2026 IEEE International Conference on Software Testing, Verification and Validation (ICST) | 979-8-3195-3308-1/26/$31.00 ©2026 IEEE | DOI: 10.1109/ICST69053.2026.00085

On the Freshness of Pinned Dependencies in Maven

Vasudev Vikram, Yuvraj Agarwal, Rohan Padhye
Carnegie Mellon University
Pittsburgh, PA, USA
vasumv@cmu.edu, yuvraj@cs.cmu.edu, rohanpadhye@cmu.edu

***Abstract*—Library dependencies in software ecosystems play a crucial role in the development of software. As newer releases of these libraries are published, developers may opt to *pin* their dependencies to a particular version. While pinning may have benefits in ensuring reproducible builds and avoiding breaking changes, it bears larger risks in using outdated dependencies that may contain bugs and security vulnerabilities. To understand the frequency and consequences of dependency pinning, we first define the concepts of *stale* and *fresh* pins, which are distinguished based on how outdated the dependency is relative to the release date of the project. We conduct an empirical study to show that over 60% of consumers of popular Maven libraries contain stale pins to their dependencies, with some outdated versions over a year old. These pinned versions often miss out on security fixes; we find that 10% of all dependency upgrades in our dataset to the latest minor or patch version would reduce security vulnerabilities.**

We prototype an approach called Pin-Freshener that can encourage developers to freshen their pins by leveraging the insight that *crowdsourced* tests of peer projects can provide additional signal for the safety of an upgrade. Running Pin-Freshener on dependency upgrades shows that just 1–5 additional test suites can provide 35–100% more coverage of a dependency, compared to that of a single consumer test suite. Our evaluation on real-world pins to the top 500 popular libraries in Maven shows that Pin-Freshener can provide an additional signal of at least 5 passing crowdsourced test suites to over 3,000 consumers to safely perform an upgrade that reduces security vulnerabilities. Pin-Freshener can provide practical confidence to developers by offering additional signal beyond their own test suites, representing an improvement over current practices.

I. Introduction

Modern software heavily relies on third-party libraries. Usage of these libraries can reduce software development time and cost by reusing existing functionality of software [1], [2]. This process has been integrated into many software ecosystems—such as Apache Maven for Java, NPM for JavaScript, and PIP for Python—for which building and installing library dependencies is a natural step for software developers. The Maven Central Repository demonstrates the popularity of this practice for Java applications, with an index containing millions of Java packages [3]. An example of the dependency network of the Maven `gemini` library is shown in Figure 1, showing many dependencies than can span multiple edges.

While the dependence on third-party libraries assists the development of new software applications, managing these dependencies can be challenging. New releases of dependencies are constantly published to the ecosystem and developers must decide whether to upgrade them to a newer version. However,

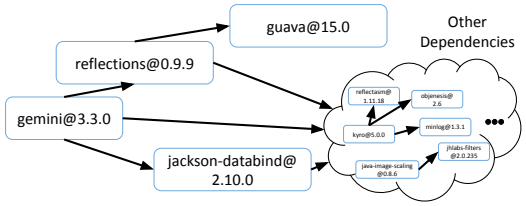

Fig. 1: Dependency tree of `gemini@3.3.0`. Arrows denote dependencies. `gemini@3.3.0` has a direct dependency on `jackson-databind@2.10.0` and an indirect dependency on `guava@15.0`.

software bugs or unexpected behavior—referred to as breaking changes—can be introduced in these new versions [4], [5], [6]. Third-party library maintainers sometimes even *knowingly* deploy breaking changes due to the build up of technical debt and pressure to release new functionality [7].

Thus, upgrading a dependency can always be risky for consumers of these libraries. They may be wary of the possibility that their project might break or even that new security vulnerabilities are introduced [8]. Developers are often faced with the choice of *floating* their dependencies to a range of versions (e.g. `2.1+`), or *pinning* their dependencies to a single specific version (e.g. `2.1`). Software ecosystems also adopt different conventions for specifying dependency versions; the default convention in NPM is to specify a floating version, whereas the default in Apache Maven is to specify a fixed pinned version. In the case of Maven, this is largely due to the packages not necessarily following the semantic versioning practices when new versions are released [4].

Thus, pinning has certain benefits such as providing reproducible builds [9] and avoiding breaking changes, but it can also bear a significant cost! New library versions often include new features, performance improvements, and crucial security patches. The high-profile 2017 Equifax data breach, in which a vulnerability in the open source Apache Struts library was exploited for leaking sensitive data of over 140 million consumers, demonstrates this drawback of pinning [10]. A patch for Apache Struts was available, but was not adopted by Equifax for over *two months*. Nowadays, tools like *Dependabot* and others [11], [12], [13], [14] help warn developers about known security vulnerabilities in outdated dependencies, though this approach is reactive rather than proactive.

So, we ask: is dependency pinning actually worth it? We first distinguish between two types of pins: *stale* pins and

979-8-3195-3308-1/26/$31.00 ©2026 IEEE
DOI 10.1109/ICST69053.2026.00085

fresh pins. A *stale* pin captures an explicit instance in which a project, at the time of release, was pinned to an outdated dependency even though a newer one was available. On the other hand, a *fresh* pin is one in which a project is pinned to the latest version of the dependency available. Using these definitions, we conduct an empirical study on the Maven ecosystem to understand the how common the practice is and its broader security implications. We use the Open Source Insights dataset [15], published by Google, containing data about dependencies, consumers, and security vulnerabilities for over 569,000 Maven packages. We construct a dataset from a targeted sample of the most popular Maven libraries from the Open Source Insights dataset and find that *over 60%* of the consumers of these libraries contain stale pins to outdated versions.

Given that stale dependency pinning is a fairly common practice in Maven, we next explore the security risks it poses. In a retrospective analysis using today's vulnerability databases, we find that 10% of library upgrades could have fixed security vulnerabilities had they adopted *fresh pinning* for their dependencies when publishing their library. In contrast, only 3% of the upgrades would have introduced new vulnerabilities. This corresponds to over 22,000 consumers in our dataset that potentially could have fixed vulnerabilities (a majority of which having high or critical severity levels) had they performed upgrades to fresh pins. Hence, we believe that developers should follow a more proactive pinning strategy since they are far likelier to fix vulnerabilities by upgrading their outdated dependencies. We acknowledge that pinning has benefits such as maintaining reproducibility of builds and avoiding unexpected breakage; however, the prevalence of stale pins and their security implications suggests developers should adopt a fresh pinning strategy.

While the overall security benefit of upgrading stale pins is clear, we must still consider the aspect of evaluating whether performing a specific upgrade is safe. In Maven, breaking changes even across minor version updates are fairly common [4]. So, how can we provide more confidence for consumers upgrade their stale pins to fresh pins? Our key insight is that the test suites of other consumers in the ecosystem can help provide additional signal about the upgrade and more confidence to the developer. To this end, we develop a prototype of *Pin-Freshener*, an approach that *crowdsources* test suites of peer consumers of the dependency to evaluate the safety of an upgrade to a fresh pin. We specifically leverage the existence of *test-JARs* in the Maven ecosystem, which contain projects' compiled tests, in order to streamline the execution of consumer test suites. By executing these additional test suites against both the pinned version and upgraded version, we can characterize the impact of the upgrade on multiple projects. Pin-Freshener reports a *confidence score* of a particular upgrade determined by the number of consumer test suites that are able to successfully run when using the upgraded dependency version.

In an experiment of executing Pin-Freshener on our dataset of these upgrades, we first find that crowdsourcing just five consumer test suites is able to provide an average improvement of almost 100% in terms of test coverage of a dependency over that of a single consumer. Pin-Freshener is able to provide an additional signal of at least five passing crowdsourced test suites to over *3,000* consumers (15%) performing upgrades that would fix security vulnerabilities. Although this signal does not provide *guarantees* of upgrade safety, it can provide developers with additional practical confidence (similar to code coverage metrics) beyond what they would get from relying solely on their own tests–the current standard practice for judging whether to perform an upgrade.

In summary, we ask the following research questions:

RQ1: To what extent are libraries in the Maven ecosystem pinning to stale dependencies?

RQ2: What is the security impact of pinning to stale dependencies?

RQ3: How much can crowdsourced test suites improve coverage of the pinned dependency?

