

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/378016543>

Predicting open source contributor turnover from value-related discussions: An analysis of GitHub issues

Conference Paper · February 2024

DOI: 10.1145/3597503.3623340

CITATIONS

17

READS

135

3 authors, including:



Jack Jamieson

NTT

26 PUBLICATIONS 186 CITATIONS

SEE PROFILE



Naomi Yamashita

Kyoto University

132 PUBLICATIONS 2,196 CITATIONS

SEE PROFILE

Predicting open source contributor turnover from value-related discussions: An analysis of GitHub issues

Jack Jamieson
jack@jackjamieson.net
NTT Communication Science Labs
Kyoto, Japan

Naomi Yamashita
naomiy@acm.org
NTT Communication Science Labs
Kyoto, Japan

Eureka Foong
eureka.foong@gmail.com
Tokyo College, the University of
Tokyo
Tokyo, Japan

ABSTRACT

Discussions about project values are important for engineering software that meets diverse human needs and positively impacts society. Because value-related discussions involve deeply held beliefs, they can lead to conflicts or other outcomes that may affect motivations to continue contributing to open source projects. However, it is unclear what kind of value-related discussions are associated with significant changes in turnover. We address this gap by identifying discussions related to important project values and investigating the extent to which those discussions predict project turnover in the following months. We collected logs of GitHub issues and commits from 52 projects that share similar ethical commitments and were identified as part of the DWeb (Decentralized Web) community. We identify issues related to DWeb's core values of respectfulness, freedom, broadmindedness, opposing centralized social power, equity & equality, and protecting the environment. We then use Granger causality analysis to examine how changes in the proportion of discussions related to those values might predict changes in incoming and outgoing turnover. We found multiple significant relationships between value-related discussions and turnover, including that discussions about respectfulness predict an increase in contributors leaving and a decrease in new contributors, while discussions about social power predicted better contributor retention. Understanding these antecedents of contributor turnover is important for managing open source projects that incorporate human-centric issues. Based on the results, we discuss implications for open source maintainers and for future research.

CCS CONCEPTS

• **Software and its engineering** → **Open source model**; • **Human-centered computing** → **Empirical studies in collaborative and social computing**.

KEYWORDS

Human Values, turnover, open source, GitHub

ACM Reference Format:

Jack Jamieson, Naomi Yamashita, and Eureka Foong. 2024. Predicting open source contributor turnover from value-related discussions: An analysis of GitHub issues. In *2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE '24)*, April 14–20, 2024, Lisbon, Portugal. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3597503.3623340>

1 INTRODUCTION

Engaging with values during development is important for engineering software that meets diverse human needs and positively impacts society. One important consideration is that engaging with values may affect contributor onboarding and turnover, particularly in open source software [25, 75]. In fact, calls to support shared values have been used to motivate collaboration and instill a sense of shared purpose among like-minded developers [35, 69]. At the same time, the open source community has been experiencing a “culture war” about the role of ethics in open source software, which has resulted in high-profile conflicts and departures [63]. Consequently, there are important questions about how increasing attention to values may simultaneously attract contributors who share project values, and drive away contributors through conflict and value clashes.

Recent research has identified that open source developers discuss a variety of values during development [33, 37, 43]. Yet, it has also advocated that some values should be more deeply engaged with in GitHub issues [37] and argued for greater attention to values in software engineering generally [48, 78]. However, since project values may shift over time and contributors may disagree on their details, discussions about values can also result in conflicts [18, 33], which can disrupt collaborative software development. Particularly, Filippova and Cho [19] found that conflicts about project values (i.e., *normative* conflicts) were negatively associated with intentions to continue contributing to open source projects, and ideological disagreements have been identified as an antecedent of discussion toxicity [42], which may also drive some contributors away [23, 51]. Prior evidence about the relationship between value-related discussions and turnover is based on self-report surveys and interviews [19, 33], which makes it difficult to determine the scale of potential impacts and to disentangle what people say they will do from their actual behaviour. This paper addresses that gap by identifying value-related discussions in a collection of projects hosted on GitHub and examining their effect on project turnover in the months following those discussions.

Our overall research question is whether changes in the degree to which values are discussed in GitHub issues lead to significant changes in contributor turnover. The first step toward addressing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICSE '24, April 14–20, 2024, Lisbon, Portugal

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0217-4/24/04...\$15.00

<https://doi.org/10.1145/3597503.3623340>

this question is to identify and summarize references to values in GitHub issues and to summarize their basic features. Identifying these issues allows us to quantify the rate of value-related issues over time. We then examine the content of value-related issues through our first research question:

- **RQ1:** *What are the qualities of value-related discussions in GitHub issues?*

The purpose of this question is to understand how and in what contexts values were discussed, which is necessary for explaining any relationships we discover between value-related discussions and turnover. Specifically, we examine potential causes of dysfunction, such as whether discussions focus on the collaborative process rather than the product being built [19] and whether a comment complains that a value has been violated [44, 45].

After addressing that question, we move to our main question:

- **RQ2:** *To what extent do proportions of value-related discussions predict contributor turnover in open source projects?*

To answer this question, we calculated two sets of time-series data for each project in our data. First, we created monthly time series of both incoming and outgoing contributor turnover for each project. The second set of time-series data describes the proportion of issues related to several values each month. Using these two sets of time series, we employ Granger causality analysis and impulse response functions to identify the extent and valence with which changes in the frequency of value-related issues dynamically predict changes in turnover.

