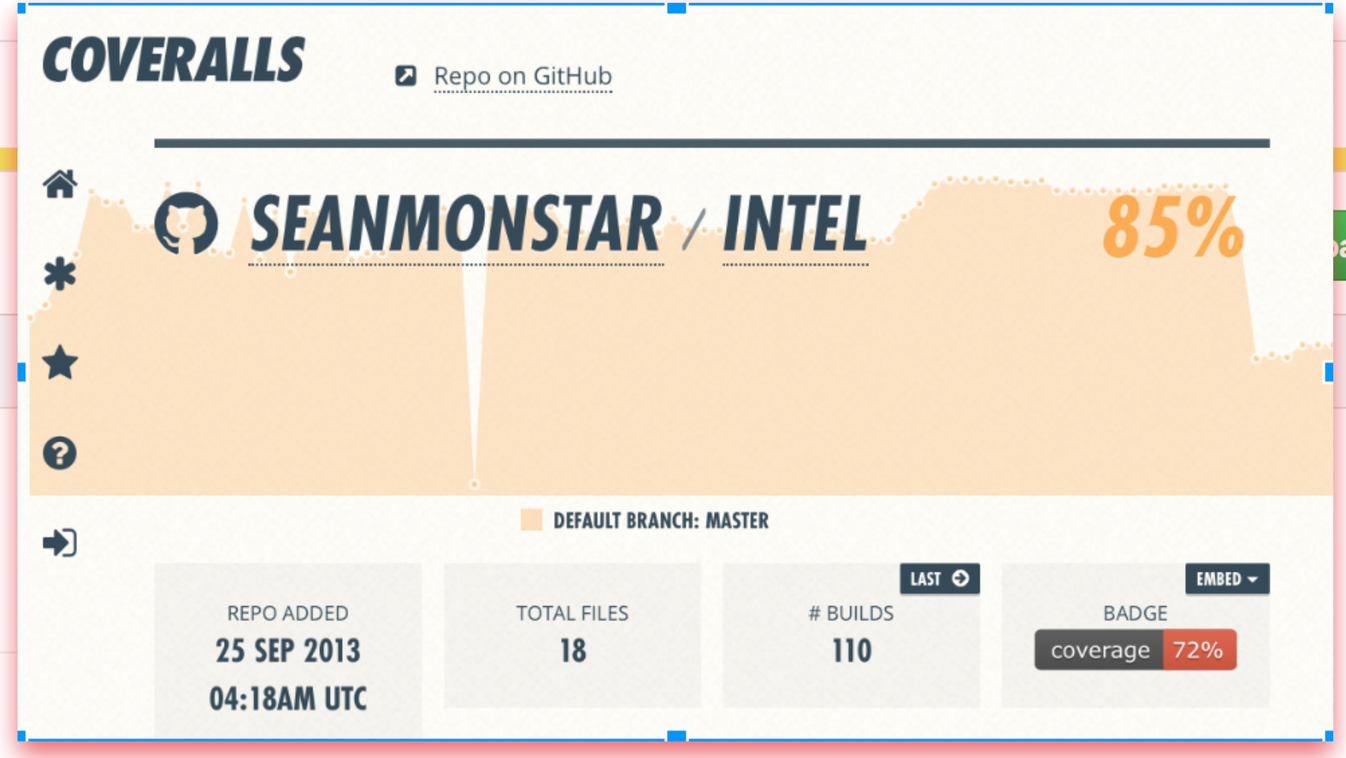

Mixed methods study

+

- 32 maintainers, 57 contributors (15% resp. rate)
- Maintainers:
  - What do you intend to signal?
  - What effects do you expect?
- Contributors:
  - What do badges tell you?

- 294,941 *npm* packages
- Mined bad
  from READ
- Measured
  test suite c
  dependen

Recap

Analysis

If all you saw was the badge, how much would that tell you?

How much more does the badge tell you, relative to existing signals?

How do things change after adding the badge?

