

Pinning Is Futile

You Need More Than Local Dependency Versioning to Defend against Supply Chain Attacks

HAO HE, Carnegie Mellon University, USA

BOGDAN VASILESCU, Carnegie Mellon University, USA

CHRISTIAN KÄSTNER, Carnegie Mellon University, USA

Recent high-profile incidents in open-source software have greatly raised practitioner attention on software supply chain attacks. To guard against potential malicious package updates, security practitioners advocate *pinning* dependency to specific versions rather than *floating* in version ranges. However, it remains controversial whether pinning carries a meaningful security benefit that outweighs the cost of maintaining outdated and possibly vulnerable dependencies. In this paper, we quantify, through counterfactual analysis and simulations, the security and maintenance impact of version constraints in the npm ecosystem. By simulating dependency resolutions over historical time points, we find that pinning direct dependencies not only (as expected) increases the cost of maintaining vulnerable and outdated dependencies, but also (surprisingly) even increases the risk of exposure to malicious package updates in larger dependency graphs due to the specifics of npm's dependency resolution mechanism. Finally, we explore collective pinning strategies to secure the ecosystem against supply chain attacks, suggesting specific changes to npm to enable such interventions. Our study provides guidance for practitioners and tool designers to manage their supply chains more securely.

CCS Concepts: • **Software and its engineering** → *Software libraries and repositories; Development frameworks and environments*; • **Security and privacy** → *Software security engineering*.

Additional Key Words and Phrases: Software supply chain security, Dependency management

ACM Reference Format:

Hao He, Bogdan Vasilescu, and Christian Kästner. 2025. Pinning Is Futile: You Need More Than Local Dependency Versioning to Defend against Supply Chain Attacks. *Proc. ACM Softw. Eng.* 2, FSE, Article FSE013 (July 2025), 24 pages. <https://doi.org/10.1145/3715728>

1 Introduction

Open-source components have become a critical part of software supply chains in today's software development. However, recent high-profile security incidents show that integrating open-source components exposes a software project to two major types of security risks:

- **Security Vulnerabilities:** An open-source component, especially outdated versions, may contain *accidentally* introduced vulnerabilities or weaknesses, which may be exploitable by a malicious actor in certain cases (e.g., the Log4Shell vulnerability [IBM 2021]).
- **Malicious Package Updates:** An open-source component may be *intentionally* compromised by an attacker to release a malicious update and exploit downstream applications (e.g., the node-ipc [Snyk 2022] and xz incidents [Sonatype 2024b]). This scenario is often referred to, in a broader sense, as *software supply chain attacks* [Ladisa et al. 2023; Ohm et al. 2020].

Authors' Contact Information: Hao He, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, haohe@andrew.cmu.edu; Bogdan Vasilescu, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, vasilescu@cmu.edu; Christian Kästner, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, kaestner@cs.cmu.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2994-970X/2025/7-ARTFSE013

<https://doi.org/10.1145/3715728>

Table 1. The commonly argued trade-offs (✔ benefits and ✘ drawbacks) between pinning and floating

The Trade-Offs (🔒 Security and ⚙️ Maintenance)	Floating	Pinning
🔒 Attack Surface for Malicious Package Updates [e.g., Bressers 2022; Ladisa et al. 2023; OpenSSF 2023; Zahan et al. 2023a]	✘ Larger	✔ Smaller
⚙️ Effort to Handle Breaking Changes [e.g., Bogart et al. 2021; Jafari et al. 2022; Leiter 2023; Lucid 2017; Radzikowski 2021; Raemaekers et al. 2017]	✘ More	✔ Less
🔒 Attack Surface from Security Vulnerabilities [e.g., Chinthanet et al. 2021; Hacker News 2017; Jafari et al. 2022; StackOverflow 2017; Wang et al. 2023]	✔ Smaller	✘ Larger
⚙️ Effort to Update Outdated Dependencies [e.g., Jafari et al. 2022]	✔ Less	✘ More
⚙️ Effort to Resolve Dependency Conflicts [e.g., Hacker News 2017]	✔ Less	✘ More

In most ecosystems, software projects need to specify *version constraints* for their dependencies on open-source components—*these constraints influence their exposure to security vulnerabilities and malicious package updates*. For example, npm packages are expected to follow semantic versioning [Preston-Werner 2023], using minor version updates to indicate non-breaking releases. Then, downstream developers can use *floating* version constraints (e.g., “`^1.2.0`” matches $1.2.0 \leq \text{version} < 2.0.0$) to allow their project to automatically install the most recent, patched but also backward compatible version of a dependency at the time of installation. Through this mechanism, a security patch fixing a vulnerability (e.g., 1.2.1 fixing 1.2.0) can be immediately propagated to a downstream project as long as it uses floating version constraints for that dependency.

However, the same floating mechanism can also lead to the immediate propagation of malicious package updates in case a dependency gets compromised and releases a malicious version as a new patch release (e.g., the node-ipc incident [Snyk 2022]). Even if malicious updates are less common, they are arguably more concerning than security vulnerabilities because many security vulnerabilities are hard to exploit [Jacobs et al. 2019; Younis et al. 2016] whereas the injection of malicious code can cause immediate catastrophic harm and is difficult to defend against. To prevent exposure to malicious package updates, it seems necessary to *pin* dependency versions (e.g., require exactly version 1.2.0). However, this strategy comes with its own costs. To receive future updates, developers have to manually update the dependency to a newer release, ideally after auditing the changes. Failure to do so can leave security vulnerabilities to hackers even when the security patch has been available for a long time [Decan et al. 2018].

The choice between *pinning* and *floating* not only represents a *trade-off* among competing security priorities but also interacts with other dependency management concerns. Specifically, the above discussion also applies to breaking changes (that are not malicious) and outdated dependencies (a form of technical debt). What’s more, modern software supply chains form complex networks in which each node (i.e., project) can only make local decisions about its direct dependencies [Decan et al. 2019], but all pinning and floating decisions an upstream project makes have implications, *both directly and transitively*, for all its downstream projects. For example, a floating version constraint leaves an attack surface for all downstream projects, while pinning puts downstream projects at a higher risk of encountering dependency conflicts [Pinckney et al. 2023c; Wang et al. 2018]. We summarize the commonly argued benefits and drawbacks of pinning and floating in Table 1.

Importantly, the concern about *malicious package updates* is fairly recent and reached broad awareness through high-profile incidents, such as *SolarWinds* in 2020 [TechTarget 2023], *node-ipc* in 2022 [Snyk 2022] and the *xz* attack targeted at *ssh* in 2024 [Sonatype 2024b]—such recent attention may shift trade-off considerations. Prior to this recent concern, there have been many studies and discussions about the trade-off between semantic versioning and breaking changes [e.g.,

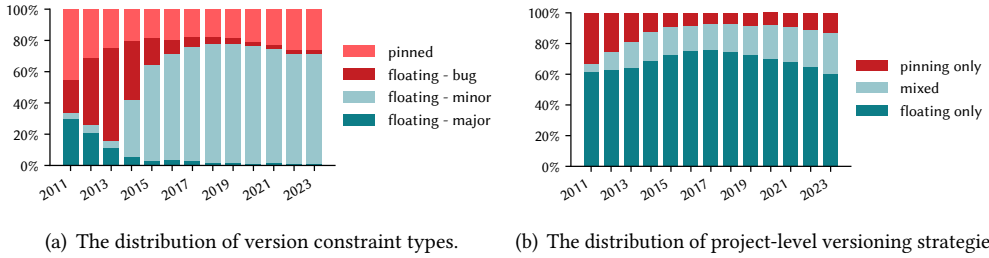


Fig. 1. The distribution of version constraint types and project-level versioning strategies in each year in the dataset we use for answering RQ1 (see Section 3.1 for details) - we observe that after years of increasing floating adoption, since about 2020 the trend reverses with a shift toward pinning. In Figure 1(a), we ignore non-floating/pinning version constraints as they only occupy 0.83% of version constraints in our dataset.

Bogart et al. 2021; Jafari et al. 2022; Raemaekers et al. 2017]. In a survey conducted in 2020, the majority of practitioners in the npm ecosystem were in favor of (semantic-versioning-based) floating [Jafari et al. 2022]. It is argued that pinning accumulates technical debt in the form of outdated dependencies [Hacker News 2018; Jafari et al. 2022], prevents security fixes from propagating downstream [Chinthanet et al. 2021; Wang et al. 2023], and increases the likelihood of dependency conflicts (or bloat in npm) [Hacker News 2017; StackOverflow 2017]. However, recent high-profile malicious package update incidents provide a new argument in favor of pinning, as pinning reduces the attack surface for malicious package updates. Given this argument, together with prior arguments that pinning increases application stability and reproducibility [Leiter 2023; Lucid 2017; Radzikowski 2021], many practitioners now argue in favor of pinning for security sensitive applications – for example, a survey of 17 domain experts and 134 software developers by Ladisa et al. [2023] ranks pinning as a top-3 safeguard against supply chain attacks in terms of utility/cost ratio; the OpenSSF Scorecard project [OpenSSF 2023; Zahan et al. 2023a] included “pinned dependencies” as one of its 18 security best practices. This discussion clearly reaches practitioners—as we show in Figure 1 (and explain in more depth later in Section 3.1), practitioners in npm have gradually shifted toward pinning more dependencies since 2020.

Overall, controversies remain on whether the benefits of pinning outweigh its drawbacks, as manifested by online discussions [Bressers 2022; Curphey 2022; Renovate Bot 2018] and evidenced by the contradicting results from the collected practitioner opinions in prior studies [Jafari et al. 2022; Ladisa et al. 2023]. However, to the best of our knowledge, *there is no quantitative evidence so far about the actual impact of pinning versus floating dependencies on the security and maintenance trade-offs in Table 1*. To gather quantitative evidence about pinning and floating, we perform a repository mining study on the npm ecosystem, the most frequently discussed ecosystem in the literature. Our study highlights an important complication that makes the discussion around pinning so difficult: Since a project can only control the version constraints of its direct dependencies, the actual impact of pinning as a local intervention is complicated and unclear, contingent upon both the decisions made upstream and the behaviors of dependency management tools. To measure this impact holistically from empirical data, we raise the following research question:

- **RQ1: (Impact of Pinning)** How does the individual choice of pinning direct dependencies affect the overall security risks and maintenance costs in one’s own dependency graph?

To answer this research question, we collect a large sample of npm packages and GitHub repositories using npm. Then, we use an experimental time-traveling feature of the npm dependency resolver

to simulate the evolution of their dependency graphs over a 12-month period. We estimate the impact of pinning on a set of five security and maintenance-related metrics using *fixed-effect panel regression* [Brüderl and Ludwig 2015; Hanck et al. 2021]. A key (and counter-intuitive) finding from our model is that for the purpose of reducing the attack surface of malicious updates, *pinning is futile and can be even harmful* for projects with ≥ 498 direct and transitive dependencies (affecting 26% of GitHub repositories in our dataset). That is, pinning is ineffective for one of its main goals, all while creating substantial costs for managing outdated, vulnerable, and conflicting dependencies.

Inspired by this finding, we argue that *ecosystem-level interventions* (e.g., investing resources to secure critical open-source projects) are more effective than (possibly futile) local interventions of pinning direct dependencies to improve the overall resilience of an open-source ecosystem against malicious package updates. To empirically test this, we explore the second research question:

- **RQ2: (Ecosystem-Level Interventions)** How effective is coordinating pinning among selected upstream packages in reducing the risk of malicious package updates ecosystem-wide?

