

Exploring the SECURITY.md in the Dependency Chain: Preliminary Analysis of the PyPI Ecosystem

Chayanid Termphaiboon[†], Raula Gaikovina Kula[§], Youmei Fan^{*},
Morakot Choetkiertikul[†], Chaiyong Ragkhitwetsagul[†], Thanwadee Sunetnanta[†], Kenichi Matsumoto^{*}

[†]Faculty of Information and Communication Technology, Mahidol University, Thailand

^{*}Graduate School of Science and Technology, Nara Institute of Science and Technology (NAIST), Japan

[§]Graduate School of Information Science and Technology, The University of Osaka, Japan

chayanid.ter@student.mahidol.ac.th, raula-k@ist.osaka-u.ac.jp, fan.youmei.fs2@is.naist.jp,
{morakot.cho, chaiyong.rag, thanwadee.sun}@mahidol.ac.th, matumoto@is.naist.jp

Abstract—Security policies, such as SECURITY.md files, are now common in open-source projects. They help guide responsible vulnerability reporting and build trust among users and contributors. Despite their growing use, it is still unclear how these policies influence the structure and evolution of software dependencies. Software dependencies are external packages or libraries that a project relies on, and their interconnected nature affects both functionality and security. This study explores the relationship between security policies and dependency management in PyPI projects. We analyzed projects with and without a SECURITY.md file by examining their dependency trees and tracking how dependencies change over time. The analysis shows that projects with a security policy tend to rely on a broader set of direct dependencies, while overall depth and transitive dependencies remain similar. Historically, projects created after the introduction of SECURITY.md, particularly later adopters, show more frequent dependency updates. These results suggest that security policies are linked to more modular and feature-rich projects, and highlight the role of SECURITY.md in promoting proactive dependency management and reducing risks in the software supply chain.

Index Terms—open-source software, security policy, software dependency

I. INTRODUCTION

Open-source software (OSS) plays a central role in modern development, with many projects depending on external packages from ecosystems like PyPI and npm. While this reuse speeds up development, it also introduces security risks. A recent study found that fewer than half of known vulnerabilities in PyPI packages are patched promptly, leaving many projects exposed [1]. This underscores the need for clear vulnerability disclosure processes. Given Python’s wide adoption and the scale of PyPI, it is a relevant setting to examine how security practices relate to dependency risks [2]. To support secure practices, platforms like GitHub have promoted the use of standardized security policies. In May 2019, GitHub introduced the SECURITY.md file as a dedicated place for security disclosure information [3]. It typically includes instructions for privately reporting vulnerabilities, often via email rather than public issue trackers. The goal is to promote trust, ensure consistent handling of vulnerabilities, and support faster response times [4]. SECURITY.md is now among the

top recommended practices for maintainers [5]. However, studies show that security-related issues are often harder to resolve, usually handled by a small group of developers, and left open longer than other issues. This highlights the value of structured policies like SECURITY.md in streamlining reports and responses [6]. Projects with SECURITY.md also tend to score higher in tools like the OpenSSF Scorecard, especially in areas such as branch protection and dependency update tools, suggesting a link between policy adoption and stronger project practices [4].

Managing dependencies is critical to maintaining open-source security. They expand the attack surface and allow vulnerabilities to spread across ecosystems, especially through transitive links. Prior studies show that many dependency vulnerabilities remain unresolved for months after disclosure, and a few maintainer accounts can become single points of failure [1], [7]. Developers’ response to these risks is influenced by how the information is presented; for example, visual dependency graphs have been shown to encourage more proactive updates than plain text reports [8]. While both dependency management and security policies are essential for secure development, past research has treated them separately, focusing either on structural risks or on reporting practices. However, intelligent software engineering aims to bring these elements together. It seeks to support developers with systems that make informed, proactive decisions based on project-level signals. Our goal is to explore how SECURITY.md, as a form of human-defined metadata, relates to dependency behavior. By connecting security policies with how projects manage and evolve their dependencies, this study provides a foundation for building intelligent approaches to support secure and maintainable software ecosystems.

To address this gap, we examine the relationship between security policies and dependency behavior in the PyPI ecosystem. This study considers both aspects together to understand how the adoption of SECURITY.md relates to the structure and evolution of software dependencies. By linking project-level policy signals with real maintenance activity, we aim to support future work in intelligent dependency management. These research questions lay the groundwork for intelligent

software engineering by exploring whether human-defined policies like SECURITY.md reflect meaningful patterns in how dependencies are organized and updated. Specifically, we investigate the following research questions:

- **RQ1:** How do the dependency trees differ between projects with and without a security policy?
- **RQ2:** Do projects with a security policy have more dependencies over time?

By answering these questions, our study makes the following contributions:

- We examine how the presence and timing of SECURITY.md adoption relate to the structure and maintenance of dependencies in open-source Python projects.
- We find that projects with a SECURITY.md file tend to have more direct dependencies, and those adopting the policy later exhibit higher rates of dependency-related commits.
- We observe that over 80% of dependency modifications occur in `setup.py` and `pyproject.toml`, highlighting their central role in managing and evolving dependencies.

