Center for
Security, Innovation,
and New Technology

Buying Security: Open Source Software Funding and Security Posture

Sara Ann Brackett and Stewart Scott with Christina Cheng

CSINT Policy Paper Series · Fall 2025

TABLE OF CONTENTS

Executive Summary

The security of open source software (OSS) has morphed from a niche technical concern to a central cybersecurity policy challenge. High-profile incidents have led to suggestions for governments to help strengthen the OSS ecosystem, including calls for funds built to support open source projects and their maintainers, such as a proposal for an EU Sovereign Tech Fund.[1] This research examines the argument that unconditional funding—namely, financial support without specific requirements for the recipient—causally improves the security posture of OSS projects.

This study empirically assesses data from 1,081 Open Collective accounts backed through the Open Collective platform and their connected 4,676 GitHub repositories over approximately 42 months—a cumulative $27.76 million in funding inflows to OSS projects—to evaluate the extent of the causal relationship between funding and security posture.

The results of this analysis reject the hypothesis that general funding alone causally improves the project security posture—here, measured as the degree to which projects follow a defined set of security practices and processes tracked by the Open Source Security Foundation's (OpenSSF) Scorecard tool.[2]

This finding challenges assumptions about the relationship between funding and OSS security and is pertinent to cybersecurity policymakers and practitioners discussing financial support to OSS. As governments and organizations seek to address insecurity in the OSS ecosystem and beyond, this research indicates that broad, general funding programs are not a panacea. This space needs further empirical study. Meanwhile, targeted support designed to strengthen security, such as subsidized code audits, training, tooling, dedicated developer time, or conditional financing, has proven more effective at improving security posture than what the current data says about general funding.

Yet, this does not mean discarding the concept of general funding for OSS projects, nor forcing security-specific interventions on the OSS community for several reasons. First, general funding serves multiple other ends, such as improving project continuity and the ability of maintainers to engage with their projects more sustainably. Second, this study only examines causality within and across funded repositories, not in comparison to similar unfunded ones, which might imply that

---

[1] Nicholas Gates, "Investing in Open Source Sustainability and Security: OFE's Proposal for an EU Sovereign Tech Fund," OpenForum Europe, April 23, 2025, https://openforumeurope.org/investing-in-open-source-sustainability-and-security-ofes-proposal-for-an-eu-sovereign-tech-fund/; Felix Reda, "We Need a European Sovereign Tech Fund," GitHub, July 23, 2025, https://github.blog/open-source/maintainers/we-need-a-european-sovereign-tech-fund/; "Fact Sheet: Biden-Harris Administration Releases Summary Report of 2023 RFI on Open Source-Software Security Initiative," The White House, August 12, 2024, https://bidenwhitehouse.archives.gov/oncd/briefing-room/2024/08/09/fact-sheet-biden-harris-administration-releases-end-of-year-report-on-open-source-software-security-initiative-2/.

[2] "OpenSSF Scorecard," Open Source Security Foundation, accessed October 21, 2025, https://openssf.org/projects/scorecard/.

whether a project receives funding at all is meaningful for security posture improvements. This research underscores the importance of evidence-based evaluation of policy interventions in the OSS ecosystem and cybersecurity policy more broadly.

About the Authors

Sara Ann Brackett is an assistant director with the Cyber Statecraft Initiative, part of the Atlantic Council Tech Programs. She focuses her work on open-source software security, software bills of materials, software liability, and software supply-chain risk management within the Cyber Statecraft Initiative's cybersecurity and policy portfolio. Brackett graduated from Duke University, where she majored in computer science and public policy and wrote a thesis on the effects of market concentration on cybersecurity. She participated in the Duke Tech Policy Lab's Platform Accountability Project and worked with the Duke Cybersecurity Leadership Program as part of Professor David Hoffman's research team.

Stewart Scott is a deputy director with the Cyber Statecraft Initiative, part of the Atlantic Council Tech Programs. He works on the Initiative's Cybersecurity, Strategy, and Policy portfolio, with focuses on software supply chain and open source software security policy. Scott earned his BA from Princeton University at the School of Public and International Affairs along with a minor in computer science. His course of study centered on misinformation, social media policy, online extremism, journalism, and American political and economic history. He joined the Atlantic Council after interning with its Cyber Statecraft Initiative.

Christina Cheng '26 was an FSI Global Policy Intern with the Atlantic Council's Cyber Statecraft Initiative in the summer of 2025, where she worked on several cybersecurity policy projects focused on open source software and the telecommunications sector. Christina is currently completing a degree in Computer Science (MS) at Stanford University.

Introduction

High-profile open source software (OSS) incidents, such as the disclosure of critical vulnerabilities within the widely used Java-based log4j framework, have become more prominent in cybersecurity policy discussions, prompting calls for enhanced security of the OSS ecosystem.[3] The 2021 log4j incident, similar in scope and severity to the 2014 Heartbleed vulnerability, led to widespread alerts, White House meetings, legislative proposals, non-profit funding initiatives, and industry consortia.[4] Both events emphasized what had long been known to cybersecurity practitioners, OSS advocates, and community members: OSS's criticality to modern digital systems remains foundational, underappreciated outside of insular technical circles, and largely without sufficient support. Since those incidents, recent near misses such as the XZ Utils compromise, as well as continuous attacks on package managers such as recent malicious backdoors targeting the Node.js package manager npm, have underlined that viewpoint, alongside policy initiatives designed to support the ecosystem.[5] Policy efforts in the United States focusing on OSS to date include funding and security investments, government working groups, and expanded partnerships. For example, in the past few years, the US government announced a series of initiatives focused the OSS ecosystem, including the Open Source Software Security Initiative (OS3I) inter-agency working group and an $11 million investment in the security of OSS.[6] In tandem, the US Cybersecurity and Infrastructure Security Agency, or CISA, published its Open Source Software Security Roadmap, an effort to harden the OSS ecosystem through government engagement and support as well as industry participation that cited log4j as a driving example.[7]

OSS is released under licenses that make the source code available for reuse, inspection, modification, and redistribution. Through the crowdsourcing of shared engineering problems, the OSS ecosystem creates reusable, modifiable, and publicly accessible software solutions, providing

---

[3] Tim Starks, "White House Hosts Open-Source Software Security Summit in Light of Expansive Log4j Flaw," *CyberScoop*, January 13, 2022, https://cyberscoop.com/white-house-log4j-open-source-software-security/.
[4] National Security Agency, "CISA, FBI, NSA, and International Partners Issue Advisory to Mitigate Apache Log4J Vulnerability," press release, December 22, 2021, https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/2881834/cisa-fbi-nsa-and-international-partners-issue-advisory-to-mitigate-apache-log4j/; Sean Lyngaas, "White House to Discuss Software Development with Tech Executives, Calling It 'Key National Security Concern,'" CNN, December 23, 2021. https://www.cnn.com/2021/12/23/politics/white-house-log4j-tech-firms-meeting; Nihal Krishan, "Senators Reintroduce Bipartisan Bill to Push for Open Source Software Security within Federal Gov," *FedScoop*, March 27, 2023, https://fedscoop.com/senators-reintroduce-bipartisan-open-source-software-bill/.
[5] Jack Cable and Aeva Black, "Lessons from XZ Utils: Achieving a More Sustainable Open Source Ecosystem," Cybersecurity and Infrastructure Security Agency, April 12, 2024. https://www.cisa.gov/news-events/news/lessons-xz-utils-achieving-more-sustainable-open-source-ecosystem; Kevin Poireault, "Open Source Community Thwarts Massive npm Supply Chain Attack," *Infosecurity Magazine*, September 9, 2025, https://www.infosecurity-magazine.com/news/npm-supply-chain-attack-averted/.
[6] "Securing the Open-Source Software Ecosystem End of Year Report: Open-Source Software Security Initiative (OS3I)," The White House, January 2024, https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/01/Securing-the-Open-Source-Software-Ecosystem-OS3I-End-of-Year-Report-MASTERCOPY.pdf; David Jones, "White House Details $11M Plan to Help Secure Open Source," *Cybersecurity Dive*, August 14, 2024, https://www.cybersecuritydive.com/news/white-house-11-million-secure-open-source/724223/.
[7] "CISA Open Source Software Security Roadmap," Cybersecurity and Infrastructure Security Agency, September 12, 2023, https://www.cisa.gov/resources-tools/resources/cisa-open-source-software-security-roadmap.

massive savings and efficiencies across the software industry and all dependent sectors.[8] These projects have proven so foundational that some studies estimate OSS comprises up to 90 percent of most modern codebases, and estimates of its aggregate market value are astronomical.[9] Previous research has argued that OSS is best understood as a public good similar to infrastructure—providing vital functionality for the world's codebases but requiring careful maintenance and support not satisfied by standard market conditions.[10] Whether the primary concern of cybersecurity policymakers is domestic critical infrastructure, military systems, consumer products, or government-procured services, open source software is at the core of all these digital systems and hardening them against compromise must involve improving the security of and around their component pieces.

This ubiquity is a significant part of the impetus behind OSS security discussions: OSS is not intrinsically less secure than proprietary code, but its wide distribution and the centrality of certain OSS projects to a massive number of critical systems multiply the potential fallout of security incidents substantially. In addition to the vast reach of OSS, its licensing schemes, usage practices, and community ethos make systematic security interventions considerably more difficult—those most impacted by OSS vulnerabilities are rarely the stewards of open source projects and seldom compensate those responsible for the day-to-day maintenance they depend on for their work, visibility into dependency on specific OSS projects and their contexts is sparse, and the community is understandably apprehensive of asks to sustain hobby projects to an enterprise security standard they, for the most part, receive no compensation for.[11] Nonetheless, security risks to OSS are security risks to innumerable dependent systems, and the potential return on investment (ROI) for OSS security improvement is massive.

OSS policy and security conversations have focused on the challenge of resourcing. Ubiquity and availability are essential characteristics of the OSS ecosystem. Yet, the resulting dynamic is one in which most software depends upon a small number of projects, often maintained by a handful of

---

[8] Manuel Hoffmann, Frank Nagle, and Yanuo Zhou, "The Value of Open Source Software," Working Paper, No. 24-038, Harvard Business School, Boston, MA, January 2024, https://www.hbs.edu/faculty/Pages/item.aspx?num=65230.

[9] Jason Perlow, "A Summary of Census II: Open Source Software Application Libraries the World Depends on," The Linux Foundation, March 7, 2022, https://www.linuxfoundation.org/blog/blog/a-summary-of-census-ii-open-source-software-application-libraries-the-world-depends-on; Rachel Layne, "Open Source Software: The $9 Trillion Resource Companies Take for Granted," Harvard Business School, March 22, 2024, https://www.library.hbs.edu/working-knowledge/open-source-software-the-nine-trillion-resource-companies-take-for-granted.

[10] Stewart Scott, Sara Ann Brackett, Trey Herr, and Maia Hamin, *Avoiding the Success Trap: Toward Policy for Open-Source Software as Infrastructure*, *Atlantic Council*, February 8, 2023, https://www.atlanticcouncil.org/in-depth-research-reports/report/open-source-software-as-infrastructure/; Nicholas Gates and Liv Marte Nordhaug, "Digital Commons and Digital Public Goods – Finding Common Ground for Policymakers," Next Generation Internet (NGI) Commons, January 31, 2025, https://commons.ngi.eu/2025/01/31/digital-commons-and-digital-public-goods-finding-common-ground-for-policymakers/.

[11] "Open Infrastructure Is Not Free: A Joint Statement on Sustainable Stewardship," Open Source Security Foundation, September 23, 2025, https://openssf.org/blog/2025/09/23/open-infrastructure-is-not-free-a-joint-statement-on-sustainable-stewardship/.

volunteers on the primary basis of goodwill.[12] Without adequate resourcing, including financial, computing, educational, or human, maintainers cannot improve, iterate, or even support their projects, jeopardizing their continuity. This dynamic has generated calls to compensate OSS maintainers better for their work, which OSS maintainers and advocates welcome.[13]

However, as is common across cybersecurity policy, empirical research on the effects of financial interventions on OSS security is relatively sparse. The small number of studies in this space focus on the relationship between donations to OSS projects and project activity.[14] Fortunately, the "open" nature of OSS, with data more publicly available and accessible than for enterprise systems, provides an opportunity to advance empirical cybersecurity research while addressing the key issue of whether general funding for OSS projects causally improves their security posture. Prior Atlantic Council research, on which this project builds, has found tentative evidence that receiving general OSS funding moderately correlates with better security posture among open source projects, with some variation across language ecosystems, funding sources, and specific security heuristics.[15] This research develops these findings in two ways: first, by looking at funding quantities rather than a strict binary variable, and second, by determining if there is a causal relationship beyond correlation between general funding and improved security posture. This is the study's core question: Does mostly unconditional funding *cause* improvements in security posture?

---

[12] Richard Speed, "Open Source Maintainers Are Really Feeling the Squeeze," *The Register*, February 16, 2025, https://www.theregister.com/2025/02/16/open_source_maintainers_state_of_open/; Open Source Security Foundation, "Open Infrastructure Is Not Free."

[13] Stephanie Glen, "Tidelift GC: Paid Open Source Can Stave off Another Log4j," TechTarget, December 13, 2022, https://www.techtarget.com/searchsoftwarequality/news/252528323/Tidelift-CEO-Paid-open-source-can-stave-off-another-Log4j; Patrick Howell O'Neill, "The Internet Runs on Free Open-Source Software. Who Pays to Fix It?" *MIT Technology Review*, December 17, 2021, https://www.technologyreview.com/2021/12/17/1042692/log4j-internet-open-source-hacking/.

