



# Lessons for Interactive Theorem Proving Researchers from a Survey of Coq Users

Ana de Almeida Borges<sup>1</sup> · Annalí Casanueva Artís<sup>2</sup> · Jean-Rémy Falleri<sup>3</sup> · Emilio Jesús Gallego Arias<sup>4</sup> · Érik Martin-Dorel<sup>5</sup> · Karl Palmiskog<sup>6</sup> · Alexander Serebrenik<sup>7</sup> · Théo Zimmermann<sup>8</sup>

Received: 28 May 2024 / Accepted: 28 January 2025  
© The Author(s) 2025

## Abstract

The Coq Community Survey 2022 was an online public survey of users of the Coq proof assistant conducted during February 2022. Broadly, the survey asked about use of Coq features, user interfaces, libraries, plugins, and tools, views on renaming Coq and Coq improvements, and also demographic data such as education and experience with Coq and other proof assistants and programming languages. The survey received 466 submitted responses, making it the largest survey of users of an interactive theorem prover (ITP) so far. We present the design of the survey, a summary of key results, and analysis of answers relevant to ITP technology development and usage. In particular, we analyze user characteristics associated with adoption of tools and libraries and make comparisons to adjacent software communities. Notably, we find that experience has significant impact on Coq user behavior, including on usage of tools, libraries, and integrated development environments (IDEs).

**Keywords** Coq · Community · Survey · Statistical analysis

## 1 Introduction

Online surveys are employed by organizations [1, 2] and researchers [3] to obtain aggregate data about software communities, e.g., with the aim of understanding community demographics, programming practices, and desired improvements to a software ecosystem. Some interactive theorem provers (ITPs) such as Coq, Isabelle/HOL, and Lean have many active users across academia and industry, and have comparable popularity to programming languages on developer platforms such as GitHub and Stack Overflow [4]. Yet, to our knowledge, the only public summary for an ITP community survey is Braibant's for the Coq Community Survey 2014 [5]. However, due to the development of Coq itself and changes in the Coq community since 2014, many of its results and conclusions are currently obsolete.

In this paper, we present and analyze the results of a new survey of Coq users, which we call the Coq Community Survey 2022. The survey was conducted publicly online during February 2022, using the LimeSurvey tool [6] on a platform hosted by Inria. The main goal of the survey was to obtain an updated picture of the Coq user community and inform future

---

Corrected By: Karl Pal

Extended author information available on the last page of the article

decisions taken by the Coq development team (*Coq Team* for short) and other stakeholders in the Coq ecosystem. In particular, the survey aimed to help the Coq Team to better prioritize issues and features for the Coq system, and enable effective decision-making about matters pertaining to the software ecosystem maintained by Coq users in academia and industry [7].

One specific Coq Team impetus for the survey was a community-wide discussion on Coq's name and logo in April 2021 on the Coq mailing list [8], centering on the name's pronunciation that is homophonous with the English word "cock", with subsequent sexual connotations. Consequently, the survey contained questions related to the name, including on respondent's views on alternative names.

The survey was designed, deployed, and managed by a working group (WG) that was formed following a public call by the Coq Team. The survey had 109 questions, some of which were conditionally visible, and was available in English and Chinese [9, 10]. Broadly, the survey asked about (1) use of Coq features, integrated development environments (IDEs), libraries, plugins, and tools, (2) views on renaming and improving Coq, and (3) demographic data, such as age, gender, and experience with Coq. The survey received 466 submitted responses.

The WG has disseminated survey results in several venues, starting with initial presentation of response summaries and basic inferred data in a series of three blog posts [11–13], followed by a Coq workshop presentation [14, 15], and more recently a conference research paper [16] with statistical analyses that unveiled characteristics of Coq users associated with different usage habits, and a final blog post [17].

This paper, authored by a subset of the WG, contains the definitive account of the survey and its results, including summaries of responses to questions and statistical analyses. Compared to the conference paper [16], we present summaries of additional multiple-choice questions, and we include a qualitative analysis component by reporting on select free-text questions, in particular on community health and contributing to Coq. We make the following contributions:

1. We present the first Coq community survey since 2014, and to our knowledge the survey of an ITP community with the largest number of respondents thus far. We highlight and discuss key survey results, focusing on data we believe is relevant to ITP researchers and practitioners, including answers to free-text questions.
2. We compare aggregate data of Coq survey respondents with data from similar surveys in related software communities (Haskell [18] and Stack Overflow [2]).
3. We perform multiple variable regression analyses on survey data to answer one main research question: **How do different categories of respondents use Coq?** In addition, we look at satisfaction and needs for improvements for the same categories of respondents. Our code and results are part of our public dataset [19].

The paper is structured as follows. Section 2 describes the survey design and deployment, the data analysis process, and the methodology for our statistical analyses, including the definition of our respondent categories of interest. Section 3 gives a descriptive analysis of some key survey results, and performs a comparison between the Coq community and the Stack Overflow and Haskell communities, based on their respective 2022 surveys. Section 4 describes responses related to community health and the experience of contributing to Coq. Section 5 compares Coq use across different categories of Coq users. Section 6 compares differing satisfaction and user needs across the same categories of Coq users. Section 7 discusses threats to validity of our results, and measures taken to mitigate them. Finally, Sect. 8 discusses our results and concludes.

## 2 Methodology

The survey was designed, deployed, managed, and analyzed by members of a working group (WG), which was formed after a public call for participation by the Coq core team. Members of this WG met weekly during a year and a half, first to design the survey, then to advertise it to the Coq community, and finally to analyze and disseminate the results, both to academic and other audiences.

### 2.1 Survey Design and Deployment

Survey questions were collaboratively created by WG members. The WG took inspiration from the previous Coq Community Survey in 2014 [5], and similar surveys for programming languages—in particular the 2020 survey of OCaml users [20]. The WG made efforts to adhere to survey best practices [21], e.g., on ordering demographic questions late to avoid their answers affecting other questions.

*Survey Language* The survey was available in English [9] and Chinese [10], having been first created in English and then translated to Chinese by a team of volunteer native Chinese speakers with knowledge of Coq. These volunteers also translated the answers to the open questions back to English. Ideally, the survey would have been translated to more languages, particularly the ones that have an existing community of Coq users who are not necessarily proficient in English. For example, there are Coq learning resources in Chinese [22] and Japanese [23], as well as dedicated categories for several languages in the Coq Discourse forum. Based on the WG's estimates about the size of Coq communities using a language and the prospect for additional responses through a survey translation, the WG contacted colleagues with the necessary knowledge to perform Chinese, Japanese, and Russian translations. However, the WG only secured a commitment for the Chinese translation. A French translation was ruled out since the WG estimated that Coq users speaking French would be willing and able to respond to the English survey.

*Survey Content and Testing* The survey had 109 questions, some of which were only visible depending on answers to previous questions. Broadly, the survey asked about (1) use of Coq features, IDEs, libraries, plugins, and tools, (2) views on renaming Coq, Coq improvements, and issues such as inclusion, and (3) demographic data, such as age, gender, and length of experience with Coq and other ITPs or functional programming languages. All the questions were optional. After a preliminary version of the survey was finalized, it was sent to a small group of beta testers, who were asked to fill out that preliminary version, time themselves, and send feedback on their experience. Some changes were made due to this feedback, and the initial prediction of the time it would take to fully complete the closed-question part of the survey (30 min) was estimated based on this beta testing. We made this estimation available to potential survey respondents on the first page of the survey and in our announcements.

*Survey Deployment* Before survey deployment, we contacted the Inria Ethics Review Board (ERB), who informed us that full ERB approval of our survey was not required and referred us to Inria's Data Protection Officer (DPO) for data protection compliance matters. Based on our communication with the DPO, we developed a GDPR conformance statement and followed best practices for data protection. In particular, we removed the raw data and only share aggregated data, including open text answers which we manually sanitized to remove any personally identifiable elements. The survey was then deployed on Inria's LimeSurvey platform and was open during the month of February 2022.

**If you have projects that are not compatible with recent versions of Coq (8.13 or more recent), why is that? Please choose all that apply:**

- It would take too much effort to port them.
- I don't care that they aren't on the latest version.
- A dependency of the project is not compatible with recent versions of Coq.
- Other: \_\_\_\_\_

**Fig. 1** Example survey question that was hidden based on answers to earlier questions; in this case, whether the respondent had used Coq in the last 6 months

*Survey Dissemination* The deployed survey was advertised to Coq users primarily via the following Coq community forums: Discourse, Zulip chat, Coq-club mailing list, Coq Twitter, Coq Reddit, and the Coq website. The WG also advertised in more general venues and other related communities, such as the Types-announce mailing list [24], Agda and Lean Zulip chats, and on language-specific forums in Chinese and in French (mailing lists of GDR IM [25] and GDR GPL [26]). These announcements asked only previous or current Coq users to answer. Still, the survey included the question “How long have you used Coq?” with the option to answer “I have never used Coq.” Those who chose that option were still asked a number of questions, but most of the survey was hidden, since most questions assumed previous Coq experience. Figure 1 shows an example of a question that was hidden when a respondent claimed never to have used Coq.

## 2.2 Response Analysis Process

After the survey closed, WG members collaboratively analyzed the response data. For GDPR compliance, the WG decided to keep the raw survey data within the EU. Therefore, the WG used Inria-hosted tools and restricted raw data access to members located in the EU. The Chinese translators translated the Chinese answers back into English, and then two authors wrote the analysis code in Python within a Jupyter notebook [27], using the `pandas` and `matplotlib` libraries. Plots were created for each closed-answer survey question (often multiple plots for each question) and automatically deployed to a static Inria-hosted website for later inclusion in summary blog posts (see Sect. 3). In order to perform the statistical analyses described in Sect. 5, we exported pre-processed data from the Jupyter notebook for a third author to perform regressions using Stata, a software package for statistical analysis [28]. We provide the Jupyter notebook and Stata code as supplementary material [19] to allow scrutiny of the source code, and also provide all the generated plots for the survey questions. However, due to privacy concerns (risk of re-identifying specific respondents), we do not provide the raw survey data, which is required to reproduce the analyses using the Jupyter notebook.