These findings shed light on how security policies affect dependency management and support future improvements in secure, maintainable open-source software. The dataset and analysis scripts are publicly available.¹

II. BACKGROUND AND RELATED WORK

A. Security Policies in Open-Source Ecosystems

Security policies guide users and contributors on how to report security vulnerabilities. On GitHub, these are typically documented in a SECURITY.md file, placed in the root, docs, or .github folder. GitHub introduced SECURITY.md² to promote responsible vulnerability disclosure [3], [9]. Security policies help open-source projects handle vulnerabilities in a clear and consistent way. They provide reporting instructions that build trust and ensure issues are addressed through private, well-defined channels [3]. They also increase accountability by making maintainers' responsibilities visible and improve response times by guiding reports to the right place, reducing the risk of public exposure. A typical SECURITY.md file specifies how to report vulnerabilities, with email being the most commonly preferred method among maintainers [4]. According to Kancharoendee et al. [4], the content of SECURITY.md files can be grouped into categories such as reporting channels, response timelines, disclosure preferences, and scope of coverage, each reflecting how project maintainers define their security practices.

B. Software Dependencies Supply Chain in Open-Source Projects

The software dependency supply chain refers to the network of software packages where projects rely on others and are also

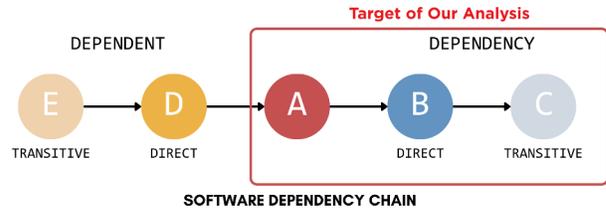


Fig. 1: An example of a software dependency supply chain, showing both dependencies and dependents

relied upon. As shown in Figure 1, project A is the center of our analysis. On the dependency side, if project A uses library B, B is a direct dependency of A. If B depends on library C, then C is a transitive dependency of A. Dependency depth measures the longest path from the focal project to any transitive dependency. On the dependent side, if project D uses project A, D is a direct dependent. If project E uses D, E becomes a transitive dependent of A. These relationships form a chain through which both functionality and vulnerabilities can spread. In this study, we examine only the dependency side. We analyze direct dependencies, transitive dependencies, and dependency depth to explore how the presence and timing of a SECURITY.md file relate to dependency structure and maintenance.

C. Related Works

Previous studies have examined security policies in open-source ecosystems, focusing on adoption and reporting practices. For example, Kancharoendee et al. [4] analyzed 679 PyPI projects, comparing declared reporting mechanisms with actual developer behavior and assessing quality using the OpenSSF Scorecard. Most projects favored private disclosure channels, such as email or GitHub advisories, although contributors often reported issues publicly. Projects with a security policy scored higher in branch protection, dependency update tools, and maintenance practices, highlighting a link between policy adoption and stronger security posture. However, the study did not examine dependency structures, leaving open the question of whether such policies relate to differences in dependency composition or evolution.

Regarding software dependencies in open-source software (OSS), several studies have highlighted their importance to project security. Prana et al. [10] showed that vulnerable dependencies can propagate risks across ecosystems. Their analysis of 450 Java, Python, and Ruby projects revealed that most vulnerabilities persist for long periods, with common types including Denial of Service and Information Disclosure. Even when fixed, vulnerabilities often take 3–5 months to resolve. Zimmermann et al. [7] further showed that many npm projects depend on vulnerable packages, that risks spread quickly, and that a few maintainer accounts act as single points of failure. Unmaintained packages were also found to cause vulnerabilities to persist for years.

To help manage these risks, visualization-based tools have been introduced to improve developers' understanding of

¹<https://github.com/morakotch/security-md-deps>

²<https://docs.github.com/en/code-security/getting-started/adding-a-security-policy-to-your-repository>