RQ4: Can crowdsourced test suites help provide additional signal for vulnerability-fixing upgrades from stale pins to fresh pins?

Our contributions are as follows:

1) We formalize the concepts of *stale* and *fresh* pins, which depend on whether a project is pinned to an outdated dependency. Using these definitions, we conduct an empirical study on the Apache Maven ecosystem using the Open Source Insights dataset to measure the frequency and retrospective security impact of dependency pinning for the top 500 most-popular libraries.
2) We design and implement a prototype of *Pin-Freshener*, an approach that crowdsources consumer test suites to better characterize the safety of an upgrade across the network and provide additional signal to developers when freshening pinned dependencies.
3) We run Pin-Freshener at a large scale on vulnerability-fixing upgrades in Maven libraries and find that Pin-Freshener is able to provide additional signal of at least 5 passing crowdsourced test suites to over 3,000 consumers that can perform an upgrade that reduces security vulnerabilities from a stale pin to a fresh pin.

II. Background and Terminology

A. Dependencies in Software Ecosystems

A software ecosystem is a collection of software libraries, each denoted by a name and a version number. We denote a library as $L@V$, where L refers to the library name and V refers to version. We define $\mathbb{L}$ as the set of all libraries in a particular software ecosystem, such as Maven for Java.

A library $L@V$ may contain a *direct dependency* to another library $L'@V'$, specified in a build system configuration file. We refer to a dependency as the specific package pulled by the build system after dependency resolution (resolving wildcards/ranges to a single version). We refer to $L'@V'$ as a *direct dependency* and $L@V$ as a *direct consumer*. A

shorthand notation for describing this direct dependency relation is $L@V \rightarrow L'@V'$. An example of a direct dependency relation can be seen in Figure 1 between `gemini@3.3.0` and `jackson-databind@2.10.0`. We define the entire dependency graph $\mathbb{G}$ as the set of all direct dependency relations (edges), and naturally define the functions *directDeps* and *directConsumers* to identify a direct dependency on D or a direct consumer C respectively as follows:

$$\begin{aligned} \textit{directDeps}(L@V) = \{&D@V' \in \mathbb{L} \mid \\ &(L@V \rightarrow D@V') \in \mathbb{G}\} \\ \textit{directConsumers}(L@V) = \{&C@V'' \in \mathbb{L} \mid \\ &(C@V'' \rightarrow L@V) \in \mathbb{G}\} \end{aligned}$$

A library dependency can also span multiple dependency edges, such as between `gemini@3.3.0` and `guava@15.0` in Figure 1. To account for these dependency relations, we define the function *allDeps* on $L@V$ to return the transitive closure of *directDeps* applied to $L@V$. We similarly define *allConsumers* as the transitive closure of *directConsumers*. These functions return the set of all dependencies and consumers of $L@V$, respectively, regardless of the number of edges. We additionally introduce the functions *indirectDeps* and *indirectConsumers* to return the sets of dependencies and consumers that are not direct.

A library has the option of *upgrading* a dependency from one version to a newer one. Continuing our example from Figure 1, the library `gemini@3.3.0` could upgrade `jackson-databind` from version `2.10.0` to `2.11.0`. We denote an *upgrade* as the pair $\langle D@V^{\alpha}, D@V^{\beta} \rangle$, where V^{α} and V^{β} are the older and newer versions of D, respectively.

B. Semantic Versioning

When performing a dependency upgrade, it's crucial for consumers to understand the types of changes being introduced in a new dependency version and whether it is backwards compatible. One practice used in many software ecosystems is *semantic versioning* [16], which defines a set of rules for assigning version numbers to new releases of libraries. When using semantic versioning, a version V is structured as `major.minor.patch[-tag]`. For example, `jackson-databind@2.10.0` has major=2, minor=10, patch=0. We define functions *major*, *minor*, and *patch* to return these components of V. This separation of version numbers also defines a total ordering between versions that compares major, minor, and patch versions numerically from left to right. We use this comparison logic throughout the paper when ordering versions (e.g. $V^{\beta} > V^{\alpha}$).

Semantic versioning is used to characterize the types of version upgrades in terms of backwards compatibility. Generally, version upgrades that include backwards *incompatible* changes increment the major version, whereas upgrades that do not break existing functionality are limited to minor or patch version increments. This allows library developers to notify consumers about the specific versions that introduce potential breaking changes, and consumers can choose which versions to adopt through a set of dependency constraints. Throughout this paper, we refer to minor and patch version upgrades as *semver-compatible*, as they should have the assurance of being backwards compatible.

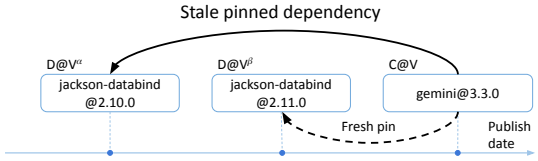

Fig. 2: Example of a direct stale and fresh pin. Since `gemini@3.3.0` is a direct consumer of `jackson-databind@2.10.0` and the latest version of the library `2.11.0` was published before the consumer, this is a stale pin. If `gemini@3.3.0` was pinned to `jackson-databind@2.11.0`, this would be a *fresh* pin, as it is the latest version available at the time.

Semantic versioning encourages consumers to perform semver-compatible upgrades on their dependencies since there should be no risk of introducing breaking changes. This can be as simple as specifying a version range for a dependency that freezes the major version, such as `[1.0.0, 2.0.0)`. However, semantic versioning is only a policy and is unenforceable throughout a software community; oftentimes new minor and patch versions may not respect the policy, resulting in unexpected breaking changes and upset consumers [17], [18]. These upgrades can even introduce accidental bugs or new security vulnerabilities, which may convince consumers to avoid semver-compatible upgrades entirely and decide to *pin* their dependencies to a single version. Package ecosystems also adopt different default approaches to versioning: NPM favors floating dependencies that follow semantic versioning conventions, automatically accepting compatible updates. Maven, on the other hand, defaults to pinning dependencies to specific versions, reflecting its ecosystem's less consistent adherence to semantic versioning principles. This is likely due to many libraries in Maven historically publishing breaking changes that disobey semantic versioning [4].

C. Dependency Pinning

The practice of specifying a single version of a dependency rather than a range is referred to as *dependency pinning*. Pinning to a specific version can provide benefits in ensuring reproducible builds and avoiding accidental breakage due to new breaking changes in an upgrade. However, the developer is then responsible for updating their pinned dependencies to the latest versions before releasing their own software. The failure to do so results in a *stale* pin, which is an explicit instance in which software is pinned to an outdated dependency version while a newer version exists. Figure 2 shows a stale pin: `gemini@3.3.0` depends on `jackson-databind@2.10.0` despite `2.11.0` being available at publication time. Depending on `2.11.0` would be a *fresh* pin.

We formally define a *stale pin* as follows: given a dependency graph $\mathbb{G}$, a stale pin is the tuple of three libraries $\langle C@V,\ D@V^{\alpha},\ D@V^{\beta}\rangle \in \mathbb{L}\times\mathbb{L}\times\mathbb{L}$ for which the following conditions hold:

1) $D@V^{\alpha} \in \mathit{allDeps}(C@V)$
2) $\mathit{publish}(V^{\alpha}) < \mathit{publish}(V^{\beta}) < \mathit{publish}(V)$
3) $(\mathit{major}(V^{\beta}) = \mathit{major}(V^{\alpha})) \wedge (V^{\beta} > V^{\alpha})$

The first condition specifies that $D@V^{\alpha}$ is a dependency of consumer $C@V$. The second condition compares publication times: if the newer dependency version V^{β} was published before the consumer version V, then $C@V$ contains a stale pin to $D@V^{\alpha}$. The final condition ensures the upgrade is semver-compatible by checking major version equality, filtering out potentially breaking major version changes.