We answer this research question by simulating coordinated efforts in the entire npm dependency network. We formulate attack models and defense strategies in the dependency network, define heuristics to intervene in selected upstream packages, and measure the overall risk of malicious package updates after the intervention. Our findings indicate that it is possible to reduce the average spread of malicious package updates in the npm dependency network by nearly 30% if we pin in 100 intentionally selected important packages. However, as we will quantify, this intervention could be substantially more effective (up to 75%) if npm implements a new *transitive pinning* feature to allow packages to pass their entire pinned dependency graphs downstream.

Summary of Contributions. Our paper contributes novel quantitative evidence on the security and maintenance impact of version constraints in a packaging ecosystem, especially in light of recent concerns about malicious package updates. Our findings align with the opinions of npm developers [Jafari et al. 2022] but contradict the views of security practitioners [Ladisa et al. 2023; Zahan et al. 2023a], suggesting that pinning direct dependencies is futile (and can be even counter-productive) as protection against malicious package updates. In addition, we contribute evidence that a new ecosystem-wide intervention, along with tooling changes and ecosystem-wide allocation of collectively pooled resources, can be an effective intervention against malicious package updates. Our findings lead to concrete advice for developers and ecosystem policy makers to improve the state of software supply chain security (Section 5).

2 Background and Related Work

Dependency Management in npm. npm is the most widely used dependency manager for JavaScript and TypeScript [npm 2023b]. It consists of a package registry that hosts open-source components

```
{ "name": "postcss-loader", "version": "7.3.3",
  "dependencies": {
    "webpack": ">=5.0.0"           ← floating - major
    "cosmiconfig": "^8.3.5",     ← floating - minor: ≥8.3.5, <9.0.0
    "jiti": "~1.21.0",           ← floating - patch: ≥1.21.0, <1.22.0
    "semver": "7.5.4"           ← pinned
    "postcss": "^7.0.0 || ^8.0.1", ← other
    "lodash": "git+ssh://git@github.com:lodash/lodash.git#v4.17", ← other
  }, "devDependencies": {...}, ... }
```

Fig. 2. An example package.json file. We modified version constraints in the original file to illustrate different types of version constraints defined in Section 3.2. Other details irrelevant to this paper are omitted.

(called *packages* in npm) and a command-line tool that helps projects reuse packages from the package registry. As of September 2024, the npm registry is the largest package registry, hosting more than 3 million packages and 25 million releases [DeBill 2024]. The npm command-line tool reads direct dependencies from a configuration file named `package.json`. In the `package.json` file (see Figure 2 for an example), developers specify package names as the *dependencies* of the project with a *version constraint* for each dependency. In addition, npm distinguishes different kinds of dependencies—usually, developers focus more on *production dependencies* [Latendresse et al. 2022], i.e., dependencies that will be included in the production build (for applications) or passed downstream (for packages), and are less concerned about security vulnerabilities if the dependencies are not used in production [Latendresse et al. 2022; Pashchenko et al. 2018].

Given a `package.json` file, npm will resolve a *dependency graph* including not only dependencies specified in the `package.json` file, but also all transitive dependencies required by dependencies. The npm ecosystem is known as particularly fragmented with many tiny packages and high levels of reuse [Abdalkareem et al. 2017; Decan et al. 2019]; a project commonly has orders of magnitude more transitive dependencies than direct ones [Decan et al. 2019]. Dependency resolution can be complicated when multiple dependencies depend on the same third dependency but with conflicting version constraints; this is known as the *diamond problem* or *dependency hell* (see Figure 7). While it is highly problematic in other ecosystems [Abate et al. 2020; Bogart et al. 2021; Fan et al. 2020; Mancinelli et al. 2006], npm can install multiple distinct versions for the same package, but it will generally try to produce a dependency graph with fewer duplicate versions (i.e., deduplication [npm 2024b]). Optimal deduplication is NP-Complete [Abate et al. 2020], so npm uses a variety of heuristics to search for a “good” dependency graph satisfying all the version constraints while having fewer duplicate versions [npm 2024b].

Importantly, *dependency graphs are volatile*. Since npm always installs the latest version inside a floating versioning constraint, the same `package.json` may result in different graphs as new package updates are released [Liu et al. 2022]. With floating, projects can keep dependencies up-to-date automatically, whereas developers often do not update pinned dependencies timely, even for those with known security vulnerabilities [Cox et al. 2015; He et al. 2023; Kula et al. 2018].

A key challenge to floating dependencies is *breaking changes*: When an update changes expected behavior, the application can break without any action by its developer. Developers then need to react and either adapt their own code or downgrade the dependency [Bogart et al. 2021]. To control breaking changes, npm advocates for *semantic versioning* [npm 2023a]: A release should increment its major version number (e.g., from 1.2.1 to 2.0.0) if it contains breaking changes, minor version number (e.g., from 1.2.1 to 1.3.0) if it contains new features, and patch version number (e.g., from 1.2.1 to 1.2.2) if it contains only bug fixes. If packages follow those conventions, automatic minor and patch updates should be safe. To support this practice, npm provides specific patterns to declare version ranges accepting minor (e.g., `^8.3.5`) and patch updates (e.g., `~1.21.0`); it also adopts floating versions for minor updates by default when a developer adds a dependency with `npm install`. Although previous research has shown that packages may fail to obey semantic versioning [Jayasuriya et al. 2023; Raemaekers et al. 2017; Raible 2019; Zhang et al. 2022] and developers may turn to pinning instead if they have lost trust in its reliability [Lucid 2017; Raible 2019], semantic-versioning-based floating is still widely accepted and used by npm developers [Decan and Mens 2021; Dietrich et al. 2019; Jafari et al. 2022; Pinckney et al. 2023a]. In fact, without discussion of supply chain attacks, the survey by Jafari et al. [2022] shows that npm developers generally consider pinned dependencies as a “smell,” indicating a compromise that will accumulate technical debt and bring negative consequences in the long term.

Finally, it is worth noting that the above discussion applies not only to floating direct dependencies but also to floating transitive dependencies even if the developer pins all their direct dependencies.

Developers usually have no control over how their dependencies float or pin transitive dependencies. As a workaround, npm (as many other ecosystems) provides functionality to *freeze* and store a project's entire dependency graph with a lock file (`package-lock.json` [npm 2023c]). A lock file ensures the reproducibility of npm builds, but it has two limitations: (1) it is restricted to applications (the root of the dependency graph) whereas components used as dependencies cannot pass their lock decisions to downstream applications; it is hence not very useful for packages; (2) it is not designed (hence very difficult) to be manually maintained, so floating version constraints in `package.json` will still cause automatic updates every time a lock file is generated.

Software Supply Chain Security. The term *software supply chain* is an umbrella term referring to all the software components, processes, and tools that a software project relies on [Sonatype 2021]. For a project using npm, all packages in its dependency graph are part of its supply chain. Most software projects nowadays are built upon a collaboratively developed open-source software supply chain [Eghbal 2016]; the 3 million packages in the npm package registry are a prominent example.

Software supply chains are fragile. Such fragility comes from the collaborative and distributed nature of open-source development: Projects are often built with a large number of open-source components, each of them maintained by different groups of developers [Schueller and Wachs 2024]. Meanwhile, popular components are directly or transitively used by a large number of downstream applications. Consequently, any single weakness in the software supply chain could have a huge impact on the entire ecosystem [Decan et al. 2019; Zimmermann et al. 2019], creating a large attack vector for malicious actors and making the software industry increasingly concerned with software supply chain security [Geer et al. 2020; Ladisa et al. 2023].

In recent years, a significant amount of practitioner attention has been drawn to the rising threat of *software supply chain attacks* [Ladisa et al. 2023; Ohm et al. 2020]. Although the attackers may attempt to trick malware installation in other ways (e.g., typosquatting [Vu et al. 2020]), the most concerning incidents of software supply chain attacks are *malicious package updates*, in which malicious actors *intentionally* compromise an open-source component to inject malware into its (often many) downstream dependents. While not every security vulnerability is exploitable in practice (leading developers to treat them as false positives) [Jacobs et al. 2019; Mirhosseini and Parnin 2017; Younis et al. 2016], malicious package updates can immediately harm the infected downstream dependents [Martínez and Durán 2021]. They may also be well obfuscated and often very difficult to detect [Sonatype 2024b]. Their high catastrophic potential and lack of controllability amplify the signal of real incidents and elevate public risk perception (e.g., in a way similar to nuclear hazards [Slovic 1987]). As a result, malicious package updates are considered one of the primary threats in today's threat landscape [ENISA 2023; The White House 2021].

As our previous discussion of the pinning versus floating debate illustrates, many dependency management practices have security implications and involve trade-offs between conflicting priorities [Pashchenko et al. 2020]. Previous studies have investigated various practices, such as dependency adoption, updates, abandonment, migration, and the use of dependency management tools [e.g., He et al. 2021, 2023; Kula et al. 2018; Miller et al. 2023; Vargas et al. 2020]. It is worth noting, however, that most prior research in the software engineering community is centered around the mitigation of security vulnerabilities. On the other hand, the security community has recently consolidated best practices to defend against supply chain attacks, such as the OpenSSF Scorecard project [Zahan et al. 2023a] and the survey by Ladisa et al. [2023] introduced in Section 1. However, these (supposed) security best practices may interact with existing *software engineering* trade-offs, among which the pinning versus floating problem is particularly complicated (Table 1); the strongly held convictions from both sides lead to a considerable amount of controversy among practitioners [Bressers 2022; Curphey 2022; Renovate Bot 2018].

Novelty of Our Research. There have been many prior studies on the use of version constraints [Decan and Mens 2021; Dietrich et al. 2019; Ladisa et al. 2023; Pinckney et al. 2023a] and the trade-offs around semantic versioning with regards to maintenance effort, breaking changes, and security vulnerabilities [Bogart et al. 2021; Jafari et al. 2022; Raemaekers et al. 2017]. However, the new threat of *malicious package updates* complicates the previously discussed trade-offs, as *malicious actors intentionally break semantic versioning*, often disguising their attacks as patch releases for maximum propagation [Snyk 2022]. The changed threat landscape and the complexity of this pinning versus floating problem eventually lead to controversies and divergence in practitioner opinions [Bressers 2022; Curphey 2022; Jafari et al. 2022; Ladisa et al. 2023; Renovate Bot 2018]. Prior studies are either collecting these opinions [Bogart et al. 2021; Jafari et al. 2022; Ladisa et al. 2023] or aggregating descriptive statistics from software repositories [Decan and Mens 2021; Dietrich et al. 2019; Pinckney et al. 2023a; Raemaekers et al. 2017]; we are not aware of any prior study that attempted to quantify the impact of version constraints using statistical modeling methods. To shed light on resolving the controversy, we believe it is vital to conduct data-driven research and quantitatively measure the actual impact of version constraints while taking all the previously discussed trade-offs into consideration (Table 1). This forms the main motivation of our study.

3 RQ1: How does the individual choice of pinning direct dependencies affect the overall security risks and maintenance costs in one's own dependency graph?

As we discussed above, dependency pinning is often acknowledged as a security best practice to reduce the attack surface for malicious package updates [Ladisa et al. 2023; Zahan et al. 2023a], but controversies remain on whether the security benefits of pinning outweigh its maintenance drawbacks [Bressers 2022; Curphey 2022; Jafari et al. 2022; Ladisa et al. 2023; Renovate Bot 2018]. Importantly, a project only has control over the version constraints of its direct dependencies, but the impact of pinning direct dependencies is complicated and unclear, contingent upon upstream practices and dependency management tools. Hence, the goal of RQ1 is to empirically measure the actual security and maintenance impact of this intervention (i.e., pinning direct dependencies).