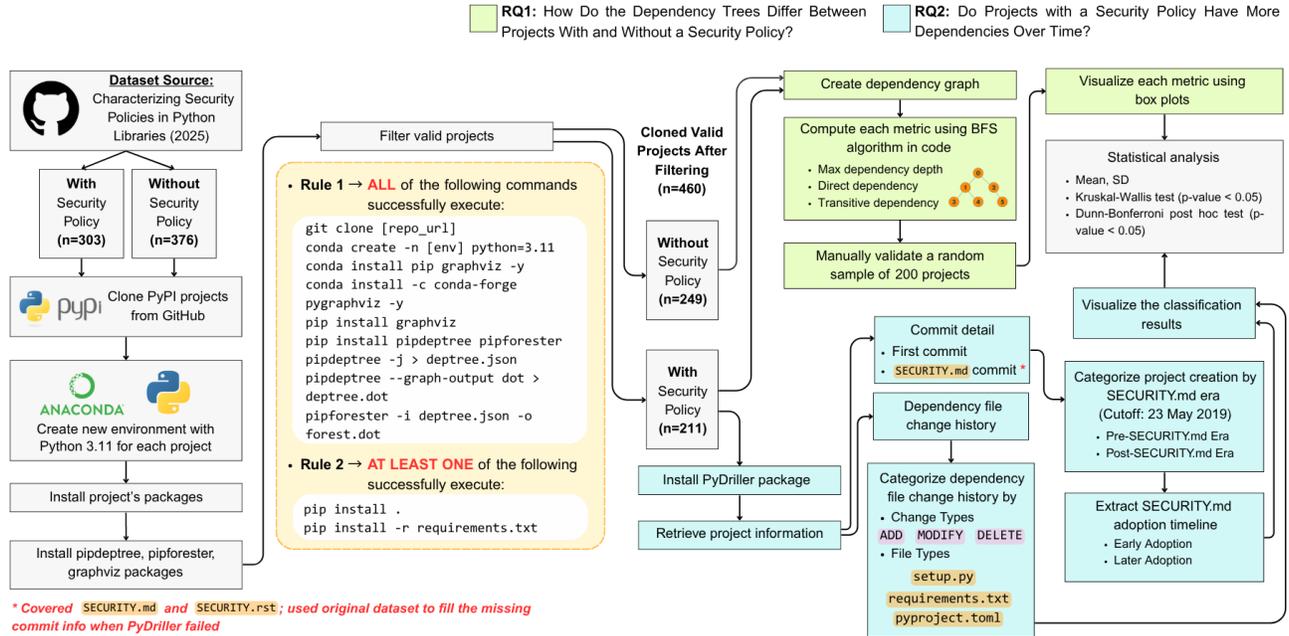


Fig. 2: Overview of the research methodology

dependency structures. V-Achilles, for example, uses interactive graphs to visualize transitive vulnerabilities in npm projects [11]. It highlights both direct and transitive risks but is limited to four levels of depth and may struggle with complex graphs. A follow-up study [8] showed that 70% of developers who used graph visualizations reprioritized updates compared to those using text-based tools. Although security policies and dependency risks have been studied, they have typically been addressed separately. Prior work has not explored whether security policies are associated with changes in the structure or evolution of a project's dependency supply chain.

III. METHODOLOGY

A. Overview

Figure 2 presents an overview of the research methodology. This study examines how dependency structures and their evolution differ between PyPI projects with and without a security policy. The dataset is derived from a prior study [4] and categorized based on the presence of a SECURITY.md file. Dependency chains are extracted and visualized using pipdeptree, pipforester, and Graphviz. Projects that do not meet installation and setup criteria are excluded—92 projects with a security policy and 127 without—resulting in a final dataset of 460 projects, with 211 having a policy and 249 without. To validate extraction accuracy, a subset of projects is manually reviewed. Dependency trees are analyzed using a breadth-first search algorithm, and project histories are examined with PyDriller to identify changes in dependency files over time. Projects are further categorized by creation date and by the period during which a security policy was adopted. All metrics are visualized, and statistical

tests are applied to assess differences between the two groups. The following sections describe each step of the methodology in detail.

B. Data Collection

This study uses a dataset from prior work [4], which includes 679 PyPI projects divided into two groups: 303 projects with a security policy and 376 without. From the provided list of projects, each repository was cloned and executed within an isolated Python environment to extract dependency information. Dependencies were installed based on the project's specification files. A dependency chain in this study refers to the sequence of required packages beginning with the main dependencies declared in the selected files. This study focuses on three primary dependency files: setup.py, pyproject.toml, and requirements.txt, as they reflect core production dependencies. Supporting tools such as pipdeptree, pipforester, and Graphviz were installed to generate and analyze dependency graphs. All installations were conducted using default commands. Projects that required additional or custom installation steps were excluded from the analysis. Development and auxiliary dependencies, such as those listed in requirements-dev.txt or requirements-tool.txt, were also excluded.

As shown in Figure 2, the specific steps for data collection and filtering are illustrated. A project was retained only if it passed both Rule 1 and Rule 2. Rule 1 checks whether all setup commands run successfully. These commands include cloning the repository, creating a Python 3.11 environment, and installing required tools such as pip, graphviz, pygraphviz, pipdeptree, and pipforester. The commands must also generate depen-

dependency graph files using `pipdeptree` and `pipforester` in DOT and JSON formats. Rule 2 checks whether at least one of the dependency installation commands succeeds. These are either `pip install .` or `pip install -r requirements.txt`. Projects that failed to satisfy both rules were excluded from further analysis. After applying these filtering rules, the final dataset consisted of 460 projects: 211 with a security policy and 249 without.

C. Data Pre-Processing and Data Extraction

This section describes how the dependency data and project metadata were processed to answer RQ1 and RQ2. For RQ1, which investigates whether dependency structures differ between projects with and without a security policy, we analyzed three key metrics: (1) maximum dependency depth, (2) number of direct dependencies, and (3) number of transitive dependencies. These metrics reflect the structural complexity and potential exposure surface of a project’s dependency supply chain. To compute these metrics, we used `pipdeptree` to extract dependency trees in JSON format, and `pipforester` to clean and structure the output. For manual validation and visualization, the structured data were converted to DOT format using `Graphviz`. A Breadth-First Search (BFS) algorithm was implemented to traverse each tree and count dependencies at each level. To validate the accuracy of the BFS implementation, we manually inspected a random sample of 200 projects.