## 2.3 Statistical Analyses

We do a multiple variable regression analysis to unveil the characteristics of Coq users associated with different usage habits and expectations. This statistical method allows us to establish the relationship between multiple regressors (or independent variables) and one dependent variable (or outcome) [29]. Including several variables at the same time is needed to disentangle the distinct effect of possibly correlated variables. If we did not include such

correlated variables simultaneously, part of the effect of the excluded variables would be captured by the coefficient of the included correlated ones. One must however keep in mind that: (1) these results cannot be interpreted causally and (2) there can still be variables correlated with the dependent or independent variables that were not included in the regression.

For each category of users, we test several hypotheses. This will lead to over-rejecting null hypotheses (i.e., some associations between variables will be considered true when in reality they are not present). To overcome this problem, it is necessary to explicitly consider the multiplicity of the testing framework. Thus, for all our estimates, we report, in addition to classic standard errors (in parentheses), the Romano-Wolf-corrected p-values for multiple hypothesis testing (in brackets) [30, 31]. This correction has a higher ability to correctly reject false null hypotheses than previous techniques such as Bonferroni and Holm [32]. Following common standards in fields with a long tradition of statistical analysis such as economics, and considering that, by controlling the family-wise error rate, the Romano-Wolf correction is more conservative than alternative approaches to multiple hypothesis testing (giving us robust levels of significance), we interpret estimates as significant as soon as the p-values are below the 0.1 threshold (instead of the 0.05 threshold which would be more common in empirical software engineering).

### 2.3.1 Defining Categories of Interest

We studied four user characteristics that we considered could be relevant in explaining differences in how respondents use Coq: their experience with Coq, their learning status, their application-domain for Coq, and their location. We explain each characteristic in turn.

*Experience Level* We define two binary variables that illustrate various stages of experience with Coq. The first variable represents whether the user has more than 2 years of experience, and the second variable whether the user has more than 5 years of experience. About 64% of our sample has more than 2 years of experience and about 40% more than 5. The omitted category represents users with less than 2 years of experience. We use these variables (corresponding to arbitrary thresholds) instead of a numeric variable representing the number of years of experience because experience does not always have a linear association with outcomes, such as the likelihood of using a particular feature. In particular, we can expect that, beyond some level of experience, additional years of experience are very unlikely to be associated with any additional changes in the outcome.

*Learners* We define a binary variable corresponding to learners. To define learner, we use the answer to the question “For what purpose have you used Coq?”. A learner is a respondent who answered any of “Learning Coq” or “Learning something other than Coq” to this (multiple-choice) question, without also responding any of “Teaching Coq”, “Teaching something other than Coq”, “Academic research”, or “Industrial research/application”. About 17% of our sample are defined as learners. The way that learners use Coq may be related to their learning context, and thus significantly differ from other Coq users. Understanding these differences could inform researchers and engineers trying to make proof assistants easier to learn.

*Application Domains for Using Coq* Coq can be used for various applications: verifying software, hardware, or theoretical systems, formalizing existing or new mathematical results, etc. The way of using Coq may differ depending on the application domain (extraction comes to mind as being specifically relevant for software verification), but also on the users’ background. This is why we decided to test whether software verification specifically (the

most common objective for Coq users) was related to differing practices. About 67% of our sample use Coq for software verification.

*Location* Various world regions have different research traditions (such as putting more or less focus on theory) and may also have different technical culture. This is why it seemed important to evaluate whether this relates to different ways of using Coq. Our baseline is Europe, where about half of our respondents live. We define two binary variables to represent the respondents' location: one for North America (which represents about 25% of our respondents' location) and one for the rest of the world (excluding both North America and Europe). This is why Europe does not appear explicitly in our tables.

### 2.3.2 Controls

In addition to the variables presented above, our regressions include additional variables as controls. In particular, we include: the age of respondents, their gender (more precisely, a binary variable for whether the respondent declared being a woman), their operating system (two binary variables for Windows and macOS), their research area (a binary variable for whether the respondent does research in math or logic), their employment (a binary variable for whether the respondent is employed by a private organization), and whether they hold a PhD. We do not analyze the results of the regressions for these variables; they are only there to avoid capturing something that could be correlated both to the independent variables and to the outcome. For example, if we include years of experience without including age, we might be capturing the effect of being of a certain generation on using a specific feature that was particularly popular during a specific period. Because years of experience is correlated with age, part of the association between age and the outcome is captured by the coefficient of years of experience. Yet, in this example, the experience would not be relevant.

### 2.3.3 Selecting Population to Analyze

Our survey received 466 responses. However, for the statistical analyses, we restrict ourselves to the 390 respondents who say that they have used Coq in the last year (as what we want to assess is current Coq practices). Furthermore, we can only perform regressions on data where our (dependent and independent) variables, including our controls, are defined. We remove 18 respondents who did not provide their age. This results in looking at a subsample of 372 responses. On one specific outcome (IDE satisfaction), we look at an even smaller subsample of 367 responses, because respondents need to have answered at least one IDE satisfaction question for this outcome to be defined.

## 3 Descriptive Observations and Comparison to Other Communities

### 3.1 Descriptive Analysis and Observations

After the survey closed, the WG performed a descriptive analysis of the data. The results were continuously shared with the community through a series of blog posts in the Coq Discourse forum. The topics of the three first articles were as follows: *Who is using Coq and in what context?* [11]; *How are people using Coq?* [12]; *How is Coq used? (features, tools, libraries)* [13]. A final blog post covered topics related to improving, contributing to, and renaming Coq, and community health [17]. The blog posts include many plots deployed from

a GitLab repository to a perennial URL,<sup>1</sup> in SVG and PNG format. Results from the blog posts were also presented and discussed at the Coq Workshop 2022 [14, 15]. In addition, the WG prepared detailed internal reports and presentations to the Coq Team on the results to the renaming questions to help the team make a decision regarding a possible renaming of Coq. At the time of writing, a decision has been made by the Coq Team to rename Coq, but the renaming has not yet been enacted.

Here, we present some key plots of descriptive response results. First, we describe the demography of respondents, including their location in Fig. 2, age in Fig. 3, education in Fig. 4, and employment in Fig. 5. Second, we describe their use of IDEs in Fig. 6 and operating system (OS) use in Fig. 7. We comment on all these plots while comparing the Coq community to two other related communities in Sect. 3.2 below.

### 3.2 Comparison of the Coq Community to Other Communities

In this section, we compare the Coq community (CC) to two distinct baseline populations. We use the 2022 Stack Overflow survey results [2] as a proxy for the “general” programming community, since Stack Overflow is a popular website used by millions of software developers. We will refer to this population as SO. We use the 2022 Haskell survey results [18] as a proxy for the functional programming (FP) community. We believe this population is interesting due to the proximity between the ITP and FP communities. We will refer to this population as HC. We chose the 2022 Haskell survey as comparison over other similar surveys, such as the OCaml 2022 survey [33], since it contains similar questions to ours and received many responses.

*Location* The United States is consistently one of the top-5 locations, as are Germany and the United Kingdom, for all three populations. As can be seen in Fig. 2, CC has a large ratio of French respondents, even outnumbering US respondents, likely for historical reasons since Coq was originally developed in France. A notable difference is that Indian respondents, numerous in the SO population, are rare in the FP and CC populations, indicating a missed opportunity for FP and ITP.

*Gender* There is no data for HC, but a comparison can be done between CC and SO. It shows the same trend of an overwhelming ratio of respondents identifying as men (about 90%), while only about 6% identify as women, and the remaining 4% as non-binary or other.

*Age* There is no data for the HC population, but SO and CC (Fig. 3) can be compared. Although the bins are slightly different between SO and CC (shifted by five years), the overall trend seems similar between the two populations. Very young programmers (less than 20 years old) are rare (about 5%), with the bulk of programmers being approximately between 20 and 40 years old. Older programmers (more than 50 years old) represent about 10% of both populations.

*Education* CC has fewer bins for education than either HC or SO (Fig. 4). We enable direct comparisons between surveys by merging bins where possible and dropping bins and corresponding answers otherwise, since bins are mutually exclusive. This analysis yields responses for no diploma for SO 3%, HC 2%, CC 1%; high school diploma for SO 26%, HC 17%, CC 5%; bachelor’s degree for SO 45%, HC 37%, CC 18%; master’s degree for SO 23%, HC 30%, CC 31%; PhD degree for SO 3%, HC 14%, CC 46%. These results place HC between

<sup>1</sup> <https://thzimmer.gitlabpages.inria.fr/coq-survey-2022-assets/asset-listing.html>.




Fig. 2 Top seven locations of respondents to the Coq Community 2022 survey

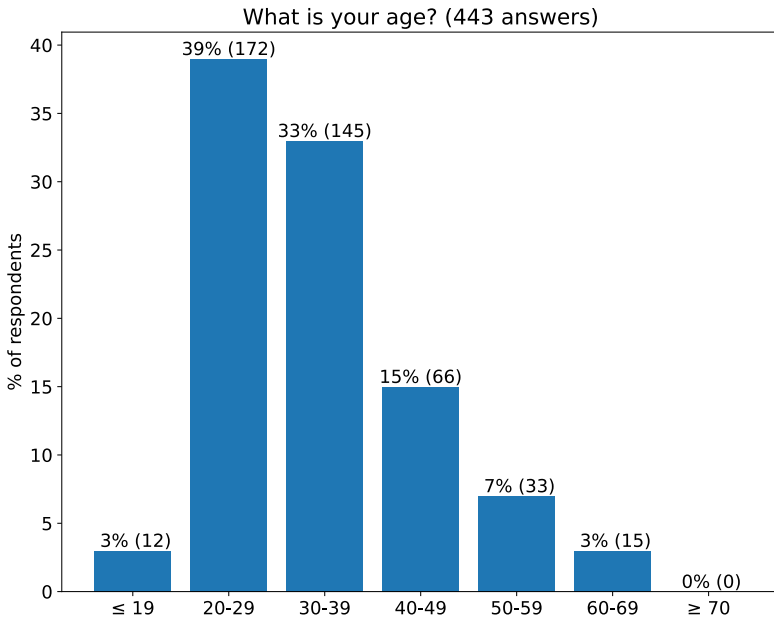



Fig. 3 Age ranges of respondents to the Coq Community Survey 2022





**Fig. 4** Highest education of respondents to the Coq Community Survey 2022

SO and CC in terms of educational achievement, but closer to SO given that CC has such a high skew towards doctoral degrees.