STRUDEL
Carnegie Mellon University

Take-aways (1)

- Open source developers rely on, and respond to, signals
  - We add both qualitative and quantitative evidence for badges

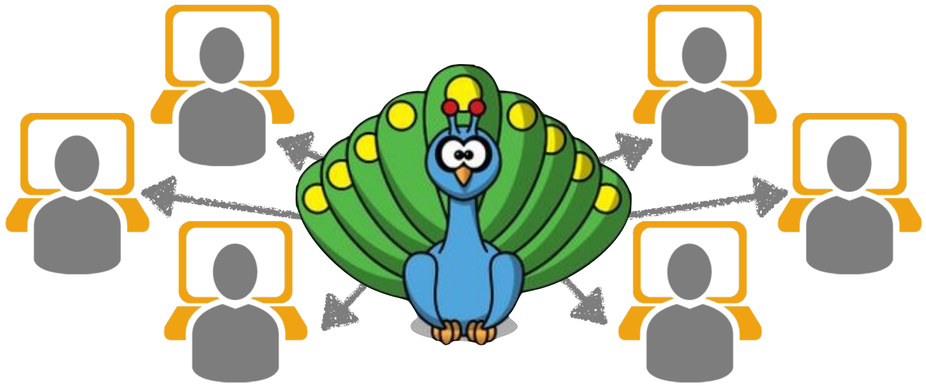

Take-aways (2)

- Harder to fake badges provide more reliable signals
  - As signaling theory predicts

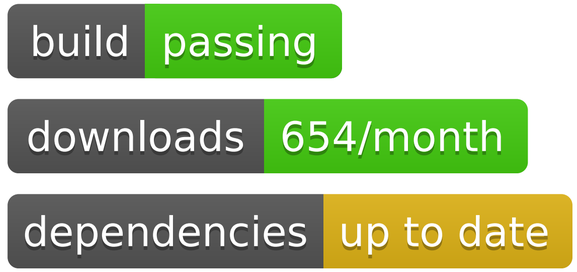

vs

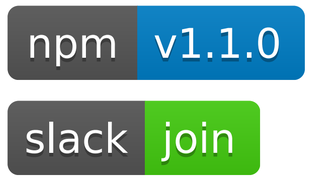

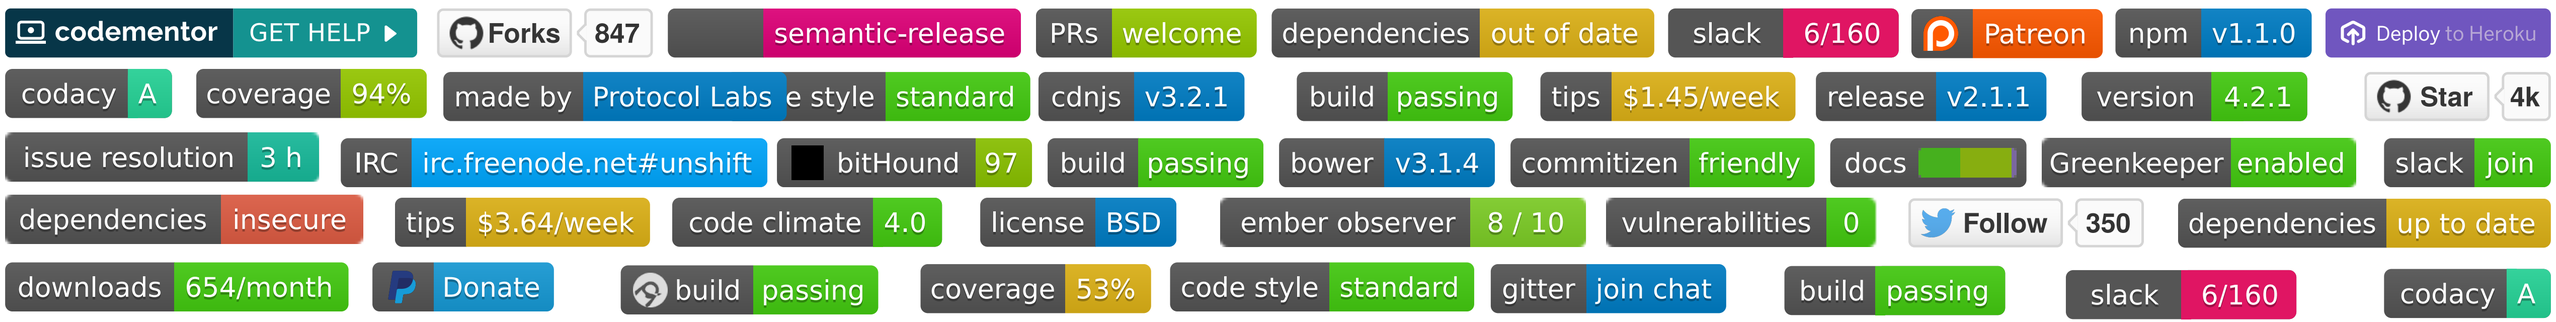

Take-aways (2)