For RQ2, which explores how dependencies evolve in projects with a security policy, we collected commit history data from the 211 valid projects using `PyDriller`. This included timestamps of the first commit, the commit introducing the security policy file, and all commits involving changes to the three primary dependency files: `setup.py`, `pyproject.toml`, and `requirements.txt`. These files were selected because they represent production-level dependencies and are the most commonly used in Python packaging. Changes were categorized as ADD, MODIFY, or DELETE to quantify the level of ongoing maintenance and dependency updates throughout the project lifecycle.

Projects were grouped based on their creation date and the timing of security policy adoption. We used 23 May 2019, when GitHub introduced support for SECURITY.md files, as the cut-off to distinguish projects created before and after this change. Each group was further divided into early and later adopters, depending on whether the policy was added within seven days of the first commit. The seven-day threshold helps identify projects that included the policy as part of their initial setup versus those that integrated it during later stages. This classification enables analysis of how adoption timing may relate to changes in dependency management. While this threshold offers a practical distinction, we acknowledge that more fine-grained analysis is possible and leave it for future work. The classification criteria are summarized in Table I, and the analysis results are presented in the following sections.

TABLE I: Project classification criteria and definitions

Classification Category	Criteria and Definition
Pre-SECURITY.md	Projects created before 23 May 2019 , the date GitHub officially introduced support for SECURITY.md.
Post-SECURITY.md	Projects created on or after 23 May 2019 , meaning they started after SECURITY.md was introduced.
Early Adoption	Projects that added a SECURITY.md file within 7 days of their first commit.
Later Adoption	Projects that added a SECURITY.md file more than 7 days after their first commit.

D. Data Analysis

To address RQ1, we examined structural differences in dependency trees between projects with and without a security policy. We compared three metrics: maximum dependency depth, number of direct dependencies, and number of transitive dependencies. The Kruskal–Wallis test [12] was used to assess overall group differences. When significant differences were observed, the Dunn–Bonferroni post hoc test [13] was applied for pairwise comparisons. These nonparametric tests were chosen because the distributions of the metrics were not normal and the group sizes were unequal.

To address RQ2, we analyzed the commit history of projects with a security policy, focusing on changes to dependency specification files to understand how dependencies evolve over time. We statistically compared the total number of commits, the number of commits involving dependency files, and the distribution of modification types (e.g., ADD, MODIFY, DELETE) between early and later adoption groups. The analysis used the mean, standard deviation, Kruskal–Wallis test [12], and Dunn–Bonferroni post hoc test [13] to assess differences in maintenance activity. These analyses allowed us to examine whether projects with a security policy tend to update or expand their dependencies more actively over time.

IV. STUDY RESULT

This section presents the results of our study based on the two research questions.

A. RQ1: How Do The Dependency Trees Differ Between Projects With And Without A Security Policy?

Figure 3 shows the distribution of three dependency metrics for projects with and without a security policy. The maximum depth (a) is similar across both groups, with nearly identical medians and overlapping ranges. For transitive dependencies (c), projects with a policy have slightly higher values, but the distributions are wide and similar. The clearest difference appears in direct dependencies (b), where projects with a policy have a higher median and a broader upper range. This suggests that these projects tend to include more direct packages than those without a policy.

Table II compares dependency metrics between 211 PyPI projects with a security policy and 249 without. Among the

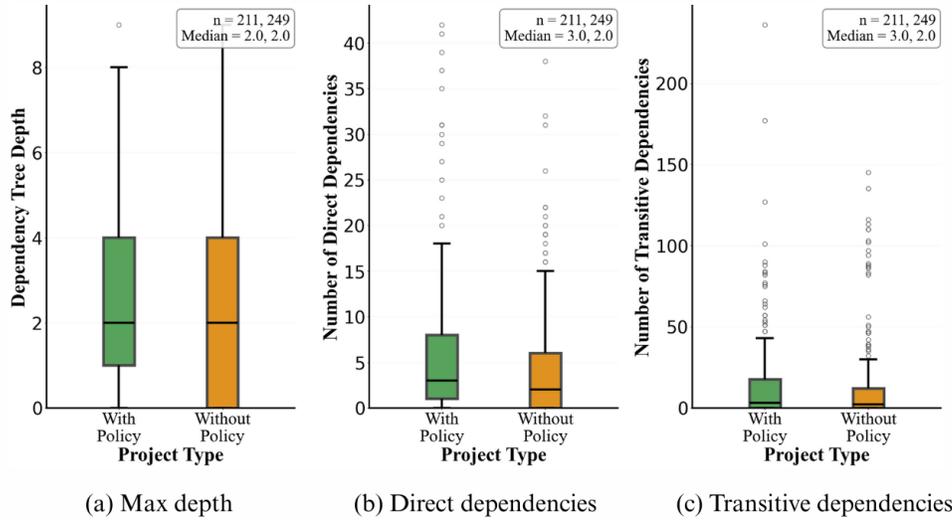


Fig. 3: Distribution of dependency metrics in projects with and without a security policy