*Academic Use* We now compare the “academic” bins of the three surveys, even if their definitions are not exactly the same. In both the CC survey and the HC survey, respondents are asked if they operate in an academic context and if they are students or not. In the SO survey, academic researcher and student are exclusive categories. In CC, as can be seen in Fig. 5, 54% of respondents are academically employed, and 32% are students. In HC, 18% of the respondents use Haskell in an academic context and 15% are students. In comparison, only 4.42% of Stack Overflow survey respondents are academic researchers and 9.13% are students. In conclusion, CC and HC respondents operate in an academic context much more frequently than SO respondents. Even in comparison with the HC respondents, CC is even more tied to academia, where it represents about half of the respondents.

*IDEs* The bins are not the same across the surveys, but we can compare the following three main IDEs: Vim, Emacs, and VS Code. See Fig. 6 for IDE use in the CC population. VS Code is used by about 70% of SO, 43% of HC and 34% of CC. Emacs is used by about 5% of SO, 30% of HC and 60% of CC. Vim is used by about 23% of SO, 40% of HC and 9% of CC. When considering only Emacs, Vi(m), and VS Code, the Coq survey answers suggest that Coq users are significantly more likely to use Emacs than either Haskell users or developers on Stack Overflow. This is partly driven by long-time Coq users selecting Emacs (and its ProofGeneral package) and VS Code support for Coq being more recent.

*Operating systems* The bins are not the same across the surveys, but we can compare the following three main Operating Systems: Linux, Windows and macOS. For the SO survey, the best data is professional usage as it reflects the coding environment. See the data for CC in Fig. 7. Windows is used by about 49% of SO, 15% of HC and 18% of CC. macOS is used

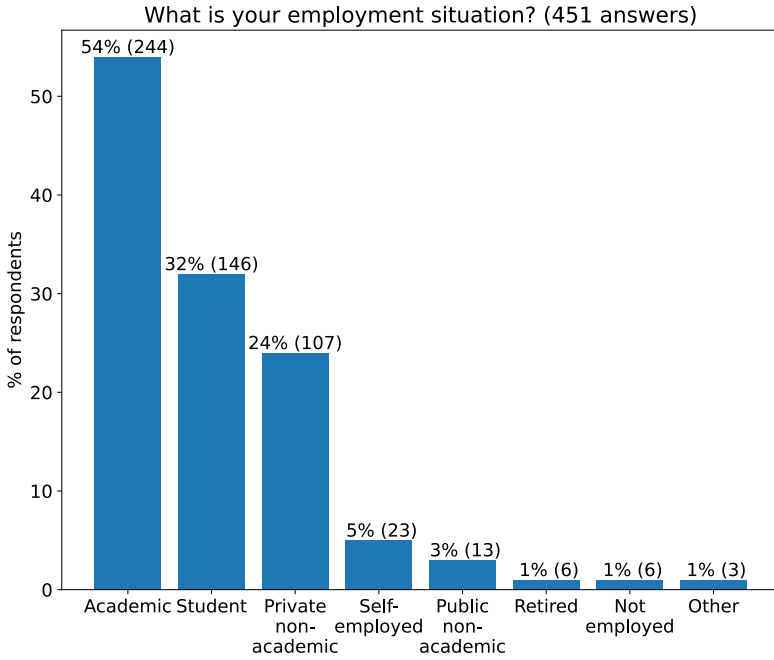


Fig. 5 Employment situation of the respondents to the Coq Community Survey 2022

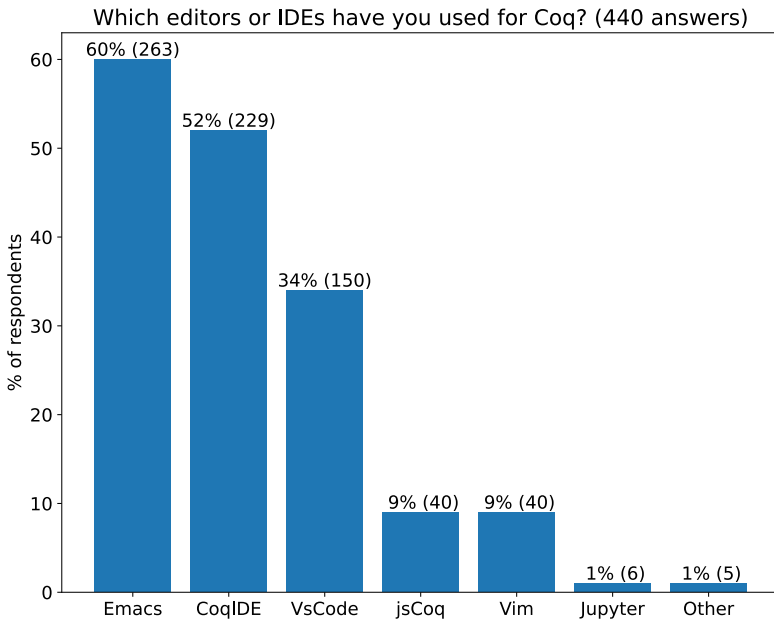
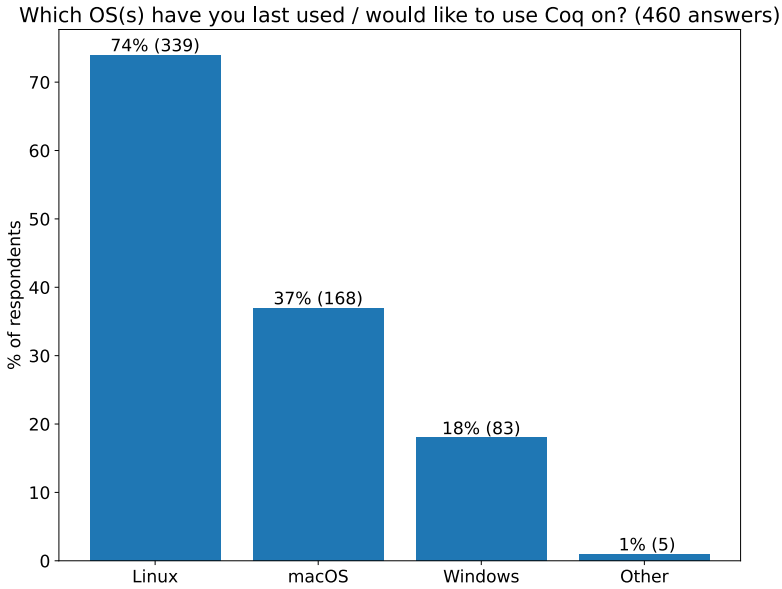


Fig. 6 IDE usage of respondents to the Coq Community Survey 2022



**Fig. 7** Operating System (OS) use of respondents to the Coq Community Survey 2022

by about 31% of SO, 33% of HC and 37% of CC. Linux is used by about 40% users of SO, 86% of HC and 74% of CC.

These percentages suggest that Linux-based operating systems are dominant among both Haskell and Coq users, with only a small minority of Windows users and a similarly-sized minority using macOS in both communities. In contrast, among Stack Overflow users, a near-majority use Windows. Windows use is slightly higher among Coq survey respondents than Haskell respondents, which is partly driven by new Coq users. For perspective on the dominance of Linux, note that Coq is embedded in the OCaml ecosystem, where Linux use is prevalent—more than 80% of respondents in the 2020 OCaml survey stated that they develop on Linux [20]. In addition, Coq and key packages in the official distribution of Coq are tested primarily on Linux [34].

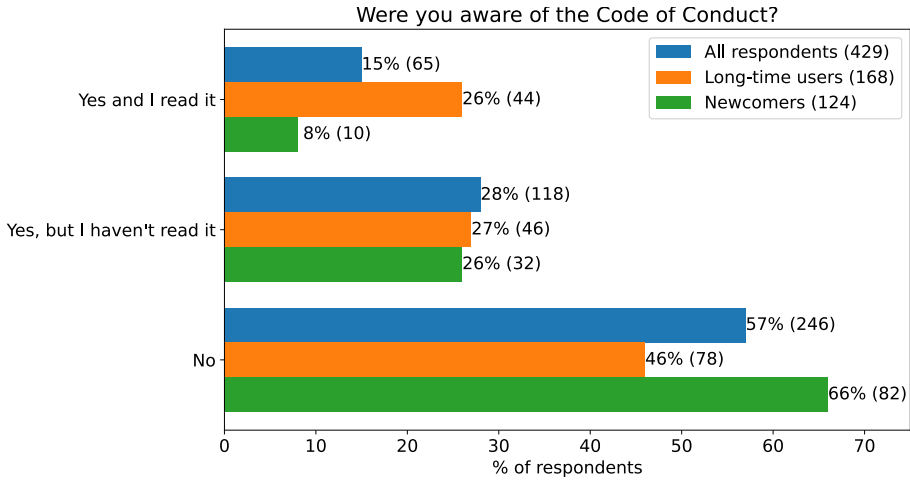
## 4 Community Health and Contributing Experience

Community health and contributing experience are vital factors for the success of any open-source project. In this section, we start by describing the state of the community health and then describe our results concerning the contributing experience.

### 4.1 Community Health

Many open source communities have adopted a Code of Conduct in recent years, and this phenomenon has received some attention of empirical software engineering research [35–38]. Coq adopted its Code of Conduct<sup>2</sup> in 2018. This code is available in the Coq repository and applies to “all spaces managed by the Coq development team” (including “the GitHub

<sup>2</sup> [https://github.com/coq/coq/blob/master/CODE\\_OF\\_CONDUCT.md](https://github.com/coq/coq/blob/master/CODE_OF_CONDUCT.md).