We then define a *fresh pin* as follows: a fresh pin is the tuple of two libraries $\langle C@V, D@V^{\gamma}\rangle$ for which the following conditions hold:

1) $D@V^{\gamma} \in \mathit{allDeps}(C@V)$
2) $\mathit{publish}(V^{\gamma}) < \mathit{publish}(V)$
3) $\nexists\, D@V^{\delta} \in \mathbb{L}$ such that:
   a) $(\mathit{major}(V^{\delta}) = \mathit{major}(V^{\gamma})) \wedge (V^{\delta} > V^{\gamma})$
   b) $\mathit{publish}(V^{\gamma}) < \mathit{publish}(V^{\delta}) < \mathit{publish}(V)$

We can further classify a pin as either *direct* or *indirect* depending on the nature of the dependency between $C@V$ and $D@V^{\alpha}$. $\langle C@V,\ D@V^{\alpha},\ D@V^{\beta}\rangle$ is direct if $D@V^{\alpha} \in \mathit{directDeps}(C@V)$ and indirect if $D@V^{\alpha} \in \mathit{indirectDeps}(C@V)$. To freshen a direct pin, a consumer would update the dependency version in the configuration file. Freshening indirect pins requires the consumer to explicitly override by introducing a direct dependency to $D@V^{\beta}$.

Freshening a pinned dependency would involve performing the upgrade from V^{α} (pinned version) to V^{β} (upgrade version) and is not necessarily a straightforward decision. Consumers may be apprehensive of incorporating changes that break their project or even introduce new security vulnerabilities. However, keeping the dependencies pinned has a risk of missing out on crucial patches for vulnerabilities that exist in the pinned version, usually fixed in minor and patch version upgrades. Without a way of characterizing the impact of these upgrades beyond semantic versioning guidelines, developers must make a difficult decision when deciding to perform these dependency upgrades.

III. Pinning in Maven

Using the Open Source Insights dataset published by Google [15], we conducted an analysis on a snapshot of the Maven ecosystem to measure the frequency and impact of pinning. We take a snapshot of the entire Maven dependency network on May 22, 2023 that includes dependencies and consumers (both direct and indirect) of ∼567,000 Maven libraries. This snapshot contains 235,959,564 dependency edges, of which 45,997,607 (19.5%) are direct dependencies. We analyze an older snapshot to provide time (roughly 1.5 years) for CVEs to be discovered and added to the vulnerability database.

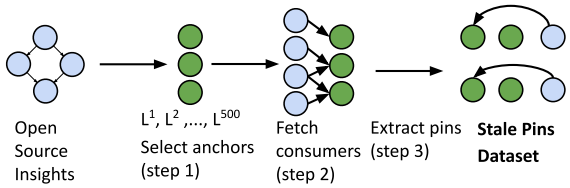

Fig. 3: Construction of stale pin dataset $\mathcal{D}$. A set of anchors are selected from Open Source Insights by popularity (consumer count). $\mathcal{D}$ contains pins from consumers to anchors.

We use Open Source Insights because it provides resolved dependency versions from Maven's resolution process rather than POM file syntax[1], handling version ranges (e.g., `2.0+`, `LATEST`) and conflict resolution for indirect dependencies. This enables finding *explicit* instances of stale pinning in the ecosystem.

A. RQ1: Frequency of Stale Pinning

In RQ1, we measure dependency pinning frequency in Maven. Since the entire ecosystem is too large to analyze completely, we target the top 500 most popular libraries (by consumer count) due to their network impact, analyzing stale pins from consumers to these libraries.

We construct a dataset of stale pins (as defined in Section II-C) using the process shown in Figure 3. The dataset uses the top 500 most popular libraries (referred to as *anchors*) as a starting point to find pins across the network. The anchors are created by selecting the library names (e.g., $L^1, L^2, \ldots, L^{500}$) with the highest number of consumers across all versions, as seen in Step 1 of Figure 3.

We can walk through an example of extracting a stale pin from a consumer to an anchor. We refer back to Figure 2 with the dependency from `gemini@3.3.0` to `jackson-databind@2.10.0`. As `jackson-databind` is one of our anchors, we would like to extract stale pins from consumers to its outdated versions. We begin by querying Open Source Insights to find all the consumers of `jackson-databind`, across all versions of the library. One such consumer is `gemini`—although there are many versions of this library, we select latest minor version (`3.3.0`) to find an up-to-date version. Since there are multiple versions of `jackson-databind` higher than version `2.10.0` published earlier than `gemini@3.3.0`, we select the highest one (`2.11.0`) and add the stale pin ⟨`gemini@3.3.0`, `jackson-databind@2.10.0`, `jackson-databind@2.11.0`⟩ to dataset $\mathcal{D}$.

The process of creating the entire dataset is formally described as follows: we first query the Open Source Insights network to find all consumers (step 2 in Figure 3) of the

[1]We considered Libraries.io [19], which stores declared POM syntax, but chose Open Source Insights for its resolved versions and recency.

TABLE I: Pinning statistics for consumers of the top 500 libraries. 148,811 (61%) contain ≥1 direct stale pin; 184,281 (83%) contain ≥1 indirect stale pin. Many consumers share the same stale pins: only 46,365 unique semver-compatible upgrades exist for direct pins and 76,317 for indirect pins.

	Consumers	**Consumers with ≥ 1 stale pin**	**Deps.**	**Upgrades**
Direct	244,819	148,811	717,705	46,365
Indirect	221,744	184,281	2,778,165	76,317

anchors libraries across all versions of each anchor, i.e.

$$anchorConsumers = \bigcup_{\substack{L^i@V^j \in \mathbb{L} \wedge \\ L^i \in anchors}} allConsumers(L^i@V^j)\}$$

We then query Open Source Insights to select the latest minor version of each consumer in *anchorConsumers*. For each consumer $C@V$, we find all of its dependencies to the anchor libraries and check whether any of them are are a stale pin. Given a dependency to an anchor $L^i@V^\alpha$, we select the highest version V^β that was published before $C@V$ and add the corresponding stale pin to $\mathcal{D}$ (step 3 of Figure 3).

Table I shows the statistics of the number of consumers, dependencies, and upgrades in $\mathcal{D}$. Interestingly, we find that *more than 60%* the direct consumers of the anchors contain at least 1 direct stale pin, and *over 80%* of the indirect consumers contain at least one indirect stale pin. Furthermore, we can see from Figure 4 that the dependency versions are outdated by a median of 370 days (7 versions) and 427 days (9 versions) for direct and indirect stale pins respectively. We see that pinning to the top 500 libraries is extremely common and features fairly outdated pinned versions! Note that there are a significantly smaller number of potential *upgrades* in $\mathcal{D}$ (as defined in Section II-A) than there are pinning consumers, suggesting that many consumers share the same stale pins to the anchors.

Finding #1: Stale pinning to the most popular libraries in Maven is a very common practice, with over 60% of consumers containing at least one direct stale pin, and 80% containing at least one indirect stale pin. The pinned versions of these libraries are fairly outdated, about 7 versions behind the upgrade version for direct pins and 9 versions for indirect pins.

B. RQ2: Security Impact of Stale Pinning

Older versions of libraries frequently contain known vulnerabilities that are patched in newer minor and patch releases. These security issues are tracked and disclosed publicly using Common Vulnerabilities and Exposures (CVEs) and other reporting mechanisms. The public Open Source Vulnerabilities (OSV) [20] database maintained by Google is a central database for CVEs and is used as a data source for the Open Source Insights dataset, which stored metadata about each

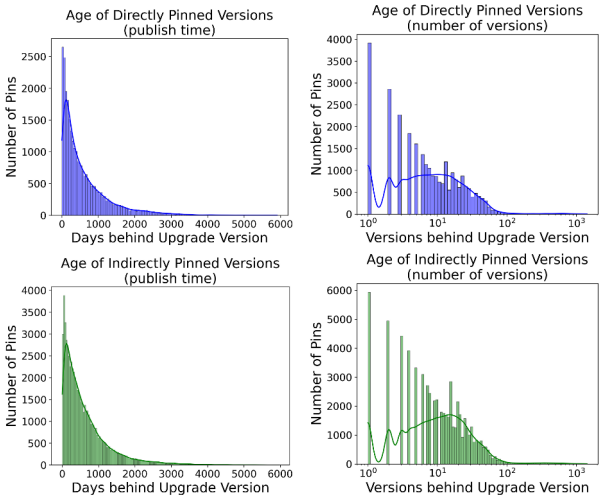

Fig. 4: Age of direct (dark blue) and indirect (light green) stale pins in $\mathcal{D}$. X-axis shows age, Y-axis shows number of stale pins (log scale for version plots).

vulnerability as an *advisory*. Each security advisory includes information about the packages and specific versions affected by the vulnerability.