three metrics—maximum depth, total direct dependencies, and total transitive dependencies—only the total number of direct dependencies shows a statistically significant difference ($p = 0.017$), based on the Kruskal–Wallis and Dunn–Bonferroni tests. Projects with a security policy have a higher mean number of direct dependencies (5.98) than those without (4.25), suggesting that these projects may be more modular or incorporate more external packages. Although the mean number of transitive dependencies is also higher in projects with a security policy (15.30 vs. 13.43), the difference is not statistically significant ($p = 0.356$), likely due to high variability across projects, as reflected in the large standard deviations.

Answer to RQ1

Projects with a security policy exhibit a significantly higher number of direct dependencies than those without. This difference suggests that these projects are more modular or make greater use of external packages. However, no significant differences are found in transitive dependencies or maximum depth. These findings imply that structural differences are mainly at the direct dependency level, possibly reflecting more deliberate dependency management practices in projects that adopt a security policy.

B. RQ2: Do Projects With a Security Policy Have More Dependencies Over Time?

Figure 4 summarizes project creation dates and SECURITY.md adoption timelines for the 211 projects with a security policy in our dataset. Projects are divided into the Pre-SECURITY.md and Post-SECURITY.md eras, using May 23, 2019—the date GitHub officially introduced SECURITY.md

TABLE II: Statistical comparison of dependency metrics between projects with and without a security policy

Metric	With Policy		Without		p-value
	Mean	SD	Mean	SD	
Max Depth	2.48	1.97	2.27	2.01	0.199 / –
Direct Deps	5.98	8.19	4.25	5.96	0.017 / 0.017
Transitive Deps	15.30	29.45	13.43	26.45	0.356 / –

Note: KW = Kruskal–Wallis; DB = Dunn–Bonferroni (reported only when $p < 0.05$).

support—as the cutoff. The figure includes a box plot (a) and three timeline subplots (b, c, and d) to show how these projects adopted SECURITY.md over time and how adoption behavior varies between older and newer projects.

As shown in Figure 4(a), the majority of projects with a security policy were created before the introduction of SECURITY.md. Specifically, 184 projects were created in the Pre-SECURITY.md era. Their median creation year is 2012.5 (Q1: 2009.4, Q3: 2016.1), meaning that half of these projects were created before mid-2012 and half after. This indicates that many projects in this group have a long development history, with some dating back to the early 2000s. In contrast, only 27 projects were created after 2019, representing the Post-SECURITY.md era. These newer projects have a median creation year of 2020.9 (Q1: 2020.1, Q3: 2022.5), which is almost a decade later than the Pre-policy group. The difference of 8.4 years between the two medians shows a substantial generational gap, highlighting that security policy adoption spans across both legacy and more recent projects. The narrower interquartile range in the Post group also reflects more concentrated development activity in recent years, which may relate to more standardized development workflows and increased awareness of platform-supported security practices.

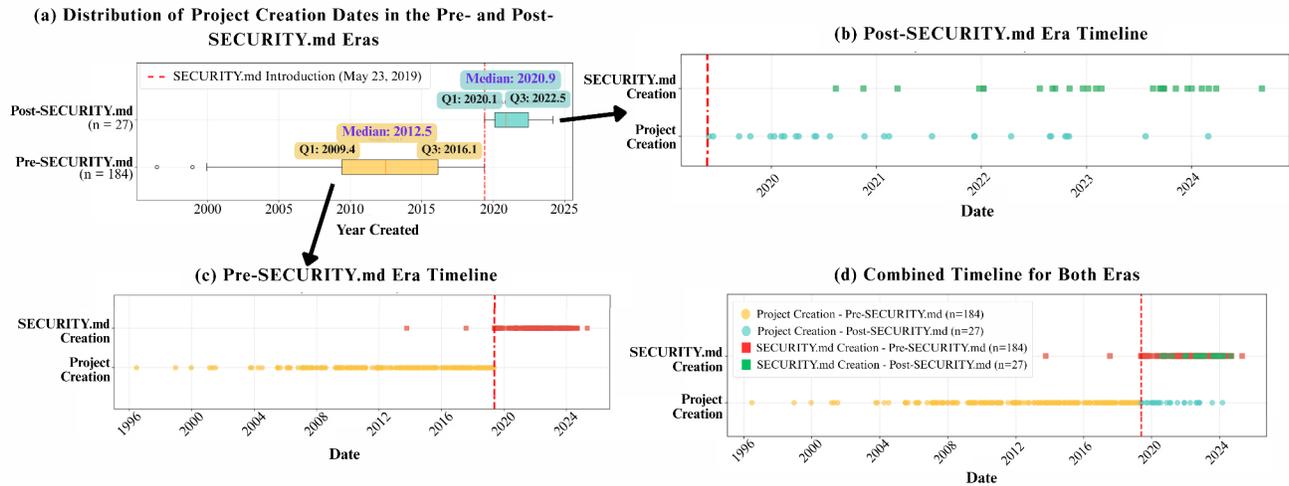


Fig. 4: Distribution and timeline of project creation and SECURITY.md adoption in projects with a security policy, categorized into Pre- and Post-SECURITY.md eras