3.1 Data Collection

We aim to build a balanced dataset of packages (e.g., libraries, frameworks) and applications (e.g., websites, tools), as both need to manage dependencies but may differ in important ways. Thus, we use two data sources: (1) *npm packages*, (2) *GitHub repositories using npm* that are not npm packages. Note that neither data source is a clear representation of libraries or applications. Some npm packages may be applications and some GitHub repositories may be libraries that have not been uploaded to the npm registry. Still, combining the two should produce a stratified sample including both representation of libraries and applications in the wild.

For the first data source, npm packages, we curate a large dataset of popular and non-obsolete npm packages. We avoid focusing only on the few most popular packages, as they are not necessarily representative of common development practices in npm (they tend to occupy an upstream position with no dependencies or a very small dependency graph). We use the following data collection process: First, we use npm-follower [Pinckney et al. 2023b] (September 2023 dump) to obtain a complete list of npm packages and releases at that time. Then, we collect year-long download statistics as a popularity measure (between September 06, 2022, and September 05, 2023) using the npm API. For each package, we also identify its corresponding GitHub repository from npm metadata. Next, using this dataset, we select the top 5% most downloaded packages (21,218+ downloads) that have links to a GitHub repository and discard the bottom 25% of these packages with the lowest number of stars and issues in GitHub (specifically, less than 10 stars and 8 issues), to ensure that the remaining are actively maintained open-source packages. To further ensure

Table 2. Descriptive statistics of the studied npm packages and GitHub repositories (excluding npm packages). The number of transitive dependencies is an estimate using successful dependency resolutions (Section 3.2).

	npm packages				GitHub repositories (excl. npm packages)			
	Min	Median	Mean	Max	Min	Median	Mean	Max
# Annual Downloads	21,234	234K	24.7M	516M			N/A	
# Stars	8	345	6,431.80	220K	14	62	694.72	330K
# Issues	10	569	6338.55	87,258	0	101	652.62	75,561
# Commits	1	1,728	30,064.11	10.1M	84	632	6215.72	7.17M
# Active Months	1	48	60.7	324	13	36	44.26	323
# Direct Dependencies (Production)	1	3	5.78	283	1	11	16.83	492
# Direct Dependencies (All)	1	10	14.58	283	1	23	32.03	502
# Transitive Dependencies (Production)	0	11	65.55	2,320	0	150	391.51	3,722
# Transitive Dependencies (All)	0	253	414.74	3,141	0	848	893.18	4,013

they are not obsolete, we discard packages that have no release in the last two years at the time of dataset cutoff (September 2023). We discard packages without any production dependencies in the latest release and end up with 41,516 npm packages.

For the second data source, we curate a dataset of GitHub repositories using npm, but are not npm packages, from World of Code [Ma et al. 2021] (version V). First, we select the JavaScript and TypeScript projects with top 1% stars (i.e., # stars ≥ 14) as a rough popularity threshold from its deduplicated GitHub metadata [Spinellis et al. 2020]. To filter out small (and likely toy) projects, we discard projects with the bottom 25% of commits and active months (specifically, ≤ 84 commits and ≤ 13 active months). Then, we filter out projects that have been inactive (no commit) in the last two years at the time of dataset cutoff (September 2023). To find GitHub repositories using npm, we retain projects with a `package.json` file in their root folder. Finally, we discard repositories without any production dependencies and repositories whose name field in `package.json` matches a package name in npm—thus removing the ones more likely to be npm packages themselves. This results in a total of 12,197 GitHub repositories.

As the two data sources are heavily imbalanced and the simulations (see Section 3.2) are computationally expensive, we downsample the two data sources to 10,000 npm packages and 10,000 GitHub repositories, respectively (uniform random sampling each). We chose 10,000 each because it is computationally feasible for our simulation (in our case, two weeks of computation) and offers very high statistical generalizability. In the remainder of RQ1, we will refer to the 10,000 npm packages and the 10,000 GitHub repositories collectively as *projects*.

Dataset Overview. We summarize the descriptive statistics from both data sources in Table 2. Most of the projects in our dataset are popular and active. As expected, GitHub repositories tend to have a much larger dependency graph than npm packages. A median GitHub project has 11 direct dependencies and 150 transitive dependencies in production (23 direct and 848 transitive if we include development dependencies). This confirms our intuition that many projects have orders of magnitude higher numbers of transitive dependencies compared with direct ones.

In Figure 1, we provide an overview of the use of version constraints in our dataset, which largely mirrors findings of prior studies [Decan and Mens 2021; Dietrich et al. 2019; Pinckney et al. 2023a], but shows a previously not observed change in trend in more recent data, where developers start to pin more dependencies since around 2020. Specifically, we aggregate dependency declarations of each project in each year, looking back at past revisions (releases for npm packages and commits for GitHub repositories). To reduce noise introduced by projects with a massive number of revisions, we sample one revision for each project in each year, as also commonly done in prior studies [e.g., Pinckney et al. 2023a; Xu et al. 2023]. For each revision, we parse the version constraints of all its production dependencies in its `package.json` file, according to the npm semver syntax [npm 2023a],

and categorize the version constraints into five different types: (1) *floating - major*, (2) *floating - minor*, (3) *floating - patch*, (4) *pinned*, and (5) *other*. Examples of these different version constraint categories can be found in Figure 2. Overall, we find that semantic-versioning-based floating is used the most in the npm ecosystem but pinning is still the second most popular choice, with a rising popularity since around 2020, even though npm does not advocate pinning [npm 2023a]. This aligns with our intuition that the rising concern of malicious package updates triggers changes in practitioner practices. To further inform our study design, we examine how individual projects demonstrate preferences for pinning or floating (Figure 1(b)). We categorize a project’s overall versioning strategy into *floating only*, *pinning only*, or *mixed* based on the version constraints of its production dependencies. Similar to the trend of version constraints, the popularity of *floating only* strategy increased from 2011 to 2018 but has fallen since then, replaced by more adoption of *pinning only* and *mixed* strategy in the last five years. Our results indicate that although projects sometimes make deliberate decisions per dependency [Jafari et al. 2023], the majority of projects make a uniform decision, motivating us to use similar alternations in our simulations (Section 3.2).

3.2 Method

Overview. At a high level, we take a *counterfactual analysis* approach [Morgan and Winship 2015]—that is, estimating the causal effect of pinning by comparing the “what-if” differences with not pinning in each project. This leads to our design of comparing metrics between a *control* condition consisting of our observed projects, and a (simulated) *treatment* condition wherein we alter each of the project’s package.json to pin all its direct dependencies. We design a set of five outcome metrics in the project dependency graph to measure the five trade-offs raised in favor or against pinning, previously summarized in Table 1. However, the comparison of these metrics is challenging as we would need to simulate how pinning would change the dependency graph of a focal project. What’s more, the values of metrics may change over time as the focal project and the ecosystem continuously evolves. To overcome the first challenge, we leverage a hack in the npm dependency resolver for accurate time-traveling dependency resolution [Pinckney et al. 2023a]. To overcome the second challenge, we organize the simulated data into *balanced panels*, wherein each project is observed and simulated at five different times under both conditions (control and treatment). Then, we use *fixed-effects panel regression* to estimate the effect of pinning separately for each of the five outcome metrics. Panel regression is a statistically sound econometric technique to model double-dimensional data (data collected from multiple units across different time periods) robustly to unobserved heterogeneity (i.e., project/time-specific characteristics that affect the outcome metric but are not directly measured or included in the model) [Brüderl and Ludwig 2015; Hanck et al. 2021]. Next, we give more details on the simulation setup and regression analysis.

Simulation Setup. To implement the above design, we get the latest revision between September 2021 and September 2022 for each project in our dataset. Then, we use September 12, 2022 as a starting time point for simulations (denoted t_0). We recreate the state of each project at this timestamp t_0 , including the package.json files, and resolve the corresponding dependency graphs (denoted as G_{t_0}). Such historical resolution is non-trivial as we will need to reconstruct the state of the npm ecosystem at those times. Fortunately, thanks to a previous report by Pinckney et al. [2023a], we use the undocumented `--before` argument in the npm command-line tool to enable *time-traveling* dependency resolution. Next, we create an altered version of each package.json file by removing the “caret” and “tilde”, i.e., pinning all direct dependencies to the minimal version specified in package.json, and resolve the simulated dependency graph (G'_{t_0}) analogously. This step results in two observations for each project, that feed into our “control” and “treatment” groups, respectively. Similarly, we check out four additional versions of the package.json file, at four evenly-spaced

fixed 90-day intervals from December 11, 2022 to September 7, 2023 (denoted $t_1 \dots t_4$). At each time point, we resolve the corresponding original and simulated (all-pinning) dependency graphs, which adds four additional observations per project into each of our “control” and “treatment” groups. Note that for each project and treatment condition, we intentionally resolve dependency graphs at different time points from the same `package.json` file; this setting, combined with our outcome metrics below, provides a framework to estimate the costs of action (or risks of not taking action) over the observation period, independent of the actual project actions taken during that period.

Outcome Metrics. For each of the five original and simulated dependency graphs resolved for a given project, we compute five security- and maintenance-related metrics, each corresponding to one dimension of trade-offs in Table 1. Their definitions are as follows:

- (1) # of floating dependencies (`n_floating`): The number of edges in the dependency graph with a floating version constraint. This metric captures *attack surface*, i.e., the number of packages where a malicious update would be installed automatically if they were attacked.
- (2) # of automatic updates (`n_auto_updates`): The number of automatic updates happened between t_0 and t for the dependency graph at time t . This metric captures the risk of *breaking changes*, assuming that breaking changes happen uniformly randomly among new releases.
- (3) # of security vulnerabilities (`n_vuln`): The number of known vulnerable dependencies in the dependency graph based on entries in GitHub Advisory [GitHub 2024] at time t . This metric captures the maintenance cost introduced by the need to remove *security vulnerabilities* in the project (usually by manually updating to a newer version).
- (4) # of outdated dependencies (`n_outdated_deps`): The number of outdated direct dependencies in the dependency graph is based on whether a newer release exists at time t . This metric captures the maintenance cost introduced by the need to manually update *outdated dependencies*.
- (5) # of bloated dependencies (`n_bloated`): The number of packages with more than one version simultaneously present in the dependency graph (i.e., bloated). This metric captures the maintenance cost introduced by *dependency conflicts* because npm resolves such conflicts by allowing different versions of the same package at different levels [npm 2024a].

Intuitively, `n_floating` and `n_auto_updates` should decrease with pinning and increase with floating; `n_vuln`, `n_outdated_deps` and `n_bloated` should decrease with floating and increase with pinning. We designed all metrics such that lower values are generally better (e.g., lower costs or risk).

Panel Regression. For each metric M , we fit the following panel regression model:

$$\ln(M + 1) \sim \text{pinning} + \ln(\text{size}(G)) + \text{pinning} \times \ln(\text{size}(G)) + \alpha + \beta \quad (1)$$

In this model specification, α and β represent project-fixed and time-fixed effects, respectively. These two variables give the model robustness to unobserved heterogeneity, allowing us to model the unmeasured project-specific or time-specific characteristics that might affect the dependent variable. The main variable of interest is the *pinning* flag, whose estimated coefficient, if statistically significantly different from zero, represents the average change in metric M relative to the default “control” condition, while holding fixed the other factors. Finally, we model an explicit interaction with the size of the dependency graph $\text{size}(G)$ (i.e., number of nodes), to test whether the effect of pinning varies with $\text{size}(G)$ —one can expect, e.g., that having more direct and transitive dependencies increases the attack surface of a downstream project. The models are fitted using the R `plm` package [Croissant et al. 2024].