[14] A study by China's National University of Defense Technology in Changsha found that donations via GitHub sponsorship only yielded short-term benefits on development activity. Similarly, an Olin College and Carnegie Mellon University study did not find strong evidence that donations increased engineering activity. See: Xunhui Zhang et al., "Who, What, Why and How? Towards the Monetary Incentive in Crowd Collaboration: A Case Study of Github's Sponsor Mechanism," in *CHI Conference on Human Factors in Computing Systems* , ed. Simone Barbos and Cliff Lamb, (New York: Association for Computing Machinery, April 29, 2022), 1–18, https://doi.org/10.1145/3491102.3501822; Cassandra Overney et al., "How to Not Get Rich: An Empirical Study of Donation in Open Source," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering* (New York: Association for Computing Machinery, October 1, 2020), 1209–21, https://doi.org/10.1145/3377811.3380410.

[15] Sara Ann Brackett, John Speed Meyers, and Stewart Scott, *O$$ Security: Does More Money for Open Source Software Mean Better Security? A Proof of Concept*, *Atlantic Council*, April 18, 2024, https://www.atlanticcouncil.org/content-series/cybersecurity-policy-and-strategy/o-security-does-more-money-for-open-source-software-mean-better-security-a-proof-of-concept/.

Background

While this study is far from the final word on whether unconditional funding improves security, its central question has more nuanced implications than what appear at first glance for two reasons. First, limited empirical evidence already exists to suggest that funding with explicit security aims boosts the OSS security posture. A 2023 study by Tidelift found that maintainers paid in exchange for making specific security improvements were more likely to implement them and improve their scores on the OpenSSF Scorecard (a tool discussed in greater detail in the analysis section of this research).[16] This is an expected result—paying for a specific outcome via grant restrictions or contract makes that outcome more likely. Notably, Tidelift's cohort received security assistance in addition to financial incentives. Similarly, a US Department of Homeland Security (DHS) program in 2006, the Open Source Hardening Project, funded access to a vulnerability scanning tool for hundreds of OSS projects, resulting in a notable increase in vulnerability detection and remediation for the enrolled projects. And most recently, the activities of GitHub's Secure Open Source Fund led to the remediation of over 1,100 vulnerabilities, the issuing of over 50 new Common Vulnerabilities and Exposures, known as CVEs, and the adoption of security practices among seventy-one projects involved in the early phases of the program.[17] More broadly, a 2024 study from Tidelift, as well as a 2024 Sonatype study, both found that project funding broadly correlated with better security practices and more formalized or codified development policies and processes.[18] This correlation matches what the Atlantic Council found in its 2024 study as well as this study's data, discussed further on.[19]

In extension of these security-specific activities and broad correlations, this study examines a more precise question: **is there empirical evidence that generally unconditional funding *causes* improvements in security practice, rather than just correlates with it.** There are several scenarios that might indicate a correlation between funding and security practice without causation. For instance, more formally managed OSS projects are likely better able to both improve project security and secure funding. Funded projects are also more likely to have maintainers able to dedicate consistent time to their projects, a non-causal precondition to making security improvements. Central to policy progress is understanding whether unconditional funding is sufficient in itself to improve security practices, and if so, in what quantities, or whether it is instead a necessary or useful precondition without direct causal power.

---

[16] "The 2023 Tidelift State of the Open Source Maintainer Report," Tidelift, April 2023, https://www.sonarsource.com/open-source-maintainer-survey-2023.pdf.
[17] Kevin Crosby and Gregg Cochran, "Securing the Supply Chain at Scale: Starting with 71 Important Open Source Projects," GitHub, August 11, 2025, https://github.blog/open-source/maintainers/securing-the-supply-chain-at-scale-starting-with-71-important-open-source-projects/.
[18] Tidelift, "State of the Open Source Maintainer Report"; John Speed Meyers and Jacqueline Kazil. "How to 'Harden' Open-Source Software," Binding Hook, November 7, 2023, https://bindinghook.com/how-to-harden-open-source-software/; "8th Annual State of the Software Supply Chain Report," 2022, Sonatype, https://www.sonatype.com/resources/state-of-the-software-supply-chain-2022/introduction.
[19] Brackett et al., *O$$ Security*.

Second, the relationships between security practices and outcomes are relatively unmeasured.[20] Security outcomes are out of scope here—this research focuses on a narrow segment of the causal chain, between funding and security posture improvements (i.e., the practices and design choices in and around a project that are largely agreed upon to improve security outcomes, regardless of the state of empirical evidence confirming that consensus). Whether these processes and practices meaningfully reduce the occurrence or severity of security incidents is a more difficult issue to resolve. Security incidents, by design, are hard to detect and their impacts challenging to quantify. Furthermore, heterogeneous codebases and use cases make comparisons at scale complicated in the absence of robust double-blind testing, absent from most cybersecurity studies for legitimate cost and feasibility reasons. As policymakers and organizations push to improve the security of the OSS ecosystem in the face of ongoing security incidents and significant potential ROI, this research seeks to further the field of quantitative study into what methods and interventions are most effective.

Methods

This project's analytical approach combines multiple data sources to create a comprehensive dataset spanning from 2021 to 2025. The core hypothesis is that general project funding causally improves project security practices. Time-series data on project funding and repository security posture are the two key inputs necessary to answer this question. This methodological approach addresses several key challenges, including the heterogeneity of funding flows, the difficulty of measuring security practices over time, and accounting for confounding variables and measurement irregularities that could affect causal analysis.

Funding data

OSS projects vary widely in their organizational structures, from individually maintained projects to large multi-contributor efforts supported by formal foundations or corporations. Such diversity complicates efforts to identify funding recipients and measure funding flows consistently across projects. This analysis focuses on funding via Open Collective, a platform with significant use throughout the OSS ecosystem that provides financial infrastructure for grassroots projects and organizations. Open Collective enables third parties, from companies and foundations to individuals, to donate to projects on a recurring or one-time basis and serves as a host for funds that flow through other mechanisms, such as (some) GitHub Sponsors payments or corporate donations, similar to a publicly viewable bank account. Open Collective offers several advantages for research purposes: standardized, public transaction records, consistent reporting formats, and clear linkages between funding recipients (called "collectives") and specific OSS repositories. Its relationship with the Open Source Collective, a group that provides legal and financial infrastructure to OSS projects, making it easier to use the Open Collective platform, helps reduce

[20] Stewart Scott, *Counting the Costs: A Cybersecurity Metrics Framework for Policy*, *Atlantic Council*, May 6, 2025, https://www.atlanticcouncil.org/in-depth-research-reports/report/counting-the-costs/.

selection bias from using the latter as a primary data source related to barriers to setting up a collective to receive funding for OSS projects in the first place.[21]

However, this focus on Open Collective necessarily excludes other funding mechanisms, such as direct corporate sponsorship, some foundation grants, contracted security or maintenance work, or direct donations that do not flow through the Open Collective platform. It also omits non-financial contributions and in-kind donations, such as developer time provided by firms to projects they depend on, and can occasionally double-count refunds as contributions. Finally, Open Collective transaction records generally do not detail when funding is conditional, as might be the case for some transactions, without deep examination of funder programs.

The study's selection of "collectives" for OSS projects relies on ecosyste.ms, a free service that connects the collectives managing funds to their specific OSS projects in the form of links to their public GitHub repositories.[22] This approach ensures entities receiving funding are those with accurate and direct connections to their OSS projects. Open Collective also provides detailed transaction histories, including the amount, date, and source of a donation, enabling time-series analysis of funding flows. To mitigate inconsistencies in project information based on the timing of data collection, the research performed re-runs on all analyses within approximately a one-month period alongside checks for parsing issues. All transactions include conversions to US dollars to ensure consistent comparisons between projects and amounts.[23]

The research uses the GitHub repository URLs for each collective to link clearly between codebases and funding accounts (Figure 1). There are some limitations to this approach: maintainers could work across projects that funding would affect, and collectives might hold links to extraneous projects, via ecosyste.ms, that funders did not intend to affect or that they did not target to receive funding.

[21] Joe Brockmeier, "Untangling the Open Collectives," LWN.net, March 8, 2024, https://lwn.net/Articles/964402/.
[22] The ecosystem.ms is a project of the Open Source Collective. It includes other code hosts, among them GitLab, which provided both the largest number of repositories and the best insight into their security practices. The research excludes collectives without GitHub repositories, which represent a negligible fraction of the total dataset.
[23] Most transactions were already in US dollars, however.

Figure 1: Diagram of funding flow and measurement sources for this project.

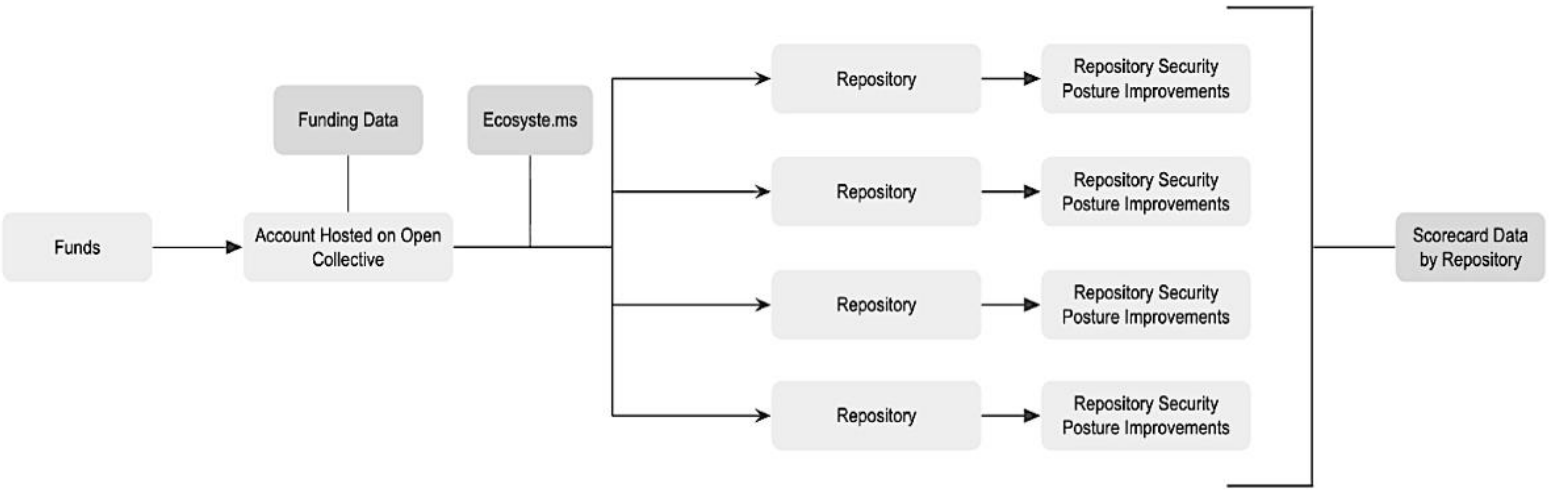

Open Collective accounts generally link to multiple repositories, and data matching funding from an Open Collective account to a specific project is sparse at best.[24] For this reason, the research tracks all affiliated projects but also flags which projects were most popular (as measured by GitHub stars) within each project cluster, as without further information, these are most likely to benefit first from funding for an organization or be the main repository associated with a funding collective.

Repository Information

This report uses GitHub's APIs to gather high-level information about the repositories identified through Open Collective and ecosystem.ms. The collection process captured dynamic and static repository data that the research hypothesizes could influence the relationship between funding and security practice changes. Static information includes overall cumulative funding, repository age, the relative composition of programming languages within repositories, and other metadata made available by the API. Dynamic information involves monthly commit and star histories as useful proxies for project activity and popularity, respectively. This data—gathered via GitHub's GraphQL API—avoids data limitations discovered in other methods, such as the Stargazers API, which only retrieves the first 40,000 stars, limiting the ability to gather monthly star data for the most popular projects in the dataset.

The usage of stars as a proxy for popularity merits additional discussion. The research views popularity to mean general public visibility and engagement rather than technical criticality and degree of usage. The hypothesis for the confounding pathway to funding is that a more visible project is likely to receive more attention, advocacy, and casual financial support. Meanwhile, the hypothesis of the relationship between popularity and scorecard improvement proposes that

[24] Very rarely, a transaction memo would provide more detail, but not often enough to produce useful data. For example, a donation marked as "for pizza party" would likely not directly contribute to security improvements, but it would be part of this dataset nonetheless.

greater visibility for a project increases the likelihood of attention paid to its Scorecard score. These two pathways are related to but distinct from the measurement of actual usage. Two other possible proxies for popularity—downloads and dependency data—face technical and causal limitations. Dependency data is generally only available for the current version of a project, with historical data contained in pull request history or updates to dependency log files, such as a requirements.txt file or a software bill of materials. Additionally, identifying critical OSS projects in terms of use is difficult, and there is no guarantee that widely critical projects are highly visible or well known.[25] Similarly, download count data tends to be limited to small timeframes or a finite number of releases or assets, and any specific cloning of a repository can indicate extremely variant types and scale of use, from individual curiosity to integration into an enterprise product. This research uses stars as a proxy for visibility, referred to throughout as popularity, based on the hypothesized causal pathways and data limitations outlined above.[26]

Security posture

This study utilizes OpenSSF Security Scorecards as the measure of security posture. Scorecards provide automated, standardized assessments of several security practices for GitHub repositories, including code review practices, dependency management, vulnerability disclosure processes, and security tooling adoption (Table 1). While this approach has limitations, it offers the advantage of a mostly consistent measurement that is trackable over several years across thousands of projects. Previous Atlantic Council research has discussed the limitations of Scorecards in greater depth, but at a high level, not all components of an overall Scorecard score are applicable to the security of all projects, and scores may not capture all the methods of satisfying the spirit of a Scorecard check.[27] Moreover, Scorecard measurements do not detail all useful security practices and behaviors. Nonetheless, it is the best population-scale tool with available public data. The study assesses Scorecard data pulled from the OpenSSF Scorecard BigQuery dataset, which provides historical score data for many OSS projects. Additionally, the rubric to assign Scorecard scores in OpenSSF's tool that generates them has changed over time. Given that the public dataset contains scores generated by different versions of the tool as it has evolved, the research employs the additional step of differencing all monthly repository scores by each month's average score across the entire dataset, controlling for tooling variance as best as possible without running a constant version of the tool on historical repository information.[28]

---

[25] Frank Nagle et al., "Census II of Free and Open Source Software – Application Libraries," The Linux Foundation, March 2022, https://www.linuxfoundation.org/research/census-ii-of-free-and-open-source-software-application-libraries.