Figure 4(b) presents the Post-SECURITY.md projects in more detail. All 27 of these projects were created after GitHub introduced SECURITY.md support in 2019. Most of them added the SECURITY.md file relatively early in their development lifecycle, typically within the first few months after project creation. This suggests that security awareness has improved among newer projects, and that maintainers are more likely to follow GitHub’s recommended practices from the outset. The adoption timing in these projects is generally more consistent, indicating that the SECURITY.md file is now part of the initial setup rather than an afterthought. This trend may reflect a shift in community norms, where including a security policy early has become part of the standard development process.

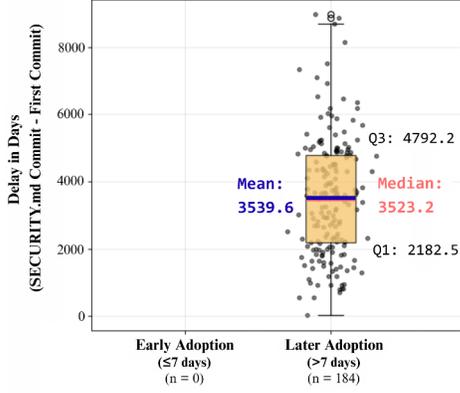
Figure 4(c) focuses on the 184 Pre-SECURITY.md projects, which were created between 1996 and 2018. Most of these older projects added a SECURITY.md file only after GitHub began officially supporting it in 2019. The timing of adoption is spread out, with many projects introducing the file months or even years later. This suggests that SECURITY.md was commonly added as part of later updates or maintenance activities, rather than during initial development. The wide range of adoption times is reflected in the quartile values: the first quartile (Q1) is around 2009.4, and the third quartile (Q3) is 2016.1. This means that 50% of Pre-SECURITY.md projects were created within that 6.7-year range, showing a long-standing base of projects that predate GitHub’s policy. These values highlight the legacy nature of many repositories and explain why the adoption of SECURITY.md was not immediate. In this study, we treat SECURITY.rst files as equivalent to SECURITY.md, as both serve the same purpose of describing the project’s security policy.

Figure 4(d) overlays both groups to provide a side-by-side comparison of adoption behavior. The figure clearly shows two distinct trends. Older projects, represented in the Pre group,

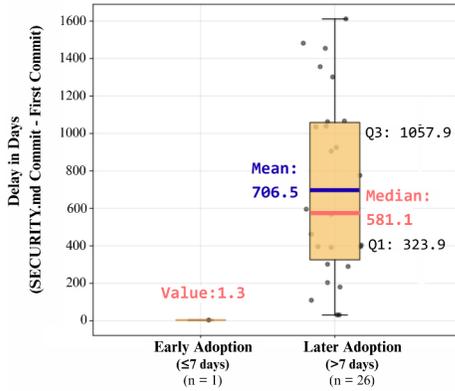
were generally slow to adopt SECURITY.md, often adding it as part of later-stage updates. In contrast, newer projects in the Post group tend to integrate SECURITY.md shortly after their creation. This contrast reflects a broader shift in development practices: whereas older repositories treated security policies as optional or delayed concerns, newer projects increasingly treat them as essential components introduced early in the project timeline. This behavioral shift may be influenced by evolving community expectations, increased tooling support, and stronger awareness of supply chain risks in open-source software. These findings highlight the need for better integration of security policy support within modern project management tools. By embedding features that encourage or even require SECURITY.md adoption from the early stages, such tools can help maintainers meet current security expectations and reduce long-term risks.

Figure 5 shows the delay between each project’s first commit and the creation of its SECURITY.md file, grouped by era and adoption timing. In the Pre-SECURITY.md era, all 184 projects were classified as later adopters, with a median delay of 3523.2 days (around 9.6 years). In contrast, Post-SECURITY.md era projects had a much shorter median delay of 581.1 days (about 1.6 years), showing a substantial gap of approximately 2942 days (around 8 years) between the two medians. This result indicates that, although delayed adoption remains common, newer projects tend to introduce SECURITY.md files much earlier than older ones. Note that using the first commit as the baseline may disadvantage Pre-SECURITY.md projects, so we also measured delay from May 23, 2019. Even then, these projects still adopted SECURITY.md later than Post-era ones.

Among the 27 Post-era projects, 26 (96.3%) were still classified as later adopters, with only one project adding the SECURITY.md file within 7 days of the first commit. In comparison, no Pre-era projects qualified as early adopters.