From RQ1, we see that a very large percentage of consumers depend on an outdated version of the most popular libraries in the Maven ecosystem. While this provides a view of how frequent dependency pinning occurs in the Maven ecosystem, we are interested in measuring the security impact of these pins: specifically, are developers avoiding introducing new security vulnerabilities into their dependencies by pinning, or they missing out on important security patches? Tools such as *dependabot* utilize these vulnerability databases to notify developers of vulnerable dependencies in the latest development snapshots; however, this data has not been used to identify the historical security impact of pinned dependencies in Maven.

To perform this analysis, we use today's vulnerability databases to *retrospectively* compare the number of known security vulnerabilities affecting the pinned version and upgrade version of the direct stale pins in dataset $\mathcal{D}$. We note that this is a retrospective analysis: many of these vulnerabilities were not yet publicly disclosed at the time the consumers made their pinning decisions, so these results represent the *exposure* that stale pinning created in hindsight rather than risk that developers could have acted on at decision time. Of the 46,365 potential upgrades (Table I), we find that 40,462 (87.2%) result in no difference in known vulnerabilities, 4,576 (9.9%) upgrades would have reduced the number of security vulnerabilities, and 1,327 (2.9%) would have introduced new ones. Thus, in retrospect, performing a semver-compatible upgrade of a pinned dependency in $\mathcal{D}$ would have been *3.45*× as likely to reduce vulnerability exposure than increase it. Figure 5 displays the histogram of the differences in vulnerabilities between the versions, excluding the upgrades

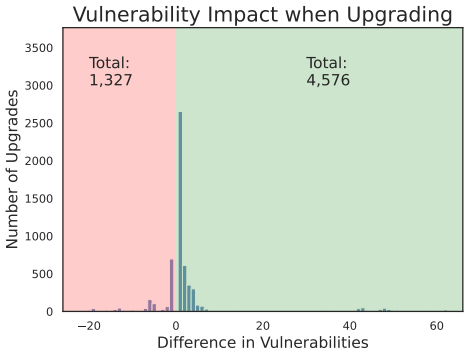

Fig. 5: Security impact of upgrading direct pins in $\mathcal{D}$. Green values reduce vulnerabilities; red values increase vulnerabilities (zero excluded). There are 4,576 upgrades that decrease vulnerabilities and 1,327 that increase vulnerabilities. Across all upgrades, vulnerabilities reduce by 20,825.

having no security impact for the sake of visualization. The majority of upgrades reduce the number of known security vulnerabilities by 1, but certain upgrades can address up to as many as 66 vulnerabilities. Across all of these upgrades, the total number of vulnerabilities would have been reduced by 20,825.

Finding #2: In a retrospective analysis using today's vulnerability databases, 10% of semver-compatible upgrades on a pinned version of a popular library would have reduced known security vulnerabilities, and only 3% would have introduced new ones. Upgrading would have been 3.45× as likely to reduce vulnerability exposure than increase it.

IV. An Approach Towards Freshening Pins

In answering RQ1 and RQ2, we have identified that stale dependency pinning to the most popular libraries in Maven is fairly common and has high security risks. However, developers of these libraries may be cautious to perform these upgrades. To upgrade a stale pin to a fresh pin, a consumer needs to be confident that the changes in the dependency upgrade are safe to introduce. One method would be to execute their own test suites against the new version of the dependency. However, even if the tests pass, they may not be comprehensive enough to thoroughly test behaviors of the new dependency version. We address this concern by developing a prototype called *Pin-Freshener* that calculates a confidence score of a given upgrade by executing *crowdsourced* test suites of other consumers of the pinned dependency and measuring their outcomes. Our insight is that consumer test suites can exercise a more thorough set of behaviors of the dependency; if multiple consumers' tests pass on both the stale pin and fresh pin, Pin-Freshener can provide additional signal for a developer to more confidently upgrade their dependency.

Pin-Freshener takes an upgrade $(D@V^\alpha, D@V^\beta)$ as input and crowdsources test suites to calculate a confidence score

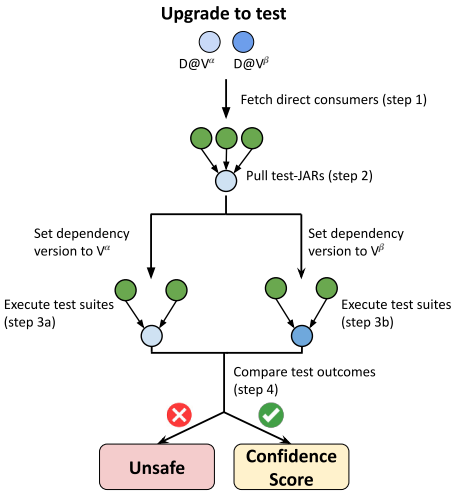

Fig. 6: Overview of Pin-Freshener. Direct consumers of $D@V^\alpha$ are fetched from Open Source Insights and their test suites executed using test-JARs. The version is then set to V^β and tests are re-executed. Test outcomes are compared: if safe, confidence equals the number of test suites executed; if unsafe, confidence is zero.

for the upgrade. Our approach follows the procedure outlined in Figure 6:

1) Query Open Source Insights to find *directConsumers*$(D@V^\alpha)$.
2) Pull the consumer test-JARs from the Maven Central Repository for each $C@V \in$ *directConsumers*. Note that not all consumers have published test-JARs; thus, we construct a set *testableConsumers* $= \{C@V \in$ *directConsumers*$(D@V^\alpha) \mid$ *testJarExists*$(C@V)\}$
3) For each consumer $C@V \in$ *testableConsumers*, execute the tests when using $D@V^\alpha$ (3a) and $D@V^\beta$ (3b) as dependencies (see Section IV-A).
4) Compare the test outcomes for each version and calculate a confidence for the upgrade $\langle D@V^\alpha, D@V^\beta \rangle$ (see Section IV-C).

Steps (1) and (2) query the Open Source Insights dataset and the Maven Central Repository respectively to fetch test-JARs of the direct consumers of the pinned version. In the following sections, we go into detail to describe Steps (3) and (4).

A. Executing Consumer Test Suites

One option to execute a consumer test suite is to download and build the source code of the repository and invoke the tests by running `mvn test`. Unfortunately, the source code for these consumers may not be publicly available. Additionally, resolving the specific version V in the repository can be a nontrivial task, as version naming conventions may differ between the source code and the Maven package.

The strategy we chose was to leverage the Maven Central Repository for *test-JARs* of the consumer, which contains

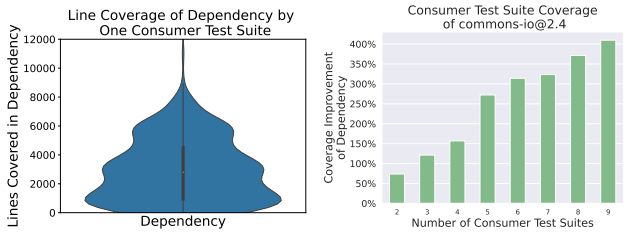

Fig. 7: Left: distribution of lines covered for a dependency from a *single* consumer test suite. Right: line coverage *improvement* of `commons-io@2.4` from additional consumer test suites (measured via random sampling).

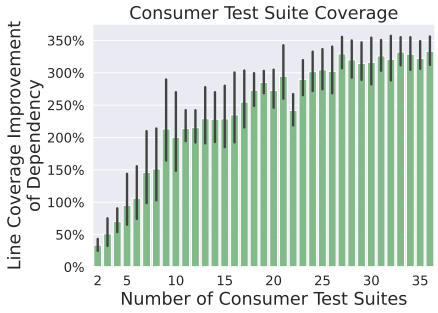

Fig. 8: Average coverage improvement by Pin-Freshener over a single consumer test suite (higher is better). X-axis: number of crowdsourced test suites; Y-axis: geometric mean line-coverage improvement across all libraries. Four additional test suites achieve nearly 100% more coverage than a single suite.

compiled classes of the test files. Test-JARs are unique to the Maven ecosystem and provide a streamlined approach of fetching and executing project test suites. While test-JARs are optional to upload to the Maven Central Repository and do not exist for certain consumers, this approach still provides a straightforward method of crowdsourcing test suites. Among the consumers in $\mathcal{D}$ with direct stale pins, we found that about 12% of projects had uploaded test-JARs to the Maven Central Repository; while we would have liked this percentage to be higher, this is still a significant number of tests available for Pin-Freshener to use to test upgrades.