[26] "REST API Endpoints for Repository Trraffic," GitHub Docs, version November 28, 2022, https://docs.github.com/en/rest/metrics/traffic?apiVersion=2022-11-28; "REST API Endpoints for Dependency Review," GitHub Docs, version November 28, 2022, https://docs.github.com/en/rest/dependency-graph/dependency-review?apiVersion=2022-11-28.

[27] Brackett et al., *O$$ Security*.

[28] "OpenSSF Scorecard Readme – Default Scorecard Checks," GitHub, accessed October 20, 2025, https://github.com/ossf/scorecard/blob/main/README.md#scorecard-checks.

Table 1: OpenSSF Scorecard subchecks and rubrics.

Name	Description	Risk Level
Binary-Artifacts	Is the project free of checked-in binaries?	High
Branch-Protection	Does the project use Branch Protection?	High
CI-Tests	Does the project run tests in continuous integration (CI), e.g., GitHub Actions, Prow?	Low
CII-Best-Practices	Has the project earned an OpenSSF Best Practices Badge—formerly part of the Core Infrastructure Initiative (CII)—at the passing, silver, or gold level?	Low
Code-Review	Does the project practice code review before code is merged?	High
Contributors	Does the project have contributors from at least two different organizations?	Low
Dangerous-Workflow	Does the project avoid dangerous coding patterns in GitHub Action workflows?	Critical
Dependency-Update-Tool	Does the project use tools to help update its dependencies?	High
Fuzzing	Does the project use fuzzing tools, e.g., OSS-Fuzz, QuickCheck, or fast-check?	Medium
License	Does the project declare a license?	Low

Name	Description	Risk Level
Maintained	Is the project at least 90 days old, and maintained?	High
Pinned-Dependencies	Does the project declare and pin dependencies?	Medium
Packaging	Does the project build and publish official packages from CI/CD, e.g., GitHub Publishing?	Medium
SAST	Does the project use Static Application Security Testing (SAST) code analysis tools, e.g., CodeQL, LGTM (deprecated), or SonarCloud?	Medium
Security-Policy	Does the project contain a security policy?	Medium
Signed-Releases	Does the project cryptographically sign releases?	High
Token-Permissions	Does the project declare GitHub workflow tokens as read only?	High
Vulnerabilities	Does the project have unfixed vulnerabilities? Uses the OSV service.	High
Webhooks	Does the webhook defined in the repository have a token configured to authenticate the origins of requests?	Critical

Aggregation and consolidation

This project's time-series analysis limits itself to projects appearing in both the Open Collective funding and OpenSSF BigQuery datasets with sufficient available data (generally a minimum of 24 months).[29] The research matches GitHub repository links associated with Open Collective collectives to repositories in the BigQuery Scorecard dataset and aggregates dynamic score and funding information, aligning monthly donation patterns with the pace of score measurement (approximately every other week) based on an average differenced score per month. A handful of months with missing score data for all repositories due to an issue with the public dataset tooling were backfilled from the most recent score, which sample figures indicate graphically with unfilled circle datapoints. While this interpolation introduces some uncertainty, the monthly aggregation enables longer-term trend analysis, minimizing the impact of short-term gaps. Moreover, the tooling gap was early in the dataset's time frame, allowing for a long series of uninterrupted, more recent (and more stable against tool variance) score data.

Finally, the research disregards repositories without 24 or more consecutive months of score and funding data and drops funding datapoints outside of the Scorecard data window of approximately 42 months, as well as a handful of repositories with name, ownership, and visibility changes that impeded data collection and interpretation, and repositories with no commit activity (i.e. no possibility of score change independent of measurement tool change). The result is a list of repositories with monthly funding and Scorecard data for the core analysis, along with several monthly and static datapoints for more sophisticated modelling. With this time-series data in hand, the research employs three methods of causal analysis—Granger tests, vector autoregression models, and panel tests—in addition to population statistics discussed in the findings section.

Granger tests

Granger tests are a simple approach to time-series causality that only look at whether a change in the independent variable (funding) predicts a change in the dependent variable (score) after a specific lag period. A statistically significant result is referred to throughout as funding Granger-causing security posture changes because no actual causal mechanism is observed, only a predictive relationship—a good first step to identifying causality, but not a comprehensive one. To select the lag period, the researchers apply a Bayesian Information Criterion (BIC) with an upper limit of six months to avoid overfitting, and for repositories failing to fit that model, they employ a fallback lag of one month.

Granger tests rely on stationary data inputs (i.e., data that does not have regular changes over time). An augmented Dickey-Fuller test (ADF) determined stationarity on the input data, which, if

[29] This time frame derives from the necessary sample size for the statistical analyses used, building into the study an assumption that a security improvement caused by funding will occur no later than 24 months after that funding is available to the project.

not stationary, was differenced until it was (i.e., changing from month-to-month scores and funding to month-to-month differences in scores and funding, and so on), before using z-score standardization to deal with the wide variety of scaling between variables (scores being between one and ten and funding, stars, and commits ranging up to the order of magnitude of hundreds of thousands). For each repository, a Granger test produces a p-value and a series of coefficients that allow for interpretation of the direction of Granger-causal effects (e.g., determining whether more funding Granger-causes positive or negative score changes in a certain repository).

Figure 2: Hypothesized causal relationship between funding and security posture improvements.

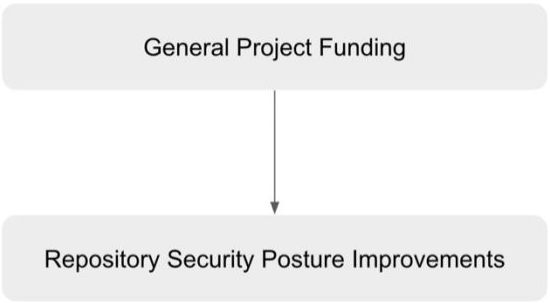

Vector autoregression

The vector autoregression analysis (VAR) extends the Granger tests by enabling the addition of confounders—variables that, might reasonably affect both funding and Scorecard score at the same time. Selected monthly GitHub stars and commit activity serve as proxies for popularity and repository activity, under the reasoning that a popular project is more likely to receive both funding and security posture improvements, and an active project is more likely to have advocates bring in funding and make security posture improvements (Figure 3). Otherwise, the process between the two tests is largely the same, including selection of lag time, differencing input data, and z-score standardization. Both analyses drop repositories without significant funding data, and, in the VAR analysis, for the repositories lacking significant confounder activity, the model leaves out that confounder variable. Under some conditions, VAR models can identify the effects of funding shocks and provide more in-depth analysis. However, given the inconclusive results and insufficiently long time-series data, this step was not taken as it would not have produced meaningful analyses.

Figure 3: Hypothesized causal relationship between funding and security posture improvements, including potential confounders of Repository Popularity and Activity.

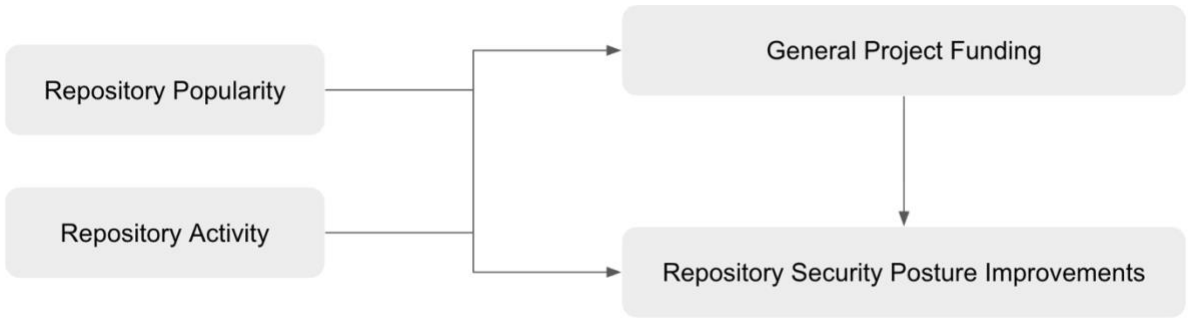

Panel

The panel analysis looks for a funding effect on scores across all repositories at once, preserving their intrinsic differences with a repository-specific variable (that would also account for differences in repository age). This approach includes stars and commits as confounders (Figure 4). An extension of this analysis also evaluates project age as an interaction variable, reasoning that the same amount of funding might have a different effect on older projects than on newer ones. For example, an older project might have a more robust maintenance infrastructure better able to optimize the funding, or first-time funding for a newer project might catalyze more robust and immediate security posture improvements (the research hypothesizes the latter—that newer projects would benefit more from funding, but the variable could reveal the inverse relationship or no relationship just as well). In addition, the research tests at both one- and two-month lag periods, as longer lags would overfit the model, with similar analyses on a subset of repositories that received at least $1,000 over their lifetime (reasoning that any less was likely insufficient to causally change security posture), producing eight total panel results.

Figure 4: Hypothesized causal relationship between funding and security posture improvements, including potential confounders of repository popularity and activity, as well as an age-funding interaction variable.

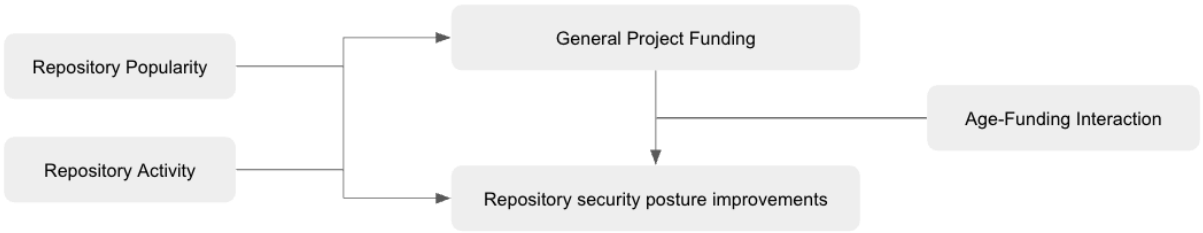

Findings

Population Statistics

This project analyzes 1,081 total collectives with a corresponding total of 4,676 GitHub repositories (Figure 5). These OSS projects had a total of $27,755,219.30 of funding flowing through Open Collective in a highly skewed distribution, with only 50 collectives representing 432 repositories receiving single transactions over $10,000 and just 633 collectives representing 2891 repositories receiving more than $1,000 total, indicating that while many projects receive small amounts of funding, a much smaller subset attracts substantial financial support. This pattern reflects broader trends in the OSS ecosystem, where a relatively small number of projects capture the majority of financial resources, and where a few projects are massively depended upon while most are less critical. This research does not examine the overlap between critical projects and well-funded ones. Figure 5 illustrates the long tail distribution of funding.

Figure 5: Funding distribution for Open Collective.

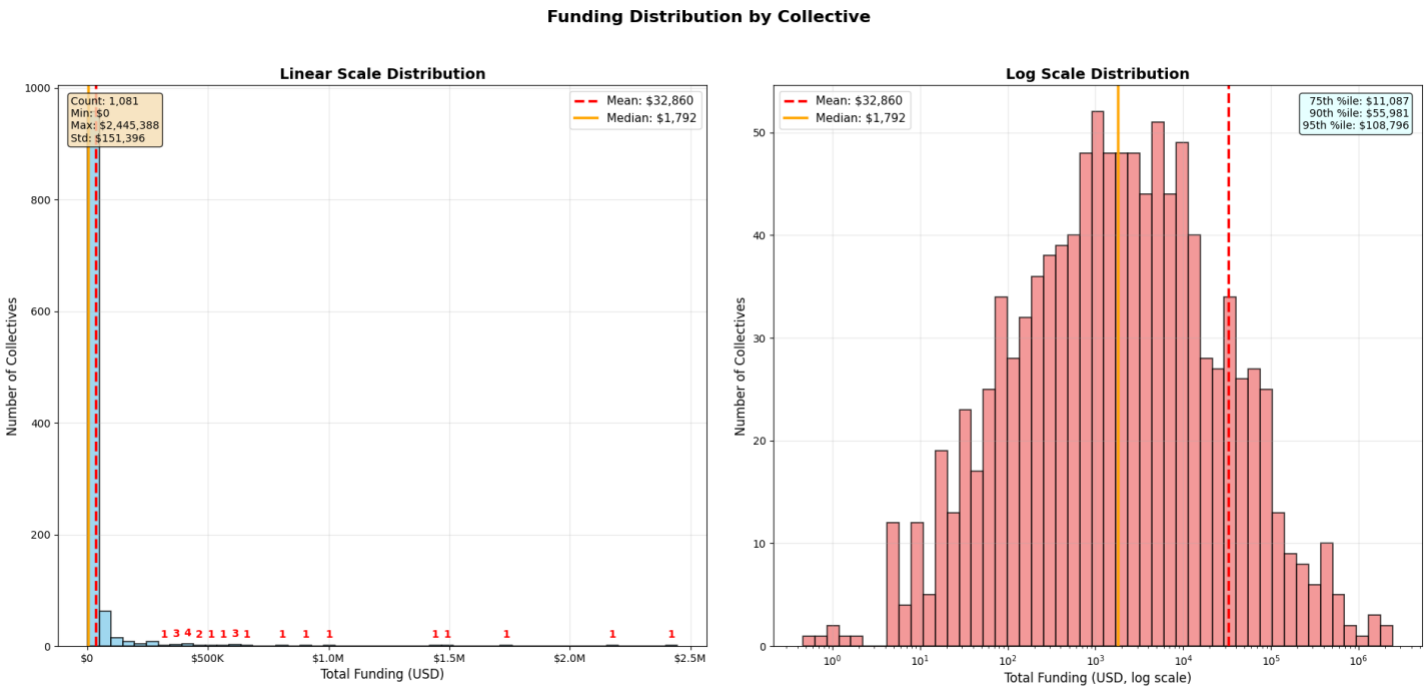

Repository popularity, as measured by GitHub stars, shows similar distribution patterns to funding, indicating the dataset captures a mix of highly popular projects and smaller efforts (Figure 6). Monthly funding patterns show considerable variation both within and across projects.[30] Some repositories receive steady monthly contributions, while others experience irregular funding spikes, and some experience both. This variability in funding patterns provides natural experiments for analyzing the relationship between funding changes and security improvements.