- Harder to fake badges provide more reliable signals
  - As signaling theory predicts
- Redesign badges as assessment signals

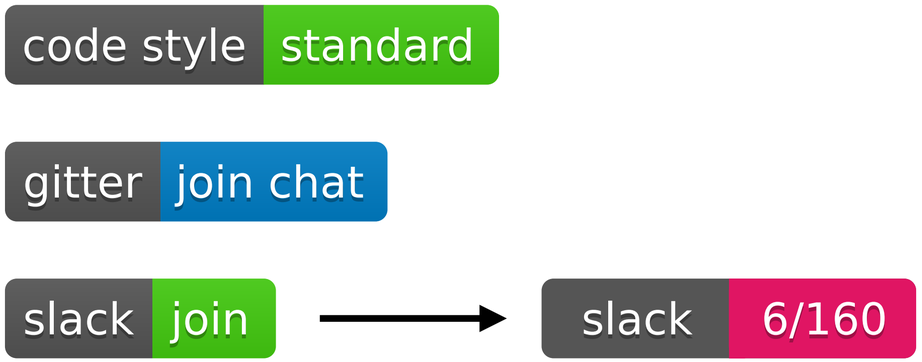

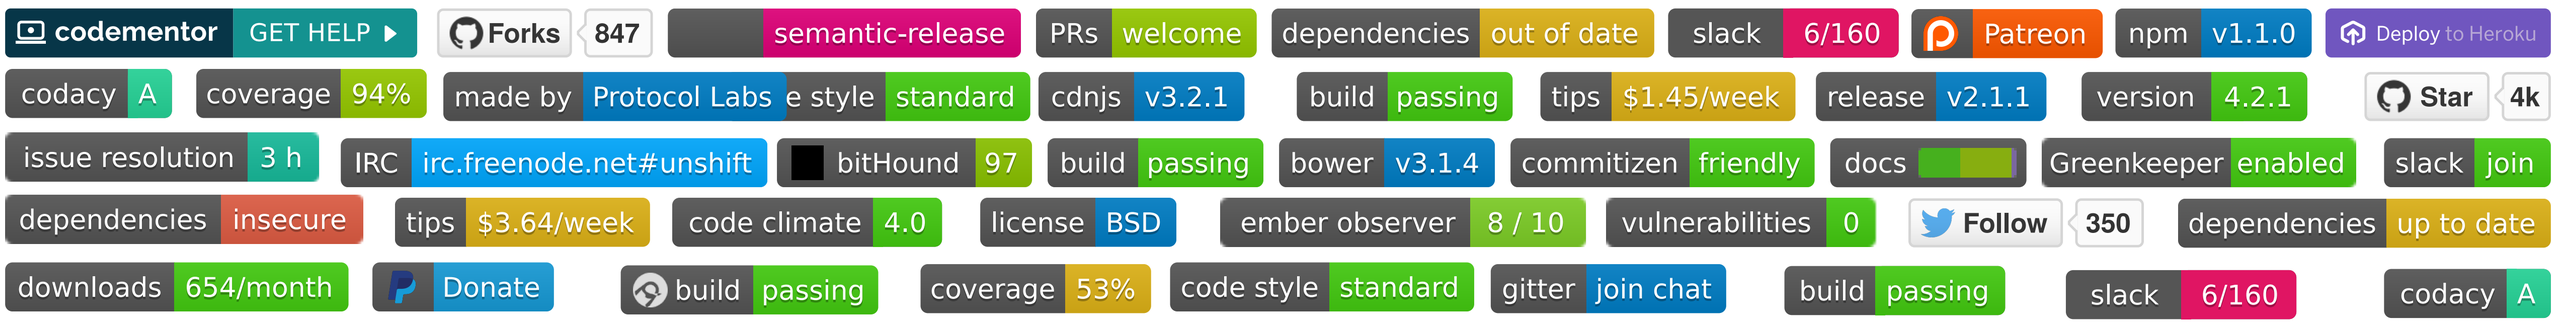

Take-aways (3)