To walk through this process, we refer to our original example of a pinned dependency from `gemini@3.3.0` to `jackson-databind@2.10.0`. The consumer `gemini@3.3.0` would use Pin-Freshener to test the upgrade of `jackson-databind` from `2.10.0` to `2.11.0`. Pin-Freshener first finds all consumers of the pinned dependency `jackson-databind@2.10.0` and pulls all consumer test-JARs that are available on the Maven Central Repository. In the case that there are multiple consumers with the same library name, we select the highest version.

For each of the consumers, Pin-Freshener first executes each of the test suites against the pinned dependency version of `jackson-databind` (`2.10.0`). Some tests may produce non-deterministic outcomes due to *flakiness* [21], [22]. Pin-Freshener executes each test with $r = 5$ repetitions to account for this flakiness. Since the tests are executed directly from the test-JARs, it also is possible that tests may have errors or fail due to missing resources. We save the test outcomes produced by Maven of each of the consumer tests to a database.

Next, Pin-Freshener upgrades the dependency version of `jackson-databind` to `2.11.0` for each of the consumer test suites. Once again, the execution of the test suites are repeated five times, and the test outcomes are saved.

B. RQ3: Coverage Improvement of Crowdsourced Consumer Test Suites

A natural question is whether using additional consumer test suites has any advantages in terms of exercising code, such as improved coverage, of the dependency. To characterize the coverage benefit of crowdsourced consumer test suites, we use the Jacoco library [23] to collect the coverage of the dependency classes only. Figure 7 (left) shows the coverage distribution of dependencies by a single consumer test suite (on average). We observe that consumer test suites cover a substantial amount of dependency code, with some even covering up to 10,000 lines. Figure 7 (right) shows the coverage *improvement* from executing additional consumer test suites on one pinned dependency `commons-io@2.4`—Pin-Freshener finds nine consumers of this dependency whose test JARs could be executed. From the figure, we see that the union line coverage of these nine test suites provided over a 400% increase in coverage of `commons-io@2.4` than if we only executed a single consumer's test suite (on average). To understand how coverage increases with the number of crowdsourced test suites, we calculate the union of the dependency coverage for each value n below 9 by randomly sampling a subset of n consumer test suites without replacement (up to 50 times) and calculating the average.

Generalizing this methodology, Figure 8 shows the average coverage improvement, across all popular libraries, with respect to the number of consumer test suites. With just a single additional consumer test suite, we can achieve an average of 40% additional coverage of the dependency; with four additional test suites, this number rises to almost 100%! Overall, we find that the crowdsourced test suites from Pin-Freshener are able to achieve a significant coverage boost of the pinned dependency over a single consumer's test suite alone.

C. Computing Confidence Score

We next explain how Pin-Freshener uses the outcomes from the consumer test suites to test an upgrade. Based on the results of the crowdsourced test suites, Pin-Freshener calculates a *confidence* score for each upgrade. We walk through our example of upgrading `jackson-databind` from version `2.10.0` to `2.11.0`. There are seven available consumer test-JARs, which Pin-Freshener executes on the pinned version `2.10.0` and upgrade version `2.11.0`. Tests that are flaky or fail in the pinned version are filtered out, and all remaining test

outcomes are compared between versions. Each of the seven consumers *vote* on whether the upgrade is safe or unsafe. If all tests in the consumer test suite pass on both dependency versions, then that consumer votes *safe*; otherwise, there exists a test that passes in the pinned version but fails in the upgrade version, indicating the presence of a breaking change. Since all seven consumers vote safe, the confidence returned by Pin-Freshener is seven, indicating the presence of seven passing consumer test suites. Had there been a consumer that voted unsafe, then the confidence returned by Pin-Freshener would be zero.

More formally, we determine confidence as follows. We define *outcome* as a function that takes in a test method t, a consumer $C@V$, and a dependency $D@V$. From Section IV-A, each test has been executed with r repetitions.

$$\mathit{outcome}(t, C@V, D@V) = \begin{cases} \mathit{pass} & \text{if } r \text{ repetitions pass} \\ \mathit{fail} & \text{if } r \text{ repetitions fail} \\ \mathit{flaky} & \text{otherwise} \end{cases}$$

Each consumer provides a *vote* for whether the upgrade is safe or unsafe depending on the results of its test suite. If all passing tests with dependency version V^{α} also pass when the dependency version is upgraded to V^{β}, then the consumer vote is *safe*. If there is a test that consistently passes with V^{α} but always fails with V^{β}, then the consumer vote is *unsafe*—this condition indicates that the upgrade has broken some functionality. In all other cases (e.g., all tests were flaky or failed in V^{α}), the consumer vote is ignored.

$$\mathit{vote}(C@V, D@V^{\alpha}, D@V^{\beta}) = \begin{cases} \mathit{safe} & \text{if } \forall\, t \in \mathit{consumerTests}(C@V): \\ & \quad \mathit{outcome}(t, C@V, D@V^{\alpha}) = \mathit{pass} \implies \\ & \quad \mathit{outcome}(t, C@V, D@V^{\beta}) = \mathit{pass} \\ \mathit{unsafe} & \exists\, t \in \mathit{consumerTests}(C@V): \\ & \quad \mathit{outcome}(t, C@V, D@V^{\alpha}) = \mathit{pass} \wedge \\ & \quad \mathit{outcome}(t, C@V, D@V^{\beta}) = \mathit{fail} \\ \mathit{ignore} & \text{otherwise} \end{cases}$$

where $\mathit{consumerTests}(C@V)$ returns the set of all test methods in the test-JAR for $C@V$.

Finally, Pin-Freshener accumulates all votes of the consumers to calculate a *confidence* score for the upgrade. If any consumers vote that the upgrade is unsafe, then the confidence is 0, since the upgrade appears to be a breaking change. Otherwise, the confidence is equal to the number of consumers that voted *safe*—higher is better.

The confidence score calculated by Pin-Freshener reports the number of consumers that had consistent test results between dependency versions. This score does not provide any guarantees about the safety of the upgrade—it is possible that the executed consumer test suites did not catch a breaking change. However, even without guarantees, this score is an improvement over the standard technique of just using one's own tests. We also note that Pin-Freshener provides a *conservative* estimate of safety by reporting a score of 0 if any of the consumer test suites fail. However, the interpretation of the score depends on the preferences of the developer looking to perform the upgrade. The confidence scores reported Pin-Freshener will also increase with more testable consumers and more available test-JARs.

TABLE II: Pin-Freshener confidence on vulnerability-reducing upgrades in $\mathcal{D}$. Of 4,576 upgrades, Pin-Freshener crowdsourced $\geq$1 test-JAR for 29% (positive + zero confidence). Pin-Freshener returns positive confidence for 9,194 (41%) consumers that could perform these upgrades.

Confidence score from Pin-Freshener	Consumers	Upgrades
Positive (potentially safe)	9,194 (41%)	850 (19%)
Zero (potentially unsafe)	3,134 (14%)	458 (10%)
Untested (no test-JARs)	10,119 (45%)	3,268 (71%)
Total	22,447 (100%)	4,576 (100%)

D. RQ4: Providing Additional Signal for Freshening Pins

A key question is how much additional signal Pin-Freshener can provide to consumers to upgrade one or more of their dependencies. We answer this RQ by running Pin-Freshener on the upgrades of direct stale pins in $\mathcal{D}$ that fix security vulnerabilities. Table II reports the distribution of upgrades that had a positive and zero confidence returned by Pin-Freshener. About 29% of all upgrades were able to be tested with at least 1 test-JAR crowdsourced from the Maven Central Repository. Out of these tested upgrades, Pin-Freshener reported a positive confidence score for 850 (65% of tested upgrades, 19% of all upgrades). In total, 9,194 (41%) consumers contained these stale pins and would have a positive signal from Pin-Freshener to perform these vulnerability-fixing upgrades.