Limitations and Threats to Validity. Our model design should allow reliable estimates of the impact of pinning, but it relies on accurate dependency resolution at past points in time. While we believe that our solution is accurate, simulations can still deviate from the realities in the past due to npm

behavior changes, configuration differences, etc. The dependency resolution attempts may also fail with an error, upon which we discard the data point: In total, we have a success ratio of 83.6% for npm packages and 87.7% for GitHub repositories. For the failed resolutions, npm’s error logs indicate the failures can be due to *bad version constraints* (41.9%), *missing packages* (37.0%), *git errors* (8.9%), *bad URLs* (5.7%), *missing local files* (2.6%), *incompatible OS* (2.4%), and others. Resolution failures can happen for many reasons, such as a release being taken down by npm and a project depending on a package not present in the npm registry (i.e., likely proprietary). As such software builds are inherently volatile [Seo et al. 2014; Tomassi et al. 2019], we believe these errors are difficult, if not impossible, to fix in an automated simulation pipeline. Thus, we conclude that the success ratio is acceptable and the scale of our data is sufficient for our study design, being comparable with or larger than datasets in the medical and econometrics literature where similar designs are adopted (e.g., [Crouse et al. 1999; Gerstorf et al. 2010; Stutzer and Frey 2006]).

Our simulation setup enables a holistic approach to measure the costs of taking action (or the risks of not), but does not fully take into account the complexity of project actions that may happen over a year. For example, a project that opts for pinning can invest no effort in updating its pinned dependencies, rapidly update upon every new release of a dependency, or take an intermediate approach (e.g., update every 30 days). In other words, for at least some of our outcome metrics (e.g., security vulnerabilities), developers using a pinning strategy can choose how much effort to invest in manual actions to further trade-off costs and risks (Figure 3). In this paper, all outcome metrics are estimated from the conservative case of no manual developer actions; we leave the exploration of other cases to future work.




Fig. 3. An illustration of the trade-off between cost and risk after pinning.

In terms of the metric design, we implicitly assume that malicious package updates and breaking changes follow a uniform random distribution, counting potential rather than actual events. For malicious package updates, real attackers are generally opportunistic, and existing research does not show any noticeable statistical patterns in attack target selection [Ladisa et al. 2023; Ohm et al. 2020]. Thus, we believe it is reasonable to assume attacks to happen uniformly randomly among packages. For breaking changes, it could be possible to get more accurate estimations from static analysis or regression testing, but adopting either would severely limit and bias the amount of data available for modeling purposes. Specifically, (1) static analysis only works well on TypeScript packages, (2) regression testing is only applicable to projects with a running test suite and prior regression testing studies collected at most a few hundred JavaScript projects with breaking changes [Kong et al. 2024; Mezzetti et al. 2018]. In fact, we believe that breaking changes are so hard to avoid—even in well-managed projects (see Hyrum’s Law [Wright 2024])—and well-tolerated in the npm ecosystem [Bogart et al. 2021] that the uniformly random assumption is actually a reasonable approximation to quantify the expected risk (or cost) of breaking changes on a large scale npm based dataset.

In terms of external validity, it is important to note that our study are specific to the npm ecosystem and the results may not generalize to other ecosystems that adopt different baseline practices. For example, the Maven ecosystem (Java) is smaller [DeBill 2024], conducts pinning extensively [Dietrich et al. 2019], and handles dependency conflicts differently [Wang et al. 2018], so the trade-off between pinning and floating could manifest in a different manner. In addition, generalizations to closed-source applications should also be careful as they may have different characteristics compared with our dataset.

Table 3. The descriptive statistics (aggregated over all time points) and the longitudinal evolution of the studied security and maintenance related metrics between the control and treatment group (pinning).

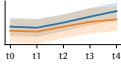
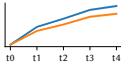
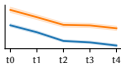
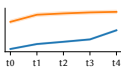
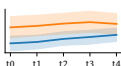
Metric	Condition	Min	25%	Median	75%	Max	Mean	Evolution
n_floating	— Control	0	8	65	335	9,191	432.35	
	— Pinning	0	5	60	327	8,909	427.83	
n_auto_updates	— Control	0	0	3	21	1,354	30.30	
	— Pinning	0	0	2	14	662	24.86	
n_vuln	— Control	0	0	0	2	85	2.13	
	— Pinning	0	0	0	2	97	2.70	
n_outdated_deps	— Control	0	1	2	7	225	5.96	
	— Pinning	0	1	4	11	302	8.78	
n_bloated	— Control	0	0	1	13	650	22.38	
	— Pinning	0	0	1	15	639	23.51	

Table 4. The fixed-effect panel regression results. We obtain a high goodness-of-fit for n_floating ($R^2 \geq 0.6$), acceptable goodness-of-fit for n_vuln, n_outdated_deps, and n_bloated ($R^2 \geq 0.1$), and low goodness-of-fit for n_auto_updates ($R^2 < 0.1$). The R^2 values are comparable with similar prior studies [see, e.g., Valiev et al. 2018; Zahan et al. 2023b]. Note that since our goal is explanatory (not predictive), we do not evaluate models based on R^2 values, but conduct a series of robustness checks to see whether the findings are robust over different settings (Section 3.3). We report local effect sizes computed using Cohen's f^2 [Selya et al. 2012].

	$\ln(\text{Dependent Variable}) + 1$				
	n_floating	n_auto_updates	n_vuln	n_outdated_deps	n_bloated
<i>pinning</i>	-0.519 ₊₊₊ ^{***} (0.006)	-0.274 ₊ ^{***} (0.005)	0.011 ₊ ^{***} (0.005)	0.302 ₊₊₊ ^{***} (0.005)	0.029 ₊ ^{***} (0.003)
$\ln(\text{size}(G))$	1.084 ₊₊₊ ^{***} (0.014)	0.794 ₊ ^{***} (0.036)	0.283 ₊ ^{***} (0.021)	0.068 ^{***} (0.011)	0.631 ₊₊ ^{***} (0.027)
<i>pinning</i> \times $\ln(\text{size}(G))$	0.084 ₊₊₊ ^{***} (0.001)	0.003 (0.001)	0.028 ₊ ^{***} (0.001)	0.019 ^{***} (0.001)	0.007 ^{***} (0.001)
R^2	0.715	0.053	0.154	0.330	0.227
Adjusted R^2	0.684	-0.052	0.059	0.255	0.141

*** $p < 0.001$, +++ large effect size, ++ medium effect size, + small effect size

3.3 Results

In this section, we report results based on *production dependency graphs* (i.e., excluding development dependencies) from all projects in all time points (Table 4). Wald tests show that the full model (with pinning and the interaction) is preferable for all five outcome variables ($p < 0.001$) compared to a null model with only $\text{size}(G)$ as an independent variable. Note while development dependencies are generally considered irrelevant for security vulnerabilities [Latendresse et al. 2022; Pashchenko et al. 2018], they form an important attack surface for malicious package updates that compromise build systems [Ladisa et al. 2023; Ohm et al. 2020]. Therefore, we have also explored the robustness of our findings: (1) in dependency graphs with development dependencies, (2) separately on npm packages and GitHub repositories, and (3) using only data collected at t_0 . The results are highly consistent and available in the replication package (Section 6).




Fig. 4. The effect of pinning direct dependencies on malicious package update attack surface (as measured by the number of floating dependencies in the entire dependency graph). Our model shows that pinning reduces attack surface for small dependency graphs, but the positive effect decreases as the dependency graphs grow larger and even flips to negative for dependency graphs with more than 498 dependencies.

For the effect of pinning, we find that pinning all direct dependencies has a significant influence on the three measures of cost—higher exposure to security vulnerabilities (n_vuln), number of outdated dependencies that need to be manually updated ($n_outdated_deps$), and higher exposure to dependency conflicts ($n_bloated$). Those results are supported by the differences in the descriptive results over different time periods (Table 3) and the significant coefficients in our panel regression models (Table 4); the interaction term shows that all three costs of pinning increase with dependency graph sizes; the exposure to security vulnerabilities has the strongest interaction effect.

The effect of pinning on the two metrics expected to show benefits of pinning (reduced attack surface and fewer breaking changes) is dim. For the number of automatic updates, the regression model fits poorly and provides little explanatory power (other factors like picking dependencies with fewer releases likely have a stronger effect than pinning). For the number of floating dependencies, the interaction effect with the size of the dependency graph (see Table 4 and illustration in Figure 4) shows that the benefit of pinning on reducing the attack surface is smaller with larger dependency graphs and even (counter-intuitively) negative for very large dependency graphs.

4 RQ2: How effective is coordinating pinning among selected upstream packages in reducing the risk of malicious package updates ecosystem-wide?

As the results from RQ1 (Section 3.3) indicate, pinning direct dependencies is ineffective and possibly even harmful for preventing malicious package updates if a project has a large and complex dependency graph (we will return to this issue in Section 5). Then, a natural follow-up question arises: Are there better, alternative interventions that can effectively reduce the risk of malicious package updates in the npm ecosystem? Inspired by previous research that reveals the structures of dependency networks [Decan et al. 2019; Zimmermann et al. 2019], we hypothesize that *the risk of malicious package updates can be greatly reduced if a small set of “core” packages pin and promptly update all their dependencies after proper auditing* (ensuring there is no outdated dependency but also no malicious update). That is, instead of all developers redundantly performing security actions, a few package maintainers (or a dedicated security team for the ecosystem) might

effectively take on the cost of (some) security work by pinning and actively managing dependencies, so that all developers benefit. This would enable a vision where pooling some community resources would be much more efficient and effective to achieve security outcomes than trying to convince every single maintainer to change their practices (e.g., in a way similar to how the Linux Foundation starts the Core Infrastructure Initiative in response to HeartBleed [Walden 2020] or how Google provides the Assured Open Source Software service [Google 2024]). This leads to our RQ2 of exploring how coordinated pinning in select upstream packages may reduce the overall risk of malicious package updates at an ecosystem level.

4.1 Method

Overview. In a nutshell, we quantify the effect a few maintainers can have on the risk of (simulated) malicious package updates on the entire npm ecosystem. We assume that we have the ability and budget to convince n package maintainers of our choice to change their dependency management practices as follows: They pin each of their dependencies, review all dependency updates to catch malicious updates, and rapidly update all non-malicious updates. This way, we would expect that some malicious updates are effectively detected by these n maintainers, preventing them from affecting other developers in the ecosystem. Focusing on whether this approach is feasible at all, we make the following optimistic, simplifying assumptions: The n maintainers review and update all dependencies immediately (without adding latency) and they recognize all malicious package updates when they audit an update. We also assume that reporting and taking down a malicious package takes some time (Ohm et al. [2020] report a mean of 209 days), so even if some maintainers detect a malicious update and protect their users, others may still be exposed. Finally, we assume all other developers are floating all their dependencies, as they already mostly do.

Dependency Network. We simulate attacks on a snapshot of the package-level npm dependency network constructed from dependency relationships specified in the latest version of each package (cut-off in September 2023), using the same npm-follower dataset [Pinckney et al. 2023b] in Section 3.1. The network has 1,566,434 packages (nodes) and 7,873,239 dependency edges.