- Too much of a good thing

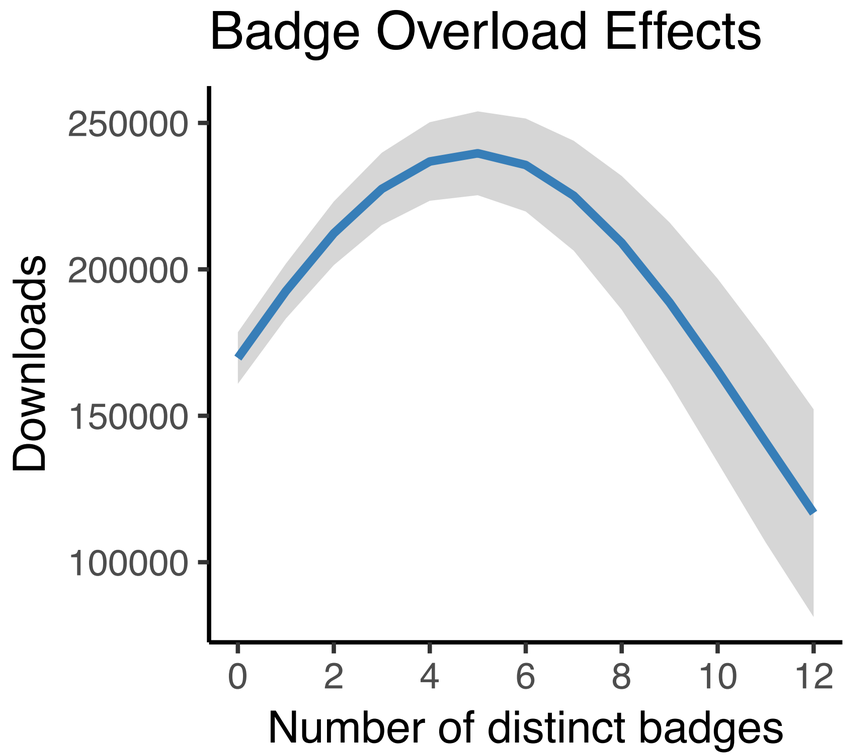

codementor GET HELP ▸ | Forks 847 | semantic-release | PRs welcome | dependencies out of date | slack 6/160 | Patreon | npm v1.1.0 | Deploy to Heroku

codacy A | coverage 94% | made by Protocol Labs | e style standard | cdnjs v3.2.1 | build passing | tips $1.45/week | release v2.1.1 | version 4.2.1 | Star 4k

issue resolution 3 h | IRC irc.freenode.net#unshift | bitHound 97 | build passing | bower v3.1.4 | commitizen friendly | docs | Greenkeeper enabled | slack join

dependencies insecure | tips $3.64/week | code climate 4.0 | license BSD | ember observer 8 / 10 | vulnerabilities 0 | Follow 350 | dependencies up to date

downloads 654/month | Donate | build passing | coverage 53% | code style standard | gitter join chat | build passing | slack 6/160 | codacy A

Take-aways (4)