For these 9,914 consumers with positive confidence scores from Pin-Freshener, what is the distribution of these scores? Figure 9 visualizes these consumers against values of the confidence score. The X-axis value of 1 is excluded for the sake of visualization and because we believe a minimum score of 1 is too low to provide enough signal for an upgrade. Overall, we find that over 3,000 (14%) of consumers would be provided a confidence score of at least 5 using Pin-Freshener. This number of consumers increases to almost *6,000* with a confidence score of at least 2. This is a significant number of consumers that would be provided additional signal to upgrade their pinned dependencies with these consumers validating the upgrade. We believe this number can be increased even further with more Maven libraries adopting the practice of publishing their test-JARs.

V. Discussion

In this section, we discuss our findings and their broader implications to practitioners and researchers.

a) Stale dependency pinning is common in Maven: Our findings confirm that stale pinning is pervasive among consumers of popular Maven libraries, consistent with prior

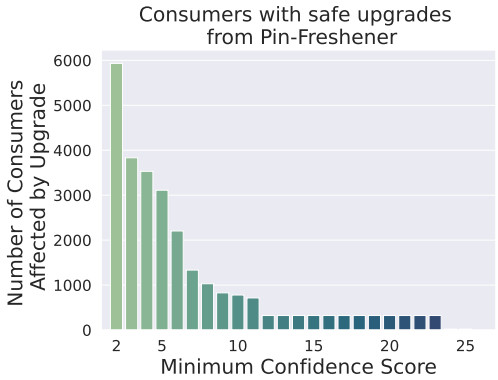

Fig. 9: Consumers with positive-confidence upgrades by Pin-Freshener confidence score. X-axis: minimum passing test suites (1 excluded); Y-axis: consumers able to upgrade at that confidence. Over 3,000 consumers have upgrades with $\geq$5 passing test suites; nearly 6,000 with $\geq$2.

studies on developer reluctance to upgrade dependencies [24], [25], [26], [27]. Maven's default of pinning to specific versions, combined with the frequent releases of popular libraries, makes it challenging for consumers to stay current.

b) Fresh pinning is safer than stale pinning: Our retrospective analysis shows that upgrading pinned dependencies would have had a net positive security impact across the ecosystem, consistent with studies linking outdated dependencies to vulnerabilities [28]. We hope developers are encouraged to adopt a more proactive strategy of keeping dependencies up to date.

c) Coverage of a dependency improves with crowdsourced test suites: It is challenging for a consumer to evaluate how their project will be affected by a dependency upgrade. While their own test suites may be able to catch certain issues, we see that *crowdsourcing* test suites from other consumers can provide a substantial boost in coverage. These test suites may be exercising different parts of the dependency, and a consumer may only care about a certain functionality that they use; nevertheless, we feel each additional test suite can only help in increasing confidence for an upgrade. Prior work has shown the potential for consumer tests [29], [30], [31], [32] in achieving reasonable coverage and fault detection capabilities in dependencies.

d) Ecosystems should encourage developers to publicize executable test suites: Pin-Freshener leverages test-JARs published to the Central Maven Repository to streamline automatic execution of consumer tests. We hope to see this practice adopted more widely, both within Maven and in other ecosystems. Our approach is analogous to how *monorepo* environments in large companies [33] select and run tests from external modules to validate code changes—Pin-Freshener applies this idea to the open source world, providing something akin to a "monorepo for the masses".

e) Limitations: While Pin-Freshener is able to provide additional signal of crowdsourced test suites to measure safety of upgrades, we acknowledge that this signal does not provide any *guarantee* of safety for a specific consumer. We believe that this type of signal, however, would still be useful for developers to understand how the upgrade impacts the rest of the ecosystem and would be beneficial in making a decision about whether to upgrade from a stale to fresh pin. This is similar to how test coverage is not a guarantee of the absence of bugs but is still a widely used measure of confidence. More coverage is always good; confidence increases monotonically even if it never reaches 100%. Our score provides a similar rating for developers familiar with such a notion of test confidence.

VI. Threats to Validity

a) Threats to Construct Validity: The validation performed by Pin-Freshener on an upgrade is dependent on the consumer tests that are executed. If there is any noise or nondeterminism affecting the test outcome, then Pin-Freshener may improperly classify certain upgrades as safe or unsafe. This can arise from flakiness [21], [34], [35] in tests. We aim to mitigate this threat through repeated execution of the tests five times (Section IV-A) on both the pinned version and the upgrade version. Pin-Freshener only compares tests that produce a consistent passing or failing outcome across all repetitions, which should filter out a majority of flaky tests.

b) Threats to Internal Validity: Pin-Freshener's approach of crowdsourcing test suites and validating upgrades assumes that consumer test suites are a valuable source testing a dependency. Since library test suites are generally focused on testing functionality of the library and not the dependencies, it may be the case that consumer tests do not exercise much behavior of dependencies. Nevertheless, Pin-Freshener executes as many consumer test suites as are available in the Maven Central Repository. We hope that publishing test-JARs becomes a more widely adopted practice in Maven, as this would increase the overall coverage of the dependency.

c) Threats to External Validity: We specifically focused on the Maven ecosystem for our analysis, and we do not know if our conclusions about dependency pinning and its security implications will generalize to other ecosystems. Additionally, Pin-Freshener depends on a central repository of crowdsourced tests that can be automatically executed; this data may not always be available in other platforms.

VII. Related Work

A. Dependencies in Software Ecosystems

The challenge of evolving and maintaining software in ecosystems is a well-researched topic [36], [37], [38], [39], [40], [41], [42], [43], [44], [45], [46], [47]. Bavota et al. [48] explore the Apache ecosystem and highlight the exponential growth in the number dependencies. They also found that application developers are reluctant to upgrade their dependencies due to the risk of API breaking changes. This issue is further quantified by Kula et al. (2015) [25], sampling 4.6K Github projects and finding that more than 80 percent of them have outdated Maven dependencies. Additional studies [49]

validate this finding for other ecosystems such as NPM by measuring technical lag in dependencies. Dietrich et al. [27] demonstrate that 85.7% of Maven libraries specify a fixed version in dependencies—our definition of stale pinning is more precise as it compares the resolved version to the latest dependency version available at the time of publishing. Cox et al. [24] conduct a mixed methods study on outdated dependencies in Maven and propose a system-level metric for dependency freshness; our analysis of stale pins confirms that outdated dependencies are frequent even in recent snapshots of the Maven ecosystem and we provide Pin-Freshener as a method for encouraging developers to upgrade to fresh pins.

Prior work [50], [51] has also measured the impact of vulnerabilities in dependencies in the NPM ecosystem. Kula et al. (2018) [26] extend their work to study the extent to which developers upgrade their dependencies and the reasons behind their reluctancy [26]. In a survey of developers, they find that 69% claimed to be unaware of vulnerabilities in their dependencies. He et al. [52] perform an empirical study on the npm ecosystem and show that dependency pinning can increase the attack surface of malicious package updates [53]. Automated dependency management bots like *Dependabot* [11] are able to address this issue by automatically notifying and creating pull requests for developers to upgrade their vulnerable dependencies. Analysis on Dependabot in practice shows that it does reduce technical lag in projects; however, its compatibility score does not reduce developer suspicion when performing upgrades [13]. Our approach can provide additional signal through the execution of consumer test suites.

B. Detection of Breaking Changes

Prior research has studied [7], [5], [54], [55], [56] and developed numerous techniques for the detection of breaking changes [57], [58], [6], [59] that can alert developers of unsafe upgrades.

a) Static Analysis Based Techniques: The majority of existing literature focuses primarily on detection of API changes between library versions. Raemaekers et al. [4] utlilize the tool `clirr` to detect API binary incompatibilities of Java code through static analysis. APIDiff is a tool developed by Brito et al. [57] that focuses on syntactic changes between Java library versions that classifies a code change as breaking or non-breaking. The more recent tool Sembid [60] locates breaking changes in Maven libraries by analyzing call chains and measuring semantic differences between versions.

b) Dynamic Analysis Based Techniques: Mostafa et al. [61] study the prevalence of *behavioral* backwards incompatibilities (BBIs) in consecutive versions of Java libraries. They find that 14 of the 15 subjects featured these types of breaking changes, with the majority of them undocumented. Prior work has also shown the effectiveness of using consumer tests to detect breaking changes and BBIs [60], [31], [29]. We highlight the main differences from our work: first, we provide a novel definition of explicit dependency pins and present a thorough empirical study on pinning in the Maven network, which is unique among related work. We also use a dataset that resolves dependency versions for old libraries at the time they were built; this is contrast to prior work that uses heuristics to resolves dependencies in older releases [31]. We focus on the security impact of pinning dependencies and validating upgrades from pins, which is unique among related work. Finally, we use crowdsourced tests from JARs published to the Maven central repository, and thus do not rely on identifying source code repositories like prior work [31], [30], [29]. A potential future extension of our work is to use generative AI and agentic techniques to automatically build repositories and run tests [62].