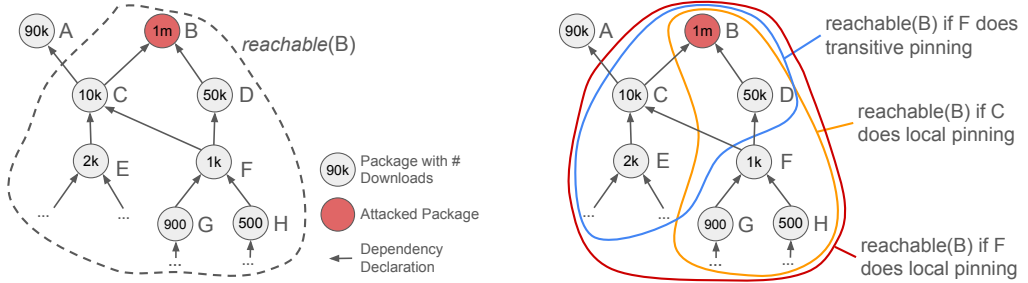
Attack Selection Strategy. In practice, attackers might target specific applications (e.g., specifically attacking cryptocurrency packages [Checkmarx 2023]) or they may try to compromise whatever package they can to opportunistically cause damage (e.g., as common for untargeted ransomware attacks). Since we do not know what exact strategy attackers may pursue and since their strategy may change over time, we make the simplifying assumption that attackers will try to compromise high-impact targets by (randomly) selecting any of the top- m packages (denoted as A) on npm as ranked by the downloads adjusted impact defined in Equation 2 (i.e., impact-based strategy).

Risk Metric. Given the set P of all packages in the network, we quantify the impact of compromising a package $a \in P$ as the download-weighted coverage of its reachable dependents within the entire npm network, where every package that depends directly or indirectly on package a is considered reachable of the edge to a in its dependency graph is floating:

$$\text{impact}(a) = \frac{\sum_{p \in \text{reachable}(a)} \text{downloads}(p)}{\sum_{p \in P} \text{downloads}(p)} \quad (2)$$

In other words, $\text{impact}(a) \in [0, 1]$ measures the fraction of npm downloads that will be affected if package a is compromised, as illustrated in Figure 5(a).

For the set of packages $A \subset P$ (i.e., the top- m most impactful packages), in which an attacker may target any of $p \in A$ based on the attack selection strategy, we define $\text{risk}(A)$ as the arithmetic mean of the impact across all possible targets in A an attacker can choose from (i.e., $\text{risk}(A) =$



(a) **Attack Impact Computation:** Packages C to H all directly or indirectly depend on B and are hence reachable if B is attacked, but package A is not. The impact of compromising package B is the sum of downloads of all reachable packages (C-H) divided by the total number of downloads across all packages. In this case, $impact(B) = (1m + 10k + 5k + 2k + 1k + 900 + 500) / (90k + 1m + 10k + 5k + 2k + 1k + 900 + 500) \approx 0.92$.

(b) **Local/Transitive Pinning:** Assume B is being attacked. If C pins its dependencies, the attack would not reach E, as C would not update to the malicious version. However, local pinning of F would have no effect, since the attack still propagates through the floating edges originating from C and D. In contrast, transitive pinning of F would protect G and H (and F's other dependents) from the attack.

Fig. 5. Examples to illustrate attack impact computation and how local/transitive pinning works.

$\frac{1}{|A|} \sum_{a \in A} impact(a)$. Intuitively, $risk(A) \in [0, 1]$ is an estimation of the expected average percentage of npm downloads that would be affected by compromising any one of the packages in A.

Pinning Mechanism. With npm's existing mechanisms, the maintainers of the n defending packages can only pin their direct dependencies (*local pinning*), but not their transitive ones. Maintainers can freeze their dependency graph with a `package-lock.json` file, which, with npm's current design, only has an effect on their own package and not on anyone importing their package.

To give maintainers more power (at a higher maintenance cost), we imagine an alternative strategy of *transitive pinning*, where maintainers can pin direct and indirect dependencies. This could be implemented with a change to npm that allows (a) package maintainers to upload their `package-lock.json` as part of the deployed package and (b) npm to consider that file in each package during dependency resolution. We illustrate the difference in Figure 5(b).

Defense Selection Strategy. Finally, we need to decide which n packages to defend with local or transitive pinning. Unfortunately, the problem of choosing an optimal package set of size n that minimizes $risk(A)$ is equivalent to the Influence Minimization problem, which is NP-hard [Xie et al. 2023]. Hence, we focus on evaluating simple heuristics to choose the defended packages and leave the optimization as future work. Specifically, we choose top- n packages with (1) highest downloads, (2) highest out degrees (i.e., the packages with the most *direct* dependencies), (3) highest betweenness centrality (i.e., the packages on most dependency *paths*), (4) highest downloads \times out degrees (i.e., popular packages with many direct dependencies), and (5) highest downloads \times betweenness centrality (i.e., popular packages on many dependency paths). Intuitively, we expect defense selection strategies based on out degree to work best for local pinning as they emphasize direct dependencies, and strategies based on betweenness to work best for transitive pinning as transitive pinning interrupts paths.

Simulation. For each combination of pinning mechanism, defense selection strategy, and number of defended packages from 0 to n , we measure the impact of the defense on the risk metric. To compute the impact, we compute $risk(A)$ with a modified *reachable* function (Equation 2, different from the standard notion of graph reachability) by breaking propagation of the attack when the final edge is pinned (local pinning) or if the path contains a defended package (transitive pinning).




Fig. 6. The reduction of malicious package update risk ($risk(A)$) by securing top n packages with five different defense selection strategies using local pinning and transitive pinning, respectively.

Limitations and Threats to Validity. To make the simulation feasible, we make several simplifying assumptions described above. Notably, we assume that attacks happen uniformly randomly on a fixed number of attack targets, while real attack selection strategies may be more targeted, more opportunistic, and may adapt in response to defense strategies. Therefore, we additionally explored two different attack selection strategies: (1) attacking any actively maintained projects randomly, and (2) attacking the most fragile or under-maintained packages (e.g., as the xz attack, which focused on a dependency with a single overworked maintainer [Sonatype 2024b]) defined as the bottom-1000 packages in our package dataset, ranked by the number of core maintainers estimated in World of Code [Ma et al. 2021]. For space and since our results are generally robust under the different attack selection strategies, we only report results under the impact-based attack selection strategy (the remaining are in the replication package, Section 6). What’s more, the dependency network is only an approximation of reverse dependency relationships. However, similar kinds of dependency networks are frequently used in prior research [Decan et al. 2019; Zimmermann et al. 2019] because it is generally impossible to compute precise reverse dependencies as would be resolved by npm due to its deduplication optimizations (Section 2). For defense selection strategies, we have to rely on simple heuristics, leaving potentially better strategies (e.g., more computationally expensive optimization attempts) to future work. Finally, due to a lack of access to real-world application data, our simulation does not accurately quantify the impact of malicious package updates on applications, whether open-source repositories or closed-source projects outside of npm. We mitigate this by considering downloads in attack impact estimation (Equation 2), but npm download data can be also noisy and inconsistent. Due to these simplifying assumptions, readers should take our results as a first exploration of broad defense strategies and their ecosystem-wide effects, rather than specific implementation recommendations.

4.2 Results

In this section, we report results for $n = 100$ and $m = 1,000$ (i.e., defending up to 100 packages to minimize the risk of attack among the top 1000 most impactful packages; the results are robust under other n, m variations which are available in the replication package, Section 6). As visible in Figure 6, one would expect 1.93% of npm downloads to be affected on average upon an attack without any intervention, which is enormous considering that npm has ~ 12 billion downloads per day in 2024 [Sonatype 2024a]. Most defense strategies can reduce ecosystem-wide risk (from doing nothing at $n = 0$) with targeted defenses at 10 to 100 packages, but their effectiveness varies widely. The best local pinning strategy can reduce ecosystem-wide risk against the 1000 most popular




Fig. 7. A minimal example illustrating how pinning direct dependencies can cause dependency conflicts and lead to a bloated dependency graph with more transitive floating dependencies. Note that this is only an artificial example for simplicity; real-world cases happen mostly in much larger dependency graphs.

packages by 12.21% with defending only 20 packages and 30.09% when defending 100 packages. The best transitive pinning strategy can reduce ecosystem-wide risk against the 1000 most popular packages by 51.34% when defending only 20 packages and 75.81% when defending 100 packages. For both pinning mechanisms, selecting packages with betweenness weighted by downloads is most effective, but transitive pinning is usually two to three times more effective than local pinning.

5 Discussion

5.1 Why can pinning direct dependencies lead to a larger total number of floating dependencies (i.e., a larger attack surface for malicious package updates)?

This surprising effect is explained by a detail of how the npm dependency resolver handles dependency conflicts [npm 2024a] (cf. Section 2): If two or more nodes in the dependency graph declare a dependency on the same package, npm will first check (with heuristics) whether there is a version of that package that meets all declared version constraints (as in our example in Figure 7(a)). If there are conflicting version constraints in dependency declarations, npm will install multiple versions of the package to satisfy all version constraints (as in our example in Figure 7(b)). This deduplication (or debloating) process is designed to download fewer, often redundant, copies of packages in the installation process. In fact, there are many possible deduplication heuristics and researchers have explored how to improve this in various ways (e.g., using an SMT solver [Pinckney et al. 2023c]).

The deduplication process interacts with pinning: If more dependencies are pinned, the dependency graph resolution has fewer opportunities to deduplicate a specific package and will be more likely to create larger dependency graphs with multiple versions of the same package. Those different versions may also then include additional dependencies—as visible in our example in Figure 7. Both effects increase the attack surface, in that attackers have possibly more version ranges they can attack for the same packages and there could be even additional packages they can attack. This leads to our conclusion that *the practice of (local) pinning is futile and even counter-productive for the purpose of reducing floating dependencies (aka attack surface) in one’s dependency graph.*

Note that the deduplication heuristics introduced here and the ability to install multiple versions of the same package are unique to npm [Bogart et al. 2021]. In most other dependency managers, the resolver must find a single version for each package and pinning can make this much harder (as captured also in our `n_bloated` metric for pinning). In those ecosystems, pinning would arguably still have only marginal effects on attack surface reduction, but not negative ones. However, the costs for resolving dependency conflicts would be much more substantial [Wang et al. 2020, 2018].

5.2 Implication for Practitioners

Dependency pinning is (as expected) an expensive intervention. Regarding the drawbacks of pinning, our RQ1 results align with the argued trade-offs in Table 1: Pinning direct dependencies substantially increases the effort for developers to manually update (and possibly review) dependencies if they want to stay up to date, or face many outdated dependencies if they do not. Also, projects that pin their dependencies are, as expected, much more likely to be exposed to known vulnerabilities in their dependency graph if they do not constantly update to available patches. While bots and notification tools like *Dependabot* can help with some of that work and even automatically run tests and merge updates (as covered in prior work [e.g., Alfadel et al. 2021; Mirhosseini and Parnin 2017]), they also have their limitations, creating additional complexity (e.g., notification fatigue [He et al. 2021]).