[30] This research does not directly measure correlations between funding and stars.

Figure 6: Repository popularity measured by GitHub stars.

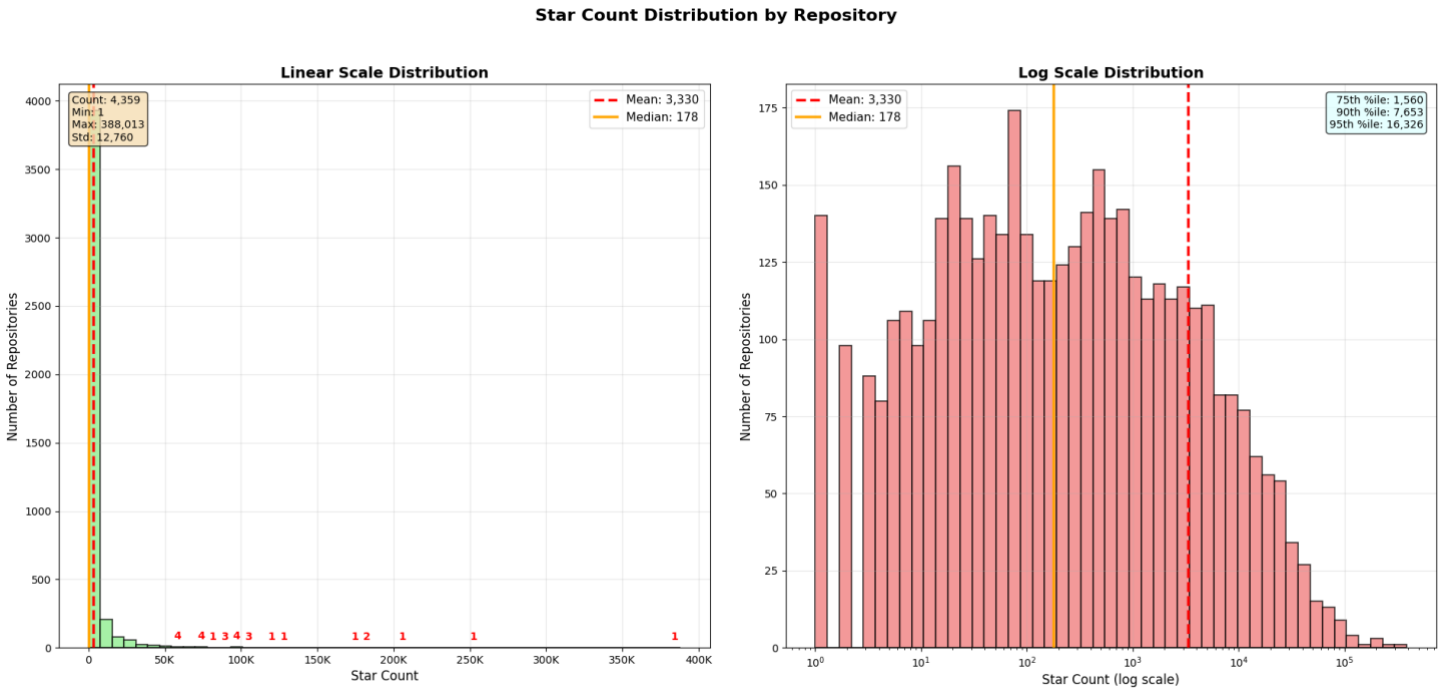

Scores distribute somewhat normally around a mean of 4 out of 10, with local clustering reflecting (Figure 7), at least in part, restrictions to the possible values of underlying subscores trickling into overall scores (e.g., some subscores are binary 0 or 10, and some are elements of limited integer ranges while each has a specific weight assigned to it used to calculate the overall score).[31] Over time, as the Scorecard tool evolved, average scores fell over the study range's first two years and stabilized for the remainder of the period (Figure 8).

[31] "OSSF/Scorecard: Check Documentation – Binary-Artifacts," GitHub, accessed October 20, 2025, https://github.com/ossf/scorecard/blob/main/docs/checks.md.

Figure 7: Score distribution showing mean of 4 out of 10.

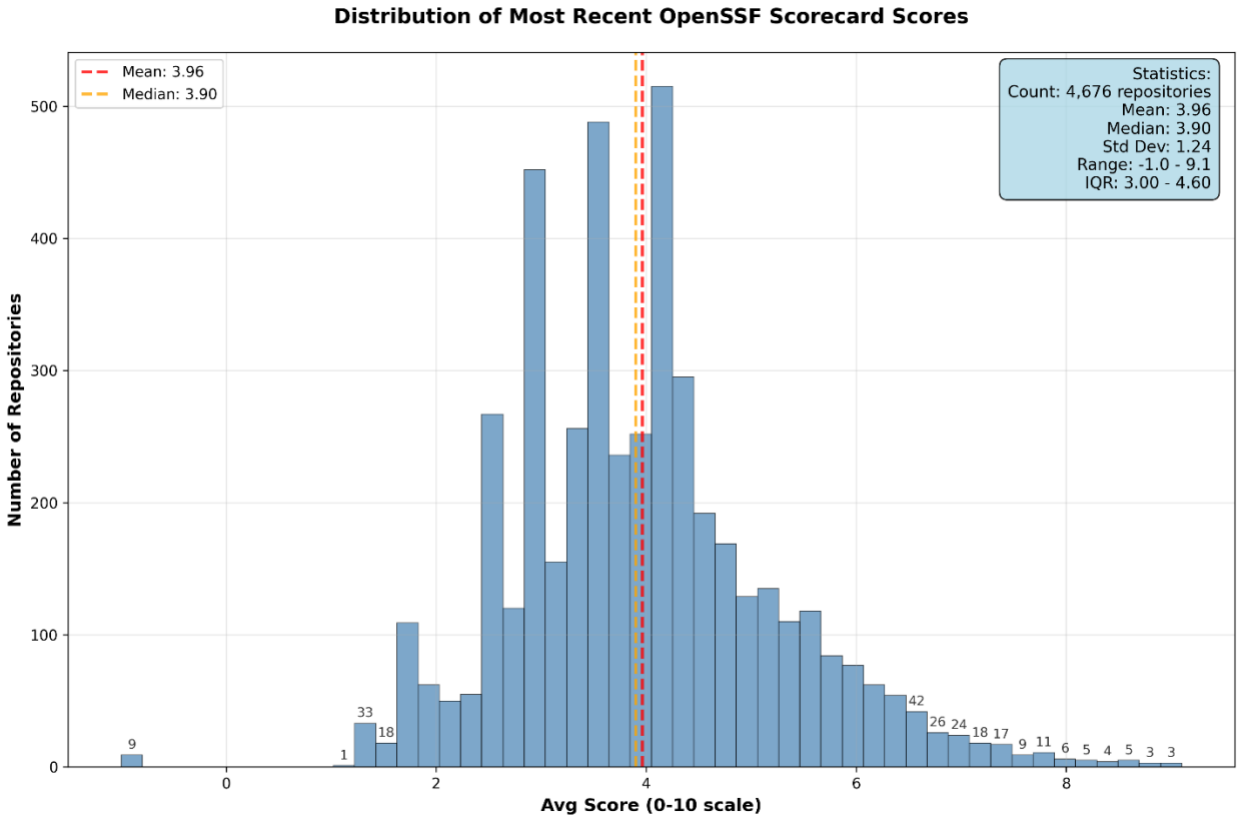

Figure 8: Scorecard trends with stabilization beginning in 2023.

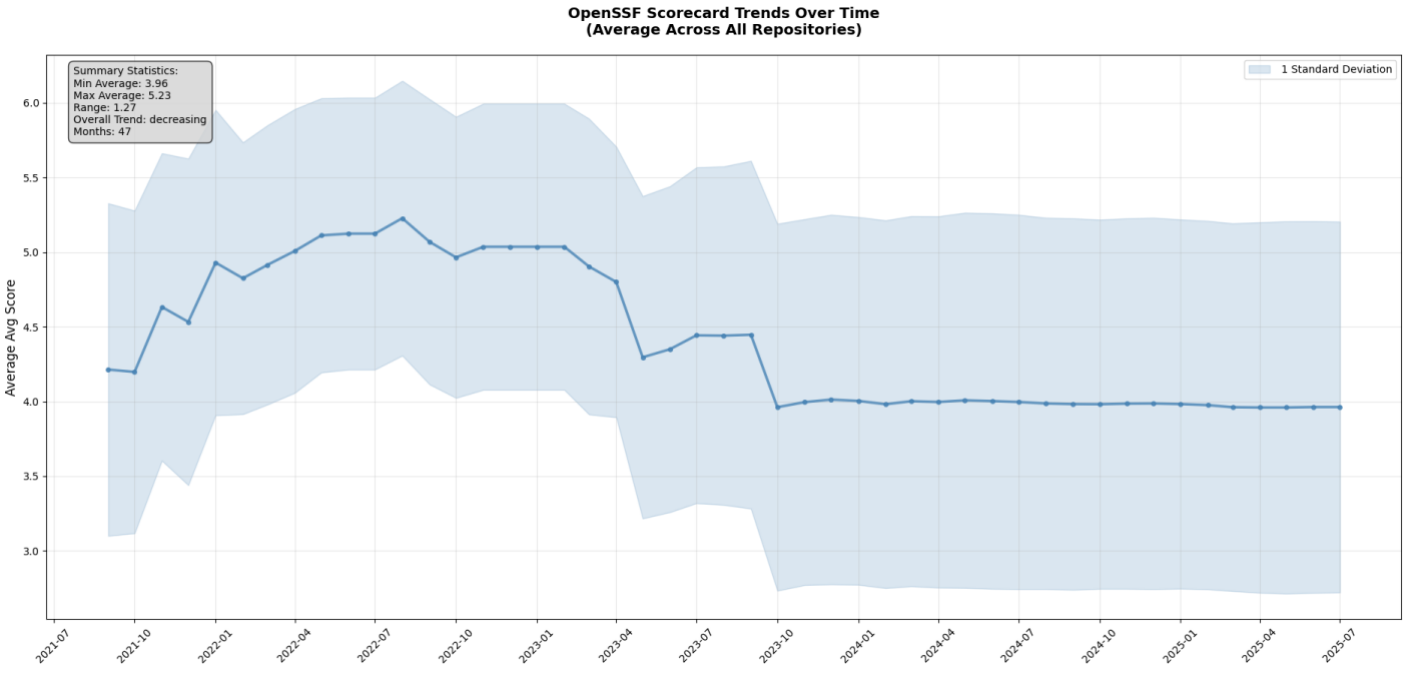

Importantly, the gathered data replicates earlier-mentioned findings that more funding generally *correlates* with better security posture. To measure this, the researchers apply a Spearman correlation between the two variables: collectives' total funding and the average of all their repositories' most recent scores (to ensure score assessment after the release of all measured funding, and at the most stable period of the Scorecard tool) (Figure 9).[32] The correlation between funding and score measured in this manner was a modest 0.2988 from a possible range of 1 to -1, with an extreme degree of statistical certainty.

Figure 9: Correlation between collective funding and repository scores.