(a) Pre-SECURITY.md Era Projects



(b) Post-SECURITY.md Era Projects

Fig. 5: Timeline of SECURITY.md Adoption Relative to Initial Commit in Pre- and Post-SECURITY.md Projects

TABLE III: Change Types Across Packaging Files in 211 Projects with Security Policies

Type	setup.py	req.txt	pyproj.toml	Total
MODIFY	33,624 (46.3%)	11,406 (15.7%)	22,923 (31.6%)	67,953 (93.6%)
ADD	989 (1.4%)	1,267 (1.7%)	1,178 (1.6%)	3,434 (4.7%)
DELETE	562 (0.8%)	450 (0.6%)	192 (0.3%)	1,204 (1.7%)
Total	35,175 (48.5%)	13,123 (18.1%)	24,293 (33.5%)	72,591 (100%)

This pattern suggests that while the overall speed of adoption has improved in recent years, immediate integration of a security policy at the beginning of a project remains rare. In short, most projects took more than seven days to add a SECURITY.md file after they started.

In the activities investigated in this study, Table III shows the distribution of change types (MODIFY, ADD, DELETE) across three packaging files in 211 projects with a security policy. Regarding the change type, MODIFY operations account for the vast majority (93.6%) of all changes, indicating

TABLE IV: Summary Statistics of Dependency File Changes by Group

Era	Adoption	Commits		Dep. Changes*		Dep. Commits			% Rate**	
		Total	Sum	Mean	SD	Sum	Mean	SD	Mean	SD
Pre	Early	0	0	-	-	0	-	-	-	-
Pre	Later	>1M	52,198	283.7	407.1	45,300	246.2	286.6	8.9	8.1
Post	Early	53	36	36.0	0.0	36	36.0	0.0	67.9	0.0
Post	Later	101,398	20,357	783.0	1,873.8	11,220	431.5	642.4	16.4	12.7

* Dependency changes include edits to requirements.txt, setup.py, and pyproject.toml.

** Percentage of commits with dependency changes; SD = standard deviation.

TABLE V: Pairwise Group Comparison of Dependency Commit Rates Using Dunn-Bonferroni Test

Group (I)		Group (II)		p-value*	Sig.**
Era	Adopt	Era	Adopt		
Pre	Later	Post	Early	0.087	No
Pre	Later	Post	Later	0.001	Yes
Post	Early	Post	Later	0.108	No

* Based on Dunn-Bonferroni Post Hoc ($p < 0.05$),

Kruskal-Wallis $p = 0.001667$

** Significant if $p < 0.05$

that projects typically focus on refining or updating existing dependency specifications rather than introducing new dependencies or removing outdated ones. ADD and DELETE changes are relatively limited, comprising only 4.7% and 1.7%, respectively. In terms of file usage, setup.py is the most actively modified file, accounting for 48.5% of all changes. This is followed by pyproject.toml (33.5%) and requirements.txt (18.1%). Together, modifications in setup.py and pyproject.toml represent over 80% of all dependency-related file changes, suggesting that these two files serve as the primary configuration points for managing dependencies in Python projects with a security policy.

Table IV summarizes dependency-related commit activity across project groups. The Pre-Early group had no valid samples. The Post-Early group included only one project, which showed a high dependency commit rate (67.9%), but this is not generalizable due to the limited sample size. Among the larger groups, the Pre-Later projects ($n = 184$) had a lower average rate of commits involving dependency changes (8.9%), while Post-Later projects ($n = 26$) showed a higher mean rate (16.4%). This difference suggests that projects initiated after GitHub's introduction of SECURITY.md may be more active in updating their dependency specifications.

The Kruskal-Wallis test revealed a significant difference in dependency commit rates across the project groups. Follow-up analysis using the Dunn-Bonferroni post hoc test (Table V) showed a statistically significant difference between the Post-Later and Pre-Later groups ($p = 0.001$). This suggests that projects created after the introduction of SECURITY.md and adopting it later (after 7 days) tend to make more frequent changes to their dependency files. In contrast, no significant differences were observed between the other group pairs. These findings highlight that project creation era, rather than the timing of SECURITY.md adoption alone, may influence

the frequency of dependency updates.

These findings raise a future research question: how do these dependency changes relate to broader project evolution, such as the introduction of new features, major code changes, or version upgrades? Investigating this relationship could help explain whether dependency specification updates are reactive (e.g., responding to security needs) or proactive (e.g., accompanying functional development), and how they align with software maintenance and release cycles.

Answer to RQ2

Dependency activity differs significantly across groups. Projects created after the introduction of SECURITY.md and adopting it later (Post-Later) show significantly higher rates of dependency file changes than those created earlier (Pre-Later). Most changes occur in `setup.py` and `pyproject.toml`, which together account for over 80% of edits. These patterns suggest that newer projects not only adopt security policies more promptly but also update their dependencies more actively, indicating stronger maintenance practices and increased attention to security.