**Fig. 8** Code of conduct awareness, when separating responses from long-time users and from newcomers

repository, the Discourse forum, the Zulip chat, the mailing lists, physical events like Coq working groups and workshops”).

Yet, Tourani et al. [38] noted that, among communities where such a Code of Conduct has been adopted, awareness of this code varies quite a lot. Figure 8 shows the answers to the question *Where you aware of the Code of Conduct?*. A majority of respondents to the survey were not aware of its existence. This is particularly true of newcomers. While a small majority of long-time users were aware of the Code of Conduct, many of them had not read it. While the Code of Conduct was initially created to respond to a need that arose in the Coq GitHub project, its scope is supposed to be larger, but the answers to this question clearly show that it is not advertised broadly enough (e.g., until after the survey took place, there was no direct link to the Code of Conduct from the Coq website). An important lesson is that merely having a Code of Conduct is not sufficient: advertisement efforts have to be made so that the users know and understand it.

For the question *Have you ever, while contributing to Coq or participating in the Coq community (online or offline, on official or non-official channels), encountered situations in which you have felt unwelcome or treated unfairly?*, we obtained the answers depicted in Fig. 9. Very few respondents (16) have encountered such situations. Women and non-binary people are overrepresented among them (4), compared to their overall share among the respondents. Yet, we cannot hope to draw reliable conclusions from such small numbers. However, particular care should be undertaken to ensure the well-being of these under-represented groups. The survey continued by asking the respondents who had encountered such situations to provide some additional details about said situations. However, it didn’t include any open-text questions to avoid risks of re-identification of respondents. Instead, it advised that respondents could write to the Code of Conduct enforcement team to report these situations (even if they are ancient). However, this team did not receive any such report following the survey.

Finally, for the question *Where did these situations occur?* we obtained 14 answers among the 16 respondents who signaled being subjected to unfair treatment. These answers are depicted in Fig. 10. One important result to notice is that most of the situations do not happen inside the project repository but rather during in-person events or in forums.

Have you encountered situations in which you have felt unwelcome or treated unfairly?




Fig. 9 Users having encountered unwelcoming/unfairness situations

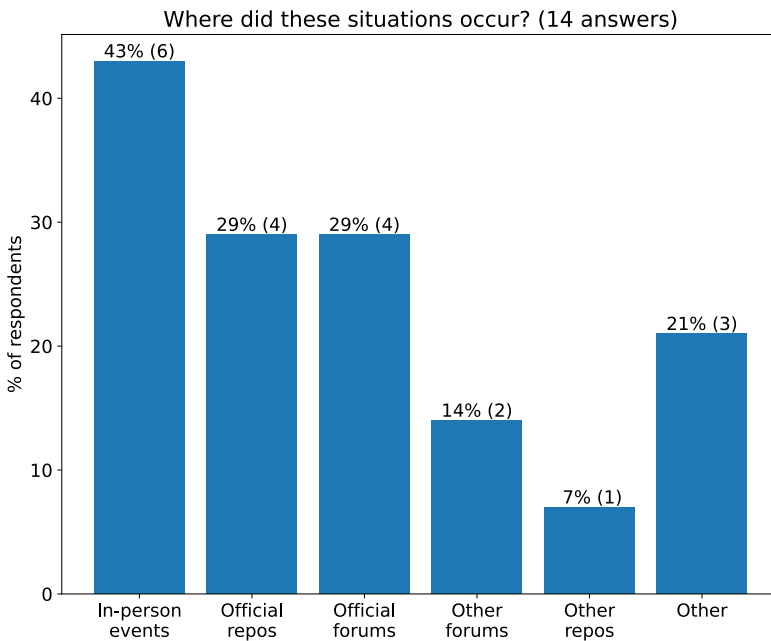


Fig. 10 Where users encountered unfairness situations

### 4.2 Contributing Experience

Figure 11 shows the answers to the question *Have you ever tried contributing changes to the Coq GitHub project, such as to libraries, documentation or code for bug fixes or improvements?* Unsurprisingly, newcomers have only rarely tried contributing to the Coq GitHub project, while a large minority of long-time users have tried contributing.

For the users who tried contributing, we asked the follow-up question *Did you accomplish what you wanted?*, to which 87% of respondents answered “Yes”. This high percentage is promising, but one should keep in mind that people who did not accomplish what they wanted might not have disclosed having tried.

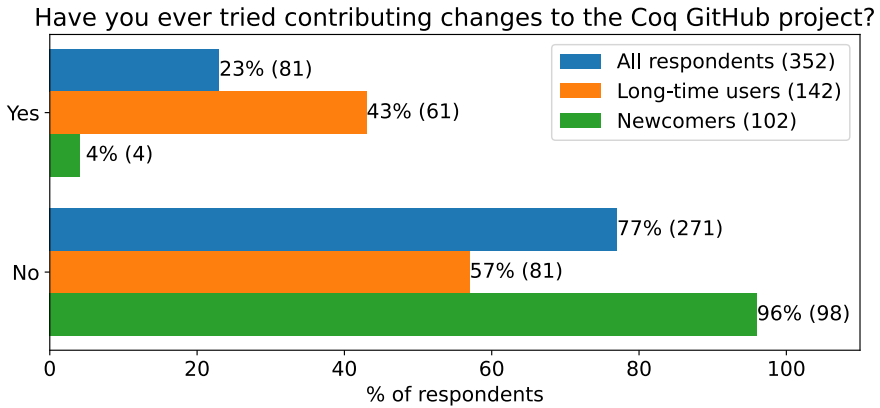


Fig. 11 Respondents having tried contributing changes to Coq, separated by categories of experience

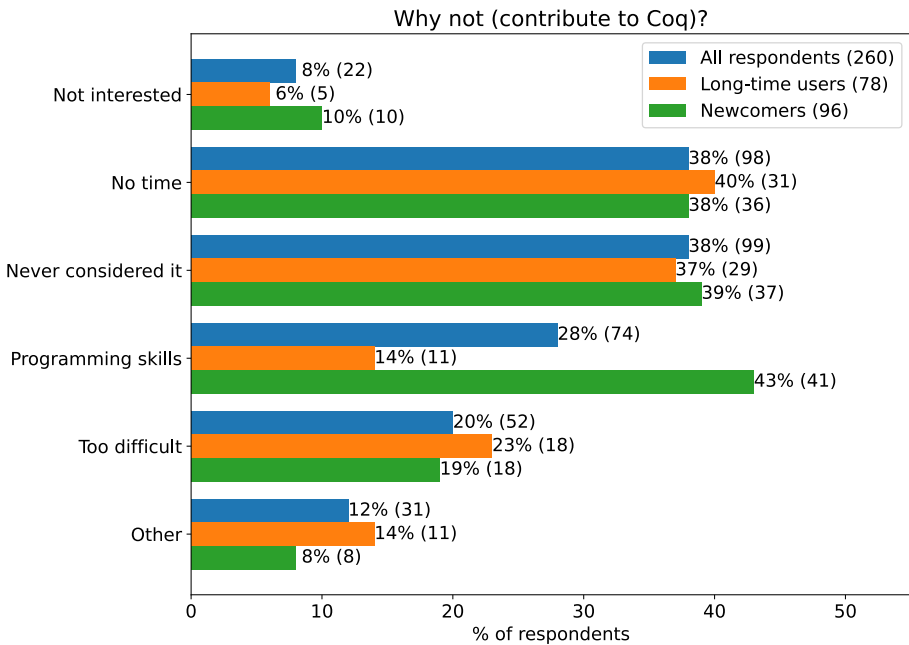
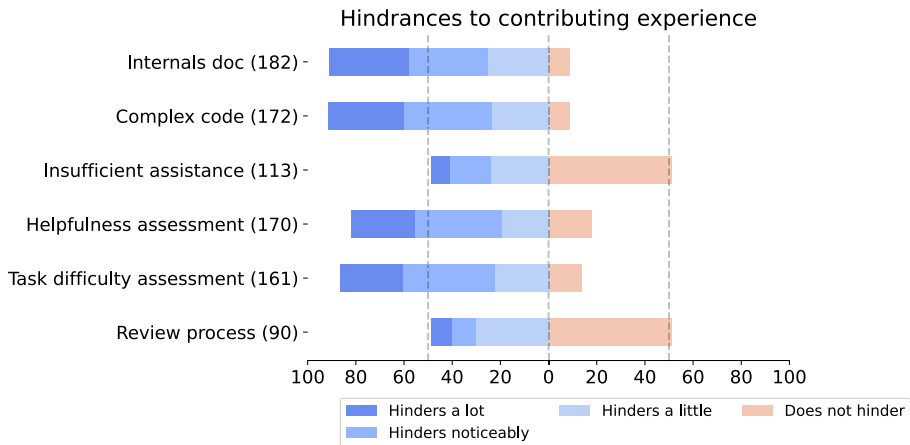


Fig. 12 Reasons for choosing not to contribute, separated by categories of experience

To the users that never tried to contribute, we asked the question *Why not?*, whose answers are depicted in Fig. 12. One of the main cited reasons is lack of time. An important phenomenon to notice is that many users did not try to contribute because they had never considered it, which suggests that the different ways of contributing to the project should be advertised more. The only reason not to contribute where we can see a large difference between newcomers and long-time users is “programming skills”, which could indicate that Coq has a steep learning curve.

Finally, we asked all users *How much do the following hinder your ability or desire to contribute changes to Coq?* The answers are depicted in Fig. 13. Of the six proposed answers,



**Fig. 13** Divergent bar plots representing the proportions of respondents answering whether each of the various proposed answers hinders contribution or not. The numbers shown in parentheses are the number of respondents for each of the proposed answers. We can see that the answers where the most respondents say “Does not hinder” are also those where the most respondents skip answering

only two (insufficient assistance and review process) have a (small) majority of respondents answering “Does not hinder”. On the other hand, the lack of internal documentation, the complexity of the code, and the difficulty in assessing which contribution would be most helpful and least difficult are all cited as noticeably hindering the contributing experience by a majority of respondents.