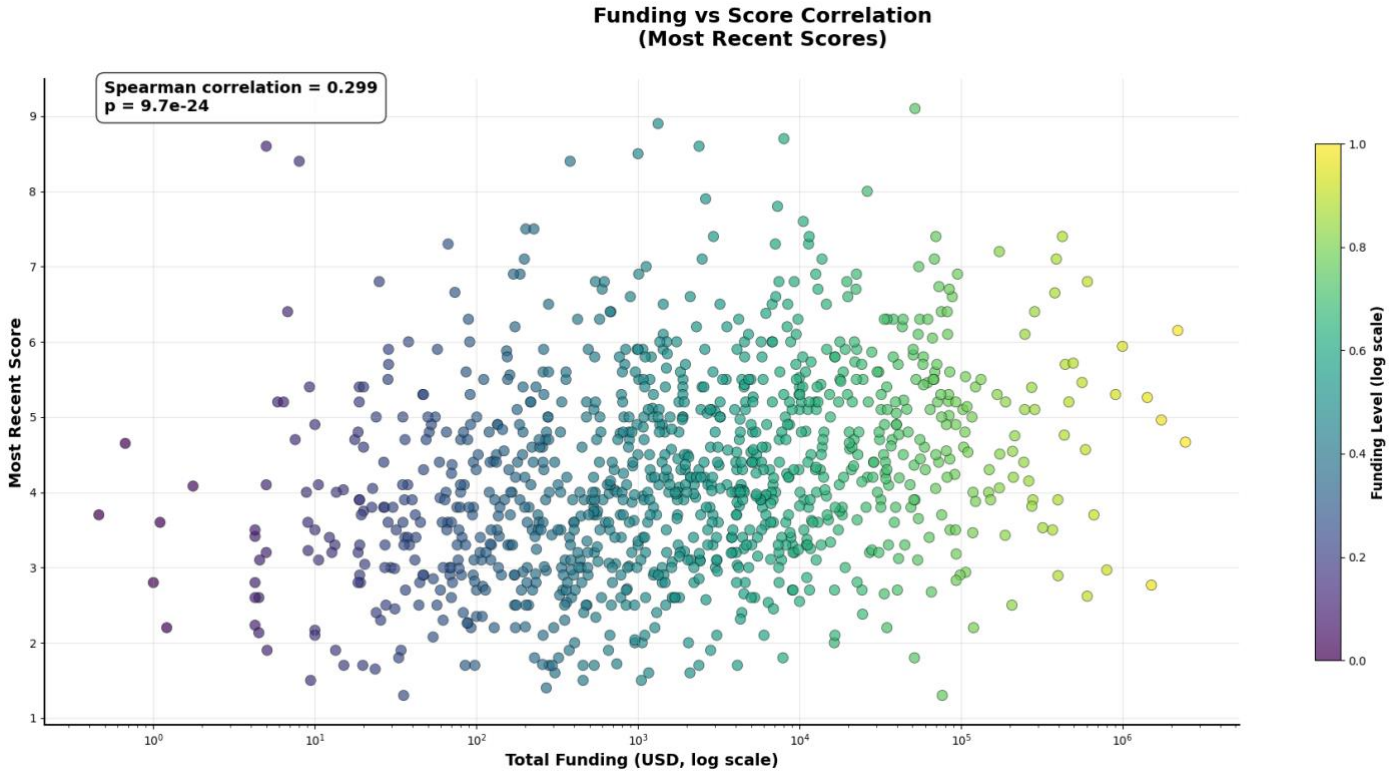

Granger Causality

Of the 4,676 repositories analyzed, 470 repositories demonstrate evidence of funding Granger-causing security score changes. However, while this is a small minority of the sampled

[32] The study uses the Spearman correlation to account for the extreme skew in funding, where most collectives had only a little funding, while a small number had massive amounts of financial support. Collectives with more funding tended to have more repositories associated with them, so grouping and averaging scores across repositories by collective prevented single well-funded collectives with many projects from skewing the results. For an explanation of the statistical tool, see: Elliot McClenaghan, "Spearman Rank Correlation," Technology Networks, https://www.technologynetworks.com/tn/articles/spearman-rank-correlation-385744.

repositories, there are important additional considerations to understand from the Granger test results. First, testing those Granger-causal repositories for "reverse causality" (i.e., score changes predicting funding changes, at the same time as funding changes predicting score changes) helps rule out repositories where the two variables merely correlate when both predict each other. Second, testing the direction of that prediction is critical, as "negative causality" is possible—when an increase in funding Granger-causes a *decrease* in score. Only 175 repositories demonstrate funding positively Granger-causing score changes. Filtering for repositories with $1,000 or more in total funding, or that were the most starred repositories associated with a collective reduces these total numbers further, but, in general, the Granger tests reveal that funding increases were only slightly more likely to predict score increases than score decreases, suggesting correlation across the sample more than Granger-causality.

Table 2: Granger Causality Analysis results–count and percentage.

Granger Causality Analysis: Funding vs OpenSSF Scorecard

Count of Repositories Demonstrating Evidence of Granger Causality

	Any Causality	Positive Causality	Negative Causality
Funding Causes Score	470	238	232
No Reverse Causality	324	175	149
Project has >$1k Funding	216	123	93
Most Starred Repository of Project	75	43	32

Granger Causality Analysis: Funding vs OpenSSF Scorecard

Percentage of Repositories Demonstrating Evidence of Granger Causality

	Any Causality	Positive Causality	Negative Causality
Funding Causes Score	10.05%	5.09%	4.96
No Reverse Causality	6.93%	3.74%	3.19
Project has >$1k Funding	4.62%	2.63%	1.99
Most Starred Repository of Project	1.6%	0.92%	0.68

Moreover, close examination of specific repositories reveals that several only show causality because of decreases in funding or that, while some funding fluctuations predict large score changes, other larger score changes occur with no funding fluctuations and some funding spikes had no effect on score changes (note that the score axis is not absolute score but variance from the standardized monthly average).

Figure 10: Granger Causality Analysis—sveltejs/svelte-loader net coefficients and data differenced.

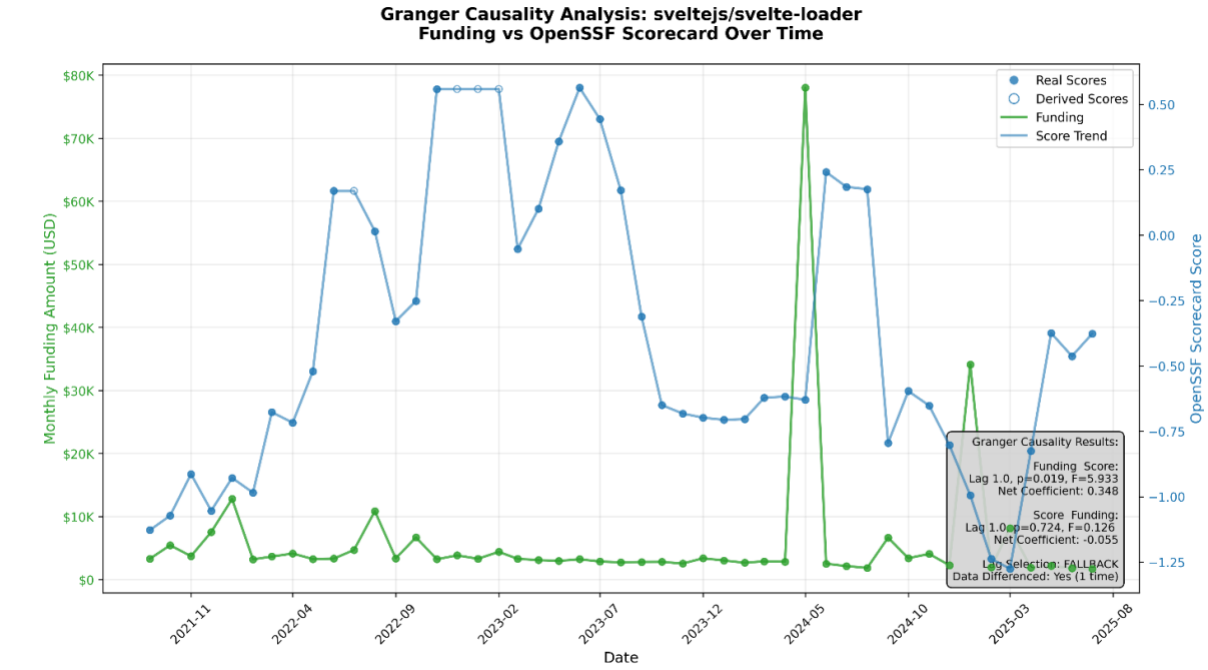

Figure 11: Granger Causality Analysis—vue-stripe/vue-stripe net coefficients and data differenced.

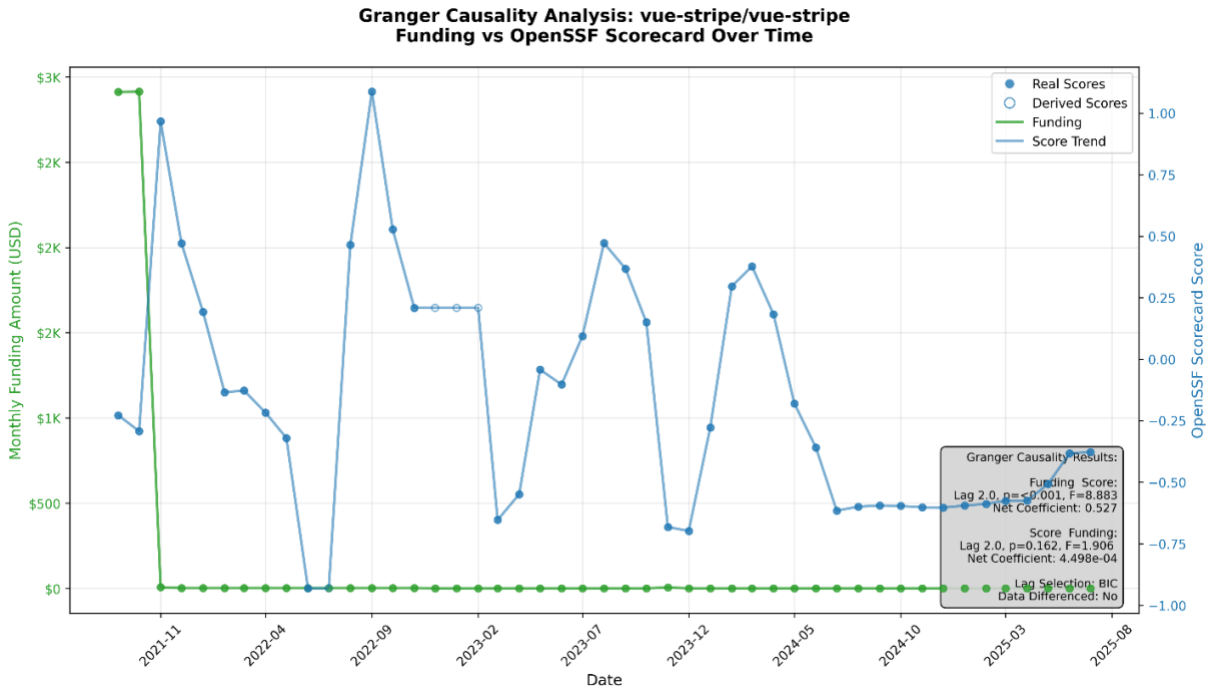

These results are generally not promising for demonstrating the hypothesized funding-score causality—the number of repositories demonstrating causality in the hypothesized direction and without concurrent reverse causality is barely above what one would expect from random noise (more on this later). Even more useful to note, too, is that some of the most statistically robust repositories are not actual codebases, but usage of GitHub as a data storage platform, with one example containing a simple list of platforms supported by Node.js under a collective with hundreds of repositories and receiving more than $120,000 in total funding (most of the collective's projects were more traditional "software" than this example). This conveys the predictive rather than strictly causal nature of Granger tests—the collective had one significant funding spike and many projects, so any well-timed fluctuation in score for any of its more than 700 repositories resulted in a statistically valid result, regardless of the underlying content and context of the project. Moreover, the volume of repositories across the entire dataset made manual cleaning of non-code repositories or similar contexts impractical.[33]

VAR

The VAR analysis controls for two confounding variables, project popularity and activity levels. It shows similar relationships between funding and security as the simple Granger tests: similar ratios of directly and inversely predictive results, a similar portion of predictive results displaying bidirectional causality, and similar proportional cuts as more constraints (minimum net funding and most starred projects) were added. This suggests that modelling in confounders weeds out a few repositories from the Granger results while at the same time adding others by overfitting a model at the lower end of statistical power (a maximum of 42 observable months). Notably, the overall decrease in causal results from the VAR model, which is generally more prone to overfitting, further suggests a null finding in terms of funding causing score improvements. Importantly, all the models containing confounders were unstable to some degree, precluding deeper analysis, and many could only model one or neither of the confounders due to sparse data (e.g., too few stars or commits or both), stationarity, or general null results. Table 3 summarizes the results, which again show only a small fraction of the overall sample displaying even statistical causality, let alone true causality.

[33] Future research could automate much of this by checking for repository language statistics and dependency data to automate identifying non-code repositories, though there are some edge cases—for example, the Node.js supported platforms list repository, which returns JavaScript as the primary language field. Similarly, some repositories returned blank primary language fields despite being large code bases.

Table 3: VAR Confounder Analysis results–count and percentage.

Vector Autoregression Analysis: Funding vs OpenSSF Scorecard

Count of Repositories Demonstrating Causal Results from VAR Model

	Any Causality	Positive Causality	Negative Causality
Funding Causes Score	363	186	177
No Reverse Causality	241	134	107
Project has >$1k Funding	204	115	89
Most Starred Repository of Project	40	22	18

Vector Autoregression Analysis: Funding vs OpenSSF Scorecard

Percentage of Repositories Demonstrating Causal Results from VAR Model

	Any Causality	Positive Causality	Negative Causality
Funding Causes Score	5.15%	2.87%	2.29
No Reverse Causality	4.36%	2.46%	1.90
Project has >$1k Funding	0.86%	0.47%	0.38
Most Starred Repository of Project	1.6%	0.92%	0.68

Overlap

Examining the overlap between repositories showing significant results in both Granger and VAR analyses provides the most conservative estimate of funding effects (Table 4). Only 265 repositories present funding-to-security causation in both analyses, with 165 showing this relationship without reverse causation, and 89 in the hypothesized direction (i.e., more funding predicting better scores after the lag period). When applying the same funding and popularity filters, these numbers drop to 139 and 16 repositories, respectively. This overlap analysis suggests that robust evidence for funding improving security exists for less than 4 percent of repositories in the sample.

Table 4: Granger and VAR Overlap Analysis results—count and percentage.

Overlap between Granger and VAR Analysis: Funding vs OpenSSF Scorecard

Count of Repositories Demonstrating Evidence of Granger Causality and Causal Results from VAR Model

	Any Causality	Positive Causality	Negative Causality
Funding Causes Score	265	129	136
No Reverse Causality	165	89	76
Project has >$1k Funding	139	76	63
Most Starred Repository of Project	16	8	8

Overlap between Granger and VAR Analysis: Funding vs OpenSSF Scorecard

Percentage of Repositories Demonstrating Evidence of Granger Causality and Causal Results from VAR Model

	Any Causality	Positive Causality	Negative Causality
Funding Causes Score	5.67%	5.09%	4.96
No Reverse Causality	3.53%	3.74%	3.19
Project has >$1k Funding	2.97%	2.63%	1.99
Most Starred Repository of Project	0.34%	0.92%	0.68

While these effects may still be practically significant for the affected repositories, the results do not support a broad claim that general funding systematically improves OSS security across the ecosystem. Where there is a causal effect, it is likely highly contextual. Moreover, the distribution of the repositories demonstrating statistical causality does not imply any underlying connection to funding totals or transaction frequency. In other words, statistically, the number of causal repositories is well within the comfort level of noise, with two exceptions. The figures in the appendix illustrate this by highlighting Granger causal repositories (i.e., without reverse causality or an inverse predictive relationship) over the log-distributions of total funding and transactions for the entire dataset. Using a log-distributed base helps adjust for the strong long-tail skew of funding and transactions, showing that a relatively consistent portion of any bin of either funding or transactions demonstrates causality.