V. DISCUSSION AND IMPLICATIONS

Our findings provide evidence that the adoption of a SECURITY.md file in Python projects is not only a policy marker but also an indicator of a diligent, security-conscious maintenance team, reflecting real maintenance behavior rather than being a direct cause of improved practices. Specifically, projects with a SECURITY.md file tend to have more direct dependencies, suggesting greater modularity or functionality. Moreover, projects that adopt SECURITY.md later in their lifecycle show higher levels of dependency-related commits, indicating more active or frequent maintenance on dependency configurations. This pattern suggests that projects which take longer to adopt SECURITY.md may be larger or more mature, thus requiring more effort to manage dependencies. Alternatively, delayed adoption may indicate greater awareness of dependency risks and a subsequent push to improve maintenance practices. The high concentration of edits in `setup.py` and `pyproject.toml` highlights the central role of these files in Python dependency management and presents them as key targets for monitoring and tooling. These results offer several implications:

- For researchers, the observed link between policy adoption and dependency activity opens new opportunities to study software evolution by combining social signals (e.g., policy files) with technical data (e.g., commits).
- For tool developers, the findings highlight `setup.py` and `pyproject.toml` as priority targets for supporting dependency tracking and automation.
- For maintainers, adopting SECURITY.md is associated with higher maintenance activity, reinforcing its role as a practical indicator of commitment to project upkeep.

VI. CONCLUSION

This study examined the relationship between SECURITY.md adoption and dependency management in PyPI projects. We found that projects with a SECURITY.md file tend to have more direct dependencies, and those adopting the policy later show higher rates of dependency-related commits, especially in `setup.py` and `pyproject.toml`. These results suggest that lightweight security policies may signal stronger maintenance practices. Future empirical work could explore how adoption patterns vary across ecosystems and project characteristics, and also delve into the qualitative nature of dependency file changes. For practice, lightweight and automated tools that incorporate policy awareness could help maintainers manage dependencies more effectively, identify update needs, and sustain long-term project health.

REFERENCES

- [1] M. Alfadel, D. E. Costa, and E. Shihab, "Empirical analysis of security vulnerabilities in python packages," *Empirical Software Engineering*, vol. 28, June 2023.
- [2] "Stack overflow 2025 developer survey." [Online]. Available: <https://survey.stackoverflow.co/2025/technology#1-programming-scripting-and-markup-languages>. (Accessed on 23/08/2025).
- [3] "Adding a security policy to your repository." [Online]. Available: <https://docs.github.com/en/code-security/getting-started/adding-a-security-policy-to-your-repository>. (Accessed on 23/08/2025).
- [4] S. Kancharoendee, T. Pichitphanphong, C. Jongyingyos, B. Reid, R. G. Kula, M. Choetkiertikul, C. Ragkhitwetsagul, and T. Sunetnanta, "On Categorizing Open Source Software Security Vulnerability Reporting Mechanisms on GitHub," in *2025 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, (Los Alamitos, CA, USA), pp. 751–756, IEEE Computer Society, Mar. 2025.
- [5] B. Vermeer, "10 github security best practices." [Online]. Available: <https://snyk.io/blog/ten-git-hub-security-best-practices/>, Feb. 2024. (Accessed on 23/08/2025).
- [6] N. Bühlmann and M. Ghafari, "How do developers deal with security issue reports on github?," in *Proceedings of the ACM Symposium on Applied Computing*, pp. 1580–1589, Association for Computing Machinery, Apr. 2022.
- [7] M. Zimmermann, C.-A. Staicu, C. Tenny, and M. Pradel, "Small world with high risks: A study of security threats in the npm ecosystem," in *28th USENIX Security Symposium (USENIX Security 19)*, (Santa Clara, CA), pp. 995–1010, USENIX Association, Aug. 2019.
- [8] C. Ragkhitwetsagul, V. Jarukitpipat, R. G. Kula, M. Choetkiertikul, K. Chhun, W. Wanprasert, and T. Sunetnanta, "See to believe: Using visualization to motivate updating third-party dependencies," in *Proceedings - 21st International Joint Conference on Computer Science and Software Engineering, JCSSE 2024*, pp. 618–625, Institute of Electrical and Electronics Engineers Inc., 2024.
- [9] "Security policy - github changelog." [Online]. Available: <https://github.blog/changelog/2019-05-22-security-policy/>, May 2019. (Accessed on 23/08/2025).
- [10] G. A. A. Prana, A. Sharma, L. K. Shar, D. Foo, A. E. Santosa, A. Sharma, and D. Lo, "Out of sight, out of mind? how vulnerable dependencies affect open-source projects," *Empirical Softw. Engg.*, vol. 26, July 2021.
- [11] V. Jarukitpipat, K. Chhun, W. Wanprasert, C. Ragkhitwetsagul, M. Choetkiertikul, T. Sunetnanta, R. G. Kula, B. Chinthanet, T. Ishio, and K. Matsumoto, "V-achilles: An interactive visualization of transitive security vulnerabilities," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering, ASE '22*, (New York, NY, USA), Association for Computing Machinery, 2023.
- [12] W. H. Kruskal and W. A. Wallis, "Use of ranks in one-criterion variance analysis," *Journal of the American Statistical Association*, vol. 47, pp. 583–621, Dec. 1952.
- [13] O. J. Dunn, "Multiple comparisons using rank sums," vol. 6, pp. 241–252, Aug. 1964.