Pinning direct dependencies is futile and (surprisingly) can be even counter-productive as a security intervention if floating is the mainstream practice in the ecosystem. Our RQ1 results regarding the expected security benefits of pinning are surprising and counter-intuitive to practitioners (based on our experience when sharing the results). The first reason is that in many projects, direct dependencies are far outnumbered by transitive dependencies for which version declarations are controlled by upstream packages (Table 2). Thus, the action of a single project has a relatively small impact on the total number of floating dependencies in its dependency graph, explaining the small effective size and the diminishing return of pinning as the dependency graph grow larger (as confirmed by the interaction term $pinning \times \ln(size(G))$ in our panel regression model and illustrated visually in Figure 4). Importantly, the security benefit of pinning does not only diminish, but it *reverses* and pinning becomes actually counterproductive in large dependency graphs (≥ 498 nodes according to our model, also visible as a crossover point in Figure 4), to the point that *pinning direct dependencies leads to more floating dependencies in the entire dependency graph*. In other words, an attacker may be able to attack the same dependencies in more version ranges or may even be able to attack additional dependencies that are only (indirectly) included in some versions of some dependency. More generally speaking, the perceived security trade-off between pinning (to reduce the risk of malicious updates) and floating (to get fewer security vulnerabilities) does not hold in practice if the upstream packages are mainly floating (e.g., in npm and PyPI).

Following the discussion above, **in an ecosystem where floating is the mainstream practice, pinned dependencies do not necessarily signal that a project is more secure from malicious package updates, especially for projects with large dependency graphs.** In such contexts, pinning can create substantial maintenance costs, for limited security benefits, if any. Therefore, we argue that pinning should not be considered when scoring the security practices of packages.

Instead of pinning, application developers should freeze their entire dependency graph with a shared lock file. In the npm ecosystem, it means using and *committing* the `package-lock.json` file (Section 2) to the version control repository (44.21% of GitHub repositories in our dataset did this). This practice enables the effective pinning of all direct and transitive dependencies and more importantly, prevents updates from being installed automatically in a new machine or in the CI/CD environment. However, *lock files are intended only for applications (not packages)*, and they are not designed to allow easy selective updates or audits of dependency updates—hence every developer would need to individually judge when to update which (possibly indirect) dependencies rather than delegating some of that judgment to package maintainers and ecosystem moderators.

To secure an ecosystem from supply chain attacks, package developers may consider community-wide coordination from a collective pool of resources. For example, our RQ2 results indicate that collective pinning actions could indeed be fairly effective at reducing the risk of malicious package updates and *localize* the cost to a few (20–100) carefully selected packages. We

believe it is much more achievable to convince (and raise labor or funds for) changing practices in 100 packages for the good of the entire ecosystem, than trying to convince almost every maintainer in npm individually to change their practices. However, our results also show clearly that npm's current local pinning is far less effective than transitive pinning could be. While transitive pinning is far more expensive for the maintainer defending a package (i.e., they need to audit updates of transitive dependencies too, not just of direct ones), it also provides a bigger lever for security. Transitive pinning is also easier to explain to developers, whereas local pinning can provide a false sense of security with counter-intuitive consequences (as shown in our RQ1). Therefore, **we suggest changes to npm infrastructure and tooling to enable transitive pinning (e.g., publishing a package-lock.json file for downstream consumption).**

However, we must note that these kinds of interventions would probably still experience pushback and be costly to implement from such an established ecosystem (as of pushing any organizational change [Schein 2010]). To make the process easier, **we encourage foundations, security companies, or companies using open source to dedicate resources (labor, funds) to such efforts.** Even if only a small number of maintainers opt to pursue this (i.e., transitive pinning and extensively review their updates), it can have a large impact on the ecosystem.

Finally, it is worth noting that the whole debate around version constraints relies on one critical assumption: *malicious package updates will go undetected in the ecosystem for a significant amount of time*, which is probably still the case as of now (as evidenced by the xz attack [Sonatype 2024b]). If malicious package updates are taken down faster or even stopped before being published in the package registry, their impact will be much smaller. Thus, **efforts are needed to make malicious package updates as short-lived as possible.** We envision that if a small number of core ecosystem packages can pin and *promptly review their updates*, they can also help in such early detection, for example, by leveraging the ongoing research about detecting malicious packages [e.g., Sejfia and Schäfer 2022; Zahan et al. 2024].

5.3 Implication for Researchers

Our study leaves a few open research questions. First, our study only focuses on the npm ecosystem, and other ecosystems adopt different baseline dependency versioning practices [Bogart et al. 2021; Dietrich et al. 2019]. Future research is necessary to explore the impact of pinning/floating under those variations. Second, while prior research provides evidence that floating jeopardizes build stability [e.g., Mukherjee et al. 2021], the `n_auto_update` model in RQ1 only explains a small amount of variance in the data ($R^2 < 0.1$), so we expect future work to better establish the causal relationship between pinning and application stability. Third, our RQ2 only takes a first step to explore possible ecosystem-level strategies, and we encourage more research in this direction. We suggest exploring better strategies than download-weighted betweenness and exploring how to consider users of packages outside npm. Researchers may also explore alternative simulation approaches going beyond our simplifying assumptions, and, for example, consider the specific existing choices of pinning and floating in the dependency graph when suggesting which packages to defend. Future research could also explore whether these strategies hold when baseline pinning and floating practices are different in other ecosystems [Bogart et al. 2021; Dietrich et al. 2019]. Another remaining question is whether a maintainer's decision to pin their package's dependencies produces unintended non-local costs for package users. For example, it is possible that the pinning in some parts of the dependency graph can affect users negatively in terms of causing them dependency conflicts (cf. `n_bloated` in Section 3.3). Researchers may explore whether different defense selection strategies (e.g., explicitly coordinating dependency updates across packages) or explicitly advocating for floating outside of a few defended packages can reduce unintended dependency conflicts. Again, we expect all these concerns to manifest differently in different

ecosystems and we envision a broad range of future research to establish the security impact of versioning practices in a wide variety of contexts.

6 Conclusion

In this paper, we have presented a counterfactual analysis and simulation study on the impact of version constraints in the npm ecosystem. We have explored the security and maintenance impact of local pinning direct dependencies and the effectiveness of pinning on “important” packages against possible malicious package updates. Our study provides counter-evidence for the use of local pinning as a security best practice, but supports the effectiveness of coordinated pinning in critical open-source packages. Our study leads to a series of implications for the design of dependency management tools and ecosystem-level coordination to enhance software supply chain security.

Data Availability. We provide a replication package at <https://doi.org/10.5281/zenodo.14693233>.

Acknowledgement. He’s and Kästner’s work was supported in part by the National Science Foundation (award 2206859). Bogdan’s work was supported in part by the National Science Foundation (award 2317168) and research awards from Google and the Digital Infrastructure Fund. We would like to thank James Herbsleb, Narayan Ramasubbu, Rohan Padhye, and the participants of S3C2 quarterly meetings for their valuable feedback to the earlier versions of this work.

References

- Pietro Abate, Roberto Di Cosmo, Georgios Gousios, and Stefano Zacchiroli. 2020. Dependency Solving Is Still Hard, but We Are Getting Better at It. In *27th IEEE International Conference on Software Analysis, Evolution and Reengineering, SANER 2020, London, ON, Canada, February 18-21, 2020*. IEEE, 547–551. <https://doi.org/10.1109/SANER48275.2020.9054837>
- Rabe Abdalkareem, Olivier Nourry, Sultan Wehaibi, Suhaib Mujahid, and Emad Shihab. 2017. Why do developers use trivial packages? An empirical case study on npm. In *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4-8, 2017*. ACM, 385–395. <https://doi.org/10.1145/3106237.3106267>
- Mahmoud Alfadel, Diego Elias Costa, Emad Shihab, and Mouafak Mkhallalati. 2021. On the Use of Dependabot Security Pull Requests. In *18th IEEE/ACM International Conference on Mining Software Repositories, MSR 2021, Madrid, Spain, May 17-19, 2021*. IEEE, 254–265. <https://doi.org/10.1109/MSR52588.2021.00037>
- Chris Bogart, Christian Kästner, James D. Herbsleb, and Ferdian Thung. 2021. When and How to Make Breaking Changes: Policies and Practices in 18 Open Source Software Ecosystems. *ACM Trans. Softw. Eng. Methodol.* 30, 4 (2021), 42:1–42:56. <https://doi.org/10.1145/3447245>
- Josh Bressers. 2022. *Fact vs. Feelings*. Retrieved December 11, 2023 from <https://opensourcesecurity.io/2022/03/21/facts-vs-feelings/>
- Josef Brüderl and Volker Ludwig. 2015. Fixed-effects panel regression. *The Sage Handbook of Regression Analysis and Causal Inference* 327 (2015), 357.
- Checkmarx. 2023. *NPM Account Takeover Results in Crypto Supply Chain Attack*. Retrieved April 28, 2024 from <https://checkmarx.com/blog/npm-account-takeover-results-in-crypto-supply-chain-attack/>
- Bodin Chinthanet, Raula Gaikovina Kula, Shane McIntosh, Takashi Ishio, Akinori Ihara, and Kenichi Matsumoto. 2021. Lags in the release, adoption, and propagation of npm vulnerability fixes. *Empir. Softw. Eng.* 26, 3 (2021), 47. <https://doi.org/10.1007/S10664-021-09951-X>
- Joel Cox, Eric Bouwers, Marko C. J. D. van Eekelen, and Joost Visser. 2015. Measuring Dependency Freshness in Software Systems. In *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015, Florence, Italy, May 16-24, 2015, Volume 2*. IEEE Computer Society, 109–118. <https://doi.org/10.1109/ICSE.2015.140>
- Yves Croissant, Giovanni Millo, Kevin Tappe, Ott Toomet, Christian Kleiber, Achim Zeileis, Arne Henningsen, Liviu Andronic, and Nina Schoenfelder. 2024. *plm: Linear Models for Panel Data*. Retrieved April 28, 2024 from <https://cran.r-project.org/web/packages/plm>
- John R Crouse, Timothy Morgan, James G Terry, Julie Ellis, Mara Vitolins, and Gregory L Burke. 1999. A randomized trial comparing the effect of casein with that of soy protein containing varying amounts of isoflavones on plasma concentrations of lipids and lipoproteins. *Archives of Internal Medicine* 159, 17 (1999), 2070–2076. <https://doi.org/10.1001/archinte.159.17.2070>