- Gamification effects

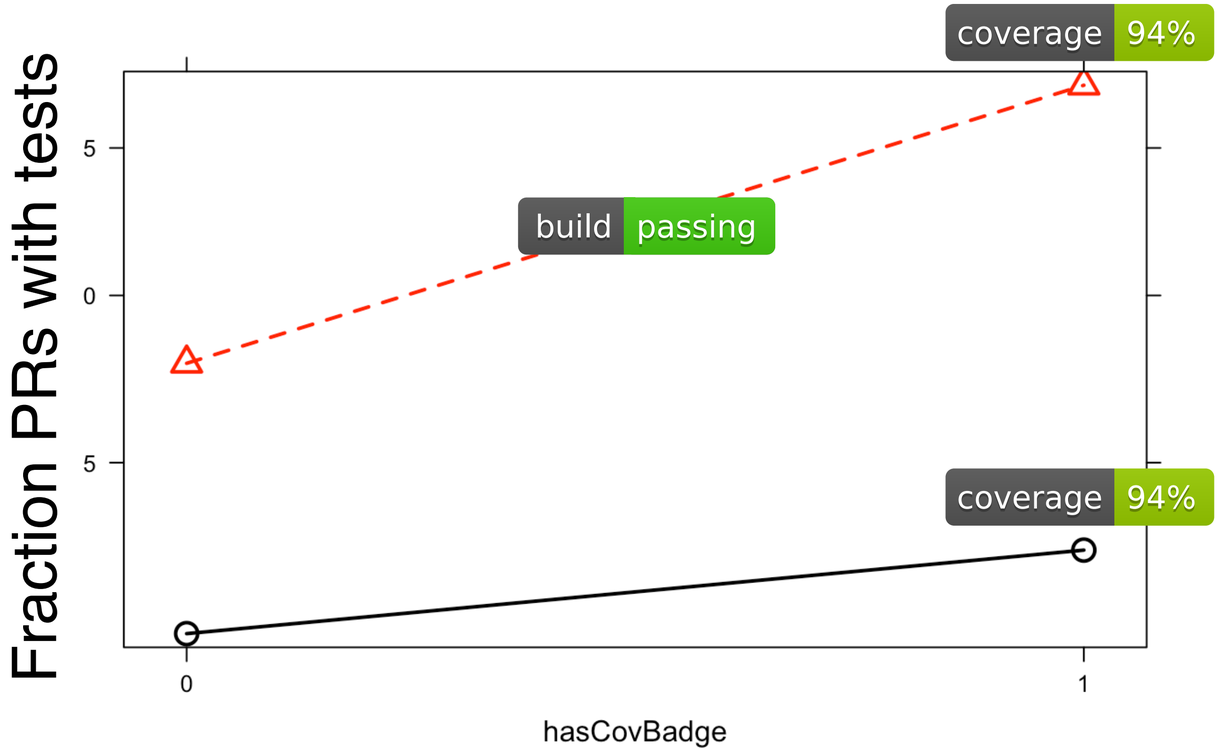

codementor GET HELP ▸ | Forks 847 | semantic-release | PRs welcome | dependencies out of date | slack 6/160 | Patreon | npm v1.1.0 | Deploy to Heroku

codacy A | coverage 94% | made by Protocol Labs | e style standard | cdnjs v3.2.1 | build passing | tips $1.45/week | release v2.1.1 | version 4.2.1 | Star 4k

issue resolution 3 h | IRC irc.freenode.net#unshift | bitHound 97 | build passing | bower v3.1.4 | commitizen friendly | docs | Greenkeeper enabled | slack join

dependencies insecure | tips $3.64/week | code climate 4.0 | license BSD | ember observer 8 / 10 | vulnerabilities 0 | Follow 350 | dependencies up to date

downloads 654/month | Donate | build passing | coverage 53% | code style standard | gitter join chat | build passing | slack 6/160 | codacy A

Take-aways (5)

- Could we design signals (**badges**?) to help **balance supply and demand for labor** in open source ecosystems?
  - Maybe