C. Client-Specific Testing

Previous techniques have aimed to measure the effect of changes in dependency components on clients. Mora et al. [63] develop an automated technique to explore behaviors of the client through symbolic execution. Zhu et al. [64] develop Compcheck, which leverages an offline knowledge base of incompatibility issues to find similar library usages and generate tests revealing client incompatibility issues [65], [66]. We believe these techniques can further enhance the use of Pin-Freshener, as they can provide more detailed analysis on test execution results; however, they require heavier-weight analysis on test executions and source code. Our approach provides a pragmatic middle-ground to execute all client test-JARs for a given library version with coverage analysis as a signal for developers. Prior automated test generation techniques to generate exploits for vulnerabilities in dependencies [67], [68] can be for additional testing in the ecosystem that could be crowdsourced.

VIII. Conclusion

In this work, we studied dependency pinning in the Maven ecosystem. We formalized the concepts of *stale* and *fresh* pins and found that a significant portion of consumers are pinned to older versions of popular libraries. A retrospective analysis shows that upgrading stale pins is more likely to fix security vulnerabilities than introduce new ones. To encourage upgrades, we prototyped Pin-Freshener, which crowdsources consumer test suites to measure upgrade safety, providing additional signal to over 19% of consumers who could fix vulnerabilities. We argue that more libraries and platforms should adopt the practice of publishing executable test binaries to enable tools like Pin-Freshener.

IX. Data Availability

We have included evaluation data in the repository at: https://doi.org/10.5281/zenodo.15021893. This data contains dependency data for each of the datasets, coverage data for consumer test suites, and test outcome data from Pin-Freshener.

X. Acknowledgements

This research was funded in part by NSF grant CCF-2120955 and the Future Enterprise Security Initiative at CyLab.

REFERENCES

[1] C. R. de Souza and D. F. Redmiles, "An empirical study of software developers' management of dependencies and changes," in *Proceedings of the 30th international conference on Software engineering*, 2008, pp. 241–250.

[2] M. Cataldo, A. Mockus, J. A. Roberts, and J. D. Herbsleb, "Software dependencies, work dependencies, and their impact on failures," *IEEE Transactions on Software Engineering*, vol. 35, no. 6, pp. 864–878, 2009.

[3] "The maven central repository," https://mvnrepository.com/repos/central, accessed: 2022-11-21.

[4] S. Raemaekers, A. Van Deursen, and J. Visser, "Semantic versioning versus breaking changes: A study of the maven repository," in *2014 IEEE 14th International Working Conference on Source Code Analysis and Manipulation*. IEEE, 2014, pp. 215–224.

[5] L. Xavier, A. Brito, A. Hora, and M. T. Valente, "Historical and impact analysis of api breaking changes: A large-scale study," in *2017 IEEE 24th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2017, pp. 138–147.

[6] L. Ochoa, T. Degueule, J.-R. Falleri, and J. Vinju, "Breaking bad? semantic versioning and impact of breaking changes in maven central: An external and differentiated replication study," *Empirical Software Engineering*, vol. 27, no. 3, p. 61, 2022.

[7] C. Bogart, C. Kästner, J. Herbsleb, and F. Thung, "How to break an api: cost negotiation and community values in three software ecosystems," in *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, 2016, pp. 109–120.

[8] S. E. Ponta, H. Plate, and A. Sabetta, "Detection, assessment and mitigation of vulnerabilities in open source dependencies," *Empirical Software Engineering*, vol. 25, no. 5, pp. 3175–3215, 2020.

[9] S. Mukherjee, A. Almanza, and C. Rubio-González, "Fixing dependency errors for python build reproducibility," in *Proceedings of the 30th ACM SIGSOFT international symposium on software testing and analysis*, 2021, pp. 439–451.

[10] M. Corporation, "CVE-2017-5638," https://www.cve.org/CVERecord?id=CVE-2017-5638, 2017. [Online]. Available: https://www.cve.org/CVERecord?id=CVE-2017-5638

[11] Github, "Dependabot," https://docs.github.com/en/code-security/dependabot/working-with-dependabot/automating-dependabot-with-github-actions#about-dependabot-and-github-actions, 2021, retrieved June 1, 2023.

[12] M. Alfadel, D. E. Costa, E. Shihab, and M. Mkhallalati, "On the use of dependabot security pull requests," in *2021 IEEE/ACM 18th International Conference on Mining Software Repositories (MSR)*. IEEE, 2021, pp. 254–265.

[13] R. He, H. He, Y. Zhang, and M. Zhou, "Automating dependency updates in practice: An exploratory study on github dependabot," *IEEE Transactions on Software Engineering*, 2023.

[14] H. Mohayeji, A. Agaronian, E. Constantinou, N. Zannone, and A. Serebrenik, "Investigating the resolution of vulnerable dependencies with dependabot security updates," in *2023 IEEE/ACM 20th International Conference on Mining Software Repositories (MSR)*. IEEE, 2023, pp. 234–246.

[15] Google, "Open Source Insights," https://deps.dev/, 2023, retrieved June 1, 2023.

[16] T. Preston-Werner, "Semantic versioning 2.0.0," https://semver.org/spec/v2.0.0.html, 2013.

[17] A. Decan and T. Mens, "What do package dependencies tell us about semantic versioning?" *IEEE Transactions on Software Engineering*, vol. PP, pp. 1–1, 05 2019.

[18] E. Wittern, P. Suter, and S. Rajagopalan, "A look at the dynamics of the javascript package ecosystem," in *Proceedings of the 13th International Conference on Mining Software Repositories*, 2016, pp. 351–361.

[19] Libraries.io, "Libraries.io Open Data," https://libraries.io/data, 2020, retrieved June 1, 2023.

[20] Google, "Open Source Vulnerabilities," https://osv.dev/, retrieved June 1, 2023.

[21] Q. Luo, F. Hariri, L. Eloussi, and D. Marinov, "An empirical analysis of flaky tests," in *Proceedings of the 22nd ACM SIGSOFT international symposium on foundations of software engineering*, 2014, pp. 643–653.

[22] M. Eck, F. Palomba, M. Castelluccio, and A. Bacchelli, "Understanding flaky tests: The developer's perspective," in *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, 2019, pp. 830–840.

[23] M. R. Hoffmann, B. Janiczak, and E. Mandrikov, "Eclemma-jacoco Java code coverage library," 2011.

[24] J. Cox, E. Bouwers, M. Van Eekelen, and J. Visser, "Measuring dependency freshness in software systems," in *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, vol. 2. IEEE, 2015, pp. 109–118.

[25] R. G. Kula, D. M. German, T. Ishio, and K. Inoue, "Trusting a library: A study of the latency to adopt the latest maven release," in *2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*. IEEE, 2015, pp. 520–524.

[26] R. G. Kula, D. M. German, A. Ouni, T. Ishio, and K. Inoue, "Do developers update their library dependencies? an empirical study on the impact of security advisories on library migration," *Empirical Software Engineering*, vol. 23, pp. 384–417, 2018.

[27] J. Dietrich, D. Pearce, J. Stringer, A. Tahir, and K. Blincoe, "Dependency versioning in the wild," in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 2019, pp. 349–359.

[28] I. Pashchenko, H. Plate, S. E. Ponta, A. Sabetta, and F. Massacci, "Vulnerable open source dependencies: Counting those that matter," in *Proceedings of the 12th ACM/IEEE international symposium on empirical software engineering and measurement*, 2018, pp. 1–10.

[29] G. Mezzetti, A. Møller, and M. T. Torp, "Type regression testing to detect breaking changes in node. js libraries," in *32nd european conference on object-oriented programming (ECOOP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.

[30] L. Chen, F. Hassan, X. Wang, and L. Zhang, "Taming behavioral backward incompatibilities via cross-project testing and analysis," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 112–124.