The survey also contained an open text question *Feel free to elaborate on the contributing experience, what we can do better, or why you do not contribute*, for which we performed a manual coding of the answers to identify recurring themes (this coding is available as part of our Zenodo artifact [19]).<sup>3</sup>

Most of the answers to this question were by people with no contributing experience. Of those who had contributed, 8 responses had positive comments and 6 responses had negative comments, and most were small contributions, such as fixing typos or adding to the documentation.

*Reasons for Not Contributing to Coq Combined with the Other* (short open text) answers to the question *Why not [contribute to Coq]?*, we can highlight a few recurring themes of why people do not contribute changes to the Coq project, together with illustrative examples of partial answers:

- Feel underqualified/intimidated (9 answers to the long text question and 10 *Other* answers to the why not contribute question).  
*“Even the idea of contributing is intimidating—if I feel like I don’t understand large parts of what happens when I use Coq, then I don’t feel like I can make meaningful contributions.”*
- Unclear which changes are needed/which changes would be likely to be accepted (5 answers to the long text question and 2 *Other* answers to the why not contribute question).  
*“It’s not clear to me what the missing parts of the standard-library are (if any) that I could possibly contribute to.”*

<sup>3</sup> NOTE TO REVIEWERS: we plan to update the Zenodo artifact with the new version of our qualitative coding for the final publication.

- Finally, in the *Other* answers to the why not contribute question, some respondents clarify that they plan to contribute (3 answers) or that they are already contributing in other ways (by opening bug reports, by contributing to the ecosystem), 2 respondents feel like the project is unwelcoming, and 8 respondents have not encountered the need to improve anything yet.

*Improving the Contribution Experience* In terms of improvement suggestions, we identified the following themes, and, as before, include illustrative partial quotes:

- Improving contributing documentation: 8 participants highlighted that a more comprehensive contributing documentation, especially targeting new contributors, would greatly help.  
*“Coq’s internals take many years to fully grasp. For newcomers, a guide through the code would be very useful (attempts at this exist, but do not seem to be sufficient).”*
- Clear roadmap: 3 participants mentioned that a clear roadmap on what are the expected new features and bug fixes could help contributors to decide on which topic to work on.  
*“Based on user input (e.g. this survey), bug reports and use case analysis by developers, etc., the Coq team should create and maintain a roadmap that identifies the improvements that will be most valuable to Coq users.”*
- Mentorship: 2 participants mentioned mentorship as a potential solution to understand some peculiarities of the internals of Coq and ease contribution for newcomers.  
*“Changing Coq itself is a completely different skill from using Coq, and so it must be learned separately. Having a teacher here is very helpful, and I don’t really see how to initially contribute without a dev holding your hand (additions to the stdlib excepted).”*

## 5 Analysis of Coq Use for Different Groups

In this section, we analyze how different population groups we identified in the survey use Coq differently, by looking at various outcomes.

### 5.1 Installation of Coq, Usage of Packages and Features, and CI

In this section, we describe and analyze the results shown in Tables 1 and 2.

#### 5.1.1 Variables

*opam* A binary variable indicating whether respondents installed Coq with opam (a software package manager for OCaml [39] and for Coq as well [40]), either directly, or relying on wrapper scripts provided by the Coq Platform [34]. Because Coq has been around for so long, there are many possible ways to install it. It is packaged in many Linux distributions and generic package managers, and there are binary installers for the main operating systems. The most flexible installation method is by using opam, because it gives access to the whole ecosystem of Coq packages. This is the most common installation method (used by 76% of respondents), but it is also more complex, so the expectation is that users start relying on it only when they encounter the need.

*External General Libraries* A binary variable for whether respondents are “casual” or “advanced” users of any of the external general libraries Coq-Extlib, Coq-std++, Math Classes, Mathematical Components, or TLC. Coq is distributed with a standard library and



**Table 1** Usage of opam, ecosystem libraries, tools, and automation

Variables	(1) opam	(2) External general libraries	(3) Number of external tools	(4) Automation
> 2 years of experience	<b>0.183***</b> (0.0576) [0.008]	0.107 (0.0663) [0.253]	-0.118 (0.233) [0.818]	0.110 (0.0620) [0.253]
> 5 years of experience	0.00177 (0.0594) [0.967]	<b>0.169**</b> (0.0684) [0.045]	<b>0.945***</b> (0.241) [0.001]	<b>0.186**</b> (0.0639) [0.016]
Learner	- <b>0.282***</b> (0.0619) [0.001]	- 0.0882 (0.0712) [0.685]	0.193 (0.250) [0.872]	-0.175 (0.0665) [0.116]
Software verification	0.0470 (0.0473) [0.533]	<b>0.165**</b> (0.0544) [0.013]	<b>0.742***</b> (0.191) [0.002]	<b>0.192***</b> (0.0509) [0.002]
North America	0.00277 (0.0519) [1.000]	0.00986 (0.0597) [1.000]	0.0795 (0.210) [0.999]	-0.125 (0.0558) [0.170]
Other locations	- 0.00968 (0.0617) [0.960]	- 0.124 (0.0710) [0.358]	- 0.0873 (0.250) [0.956]	- 0.0269 (0.0663) [0.956]
Additional controls	Yes	Yes	Yes	Yes
Mean of dep. variable	0.758	0.366	0.995	0.702
Observations	372	372	372	372

Each column corresponds to a regression. Standard errors are in parentheses, Romano-Wolf adjusted p-values in brackets. Bolded text with stars indicate different levels of significance (\*\*\*for p < 0.01, \*\*for p < 0.05, and \*for p < 0.1) according to Romano-Wolf adjusted p-values. Independent variables are described in Sect. 2.3.1 and additional controls in Sect. 2.3.2. Mean of dep. (dependent) variable corresponds to the mean of the outcome of each column

**Table 2** Usage of SSReflect, continuous integration (CI), and extraction

Variables	(5) SSReflect	(6) CI	(7) Extraction
> 2 years of experience	−0.0196 (0.0644) [0.818]	0.0512 (0.0600) [0.720]	<b>0.231***</b> (0.0653) [0.005]
> 5 years of experience	<b>0.138*</b> (0.0664) [0.059]	<b>0.320***</b> (0.0619) [0.001]	<b>0.169**</b> (0.0673) [0.045]
Learner	−0.0363 (0.0691) [0.872]	− <b>0.197**</b> (0.0645) [0.042]	−0.0513 (0.0701) [0.872]
Software verification	<b>0.126*</b> (0.0528) [0.051]	0.0371 (0.0493) [0.533]	<b>0.290***</b> (0.0536) [0.001]
North America	−0.0637 (0.0580) [0.860]	0.00597 (0.0541) [1.000]	0.00654 (0.0588) [1.000]
Other locations	−0.0673 (0.0689) [0.802]	−0.122 (0.0643) [0.290]	<b>0.172*</b> (0.0699) [0.089]
Additional controls	Yes	Yes	Yes
Mean of dep. variable	0.288	0.315	0.446
Observations	372	372	372

Each column corresponds to a regression. Standard errors are in parentheses, Romano-Wolf adjusted p-values in brackets. Bolded text with stars indicate different levels of significance (\*\*for  $p < 0.01$ , \*\*for  $p < 0.05$ , and \*for  $p < 0.1$ ) according to Romano-Wolf adjusted p-values. Independent variables are described in Sect. 2.3.1 and additional controls in Sect. 2.3.2. Mean of dep. (dependent) variable corresponds to the mean of the outcome of each column

some embedded tools, but it also comes with a rich package ecosystem. Power users rely on this package ecosystem to be more productive. With this outcome and the next one, we evaluate how different types of users take advantage (or not) of the package ecosystem. Only 37% of respondents are users of a general external library.

*Number of External Tools* A natural number variable counting how many external tools respondents said they were using (“Casual user” or “Advanced user”) among a long list: Mtac2, MetaCoq, and Coq-Elpi (from the “Tactic languages” question), as well as all 16 tools from the “Ecosystem plugins and tools” question. Respondents to the survey only use 0.995 of the 19 listed external tools on average.

*Automation* A binary variable for whether respondents are “casual” or “advanced” users of any of several core automation tools (Micromega, Nsatz, Ring, and the solvers for logic and equality from the standard library, such as firstorder) or external automation tools (AAC Tactics, CoqHammer, and SMTCoq). 70% of respondents are using such a tool.

*SSReflect* A binary variable for whether respondents are “casual” or “advanced” users of the SSReflect proof language. SSReflect is a consistent set of tactics providing an alternative to

the standard tactics. It has been distributed as part of the main Coq package for more than 5 years. 29% of our respondents use SSReflect.

*CI* A binary variable for whether respondents use Continuous Integration (CI) on any of their projects. CI is the practice of automatically running tests at every change when developing software. This is a very common practice today, which helps to produce better quality software [41]. CI can also be used on mechanized proof projects, usually not to run tests, but simply to verify that the proofs are still accepted. This is particularly useful when developing Coq libraries, because CI can help ensure that a library remains compatible with several consecutive versions of Coq, even though its maintainers only test one version on their machines. Because of this main application, it is expected that respondents who do not yet have big enough projects will not use CI. 31% of our respondents declared they do.

*Extraction* A binary variable for whether respondents are “casual” or “advanced” users of any of the officially supported extraction targets. Extraction is a Coq feature allowing the automatic generation of executable code in OCaml, Haskell or Scheme from a program written in Gallina, the functional programming language embedded inside Coq. Since Coq can be used to prove a Gallina program correct with respect to its specification, extraction is a useful mechanism to produce executable code for this program, and it is expected to be strongly associated with using Coq for verifying software. 45% of our respondents said they use extraction.

### 5.1.2 Results on Experience Level

We observe statistically significant results relating experience to each of our outcomes. For most of our outcomes (external general libraries, number of external tools, automation, SSReflect and CI), the threshold to see a statistically significant difference is to reach more than 5 years of experience. The most important jump being in the use of CI, which is of the same magnitude as the mean of the variable over our population.