- Mark Curphey. 2022. *Dependency Pinning Only Works If You Actually Review the Updates*. Retrieved December 11, 2023 from <https://crashoverride.com/blog/dependency-pinning-only-works-if-you-actually-review-the-updates>
- Erik DeBill. 2024. *Modulecounts*. Retrieved May 23, 2024 from <http://www.modulecounts.com/>
- Alexandre Decan and Tom Mens. 2021. What Do Package Dependencies Tell Us About Semantic Versioning? *IEEE Trans. Software Eng.* 47, 6 (2021), 1226–1240. <https://doi.org/10.1109/TSE.2019.2918315>
- Alexandre Decan, Tom Mens, and Eleni Constantinou. 2018. On the impact of security vulnerabilities in the npm package dependency network. In *Proceedings of the 15th International Conference on Mining Software Repositories, MSR 2018, Gothenburg, Sweden, May 28-29, 2018*. ACM, 181–191. <https://doi.org/10.1145/3196398.3196401>
- Alexandre Decan, Tom Mens, and Philippe Grosjean. 2019. An empirical comparison of dependency network evolution in seven software packaging ecosystems. *Empir. Softw. Eng.* 24, 1 (2019), 381–416. <https://doi.org/10.1007/S10664-017-9589-Y>
- Jens Dietrich, David J. Pearce, Jacob Stringer, Amjed Tahir, and Kelly Blincoe. 2019. Dependency versioning in the wild. In *Proceedings of the 16th International Conference on Mining Software Repositories, MSR 2019, 26-27 May 2019, Montreal, Canada*. IEEE / ACM, 349–359. <https://doi.org/10.1109/MSR.2019.00061>
- Nadia Eghbal. 2016. Roads and Bridges. *The Unseen Labor Behind our Digital Infrastructure* (2016).
- ENISA. 2023. *Good Practices for Supply Chain Cybersecurity*. Retrieved August 27, 2024 from <https://www.enisa.europa.eu/publications/good-practices-for-supply-chain-cybersecurity>
- Gang Fan, Chengpeng Wang, Rongxin Wu, Xiao Xiao, Qingkai Shi, and Charles Zhang. 2020. Escaping dependency hell: finding build dependency errors with the unified dependency graph. In *ISSTA '20: 29th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, USA, July 18-22, 2020*. ACM, 463–474. <https://doi.org/10.1145/3395363.3397388>
- Dan Geer, Bentz Tozer, and John Speed Meyers. 2020. For Good Measure: Counting Broken Links: A Quant’s View of Software Supply Chain Security. *login Usenix Mag.* 45, 4 (2020). <https://www.usenix.org/publications/login/winter2020/geer>
- Denis Gerstorf, Nilam Ram, Guy Mayraz, Mira Hidajat, Ulman Lindenberger, Gert G Wagner, and Jürgen Schupp. 2010. Late-life decline in well-being across adulthood in Germany, the United Kingdom, and the United States: Something is seriously wrong at the end of life. *Psychology and Aging* 25, 2 (2010), 477. <https://doi.org/doi.org/10.1037/a0017543>
- GitHub. 2024. *GitHub Advisory Database*. Retrieved April 21, 2024 from <https://github.com/advisories>
- Google. 2024. *Assured Open Source Software*. Retrieved September 8, 2024 from <https://cloud.google.com/security/products/assured-open-source-software>
- Hacker News. 2017. *Package Management: The problem with using version ranges*. Retrieved December 11, 2023 from <https://news.ycombinator.com/item?id=14759329>
- Hacker News. 2018. *I have misgivings about all these version pinning...* Retrieved December 13, 2023 from <https://news.ycombinator.com/item?id=16422916>
- Christoph Hanck, Martin Arnold, Alexander Gerber, and Martin Schmelzer. 2021. Regression with Panel Data. In *Introduction to Econometrics with R*. Universität Duisburg-Essen. <https://www.econometrics-with-r.org/10-rwpd.html>
- Hao He, Runzhi He, Haiqiao Gu, and Minghui Zhou. 2021. A large-scale empirical study on Java library migrations: prevalence, trends, and rationales. In *ESEC/FSE '21: 29th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Athens, Greece, August 23-28, 2021*. ACM, 478–490. <https://doi.org/10.1145/3468264.3468571>
- Runzhi He, Hao He, Yuxia Zhang, and Minghui Zhou. 2023. Automating Dependency Updates in Practice: An Exploratory Study on GitHub Dependabot. *IEEE Trans. Software Eng.* 49, 8 (2023), 4004–4022. <https://doi.org/10.1109/TSE.2023.3278129>
- IBM. 2021. *What is the Log4j vulnerability?* Retrieved December 10, 2023 from <https://www.ibm.com/topics/log4j>
- Jay Jacobs, Sasha Romanosky, Benjamin Edwards, Michael Roytman, and Idris Adjerid. 2019. Exploit Prediction Scoring System (EPSS). *CoRR abs/1908.04856* (2019). arXiv:1908.04856 <http://arxiv.org/abs/1908.04856>
- Abbas Javan Jafari, Diego Elias Costa, Rabe Abdalkareem, Emad Shihab, and Nikolaos Tsantalis. 2022. Dependency Smells in JavaScript Projects. *IEEE Trans. Software Eng.* 48, 10 (2022), 3790–3807. <https://doi.org/10.1109/TSE.2021.3106247>
- Abbas Javan Jafari, Diego Elias Costa, Emad Shihab, and Rabe Abdalkareem. 2023. Dependency Update Strategies and Package Characteristics. *ACM Trans. Softw. Eng. Methodol.* 32, 6 (2023), 149:1–149:29. <https://doi.org/10.1145/3603110>
- Dhanushka Jayasuriya, Valerio Terragni, Jens Dietrich, Samuel Ou, and Kelly Blincoe. 2023. Understanding Breaking Changes in the Wild. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSA 2023, Seattle, WA, USA, July 17-21, 2023*. ACM, 1433–1444. <https://doi.org/10.1145/3597926.3598147>
- Dezhen Kong, Jiakun Liu, Lingfeng Bao, and David Lo. 2024. Towards Better Comprehension of Breaking Changes in the NPM Ecosystem. *arXiv preprint arXiv:2408.14431* (2024). <https://arxiv.org/pdf/2408.14431>
- Raula Gaikovina Kula, Daniel M. Germán, Ali Ouni, Takashi Ishio, and Katsuro Inoue. 2018. Do developers update their library dependencies? - An empirical study on the impact of security advisories on library migration. *Empir. Softw. Eng.* 23, 1 (2018), 384–417. <https://doi.org/10.1007/S10664-017-9521-5>
- Piergiorgio Ladisa, Henrik Plate, Matias Martinez, and Olivier Barais. 2023. SoK: Taxonomy of Attacks on Open-Source Software Supply Chains. In *44th IEEE Symposium on Security and Privacy, SP 2023, San Francisco, CA, USA, May 21-25,*

2023. IEEE, 1509–1526. <https://doi.org/10.1109/SP46215.2023.10179304>
- Jasmine Latendresse, Suhaib Mujahid, Diego Elias Costa, and Emad Shihab. 2022. Not All Dependencies are Equal: An Empirical Study on Production Dependencies in NPM. In *37th IEEE/ACM International Conference on Automated Software Engineering, ASE 2022, Rochester, MI, USA, October 10-14, 2022*. ACM, 73:1–73:12. <https://doi.org/10.1145/3551349.3556896>
- Max Leiter. 2023. *Pin your npm/yarn dependencies*. Retrieved December 13, 2023 from <https://maxleiter.com/blog/pin-dependencies>
- Chengwei Liu, Sen Chen, Lingling Fan, Bihuan Chen, Yang Liu, and Xin Peng. 2022. Demystifying the Vulnerability Propagation and Its Evolution via Dependency Trees in the NPM Ecosystem. In *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*. ACM, 672–684. <https://doi.org/10.1145/3510003.3510142>
- Lucid. 2017. *Package Management: Stop Using Version Ranges*. Retrieved December 13, 2023 from <https://www.lucidchart.com/techblog/2017/03/15/package-management-stop-using-version-ranges/>
- Yuxing Ma, Tapajit Dey, Chris Bogart, Sadika Amreen, Marat Valiev, Adam Tutko, David Kennard, Russell Zaretzki, and Audris Mockus. 2021. World of code: enabling a research workflow for mining and analyzing the universe of open source VCS data. *Empir. Softw. Eng.* 26, 2 (2021), 22. <https://doi.org/10.1007/S10664-020-09905-9>
- Fabio Mancinelli, Jaap Boender, Roberto Di Cosmo, Jerome Vouillon, Berke Durak, Xavier Leroy, and Ralf Treinen. 2006. Managing the Complexity of Large Free and Open Source Package-Based Software Distributions. In *21st IEEE/ACM International Conference on Automated Software Engineering (ASE 2006), 18-22 September 2006, Tokyo, Japan*. IEEE Computer Society, 199–208. <https://doi.org/10.1109/ASE.2006.49>
- Jeferson Martinez and Javier M Durán. 2021. Software supply chain attacks, a threat to global cybersecurity: SolarWinds' case study. *International Journal of Safety and Security Engineering* 11, 5 (2021), 537–545. <https://doi.org/10.18280/ijss.110505>
- Gianluca Mezzetti, Anders Møller, and Martin Toldam Torp. 2018. Type Regression Testing to Detect Breaking Changes in Node.js Libraries. In *32nd European Conference on Object-Oriented Programming, ECOOP 2018, July 16-21, 2018, Amsterdam, The Netherlands (LIPICs, Vol. 109)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:24. <https://doi.org/10.4230/LIPICs.ECOOP.2018.7>
- Courtney Miller, Christian Kästner, and Bogdan Vasilescu. 2023. "We Feel Like We're Winging It": A Study on Navigating Open-Source Dependency Abandonment. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2023, San Francisco, CA, USA, December 3-9, 2023*. ACM, 1281–1293. <https://doi.org/10.1145/3611643.3616293>
- Samim Mirhosseini and Chris Parnin. 2017. Can automated pull requests encourage software developers to upgrade out-of-date dependencies?. In *Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering, ASE 2017, Urbana, IL, USA, October 30 - November 03, 2017*. IEEE Computer Society, 84–94. <https://doi.org/10.1109/ASE.2017.8115621>
- Stephen L Morgan and Christopher Winship. 2015. *Counterfactuals and causal inference*. Cambridge University Press.
- Suchita Mukherjee, Abigail Almanza, and Cindy Rubio-González. 2021. Fixing dependency errors for Python build reproducibility. In *ISSA '21: 30th ACM SIGSOFT International Symposium on Software Testing and Analysis, Virtual Event, Denmark, July 11-17, 2021*, Cristian Cadar and Xiangyu Zhang (Eds.). ACM, 439–451. <https://doi.org/10.1145/3460319.3464797>
- npm. 2023a. *About semantic versioning | npm Docs*. Retrieved December 14, 2023 from <https://docs.npmjs.com/about-semantic-versioning>
- npm. 2023b. *npm*. Retrieved December 11, 2023 from <https://docs.npmjs.com/about-npm>
- npm. 2023c. *package-lock.json | npm Docs*. Retrieved December 13, 2023 from <https://docs.npmjs.com/cli/v10/configuring-npm/package-lock-json>
- npm. 2024a. *How npm3 Works*. Retrieved April 28, 2024 from <https://npm.github.io/how-npm-works-docs/npm3/how-npm3-works.html>
- npm. 2024b. *npm-install | npm Docs*. Retrieved April 17, 2023 from <https://docs.npmjs.com/cli/v8/commands/npm-install>
- Marc Ohm, Henrik Plate, Arnold Sykosch, and Michael Meier. 2020. Backstabber's Knife Collection: A Review of Open Source Software Supply Chain Attacks. In *Detection of Intrusions and Malware, and Vulnerability Assessment - 17th International Conference, DIMVA 2020, Lisbon, Portugal, June 24-26, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12223)*. Springer, 23–43. https://doi.org/10.1007/978-3-030-52683-2_2
- OpenSSF. 2023. *OpenSSF Scorecard*. Retrieved December 11, 2023 from <https://github.com/ossf/scorecard/blob/main/docs/checks.md#pinned-dependencies>
- Ivan Pashchenko, Henrik Plate, Serena Elisa Ponta, Antonino Sabetta, and Fabio Massacci. 2018. Vulnerable open source dependencies: counting those that matter. In *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement, ESEM 2018, Oulu, Finland, October 11-12, 2018*. ACM, 42:1–42:10. <https://doi.org/10.1145/3239235.3268920>