[31] S. Mujahid, R. Abdalkareem, E. Shihab, and S. McIntosh, "Using others' tests to identify breaking updates," in *Proceedings of the 17th International Conference on Mining Software Repositories*, 2020, pp. 466–476.

[32] J. Hejderup and G. Gousios, "Can we trust tests to automate dependency updates? a case study of Java projects," *Journal of Systems and Software*, vol. 183, p. 111097, 2022.

[33] A. Memon, Z. Gao, B. Nguyen, S. Dhanda, E. Nickell, R. Siemborski, and J. Micco, "Taming google-scale continuous testing," in *2017 IEEE/ACM 39th International Conference on Software Engineering: Software Engineering in Practice Track (ICSE-SEIP)*. IEEE, 2017, pp. 233–242.

[34] W. Lam, K. Muşlu, H. Sajnani, and S. Thummalapenta, "A study on the lifecycle of flaky tests," in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, 2020, pp. 1471–1482.

[35] O. Parry, G. M. Kapfhammer, M. Hilton, and P. McMinn, "A survey of flaky tests," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 31, no. 1, pp. 1–74, 2021.

[36] M. M. Lehman, D. E. Perry, and J. F. Ramil, "Implications of evolution metrics on software maintenance," in *Proceedings. International Conference on Software Maintenance (Cat. No. 98CB36272)*. IEEE, 1998, pp. 208–217.

[37] M. M. Lehman, "The role and impact of assumptions in software development, maintenance and evolution," in *IEEE International Workshop on Software Evolvability (Software-Evolvability'05)*. IEEE, 2005, pp. 3–14.

[38] M. W. Godfrey and D. M. German, "The past, present, and future of software evolution," in *2008 Frontiers of Software Maintenance*. IEEE, 2008, pp. 129–138.

[39] M. D'Ambros, "Supporting software evolution analysis with historical dependencies and defect information," in *2008 IEEE International Conference on Software Maintenance*. IEEE, 2008, pp. 412–415.

[40] V. Rajlich, "Software evolution and maintenance," in *Future of Software Engineering Proceedings*, 2014, pp. 133–144.

[41] P. Tripathy and K. Naik, *Software evolution and maintenance: a practitioner's approach*. John Wiley & Sons, 2014.

[42] N. Chapin, J. E. Hale, K. M. Khan, J. F. Ramil, and W.-G. Tan, "Types of software evolution and software maintenance," *Journal of software maintenance and evolution: Research and Practice*, vol. 13, no. 1, pp. 3–30, 2001.

[43] T. Mens, M. Wermelinger, S. Ducasse, S. Demeyer, R. Hirschfeld, and M. Jazayeri, "Challenges in software evolution," in *Eighth International Workshop on Principles of Software Evolution (IWPSE'05)*. IEEE, 2005, pp. 13–22.

[44] T. Mens, S. Demeyer, M. D'Ambros, H. Gall, M. Lanza, and M. Pinzger, "Analysing software repositories to understand software evolution," *Software evolution*, pp. 37–67, 2008.

[45] K. Manikas and K. M. Hansen, "Software ecosystems–a systematic literature review," *Journal of Systems and Software*, vol. 86, no. 5, pp. 1294–1306, 2013.

[46] R. Cox, "Surviving software dependencies," *Communications of the ACM*, vol. 62, no. 9, pp. 36–43, 2019.

[47] D. Zhou, Y. Wu, X. Peng, J. Zhang, and Z. Li, "Revealing code change propagation channels by evolution history mining," *Journal of Systems and Software*, vol. 208, p. 111912, 2024.

[48] G. Bavota, G. Canfora, M. Di Penta, R. Oliveto, and S. Panichella, "How the apache community upgrades dependencies: an evolutionary study," *Empirical Software Engineering*, vol. 20, no. 5, pp. 1275–1317, 2015.

[49] A. Zerouali, E. Constantinou, T. Mens, G. Robles, and J. González-Barahona, "An empirical analysis of technical lag in npm package dependencies," in *International Conference on Software Reuse*. Springer, 2018, pp. 95–110.

[50] C. Liu, S. Chen, L. Fan, B. Chen, Y. Liu, and X. Peng, "Demystifying the vulnerability propagation and its evolution via dependency trees in the npm ecosystem," in *Proceedings of the 44th International Conference on Software Engineering*, 2022, pp. 672–684.

[51] A. Decan, T. Mens, and P. Grosjean, "An empirical comparison of dependency network evolution in seven software packaging ecosystems," *Empirical Software Engineering*, vol. 24, no. 1, pp. 381–416, 2019.

[52] H. He, B. Vasilescu, and C. Kästner, "Pinning is futile: You need more than local dependency versioning to defend against supply chain attacks," *Proceedings of the ACM on Software Engineering*, no. FSE, pp. 1–24, 2025, fSE '25.

[53] I. Rahman, J. Marley, W. Enck, and L. Williams, "Which is better for reducing outdated and vulnerable dependencies: Pinning or floating?" *arXiv preprint arXiv:2510.08609*, 2025.

[54] C. Bogart, C. Kästner, J. Herbsleb, and F. Thung, "When and how to make breaking changes: Policies and practices in 18 open source software ecosystems," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 30, no. 4, pp. 1–56, 2021.

[55] D. Jayasuriya, V. Terragni, J. Dietrich, S. Ou, and K. Blincoe, "Understanding breaking changes in the wild," in *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2023, pp. 1433–1444.

[56] D. Jayasuriya, V. Terragni, J. Dietrich, and K. Blincoe, "Understanding the impact of apis behavioral breaking changes on client applications," *Proceedings of the ACM on Software Engineering*, vol. 1, no. FSE, pp. 1238–1261, 2024.

[57] A. Brito, L. Xavier, A. Hora, and M. T. Valente, "Apidiff: Detecting api breaking changes," in *2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 2018, pp. 507–511.

[58] X. Du and J. Ma, "Aexpy: Detecting api breaking changes in python packages," in *2022 IEEE 33rd International Symposium on Software Reliability Engineering (ISSRE)*. IEEE, 2022, pp. 470–481.

[59] L. Ochoa, T. Degueule, and J.-R. Falleri, "Breakbot: Analyzing the impact of breaking changes to assist library evolution," in *Proceedings of the ACM/IEEE 44th International Conference on Software Engineering: New Ideas and Emerging Results*, 2022, pp. 26–30.

[60] L. Zhang, C. Liu, Z. Xu, S. Chen, L. Fan, B. Chen, and Y. Liu, "Has my release disobeyed semantic versioning? static detection based on semantic differencing," in *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, 2022, pp. 1–12.

[61] S. Mostafa, R. Rodriguez, and X. Wang, "Experience paper: a study on behavioral backward incompatibilities of Java software libraries," in *Proceedings of the 26th ACM SIGSOFT international symposium on software testing and analysis*, 2017, pp. 215–225.

[62] I. Bouzenia and M. Pradel, "You name it, i run it: An llm agent to execute tests of arbitrary projects," *arXiv preprint arXiv:2412.10133*, 2024.

[63] F. Mora, Y. Li, J. Rubin, and M. Chechik, "Client-specific equivalence checking," in *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, 2018, pp. 441–451.

[64] C. Zhu, M. Zhang, X. Wu, X. Xu, and Y. Li, "Client-specific upgrade compatibility checking via knowledge-guided discovery," *ACM Transactions on Software Engineering and Methodology*, vol. 32, no. 4, pp. 1–31, 2023.

[65] Y. Xie, Z. Jia, S. Li, Y. Wang, J. Ma, X. Li, H. Liu, Y. Fu, and X. Liao, "How to pet a two-headed snake? solving cross-repository compatibility issues with hera," in *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*, 2024, pp. 694–705.

[66] G. Monce, T. Degueule, J.-R. Falleri, and R. Robbes, "Client-library compatibility testing with api interaction snapshots," in *2025 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 2025, pp. 791–796.

[67] H. J. Kang, T. G. Nguyen, B. Le, C. S. Păsăreanu, and D. Lo, "Test mimicry to assess the exploitability of library vulnerabilities," in *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis*, 2022, pp. 276–288.

[68] Z. Chen, X. Hu, X. Xia, Y. Gao, T. Xu, D. Lo, and X. Yang, "Exploiting library vulnerability via migration based automating test generation," *arXiv preprint arXiv:2312.09564*, 2023.