Sometimes, 2 years of experience are enough to see significant differences: the use of opam increases (18 percentage points (pp) more respondents using opam) at the 2 years of experience threshold, and the use of extraction increases twice, at each of the two experience thresholds (23 pp more respondents using extraction after 2 years, and 17 pp more respondents using extraction after 5 years).

### 5.1.3 Results on Learners

Learners use opam significantly less often (28 pp fewer respondents use opam when considering learners compared to non-learners with every other independent variable fixed, including experience). This may be explained by learners having a stronger need for an easy installation method and a lesser need for the access to packages that installing with opam provides. The only other statistically significant difference robust to multiple hypothesis testing is in the use of CI. 20 pp fewer learners use CI for their projects compared to non-learners, which can be explained by learners not having any project that really requires the use of CI. Other coefficients are estimated to be negative as well, but are not statistically significant or not robust to multiple hypothesis testing corrections.

### 5.1.4 Results on Software Verification

Besides experience level, the other category of users associated with many statistically significant differences in their use of Coq are respondents applying Coq to do software verification. These users are more likely to use external general libraries (16 additional pp), they use more external tools (an average of 0.74 additional tools, which is close in magnitude to the average number of external tools used by our respondents), they are more likely to use automation, more likely to use SSReflect, and more likely to use extraction. In short, we can say that these respondents (who represent about 67% of our respondents) are the ones who make the most use of the Coq package ecosystem and features. This may not come as a surprise for, e.g., the use of extraction, which is a feature particularly suited to software verification, but may be more surprising when talking about the use of external general libraries or SSReflect, given that MathComp (one of these general libraries) and SSReflect were not created to verify software, but to formalize pure mathematics.

Beyond the previously listed significant results, we can also note a significant, but difficult to interpret difference on the use of extraction by respondents from other locations (excluding Europe and North America).

## 5.2 User Interfaces

We relate the probability to have used one of four Integrated Development Environment (IDEs) for Coq (Emacs, CoqIDE, VS Code and jsCoq) to our previously defined categories. An IDE is a fundamental part of proof assistants such as Coq, as it is the interaction medium between users and the system.

Our survey included many IDE questions, meant to get a better idea of their user base, their limitations, and to be able to inform stakeholders. In Table 3, we look at the responses to the question “Which editors or IDEs have you used for Coq?” and relate the results for four of the most used IDEs to our previously defined categories of respondents. Among the four selected IDEs, the two with the most users according to the survey results are Emacs and CoqIDE (used respectively by 61 and 52% of respondents). This is explained by observing that they were the only two available options for many years. Recently, more focus has been put on developing more modern user interfaces, e.g., based on the very popular Visual Studio Code (VS Code, used by 36% of respondents), or directly in the browser, with jsCoq (used by 9% of respondents).

Unfortunately, most of the significant results obtained in this table were not robust to Romano-Wolf corrections for multiple hypothesis testing. The only two robust results are:

1. Users with more than 2 years of experience are more likely to have used CoqIDE. This can be interpreted by CoqIDE being the most standard user interface for Coq, since it is distributed along with the system, and the question being formulated as “Which editors or IDEs have you used for Coq?” and not “Which editors or IDEs are you *currently using* for Coq?”. Thus, users are very likely to have used CoqIDE at some point when they are experienced enough.
2. Users in North America are 13 pp more likely to have used Emacs, compared to the Europe baseline.

We do not get any other statistically significant results, although some are not very far from being significant. In particular, the learner category looks like it could be more frequently using jsCoq, but it is very difficult to draw reliable conclusions given the small number of jsCoq users among our respondents (40).

**Table 3** Usage of integrated development environment (IDE)

Variables	(1) Emacs	(2) CoqIDE	(3) VS Code	(4) jsCoq
> 2 years of experience	0.0164 (0.0651) [0.991]	<b>0.220***</b> (0.0679) [0.009]	0.00848 (0.0679) [0.991]	0.0135 (0.0418) [0.991]
> 5 years of experience	0.0778 (0.0672) [0.688]	-0.0318 (0.0700) [0.963]	-0.0134 (0.0700) [0.963]	-0.0146 (0.0431) [0.963]
Learner	-0.109 (0.0699) [0.265]	0.124 (0.0729) [0.265]	-0.0897 (0.0729) [0.265]	0.0963 (0.0448) [0.147]
Software verification	0.120 (0.0535) [0.118]	0.0376 (0.0557) [0.515]	-0.110 (0.0557) [0.170]	0.0431 (0.0343) [0.398]
North America	<b>0.132*</b> (0.0587) [0.077]	-0.117 (0.0612) [0.145]	-0.00667 (0.0612) [0.899]	-0.0416 (0.0376) [0.458]
Other location	-0.0743 (0.0697) [0.629]	-0.0147 (0.0727) [0.976]	-0.0140 (0.0726) [0.976]	-0.0759 (0.0447) [0.250]
Additional controls	Yes	Yes	Yes	Yes
Mean of dependent variable	0.608	0.516	0.360	0.091
Observations	372	372	372	372

Each column corresponds to a regression. Standard errors are in parentheses, Romano-Wolf adjusted p-values in brackets. Bolded text with stars indicate different levels of significance (\*\*\*for  $p < 0.01$ , \*\*for  $p < 0.05$ , and \*for  $p < 0.1$ ) according to Romano-Wolf adjusted p-values. Independent variables are described in Sect. 2.3.1 and additional controls in Sect. 2.3.2. Mean of dep. (dependent) variable corresponds to the mean of the outcome of each column

## 6 Satisfaction and Needs of Different Groups

In this section, we look at how our different population groups perceive Coq and what needs they express. This should inform researchers willing to improve proof assistants for specific use cases or target audiences.

We base our analysis on the responses to the satisfaction questions for the six IDEs that we included additional questions about, as well as one question on the importance of various improvements to be possibly made to Coq, one question asking what additional extraction targets respondents would like to have, and one question asking what additional languages to support in the documentation. Figure 14 presents the descriptive results to the question about improvements to be possibly made to Coq. We next describe the variables and results from Table 4.


### 6.1 Variables

*IDE Satisfaction* A natural number variable defined as the maximum IDE satisfaction level (from 0 to 4) for each respondent. We say that the IDEs each respondent assigned a maximum

**Table 4** Expressed satisfaction and needs

Variables	(1) IDE satisfaction	(2) Essential tech. vs comm. improvements	(3) New extraction targets	(4) Other languages
> 2 years of experience	0.116 (0.105) [0.632]	0.105 (0.185) [0.826]	-0.0154 (0.0510) [0.826]	-0.0483 (0.0406) [0.632]
> 5 years of experience	- <b>0.236*</b> (0.108) [0.085]	<b>0.484**</b> (0.191) [0.047]	-0.00963 (0.0526) [0.976]	-0.00944 (0.0419) [0.976]
Learner	-0.224 (0.112) [0.159]	-0.0625 (0.199) [0.936]	0.0519 (0.0548) [0.704]	-0.00920 (0.0436) [0.936]
Software verification	0.0429 (0.0856) [0.624]	0.187 (0.152) [0.545]	<b>0.163***</b> (0.0419) [0.002]	0.0340 (0.0333) [0.559]
North America	0.0721 (0.0942) [0.779]	-0.286 (0.167) [0.265]	-0.00620 (0.0460) [0.982]	0.000884 (0.0366) [0.982]
Other locations	0.00952 (0.113) [0.923]	-0.219 (0.198) [0.415]	0.0872 (0.0546) [0.304]	<b>0.122**</b> (0.0434) [0.038]
Additional controls	Yes	Yes	Yes	Yes
Mean of dep. variable	3.049	0.325	0.153	0.089
Observations	367	372	372	372

Each column corresponds to a regression. Standard errors are in parentheses, Romano-Wolf adjusted p-values in brackets. Bolded text with stars indicate different levels of significance (\*\* for  $p < 0.01$ , \*\*\* for  $p < 0.001$ , \* for  $p < 0.05$ , and \*\* for  $p < 0.1$ ) according to Romano-Wolf adjusted p-values. Independent variables are described in Sect. 2.3.1 and additional controls in Sect. 2.3.2. Mean of dep. (dependent) variable corresponds to the mean of the outcome of each column



**Fig. 14** Responses to the question “In order to make you more productive in Coq and to encourage others to learn and use Coq, how important are improvements in the following areas (relative to their current state)?”. On the right hand-side, the percentages correspond, in order, to the proportion of “Unwise”, “Unimportant”, “Worthwhile”, and “Essential” answers. Improvement categories are ordered according to the majority judgment methodology [42]

satisfaction level to are their favorites. To be defined, this outcome needs respondents to have answered at least one IDE satisfaction question, which is why we have 5 fewer observations for this outcome. The average satisfaction level given to the favorite IDE is 3, which corresponds to “Satisfied”.

*Essential Technical vs Community Improvements* An integer number variable defined as the difference between the number of essential technical improvements and the number of essential community improvements according to each respondent. We separate Coq improvements (from the question of Fig. 14) into two categories: technical improvements (fixing bugs, refining features, IDEs, development tools, new features, performance) and community improvements (documentation, libraries, community building, support channels). We exclude the “Compatibility/migration support” improvement because it could be perceived as belonging to any of the two categories depending on how “support” is interpreted. We count for each of these categories the number of essential improvements, and we take the difference. Thus, a respondent with a positive value for this variable is listing more technical improvements as essential, and a respondent with a negative value for this variable is listing more community improvements as essential. On average, respondents listed 0.3 additional essential technical improvements compared to community improvements, but note that the list of technical improvements contained two more items than the list of community improvements.

*New Extraction Targets* A binary variable for whether respondents provided a non-empty answer to the question on new extraction targets. 15% of the respondents asked for new extraction targets using this open text question.

*Other Languages* A binary variable for whether respondents provided a non-empty answer to the question on languages to support in the documentation. 9% of our respondents asked for documentation in languages other than English using this open text question.