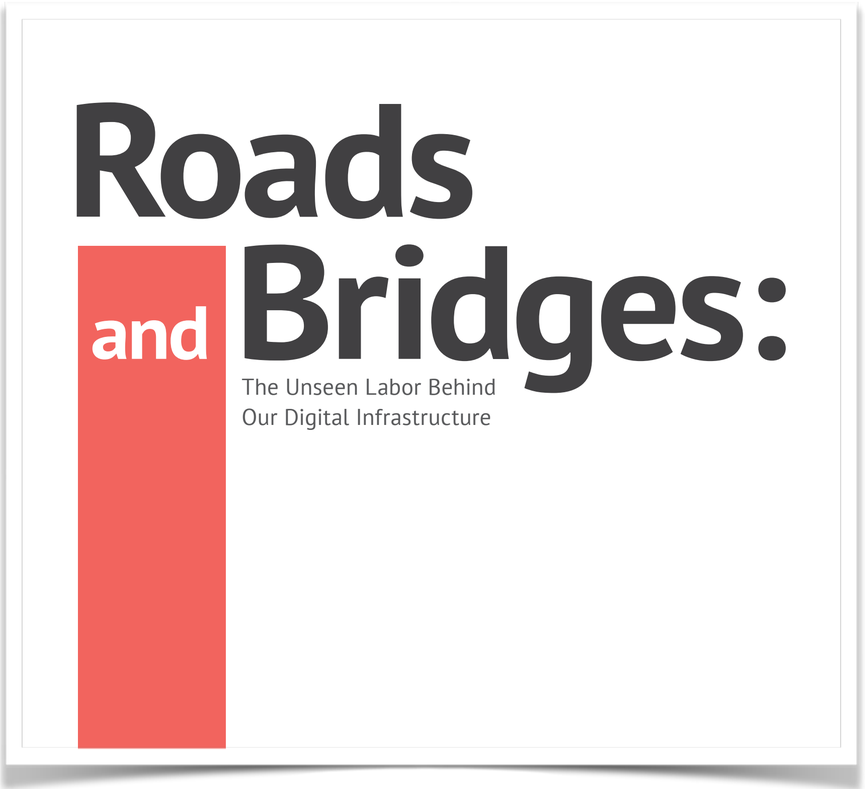

Adding Sparkle to Social Coding

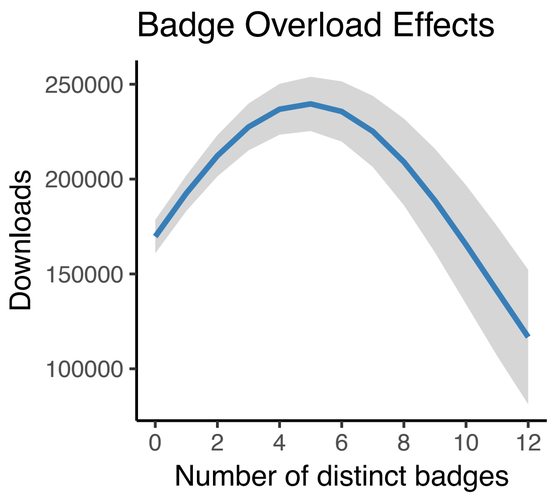

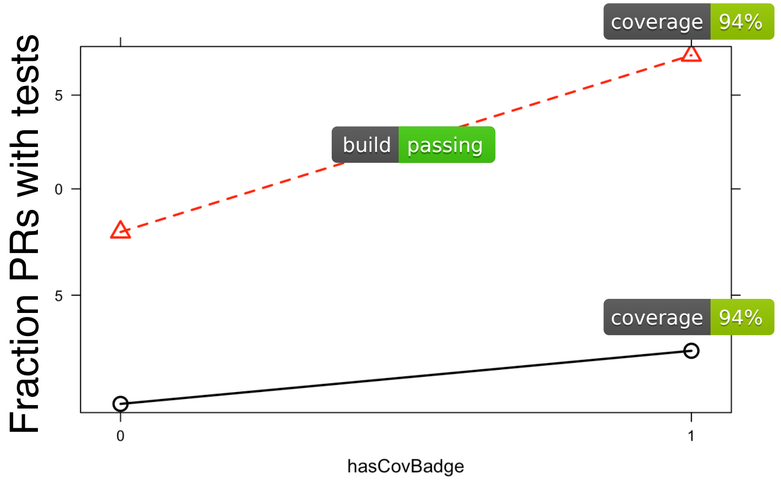

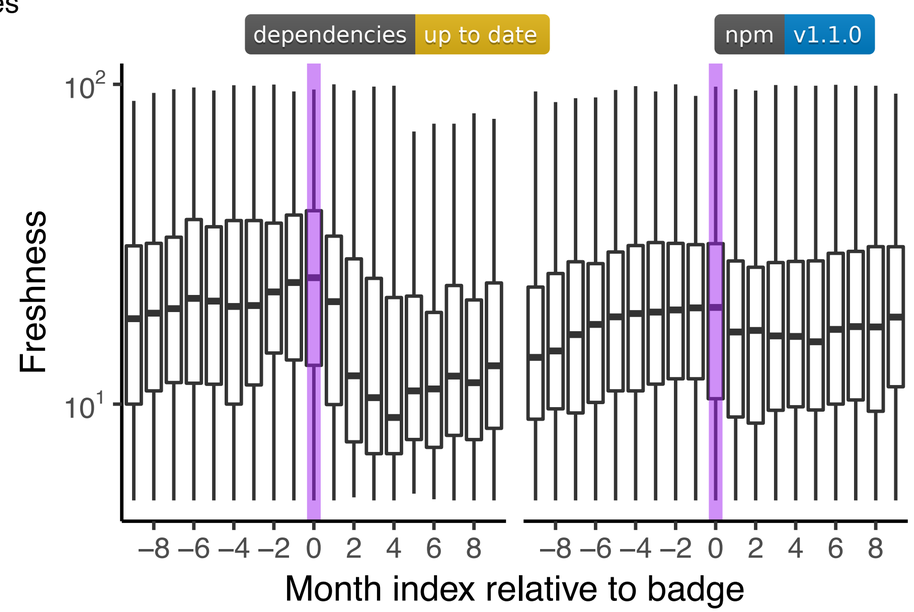

codementor GET HELP ▶ | npm v1.1.0 | Donate

build passing | code style standard | tips $3.64/week

ember observer 8 / 10 | dependencies out of date

build passing | license BSD | vulnerabilities 0

bitHound 97 | Star 4k | cdnjs v3.2.1

bower v3.1.4 | downloads 654/month | codacy A

issue resolution 3 h | coverage 53% | Follow 350

slack join | build passing | dependencies insecure

dependencies up to date | IRC irc.freenode.net#unshift

coverage 94% | gitter join chat | code climate 4.0

build passing | commitizen friendly | version 4.2.1

release v2.1.1 | Greenkeeper enabled | docs

Patreon | slack 6/160 | made by Protocol Labs

code style standard | tips $1.45/week | Deploy to Heroku

Forks 847 | semantic-release | PRs welcome

STRUDEL
Carnegie Mellon University

Badge vs. tool (1)

- 83.6% of projects using Travis also use the badge