- Ivan Pashchenko, Duc Ly Vu, and Fabio Massacci. 2020. A Qualitative Study of Dependency Management and Its Security Implications. In *CCS '20: 2020 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, USA, November 9-13, 2020*. ACM, 1513–1531. <https://doi.org/10.1145/3372297.3417232>
- Donald Pinckney, Federico Cassano, Arjun Guha, and Jonathan Bell. 2023a. A Large Scale Analysis of Semantic Versioning in NPM. In *20th IEEE/ACM International Conference on Mining Software Repositories, MSR 2023, Melbourne, Australia, May 15-16, 2023*. IEEE, 485–497. <https://doi.org/10.1109/MSR59073.2023.00073>
- Donald Pinckney, Federico Cassano, Arjun Guha, and Jonathan Bell. 2023b. npm-follower: A Complete Dataset Tracking the NPM Ecosystem. In *Proceedings of the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2023, San Francisco, CA, USA, December 3-9, 2023*. ACM, 2132–2136. <https://doi.org/10.1145/3611643.3613094>
- Donald Pinckney, Federico Cassano, Arjun Guha, Jonathan Bell, Massimiliano Culpò, and Todd Gamblin. 2023c. Flexible and Optimal Dependency Management via Max-SMT. In *45th IEEE/ACM International Conference on Software Engineering, ICSE 2023, Melbourne, Australia, May 14-20, 2023*. IEEE, 1418–1429. <https://doi.org/10.1109/ICSE48619.2023.00124>
- Tom Preston-Werner. 2023. *Semantic versioning 2.0.0*. Retrieved December 11, 2023 from <https://semver.org/>
- Maciej Radzikowski. 2021. *Pin exact dependency versions*. Retrieved December 13, 2023 from <https://betterdev.blog/pin-exact-dependency-versions/>
- Steven Raemaekers, Arie van Deursen, and Joost Visser. 2017. Semantic versioning and impact of breaking changes in the Maven repository. *J. Syst. Softw.* 129 (2017), 140–158. <https://doi.org/10.1016/J.JSS.2016.04.008>
- Matt Raible. 2019. *Semantic Versioning Sucks! Long Live Semantic Versioning*. Retrieved May 23, 2024 from <https://developer.okta.com/blog/2019/12/16/semantic-versioning>
- Renovate Bot. 2018. *Should you pin your JavaScript dependencies?* Retrieved December 11, 2023 from <https://docs.renovatebot.com/dependency-pinning>
- Edgar H Schein. 2010. *Organizational culture and leadership*. Vol. 2. John Wiley & Sons.
- William Schueller and Johannes Wachs. 2024. Modeling interconnected social and technical risks in open source software ecosystems. *Collective Intelligence* 3, 1 (2024), 26339137241231912.
- Adriana Sejfia and Max Schäfer. 2022. Practical Automated Detection of Malicious npm Packages. In *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*. ACM, 1681–1692. <https://doi.org/10.1145/3510003.3510104>
- Arielle S Selya, Jennifer S Rose, Lisa C Dierker, Donald Hedeker, and Robin J Mermelstein. 2012. A practical guide to calculating Cohen’s f^2 , a measure of local effect size, from PROC MIXED. *Frontiers in Psychology* 3 (2012), 111.
- Hyunmin Seo, Caitlin Sadowski, Sebastian G. Elbaum, Edward Aftandilian, and Robert W. Bowdidge. 2014. Programmers’ build errors: A case study (at Google). In *36th International Conference on Software Engineering, ICSE '14, Hyderabad, India - May 31 - June 07, 2014*. ACM, 724–734. <https://doi.org/10.1145/2568225.2568255>
- Paul Slovic. 1987. Perception of risk. *Science* 236, 4799 (1987), 280–285.
- Snyk. 2022. *Alert: peacenotwar module sabotages npm developers in the node-ipc package to protest the invasion of Ukraine*. Retrieved August 22, 2023 from <https://snyk.io/blog/peacenotwar-malicious-npm-node-ipc-package-vulnerability/>
- Sonatype. 2021. *Software supply chains: An introductory guide*. Retrieved May 24, 2024 from <https://www.sonatype.com/blog/software-supply-chain-a-definition-and-introductory-guide>
- Sonatype. 2024a. *10th Annual State of the Software Supply Chain*. Retrieved Jan 16, 2025 from <https://www.sonatype.com/state-of-the-software-supply-chain/2024>
- Sonatype. 2024b. *CVE-2024-3094 The targeted backdoor supply chain attack against XZ and liblzma*. Retrieved April 17, 2024 from <https://blog.sonatype.com/cve-2024-3094-the-targeted-backdoor-supply-chain-attack-against-xz-and-liblzma>
- Diomidis Spinellis, Zoe Kotti, and Audris Mockus. 2020. A Dataset for GitHub Repository Deduplication. In *MSR '20: 17th International Conference on Mining Software Repositories, Seoul, Republic of Korea, 29-30 June, 2020*. ACM, 523–527. <https://doi.org/10.1145/3379597.3387496>
- StackOverflow. 2017. *Should I pin my Python dependencies versions?* Retrieved December 13, 2023 from <https://stackoverflow.com/questions/28509481/should-i-pin-my-python-dependencies-versions>
- Alois Stutzer and Bruno S Frey. 2006. Does marriage make people happy, or do happy people get married? *The Journal of Socio-Economics* 35, 2 (2006), 326–347. <https://doi.org/doi.org/10.1016/j.socrec.2005.11.043>
- TechTarget. 2023. *SolarWinds hack explained: Everything you need to know*. Retrieved May 22, 2024 from <https://www.techtarget.com/whatis/feature/SolarWinds-hack-explained-Everything-you-need-to-know>
- The White House. 2021. *Executive Order on Improving the Nation’s Cybersecurity*. Retrieved August 27, 2024 from <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>
- David A. Tomassi, Naji Dmeiri, Yichen Wang, Antara Bhowmick, Yen-Chuan Liu, Premkumar T. Devanbu, Bogdan Vasilescu, and Cindy Rubio-González. 2019. BugSwarm: Mining and continuously growing a dataset of reproducible failures and fixes. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, QC, Canada, May*

- 25-31, 2019. IEEE / ACM, 339–349. <https://doi.org/10.1109/ICSE.2019.00048>
- Marat Valiev, Bogdan Vasilescu, and James D. Herbsleb. 2018. Ecosystem-level determinants of sustained activity in open-source projects: A case study of the PyPI ecosystem. In *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*. ACM, 644–655. <https://doi.org/10.1145/3236024.3236062>
- Enrique Larios Vargas, Mauricio Finavaro Aniche, Christoph Treude, Magiel Bruntink, and Georgios Gousios. 2020. Selecting third-party libraries: the practitioners’ perspective. In *ESEC/FSE ’20: 28th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, Virtual Event, USA, November 8-13, 2020*. ACM, 245–256. <https://doi.org/10.1145/3368089.3409711>
- Duc-Ly Vu, Ivan Pashchenko, Fabio Massacci, Henrik Plate, and Antonino Sabetta. 2020. Typosquatting and Composquatting Attacks on the Python Ecosystem. In *IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2020, Genoa, Italy, September 7-11, 2020*. IEEE, 509–514. <https://doi.org/10.1109/EUROSPW51379.2020.00074>
- James Walden. 2020. The Impact of a Major Security Event on an Open Source Project: The Case of OpenSSL. In *MSR ’20: 17th International Conference on Mining Software Repositories, Seoul, Republic of Korea, 29-30 June, 2020*. ACM, 409–419. <https://doi.org/10.1145/3379597.3387465>
- Ying Wang, Peng Sun, Lin Pei, Yue Yu, Chang Xu, Shing-Chi Cheung, Hai Yu, and Zhiliang Zhu. 2023. Plumber: Boosting the Propagation of Vulnerability Fixes in the npm Ecosystem. *IEEE Trans. Software Eng.* 49, 5 (2023), 3155–3181. <https://doi.org/10.1109/TSE.2023.3243262>
- Ying Wang, Ming Wen, Yepang Liu, Yibo Wang, Zhenming Li, Chao Wang, Hai Yu, Shing-Chi Cheung, Chang Xu, and Zhiliang Zhu. 2020. Watchman: monitoring dependency conflicts for Python library ecosystem. In *ICSE ’20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*. ACM, 125–135. <https://doi.org/10.1145/3377811.3380426>
- Ying Wang, Ming Wen, Zhenwei Liu, Rongxin Wu, Rui Wang, Bo Yang, Hai Yu, Zhiliang Zhu, and Shing-Chi Cheung. 2018. Do the dependency conflicts in my project matter?. In *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*. ACM, 319–330. <https://doi.org/10.1145/3236024.3236056>
- Hyrum Wright. 2024. *Hyrum’s Law*. Retrieved May 23, 2024 from <https://www.hyrumslaw.com/>
- Jiadong Xie, Fan Zhang, Kai Wang, Xuemin Lin, and Wenjie Zhang. 2023. Minimizing the Influence of Misinformation via Vertex Blocking. In *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*. IEEE, 789–801. <https://doi.org/10.1109/ICDE55515.2023.00066>
- Weiwei Xu, Hao He, Kai Gao, and Minghui Zhou. 2023. Understanding and Remediating Open-Source License Incompatibilities in the PyPI Ecosystem. In *38th IEEE/ACM International Conference on Automated Software Engineering, ASE 2023, Luxembourg, September 11-15, 2023*. IEEE, 178–190. <https://doi.org/10.1109/ASE56229.2023.00175>
- Awad A. Younis, Yashwant K. Malaiya, Charles Anderson, and Indrajit Ray. 2016. To Fear or Not to Fear That is the Question: Code Characteristics of a Vulnerable Function with an Existing Exploit. In *Proceedings of the Sixth ACM on Conference on Data and Application Security and Privacy, CODASPY 2016, New Orleans, LA, USA, March 9-11, 2016*. ACM, 97–104. <https://doi.org/10.1145/2857705.2857750>
- Nusrat Zahan, Philipp Burckhardt, Mikola Lysenko, Feross Aboukhadijeh, and Laurie A. Williams. 2024. Shifting the Lens: Detecting Malware in npm Ecosystem with Large Language Models. *CoRR abs/2403.12196* (2024). <https://doi.org/10.48550/ARXIV.2403.12196> arXiv:2403.12196
- Nusrat Zahan, Parth Kanakia, Brian Hambleton, Shohanuzzaman Shohan, and Laurie A. Williams. 2023a. OpenSSF Scorecard: On the Path Toward Ecosystem-Wide Automated Security Metrics. *IEEE Secur. Priv.* 21, 6 (2023), 76–88. <https://doi.org/10.1109/MSEC.2023.3279773>
- Nusrat Zahan, Shohanuzzaman Shohan, Dan Harris, and Laurie A. Williams. 2023b. Do Software Security Practices Yield Fewer Vulnerabilities?. In *45th IEEE/ACM International Conference on Software Engineering: Software Engineering in Practice, SEIP@ICSE 2023, Melbourne, Australia, May 14-20, 2023*. IEEE, 292–303. <https://doi.org/10.1109/ICSE-SEIP58684.2023.00032>
- Lyuye Zhang, Chengwei Liu, Zhengzi Xu, Sen Chen, Lingling Fan, Bihuan Chen, and Yang Liu. 2022. Has My Release Disobeyed Semantic Versioning? Static Detection Based on Semantic Differencing. In *37th IEEE/ACM International Conference on Automated Software Engineering, ASE 2022, Rochester, MI, USA, October 10-14, 2022*. ACM, 51:1–51:12. <https://doi.org/10.1145/3551349.3556956>
- Markus Zimmermann, Cristian-Alexandru Staicu, Cam Tenny, and Michael Pradel. 2019. Small World with High Risks: A Study of Security Threats in the npm Ecosystem. In *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*. USENIX Association, 995–1010. <https://www.usenix.org/conference/usenixsecurity19/presentation/zimmerman>