## 6.2 Results on Experience Level

Respondents with more than 5 years of experience are significantly more likely to assign a lower satisfaction score (0.24 lower on average) to their favorite IDE. They also list more essential technical improvements than community improvements, compared to other respondents (a 0.48 average difference).

We can explain experienced respondents being less satisfied with their IDE and more adamant about technical improvements by them having learned limitations of their IDEs or other technical limitations during their years of experience.

## 6.3 Other Results

The other two statistically significant differences (robust to multiple hypothesis testing) are:

1. Respondents using Coq for software verification are more likely to request new extraction targets (an additional 0.16 pp, which is of the same magnitude as the mean of the variable over our population). Once again, this is not surprising as extraction is particularly suited to software verification, but it is limited by the restricted number of programming languages being currently officially supported as extraction targets.
2. Respondents in other locations (outside Europe and North America) are more likely to request support for new languages in the documentation (an additional 0.12 pp, which is largely superior to the mean of the variable over our population).

## 7 Threats to Validity

As with any empirical study, the validity of our conclusions may be threatened by various factors. In the following discussion we adhere to the well-established threats to validity framework by Shadish et al. [43]. We acknowledge the ongoing debate regarding trade-offs in research design and the associated threats to validity across many scientific disciplines [44–47]. However, in the absence of a widely accepted alternative to the threats to validity framework by Shadish, Cook, and Campbell, we adhere to this framework but recognize its limitations [48, 49].

In terms of external validity (i.e., the validity of applying the conclusions of our study outside the specific context), there are two major concerns. First, it is possible that the Coq community behaves differently from other ITP communities. This can be due both to technical reasons (e.g., different proof assistants may be best suited to different use cases) and to sociological reasons (“birds of a feather flock together”), as could be observed, e.g., with the success of Lean among mathematicians. Second, even within the Coq community, the people who responded to the survey did not form a random sample of the community. It is likely that survey respondents had characteristics different from the average Coq user. In particular, and given the length of the survey, only people with enough time and motivation completed the survey in full. The results then concern this subsample of users and may not be extrapolated to the whole community. We took measures to avoid selection bias due to the use of the English language, but by only making the survey available in a second language (Chinese), and not many more, we are likely to have missed potential respondents in other linguistic communities. We tried to reach Coq users broadly by advertising the survey on many forums and mailing lists, including language-specific ones, on social networks, and on the Coq website, but each choice of platform necessarily introduced bias in the sample.



There is no way to completely avoid that beyond trying to reach even more broadly. For instance, we tried to reach learners as well by asking teachers to advertise the survey to their students, but we do not know how many did so. Our analyses of how different categories of respondents use Coq differently or express different needs reduce the issues of having a biased sample by including many controls, but we very likely missed some important ones.

Beyond external validity, there are other factors that can lead to misleading results. For example, the relatively low number of respondents may reduce the precision of the estimates, increase confidence intervals and reduce the power of the analysis (i.e., the ability to find a statistically significant effect when it exists). The order in which questions were asked may have affected the answers, which is a well-known phenomenon [50]. It was to mitigate this kind of issue that we decided, for example, to place demographics questions last. Respondents may have misunderstood some questions or questions options. We took many measures to make questions as clear as possible (using simple English, collaboratively writing and proofreading questions, using beta testers), but, despite that, we became aware after launching the survey that a few questions were still ambiguous, or did not have the originally intended meaning.

## 8 Discussion and Conclusions

We presented the design, deployment, and summaries and analyses of the responses to the Coq Community Survey 2022. Besides providing up-to-date descriptive data on the Coq community and decision support for the Coq Team, we hope our work can be useful for constructing future surveys targeted at ITP communities.

*Data Retention* Planning, running, and analyzing the survey was work intensive. The WG decided to announce a retention period of one year for the raw survey response data up front. However, this turned out to be a limiting decision in terms of the WG's time and resources. Keeping raw data is a delicate matter, since respondents and their answers may sometimes be identified by correlating survey response data with data from other sources. Nevertheless, we believe asking respondents for permission to store their data for more than one year is warranted, due to the time required for rigorous response analysis.

*Responder Retention* An important problem with recurring surveys is responder retention, as highlighted by the significant drop in the number of responses to the OCaml community survey between 2020 and 2022 [33]. Three years or more may be needed between surveys in an ITP community to ensure there will be enough responses to repeat previous analyses in a reliable way. In a small community, if filling in a survey becomes perceived as time-consuming routine work of little value, many members may not volunteer for the task. Making most or all questions optional does not address this issue.

*Renaming of Coq* This paper focuses on results from the survey relevant to ITP researchers in general. Hence, we do not discuss in detail the questions and results related to Coq's name and renaming of Coq. However, summaries on responses related to the name are available in the WG's final blog post [17]. In brief, our only significant finding is that respondents that want to change Coq's name tend to reside in English-speaking regions.

*Future Work* During analysis of tentative future Coq user surveys, changes could be tracked for recurring questions and results of our regression analyses, e.g., on users performing software verification, could be replicated. Since our analyses indicate that experience has a large impact on Coq user behavior, future research could investigate in more detail how users

gain experience, and the effectiveness of their methods to improve. Another avenue is to look in detail at methodological problems that users face when applying Coq in their domain, and how users develop workarounds and solutions to these problems.

## Supplementary information

As part of the supplementary material to our paper [19], we provide (1) the survey definition in LimeSurvey and HTML format, (2) the Jupyter Notebook and Stata code we used to analyze the response data, (3) plots of answers to all closed questions and plots of some interactions between questions, (4) manual analysis of some open-text questions, and (5) answers to open-text questions. In particular, the code in our supplementary material shows how we generated the plots in this paper. However, for GDPR conformance, we do not include the raw response data that could fully replicate plots and other analyses.

**Acknowledgements** The work was partially funded by the Swedish Foundation for Strategic Research under the project Trustfull.

**Author Contributions** All authors jointly wrote the manuscript text. All authors reviewed the manuscript. J.-R. F. and T.Z. wrote the Jupyter code to analyze the survey data. A. C. A. wrote the Stata code to perform regression analysis of survey data.

**Funding** Open access funding provided by Royal Institute of Technology.

**Data Availability** As part of the supplementary material to our paper (<https://doi.org/10.5281/zenodo.7930567>), we provide (1) the survey definition in LimeSurvey and HTML format, (2) the Jupyter Notebook and Stata code we used to analyze the response data, (3) plots of answers to all closed questions and plots of some interactions between questions, (4) manual analysis of some open-text questions, and (5) answers to open-text questions. In particular, the code in our supplementary material shows how we generated the plots in this paper. However, for GDPR conformance, we do not include the raw response data that could fully replicate plots and other analyses.

## Declarations

**Conflict of interest** The authors have no conflict of interest to declare that are relevant to the content of this article.

**Ethical Approval** The authors contacted Inria's Ethics Review Board (ERB) before survey deployment. The ERB informed the authors that full ERB approval was not required and referred to Inria's Data Protection Officer for compliance matters. The authors developed a GDPR conformance statement and share only aggregated survey data including open text answers from which information that could identify respondents was removed.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Greif, S., Burel, E.: State of JS (2022). <https://2022.stateofjs.com/en-US/>