A single Open Collective account linked to 760 repositories and with $122,806.30 of total funding drives the main exception to this finding, with much of its funding concentrated in a single month of outsized contribution in mid-2022 (Figure 12). This large funding spike and the large number of repositories tested against it during the more unstable Scorecard score period across the dataset likely returned an outsized number of false positives, particularly given that (a) this Open Collective account associates with a large number of non-code repositories (e.g., a list of recommended science fiction novels or a compendium of ASCII art depicting cows) and (b) the common score improvement throughout many of the repositories sits well outside any tested (or indeed testable) lag window and likely reflects a common change to the repository owner's default settings and practices that improved Scorecard scores at scale. Two other Open Collective accounts with dozens rather than hundreds of associated repositories demonstrate similar behavior on a much smaller scale. Otherwise, all causal results across the dataset are well within or below the limits of expected noise, though it is incorrect to assume that therefore no repository demonstrated funding causing security improvements.

Figure 12: Granger Causality Analysis—Net Coefficients and Data Differenced.

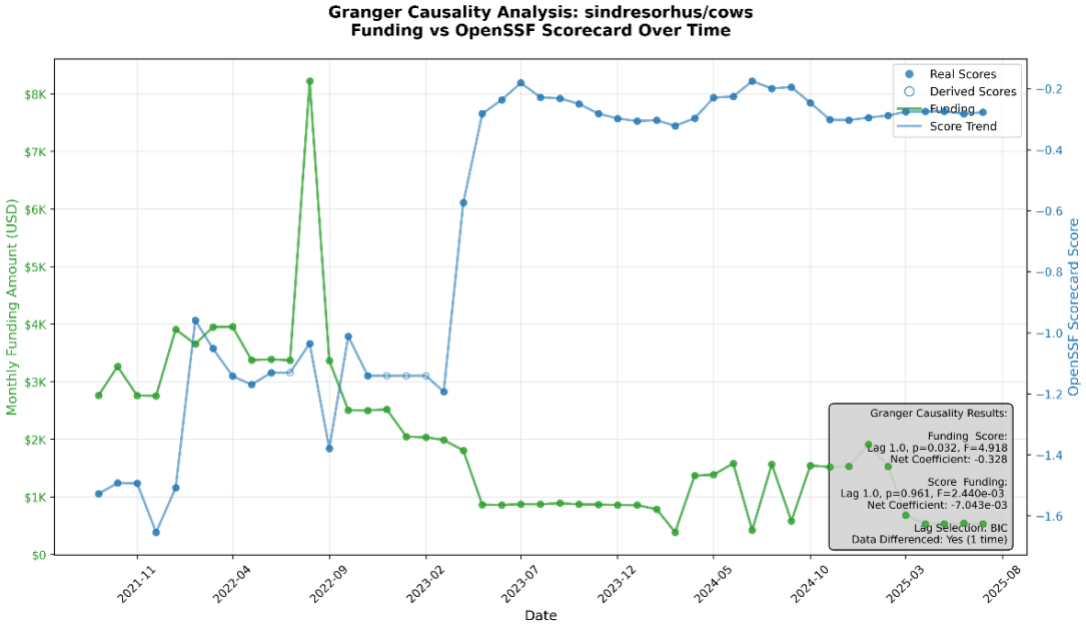

Panel

The panel analysis provides the most comprehensive test of funding effects by examining all repositories with at least 24 months of available data while controlling for heterogeneity across repositories. Several versions of the analysis were carried out. One alternately varies the monthly lag (denoted as p) between 1 and 2, one toggles the inclusion of a project-age interaction variable to see whether funding affected older and newer projects differently, and the final only includes projects with a total amount of funding over $1,000. Across these, no results suggest that funding led to meaningful, statistically significant effects for project security scores. Moreover, the analyses without the funding cap reveal statistically significant and near-zero funding effects, rather than insignificant results. Figures 13 and 14 depict the findings:

- A statistically significant yet near-zero effect of funding on security.
- A marginally larger but similarly near-zero effect from monthly changes in stars.
- No meaningful project-age funding interaction effect (either insignificant or significantly zero).
- And most importantly, a consistently strong effect from commit activity per month across panel specifications, suggesting that project development activity is the only measurable cause of change in security posture in this research.

It is useful to note that the coefficient values are in terms of standard deviation percent per additional standard deviation in funding, i.e., a standard deviation change in funding produced less than 1 percent of a standard deviation in score change after the lag, while a standard deviation in commits produced around a 10 percent of a standard deviation in score change in score in the same direction as the change in monthly commits.

Figure 13: Panel Regression Analysis.

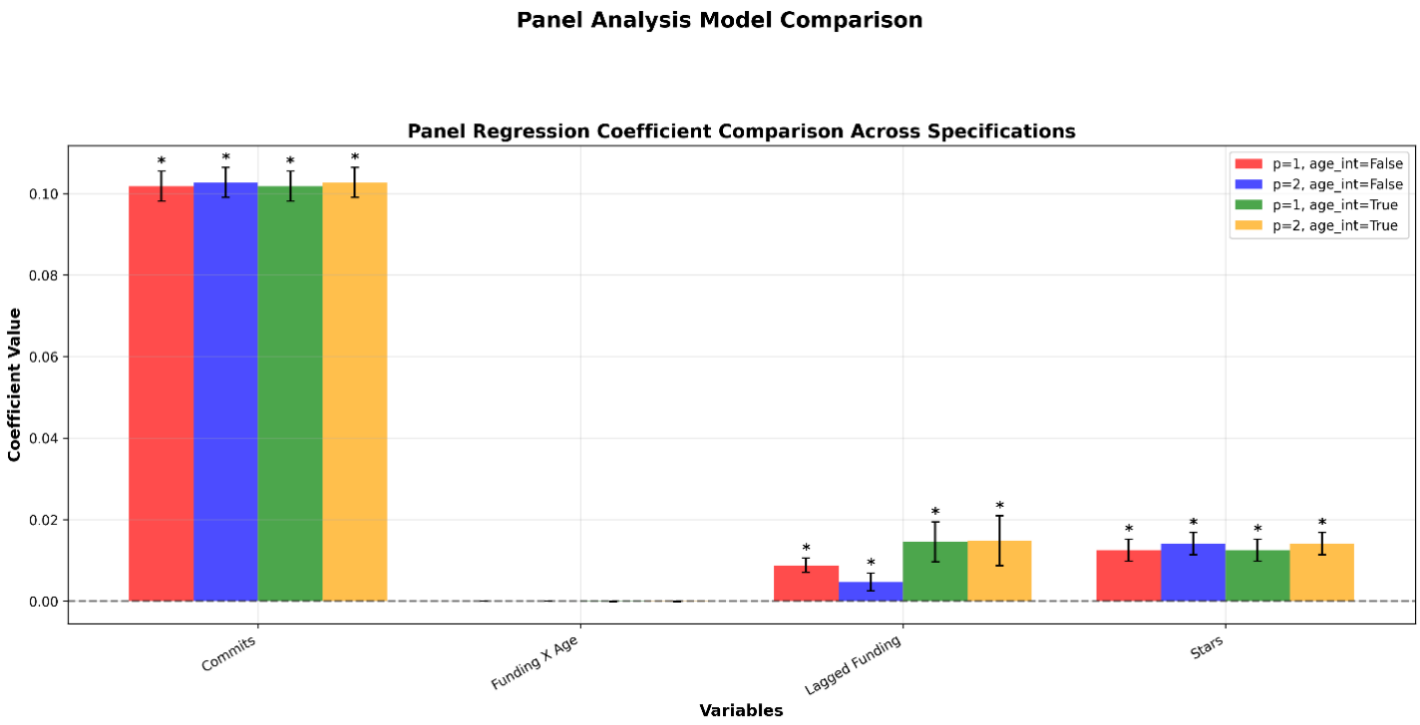

Model Specification	R^2	Observations	Entities
p=1, age_int=False	0.0108	306,532	4,675
p=2, age_int=False	0.0111	301,857	4,675
p=1, age_int=True	0.0108	306,532	4,675
p=2, age_int=True	0.0111	301,857	4,675

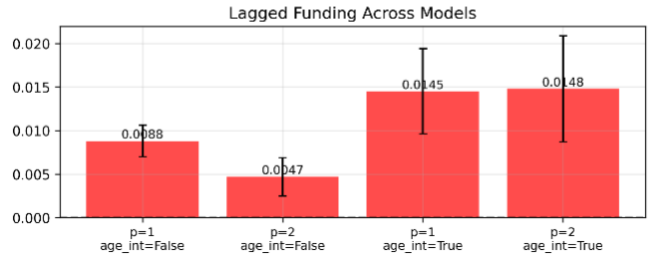

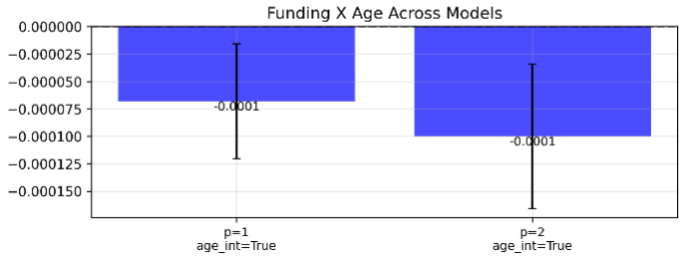

Figure 14: Statistics Comparing Model Coefficients.

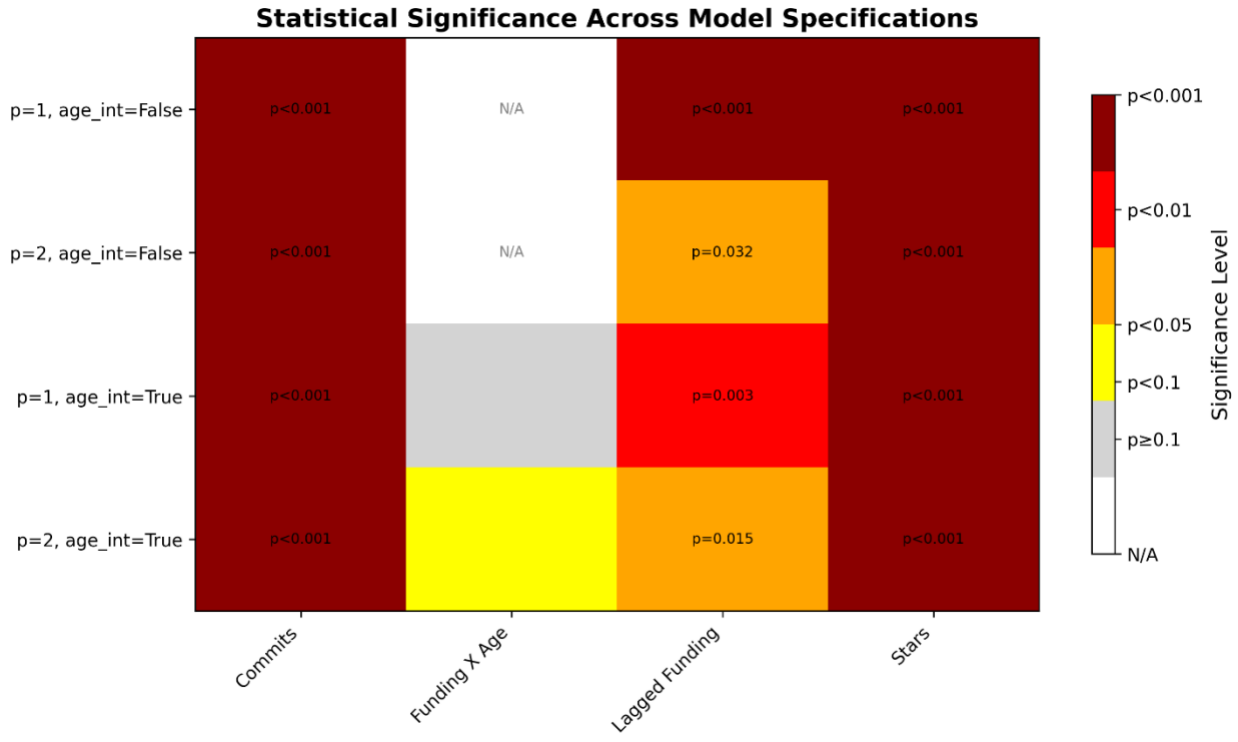

Analysis and Discussion

Granger tests of causality within repositories hold significant results in the hypothesized direction at a rate only slightly above what would be expected from random noise when distributed across net funding distributions. Further, adding confounders into VAR models slightly decrease the number of causal-significant results, the opposite of what one would expect if there were a non-null result, as the increase in model complexity would generally tend towards overfitting. The panel results provide a more definitive finding—across a variety of lag parameters and models (including a minimum $1,000 funding amount, which the appendix summarizes), they reveal a statistically significant near-zero impact of lagged funding (less than 1 percent of a standard deviation). Moreover, simple observation of score and funding charts over time, alongside the underlying content of some repositories, show a variety of circumstances that, while mathematically causal, are plainly noise. Finally, the fact that the panel result produced statistically significant relationships between commits—i.e., changing the content of repositories—and improvements in security posture adds extra heft to this finding: that the best intervention for improving security posture is to change the underlying codebase and its surrounding infrastructure directly.