- Projects with the badge tend to have more tests vs. those with just Travis

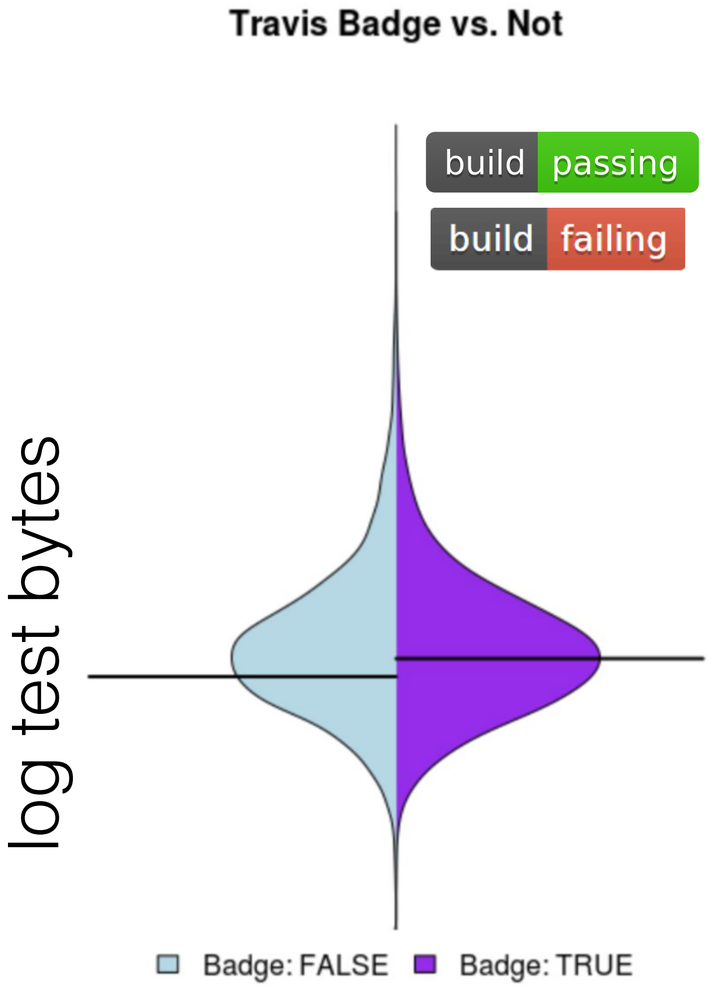

Badge vs. tool (2)

- 83.6% of projects using Travis also use the badge
- Logistic regression to predict current build status (pass/fail):

If you show the badge, it's more likely 

than 


```
Coefficients:
                                Estimate Std. Error z value Pr(>|z|)
(Intercept)                     -0.52484    0.03236 -16.219  < 2e-16 ***
scale(log(downloads + 1))       -0.05058    0.01288  -3.927 8.61e-05 ***
scale(log(active_age + 1))       0.06748    0.01307   5.165 2.41e-07 ***
scale(log(commits + 1))          0.27668    0.01633  16.941  < 2e-16 ***
scale(log(dependencies + 1))     0.14626    0.01052  13.909  < 2e-16 ***
factor(has_tests)1              -0.30313    0.03295  -9.199  < 2e-16 ***
scale(log(build_num + 1))       -0.03652    0.01607  -2.273    0.023 *
has_badge                       -1.12270    0.02432 -46.156  < 2e-16 ***
```

Has_badge implies 0.3 multiplicative diff in odds of failing

STRUDEL
Carnegie Mellon University