2. Stack Overflow: Stack Overflow Developer Survey (2022). <https://survey.stackoverflow.co/2022/>
3. Wessel, M., Serebrenik, A., Wiese, I., Steinmacher, I., Gerosa, M.A.: What to expect from code review bots on GitHub? A survey with OSS maintainers. In: Proceedings of the XXXIV Brazilian Symposium on Software Engineering. SBES '20, pp. 457–462. Association for Computing Machinery, New York (2020). <https://doi.org/10.1145/3422392.3422459>
4. O'Grady, S.: The RedMonk Programming Language Rankings: June 2022 (2022). <https://redmonk.com/sogrady/2022/10/20/language-rankings-6-22/>
5. Braibant, T.: Coq survey (2014). <https://github.com/braibant/coq-survey/blob/master/popl-coq.pdf>
6. LimeSurvey Developers: LimeSurvey (2024). <https://www.limesurvey.org>
7. Appel, A.W.: Coq's vibrant ecosystem for verification engineering (invited talk). In: Proceedings of the 11th ACM SIGPLAN International Conference on Certified Programs and Proofs. CPP 2022, pp. 2–11. Association for Computing Machinery, New York (2022). <https://doi.org/10.1145/3497775.3503951>
8. Wilcox, J.R.: Why is the Coq logo made to look like a penis? (2021). <https://sympa.inria.fr/sympa/arc/coq-club/2021-04/msg00006.html>
9. The Coq Survey Working Group: Coq Community Survey 2022 in English (2022). [https://thzimmer.gitlabpages.inria.fr/coq-survey-2022-assets/LimeSurvey/questionnaire\\_356388\\_en.html](https://thzimmer.gitlabpages.inria.fr/coq-survey-2022-assets/LimeSurvey/questionnaire_356388_en.html)
10. The Coq Survey Working Group: Coq Community Survey 2022 in Chinese (2022). [https://thzimmer.gitlabpages.inria.fr/coq-survey-2022-assets/LimeSurvey/questionnaire\\_chinese.html](https://thzimmer.gitlabpages.inria.fr/coq-survey-2022-assets/LimeSurvey/questionnaire_chinese.html)
11. The Coq Survey Working Group: Coq Community Survey 2022 Results: Part I (blog post). (Who is using Coq and in what context?) (2022). <https://coq.discourse.group/t/coq-community-survey-2022-results-part-i/1730>
12. The Coq Survey Working Group: Coq Community Survey 2022 Results: Part II (blog post). (How people are using Coq?—OS, IDEs, CI/CD) (2022). <https://coq.discourse.group/t/coq-community-survey-2022-results-part-ii/1746>
13. The Coq Survey Working Group: Coq Community Survey 2022 Results: Part III (blog post). (How is Coq used?—features, tools, libraries) (2022). <https://coq.discourse.group/t/coq-community-survey-2022-results-part-iii/1777>
14. de Almeida Borges, A., Falleri, J.-R., Fehrlé, J., Gallego Arias, E.J., Martin-Dorel, É., Palmiskog, K., Serebrenik, A., Zimmermann, T.: Coq Community Survey 2022: Summary of Results, abstract. The Coq Workshop 2022 (2022). <https://coq-workshop.gitlab.io/2022/abstracts/Coq2022-04-01-community-survey.pdf>
15. de Almeida Borges, A., Falleri, J.-R., Fehrlé, J., Gallego Arias, E.J., Martin-Dorel, É., Palmiskog, K., Serebrenik, A., Zimmermann, T.: Coq Community Survey 2022: Summary of Results, slides. The Coq Workshop 2022 (2022). <https://coq-workshop.gitlab.io/2022/slides/Coq2022-04-01-community-survey.pdf>
16. Almeida Borges, A., Casanueva Artís, A., Falleri, J.-R., Gallego Arias, E.J., Martin-Dorel, E., Palmiskog, K., Serebrenik, A., Zimmermann, T.: Lessons for interactive theorem proving researchers from a survey of Coq users. In: Naumowicz, A., Thiemann, R. (eds.) 14th International Conference on Interactive Theorem Proving (ITP 2023). Leibniz International Proceedings in Informatics (LIPIcs), vol. 268, pp. 12–11218. Schloss Dagstuhl—Leibniz-Zentrum für Informatik, Dagstuhl (2023). <https://doi.org/10.4230/LIPIcs.ITP.2023.12>
17. Coq Survey Working Group, T.: Coq Community Survey 2022 Results: Part IV (blog post). (Improvements to Coq, contributing to Coq, community health, and renaming Coq) (2023). <https://coq.discourse.group/t/coq-community-survey-2022-results-part-iv-and-ity-paper-announcement/2001>
18. Fausak, T.: State of Haskell Survey (2022). <https://taylor.fausak.me/2022/11/18/haskell-survey-results/>
19. de Almeida Borges, A., Casanueva Artís, A., Falleri, J.-R., Gallego Arias, E.J., Martin-Dorel, É., Palmiskog, K., Serebrenik, A., Zimmermann, T.: Supplementary material for the article “lessons for interactive theorem proving researchers from a survey of Coq users”. Zenodo (2023). <https://doi.org/10.5281/zenodo.7930567>
20. OCaml Software Foundation: OCaml User Survey (2020). <https://discuss.ocaml.org/t/suggestions-from-the-ocaml-survey-result/6791>
21. Mitchell, M.L., Jolley, J.M.: Research Design Explained, 7th edn. Cengage Learning, Wadsworth (2010)
22. Pierce, B.C., Amorim, A.A., Casinghino, C., Gaboardi, M., Greenberg, M., Hritcu, C., Sjöberg, V., Yorgey, B.: Logical Foundations. Software Foundations, vol. 1. Electronic textbook, Online (2022). Chinese translation, version 5.7. <https://coq-zh.github.io/SF-zh/lf-current/index.html>
23. Pierce, B.C., Casinghino, C., Greenberg, M., Sjöberg, V., Yorgey, B.: Software Foundations. Electronic textbook, Online (2011). Japanese translation. <http://proofcafe.org/sf/>. Accessed 27 May 2024
24. TYPES community: Types-announce mailing list (2024). <https://lists.seas.upenn.edu/mailman/listinfo/types-announce>

25. GDR IM: GDR IM (Informatics and Mathematics National Working Group) mailing list (2024). <https://mail.gdr-im.fr/sympa>
26. GDR GPL: GDR GPL (Software and Programming National Working Group) mailing list (2024). <https://listes.univ-grenoble-alpes.fr/sympa/info/gdr.gpl>
27. Project Jupyter: Jupyter (2024). <https://jupyter.org>
28. StataCorp: Stata (2024). <https://www.stata.com>
29. Wooldridge, J.M.: *Introductory Econometrics: A Modern Approach*, 6th edn. Cengage Learning, Mason (2015)
30. Romano, J.P., Wolf, M.: Stepwise multiple testing as formalized data snooping. *Econometrica* **73**(4), 1237–1282 (2005)
31. Clarke, D., Romano, J.P., Wolf, M.: The Romano-Wolf multiple-hypothesis correction in Stata. *Stata J.* **20**(4), 812–843 (2020). <https://doi.org/10.1177/1536867X20976314>
32. Holm, S.: A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* **6**(2), 65–70 (1979)
33. OCaml Software Foundation: OCaml Users Survey (2022). <https://ocaml-sf.org/docs/2022/ocaml-user-survey-2022.pdf>
34. Palmkog, K., Tassi, E., Zimmermann, T.: Reliably reproducing machine-checked proofs with the Coq Platform. In: *Workshop on Reproducibility and Replication of Research Results* (2022). <https://doi.org/10.48550/arXiv.2203.09835>
35. Li, R., Pandurangan, P., Fruckaj, H., Dabbish, L.: Code of conduct conversations in open source software projects on GitHub. *Proc. ACM Hum. Comput. Interact.* **5**(CSCW1), 1–31 (2021)
36. Robson, N.: Diversity and decorum in open source communities. In: *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pp. 986–987 (2018)
37. Singh, V., Bongiovanni, B., Brandon, W.: Codes of conduct in Open Source Software—for warm and fuzzy feelings or equality in community? *Softw. Qual. J.* (2021). <https://doi.org/10.1007/s11219-020-09543-w>
38. Tourani, P., Adams, B., Serebrenik, A.: Code of conduct in open source projects. In: *2017 IEEE 24th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pp. 24–33. IEEE (2017)
39. opam development team: OCaml Package Manager (2023). <https://opam.ocaml.org>
40. The Coq Development Team: opam archive for Coq (2023). <https://github.com/coq/opam-coq-archive>
41. Duvall, P., Matyas, S., Glover, A.: *Continuous Integration: Improving Software Quality and Reducing Risk*, 1st edn. Addison-Wesley Professional, Upper Saddle River (2007)
42. Fabre, A.: Tie-breaking the highest median: alternatives to the majority judgment. *Soc. Choice Welf.* **56**(1), 101–124 (2021)
43. Shadish, W.R., Cook, T.D., Campbell, D.T.: *Experimental and Quasi-experimental Designs for Generalized Causal Inference*. Houghton, Mifflin and Company (2002)
44. Capano, G., Engeli, I.: Using instrument typologies in comparative research: conceptual and methodological trade-offs. *J. Comp. Policy Anal. Res. Pract.* **24**(2), 99–116 (2022). <https://doi.org/10.1080/13876988.2020.1871297>
45. Mercer, S.L., DeVinney, B.J., Fine, L.J., Green, L.W., Dougherty, D.: Study designs for effectiveness and translation research: identifying trade-offs. *Am. J. Prev. Med.* **33**(2), 139–154 (2007). <https://doi.org/10.1016/j.amepre.2007.04.005>
46. Wolff, M., Haase, A.: Viewpoint: dealing with trade-offs in comparative urban studies. *Cities* **96** (2020). <https://doi.org/10.1016/j.cities.2019.102417>
47. Verdecchia, R., Engström, E., Lago, P., Runeson, P., Song, Q.: Threats to validity in software engineering research: a critical reflection. *Inf. Softw. Technol.* **164**, 107329 (2023). <https://doi.org/10.1016/j.infsof.2023.107329>
48. Reichardt, C.S.: Criticisms of and an alternative to the Shadish, cook, and Campbell validity typology. *N. Dir. Eval.* **2011**(130), 43–53 (2011). <https://doi.org/10.1002/ev.364>
49. Robillard, M.P., Arya, D.M., Ernst, N.A., Guo, J.L.C., Lamothe, M., Nassif, M., Novielli, N., Serebrenik, A., Steinmacher, I., Stol, K.-J.: Communicating study design trade-offs in software engineering. *ACM Trans. Softw. Eng. Methodol.* **33**(5), 1–10 (2024). <https://doi.org/10.1145/3649598>
50. Israel, G.D., Taylor, C.: Can response order bias evaluations? *Eval. Program Plan.* **13**(4), 365–371 (1990)

## Authors and Affiliations

Ana de Almeida Borges<sup>1</sup> · Annalí Casanueva Artís<sup>2</sup> · Jean-Rémy Falleri<sup>3</sup> · Emilio Jesús Gallego Arias<sup>4</sup> · Érik Martin-Dorel<sup>5</sup> · Karl Palmskog<sup>6</sup> · Alexander Serebrenik<sup>7</sup> · Théo Zimmermann<sup>8</sup>

✉ Karl Palmskog  
palmskog@kth.se

Ana de Almeida Borges  
ana.agvb@gmail.com

Annalí Casanueva Artís  
annali.casanueva@gmail.com

Jean-Rémy Falleri  
falleri@labri.fr

Emilio Jesús Gallego Arias  
emilio-jesus.gallego-arias@inria.fr

Érik Martin-Dorel  
erik.martin-dorel@irit.fr

Alexander Serebrenik  
a.serebrenik@tue.nl

Théo Zimmermann  
theo.zimmermann@telecom-paris.fr

<sup>1</sup> Dep., Universitat de Barcelona, 100190 Barcelona, Spain

<sup>2</sup> Dep., Ifo Institute, 100190 Munich, Germany

<sup>3</sup> Univ. Bordeaux, Bordeaux INP, CNRS, UMR 5800 LaBRI, 33400 Talence, France

<sup>4</sup> Université Paris Cité, CNRS, Inria, IRIF, 75013 Paris, France

<sup>5</sup> Dep., IRIT, CNRS, Toulouse INP, Université de Toulouse, 118 route, 31062 Toulouse Cedex 9, France

<sup>6</sup> Department of Computer Science, KTH Royal Institute of Technology, Lindstedtsvägen 3, 100 44 Stockholm, Sweden

<sup>7</sup> Mathematics and Computer Science, Eindhoven University of Technology, POB 513, 5600MB, North, Brabant, Netherlands

<sup>8</sup> LTCI, Télécom Paris, Institut Polytechnique de Paris, 100190 Paris, France