While this project does not analyze the motivations of the wide variety of funders that route donations through Open Collective, including corporations, foundations, portions of GitHub sponsors funding, and individuals, the goal of Open Collective itself is not to improve the security posture of projects. As discussed previously, while some research—including from the Atlantic

Council—has posited that there could be a relationship between general, unrestricted funding and improvements in security posture, the results of this study suggest that a causal relationship is not present, at least in the given sample. The correlation most studies reference, and that this one also identified, therefore, likely reflects that more mature, institutionally supported projects are more likely to both receive funding and follow Scorecard-measured security practices—particularly given that such projects often engage with the Scorecard project itself through OpenSSF.

Several factors may explain the lack of systematic funding effects observed in this analysis. First, the assumption that general financial resources are the primary, direct constraint limiting security improvements in OSS projects may be incorrect. Many security practices measured by OpenSSF Scorecards might simply be unknown or unfamiliar to maintainers, or too burdensome to implement without a level of funding rarely reached in this dataset. More realistically, funding might allow a project to continue to exist and be actively maintained—both critical preconditions to improving security—but not mechanically cause those improvements in security posture directly. The causal pathways shown in DHS's Hardening Open Source Program and Tidelift's 2023 study imply that access to tooling, security-specific knowledge, and direct security incentives are far stronger drivers of security posture improvement than plain cash.

Second, rather than simply the amount provided, the effectiveness of funding may depend critically on its allocation and use. General funding without specific requirements or guidance for security improvements might cover other project needs, such as feature development, documentation, community management, or simply helping a maintainer pay their bills. While these could reasonably have security benefits down the line, the indirect nature of that effect would introduce heterogeneity in effect lags and reasonably could extend those lags beyond what is measurable in the current dataset. Some projects may lack the security expertise necessary to directly translate any financial resources into meaningful security improvements as well.

Third, the relationship between funding and security practices may be highly contextual, depending on project characteristics, such as size, complexity, contributor expertise, and existing security practices. Funding may benefit small projects maintained by individual developers differently than large projects with multiple institutional contributors, existing security resources, and contractual agreements not conveyed through Open Collective transaction records, as this research did not attempt to exhaustively measure an individual project's capacity to use funding. For example, a small project might be more likely to distribute unconditional funding directly to its maintainers, whereas a larger project might dedicate that resourcing to direct project changes. Moreover, documentation and observability of security improvements likely vary across projects, meaning that some funding might have contributed directly to security improvements that the Scorecard tool simply didn't pick up, with the possibility of bias toward measurement of improvements in projects aware of the tool regardless of actual underlying security improvement.

This study's methods and data sources also face several methodological limitations. Open Collective accounts only capture a portion of possible OSS funding, which can flow through traditional bank accounts, other payment platforms, or indirect subsidies for tooling and developer

time. Open Collective transaction data does not commonly specify funding use for projects. The limited observational timeframe also precluded analysis of lag periods greater than six months, while in practice, funding might sit in a bank account for far longer before being deployed toward security ends or take longer to produce posture improvements. Changes or outages in the Scorecard measurement tool add noise to the data sources, and Scorecard security practices do not fully encompass the means by which a maintainer can improve a project's security posture. For example, funding spent on developer time used to discover and remediate a vulnerability is unlikely to improve a project's Vulnerability subscore, while indiscriminately ignoring all known vulnerabilities in a project as flagged by the Open Source Vulnerabilities service would technically improve that project's score by simply including an osv-scanner.toml file signaling each to be ignored. Relatedly, commit data as used here does not contextualize the *content* of commits, which can vary widely in scope and is independent from the volume of commits made. Some projects might have fewer commits precisely because their codebases are relatively stable and secure.

Emphatically, these results should *not* be interpreted as dispositive for advocates of the open source community that wish to see additional funding flow towards the OSS ecosystem. Rather, this research suggests that financial resources are not the most proximate limiting factor for open source project security posture improvements—that does not imply that unconditional funding does nothing, but rather, that alone it does not improve security. Indeed, it is particularly likely that financial support is a useful if not necessary precondition for maintainers to dedicate additional time and effort to improving project security, as well as a normatively important one if the demand signal for improved OSS security originates primarily in that software's beneficiaries more than its creators and stewards. Policy interventions aimed at improving the security posture of critical projects should not focus only on funding, but also on targeted support, such as subsidized security audits and tooling, education, improvements to the security of default repository settings, and the provision of external developer support.

Future Research

OSS's criticality to modern digital systems demands a better grasp of open source project development needs and how funding impacts the ecosystem and its security posture. Despite the limited empirical analysis available, the open nature of the OSS ecosystem provides plenty of public data, which further study can aggregate and analyze to answer additional questions about the effects of funding. Even within this study's dataset, there is room for wider analysis on both the macro-scale and into specific repositories that are known to have made security improvements using funding through Open Collective.

One subject of future analysis that follows most directly from this research would be measuring the causal impacts of project funding on specific OpenSSF Scorecard subscores. The subscores of each project give detailed information about a series of checks performed automatically on GitHub repositories. Each subscore varies in importance and weight to derive the overall Scorecard score. It's possible that funding could have distinct effects on individual subscores lost in the noise, and the overall impact of different security practices on security outcomes remains understudied. For

example, scores in the Maintained subcheck, which examines the recency of project activity, could improve with funding that allows maintainers to dedicate more time to a project. Notably, the maintained subcheck had by far the strongest and most consistent correlation to funding in previous research and offers the most intuitive causal model.[34] Future research could directly examine each subscore to determine the individual effects of funding. However, subscores all measure differently, with some checks producing binary subscores or subscores with limited ranges (e.g., only integers 1, 3, 6, 9, and 10), posing data sparsity challenges over an already limited and variable timeframe of available data.

Additionally, Scorecard scores introduce a broader practical question for researchers—What is the best way to measure OSS security? This project does not examine the connection between security posture, or which security practices a project follows and adopts, and the project's security outcomes, which would be the actual security incidents experienced by a project and their impact on end users. Additional research to connect funding to changes in project attributes, such as security posture, and then to outcomes, would further direct intervention and investment in both open source software and cybersecurity writ large. For policymakers, this is a particularly key challenge. Directly incentivizing security practice changes is relatively straightforward—however, knowing which practices to incentivize and what the appropriate value of that incentive is against the value of the improvements in security outcomes those practices generate is far more critical and understudied.

Outside of Open Collective, alternative funding sources could involve different projects and maintainers, which might change the effect of funding on security posture. General funding flowing from different sources, such as GitHub sponsors or foundation donations, or on different recurring timeframes, such as monthly or yearly donations, could have distinct effects that would require parsing out in more detail through alternate methods examining funding flows. Additional interventions outside of general funding could be examined with the same methodology for their impact on security posture. One component of this project demonstrated the relationship between commit activity and improvements in security posture, and other connections could be similarly identified, with particular focus on the degree to which commit volume meets external demand for codebase changes.[35]

The previous research this project builds off found noticeable differences across package managers and language ecosystems when measuring the impacts of funding on project security.[36] That project examined those differences because data was heterogeneously collected from

---

[34] Brackett et al., *O$$ Security*.

[35] Kaylea Champion and Benjamin Mako Hill, "Underproduction: An Approach for Measuring Risk in Open Source Software," IEEE, 2021 IEEE International Conference on Software Analysis, Evolution and Reengineering, *2021 International Conference on Software Analysis, Evolution and Reengineering (SANER)*, (Honolulu: IEEE, March 2021), 388–99, https://ieeexplore.ieee.org/document/9426043.

[36] Brackett et al., *O$$ Security*.

different package managers by download frequency, but similar analysis could be performed on data gathered from funding sources, such as the data from this project. Examining differences in language ecosystem adoption of security practices and measuring if there are different implications of funding within those ecosystems could be useful information for policymakers looking to develop targeted interventions.

Conclusion

This study set out to answer a fundamental question in OSS security policy: Does general funding for OSS projects improve their security posture? After analyzing over 4,600 repositories and nearly $28 million in funding flows, this research suggests that there is not a general, direct causal relationship. While a small percentage of projects (less than 4 percent, depending on the test) show statistical evidence of funding leading to or predicting security improvements, these effects are not robust across methodological approaches and represent a small fraction of generally funded projects, and a dataset-wide panel analysis found precisely the opposite: that funding does little to directly improve security posture, while, intuitively, changing a code base has a strong effect on its security. The majority of repositories showed no detectable relationship between funding levels and security scores, even when controlling for confounding factors such as project popularity and development activity. These findings challenge the assumption that financial constraints are the primary barrier to better securing open source projects. Instead, they suggest that knowledge, expertise, and targeted interventions are more important in determining security outcomes than general funding levels.

For policymakers and organizations committed to improving OSS security, these results point toward more targeted approaches than broad-based funding programs alone. Security audits, tailored tooling, training programs, and conditional funding with specific security requirements (and the resourcing to achieve them) may prove more effective than unrestricted financial support in driving security posture improvements, and hopefully security outcomes. This emphasis could make comprehensive OSS security improvements more achievable while using available resources more efficiently. Some entities within the open source ecosystem do focus on security-specific support through a wide variety of activities: the alpha-omega project, GitHub's nascent Secure Open Source Fund, and the Open Source Technology Improvement Fund, to name a few.[37]

Rigorous evaluation of policy interventions, using methodologies like those developed for this study, can help ensure directing limited resources toward approaches that effectively achieve their goals. This analysis captures only one funding mechanism among many that support OSS projects and focuses on a narrow set of security practices rather than outcomes. Future research examining alternative funding sources, longer-term effects, and connections to actual security outcomes will

[37] "Alpha-Omega Project," The Linux Foundation, accessed October 20, 2025, https://alpha-omega.dev/; "GitHub Secure Open Source Fund," GitHub, accessed October 20, 2025, https://resources.github.com/github-secure-open-source-fund/; "Open Source Technology Improvement Fund: Securing Open Source for the World," Open Source Technology Improvement Fund, accessed October 20, 2025, https://ostif.org/.

further refine the understanding of how financial resources can most effectively improve OSS security. The path forward requires a willingness to test and adapt interventions in the open source ecosystem, which this research contributes to with empirical evidence about the relationship between general financial inflows and security posture improvements. Securing the open source ecosystem is a pressing challenge for the modern software systems that depend upon it, and ineffective interventions will only slow progress towards that goal.

Acknowledgements

The authors would like to thank John Speed Meyers, Cailean Osborne, Lauren Hanford, Paul Sharratt, Katharina Meyer, and Kevin Crosby for their review of a draft of this report. This report was supported by the Sovereign Tech Agency and American University.

Appendix: Figures

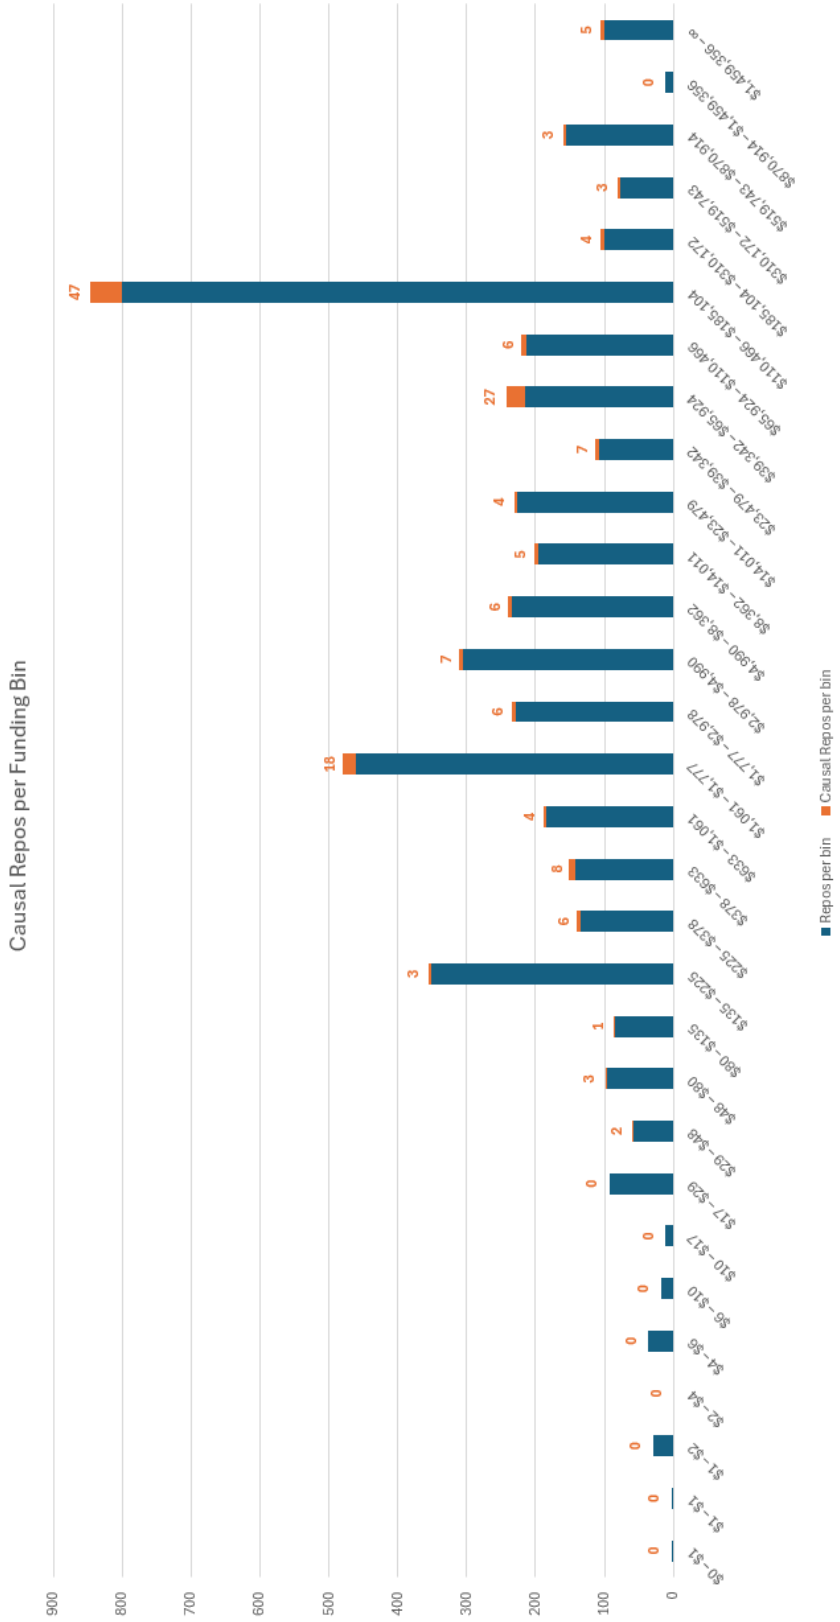

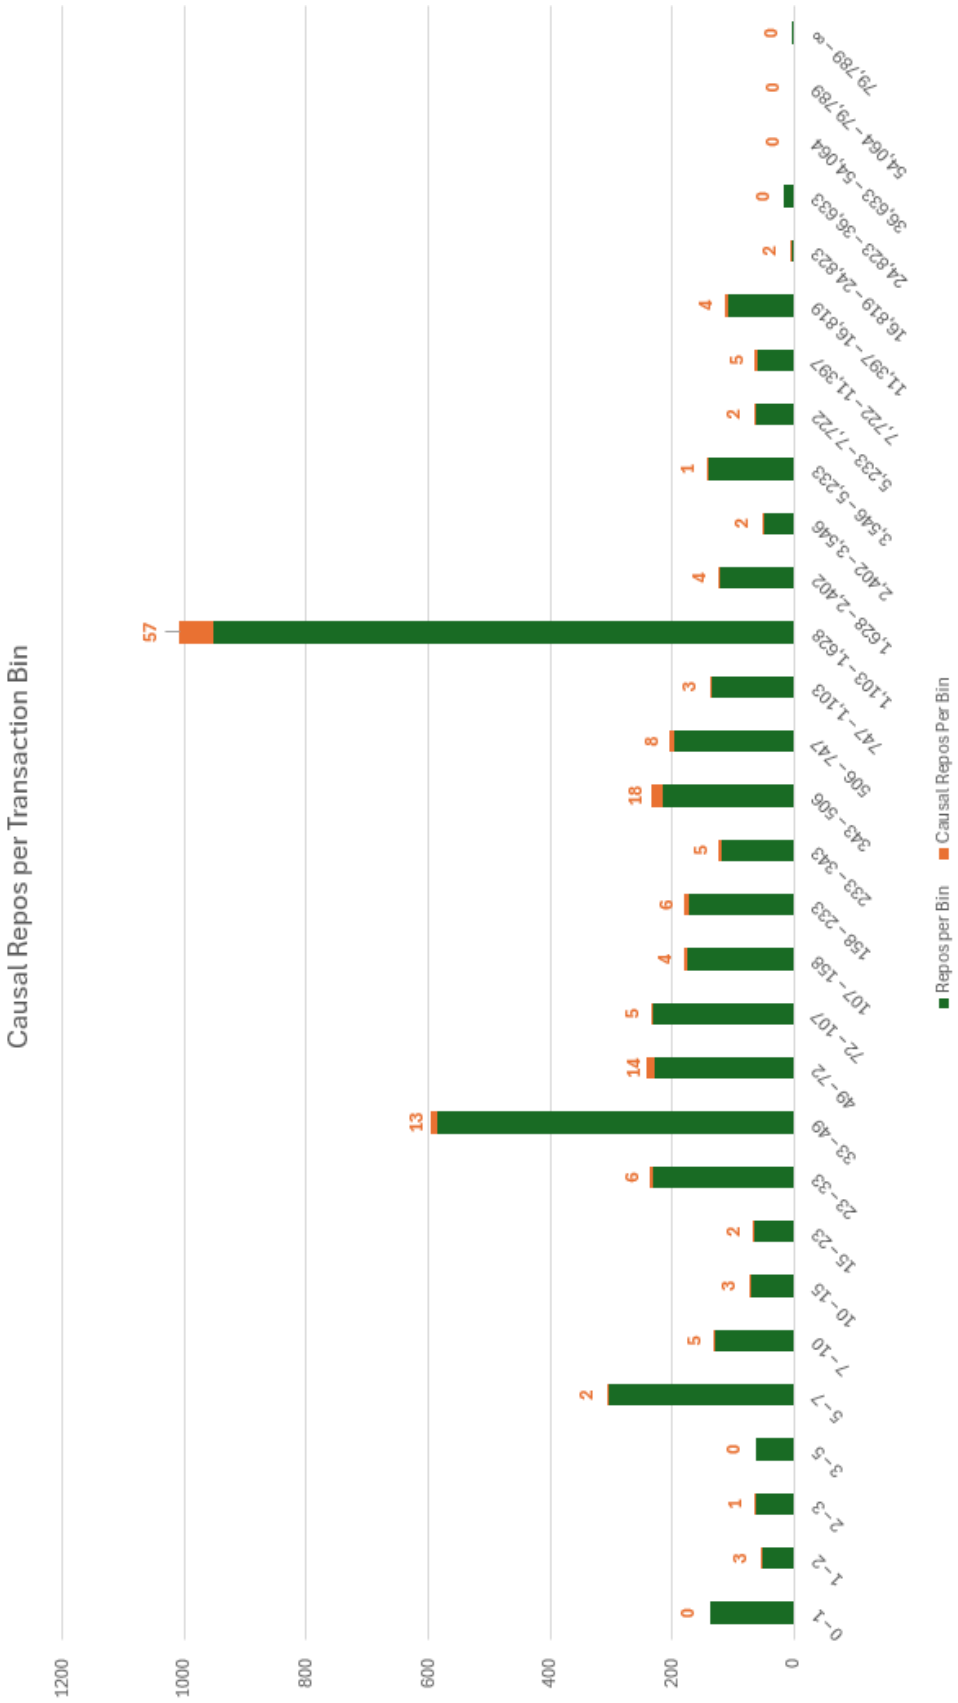
Causal Repos per Transaction Bin
Repos per Bin
Causal Repos Per Bin
0
200
400
600
800
1000
1200
0-1
1-2
2-3
3-5
5-7
7-10
10-15
15-23
23-33
33-49
49-72
72-107
107-158
158-233
233-343
343-506
506-747
747-1,103
1,103-1,628
1,628-2,402
2,402-3,546
3,546-5,233
5,233-7,722
7,722-11,397
11,397-16,819
16,819-24,823
24,823-36,633
36,633-54,064
54,064-79,789
79,789-∞

Panel Analysis Model Comparison

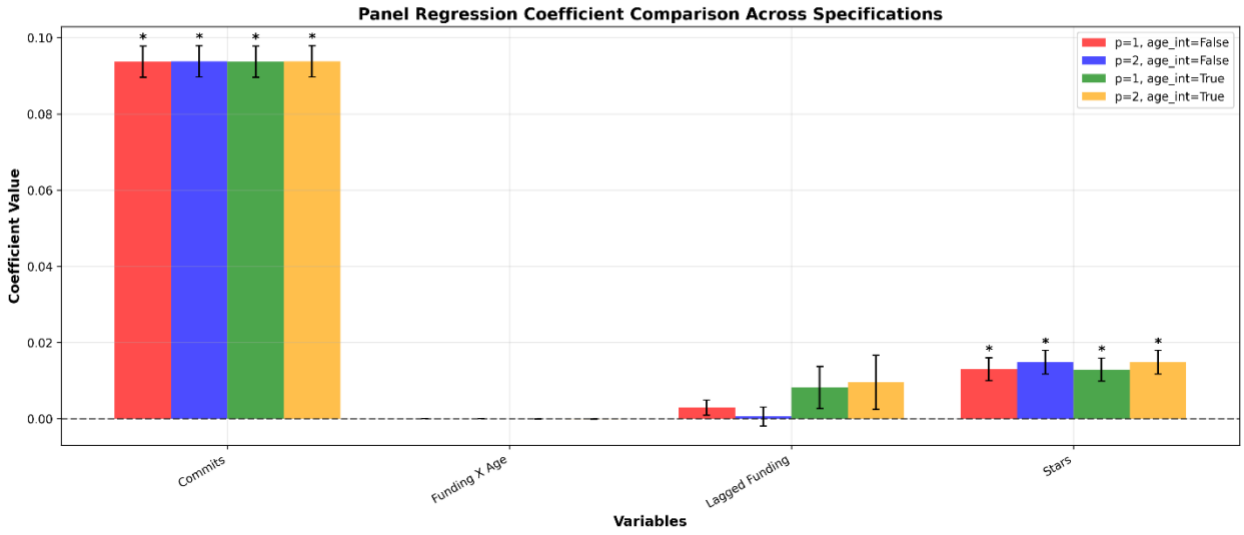

Model Specification	R^2	Observations	Entities
p=1, age_int=False	0.0094	234,422	3,364
p=2, age_int=False	0.0096	231,058	3,364
p=1, age_int=True	0.0094	234,422	3,364
p=2, age_int=True	0.0096	231,058	3,364

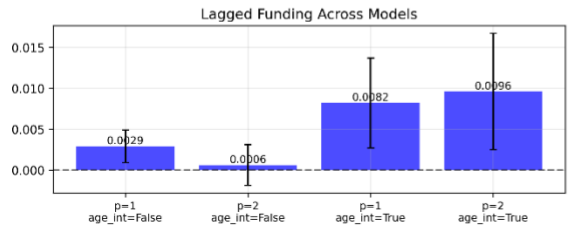

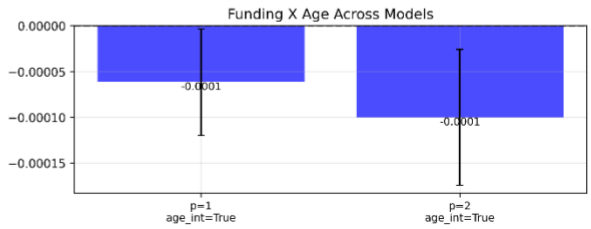

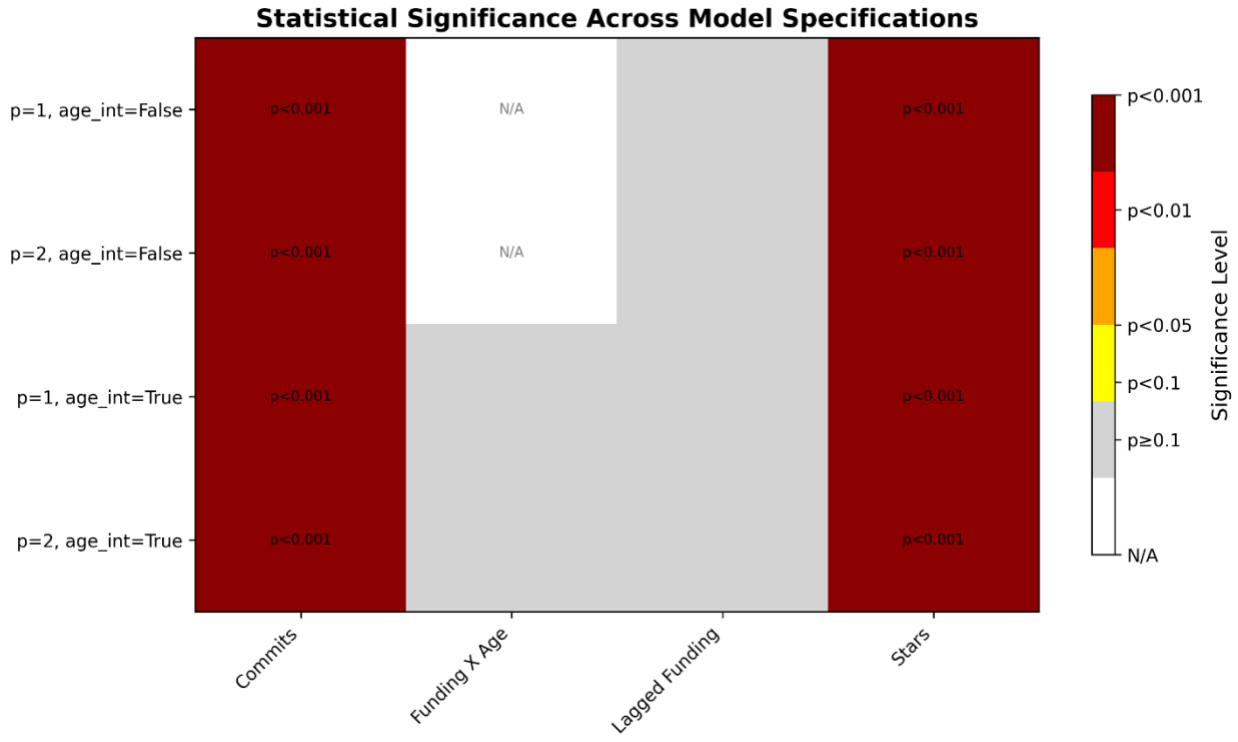
Statistical Significance Across Model Specifications
p=1, age_int=False
p=2, age_int=False
p=1, age_int=True
p=2, age_int=True
p<0.001
N/A
Commits
Funding X Age
Lagged Funding
Stars
p<0.001
p<0.01
p<0.05
p<0.1
p≥0.1
N/A
Significance Level

CENTER FOR SECURITY, INNOVATION, AND NEW TECHNOLOGY

Our mission is to examine the geopolitical implications of emerging technologies for security, democracy, and society. Based at American University's School of International Service, we conduct original research that advances theory and informs evidence-based policymaking, while also training the next generation of thought leaders and practitioners.

Visit the **Center for Security, Innovation, and New Technology** website for more information about our work and team.

For publication inquiries, please contact **csint@american.edu**