We investigate these relationships between value-related discussions and contributor turnover in the context of a community of projects for building decentralized web (DWeb) technologies, which are loosely united by shared commitments to common values [35, 70]. One benefit of focusing on these projects is that, because they share similar values and engage in related technical activities, they tend to discuss values in similar ways, which simplifies our interpretation of these discussions.

We identified that value-related discussions do have effects on turnover. The most striking result was that an increase in the percentage of issues related to ‘respectfulness’ predicts an increase in outgoing turnover and a decrease in new users in the following months. Finally, our discussion synthesizes the results of both research questions to propose reasons for this and other predictive influences between value-related discussions and turnover. We then discuss implications for open source maintainers and for future research about discussions in open source software communities.

Based on an empirical analysis, this paper identifies concrete ways that value-related discussions serve as antecedents of turnover in open source software. We provide new details about how discussions related to specific values predict turnover among different types of contributors, including scales of predicted turnover on a project level. These findings provide insights for improving retention and participation in open source projects, and for understanding the role of values in software engineering practices.

2 BACKGROUND

2.1 Value-related discussions in open source

Following Cheng and Fleischmann [5], recent software engineering research [48] has adopted a general definition of human values as “guiding principles of what people consider important in life.” Much of the research about values in software engineering has focused on techniques to integrate values into development processes [13, 49], understanding what values developers engage with during development and how [43] and mapping what values are supported or violated in software [44, 66]. Alongside the question of how developers work toward achieving values in their product, it is important to consider how values shape collaborative work itself [11].

The trace data produced through open source collaborations is a rich resource for investigating how values impact both the design of end products and the processes of collaboration. Public, logged discussions are a major part of open source development, such as mailing lists [28], GitHub issues [37], and Q&A forums [74]. These logs are a rich resource for investigating how values impact both the design of end-products and the processes of collaboration. GitHub, a popular platform used for both hosting code and communicating among open source contributors, is a particularly valuable platform for understanding how open source work is coordinated through “social coding” [9]. Using GitHub issues – discussion threads used largely for bug reports and feature requests – researchers have identified many ways in which human values are interwoven with open source development work. For example, recent work has identified and mapped value-related discussions in GitHub issues and found that human-centric issues vary significantly between projects [37], proposed software-specific definitions of values based on issues discussions [43], and observed that the proportion of GitHub issues related to human values increased after the first year of a project [33]. That value-related discussions vary by project and over time is consistent with arguments that changes in the social context around software [21] can prompt developers to revisit ethical and social challenges that were previously regarded as settled.

Other work has focused specifically on the role of human values in open source collaborative processes. For example, Li et al. [40] examined how GitHub users discuss codes of conduct – rules that codify community values – and found that they are used both proactively and reactively to govern community behaviour, such as by enforcing standards of mutual respect. Additionally, managing conflicts has long been recognized as a central feature of open source software development, and values both cause and exacerbate such conflicts [18, 19, 73].

2.2 Joining and leaving open source projects

A significant amount of earlier research about open source communities focused on people’s motivations for contributing, such as career development, learning, altruism, and personal enjoyment [e.g. 30, 38, 46, 75]. More recent research has extended upon that foundation, examining changes in motivation over time [25] and how potential contributors decide which projects to join [51], including how welcomeness is conveyed in project discussions.

Turnover is particularly important to understand because it can result in abandoned code, project instability, and knowledge

loss [57]. Indeed, recent work has explored why contributors leave open source projects, such as due to life-events [31], frustration with uncivil and toxic behaviours [23], and different experiences according to contributors' gender [12, 22, 50, 53]. Further, although conflicts, in general, have mixed effects on open source contributors' motivations, a 2015 survey concluded that conflicts about project values are associated with intentions to stop contributing [19]. And, open source contributors who were interviewed about discussing values on GitHub explained that these discussions often involved deeply held ideological beliefs, and could impact motivations to quit or continue contributing [33].

These prior studies have used surveys [19, 23, 53] and interviews [12, 31, 33] to provide evidence that conversations in open source settings affect motivations to leave projects. However, these findings are based on self-reported perceptions of open source contributors, rather than observational evidence through conversational logs. Other studies have cited concerns about turnover as partial motivations for studying toxicity and uncivil language in conversational logs, but have not directly measured turnover in their analysis [42, 56]. This leaves an open gap in knowledge about whether researchers can identify conversational antecedents of turnover using log data, which we address through an empirical analysis of relationships between value-related discussions and turnover using log data from GitHub.

3 METHOD

We address our research questions through an analysis of log data collected from GitHub's API. We collected data from a community of projects hosted on GitHub with commitments to a common set of values, in order to ensure that value-related discussions would have common characteristics across the data set, aiding comparability. The overall research process is depicted in Figure 1. We identified and summarized value-related discussions. Then, we built two sets of monthly timelines for each project, measuring the proportion of issues related to values and turnover rates. Finally, we identified predictive relationships between these timelines using Granger causality analysis. Details of this procedure are described in this section. Analysis scripts are included in a replication package.¹

3.1 Data collection

We focus on GitHub projects related to the Decentralized Web (DWeb), a community of related projects building protocols, software, and related technologies to support decentralized online networks. In doing so, DWeb prioritizes values-driven, commons-based efforts to reshape the political economy of the web [32, 71].

The DWeb community is anchored by events and resources hosted by the Internet Archive, whose founder Brewster Kahle published a call to "bake our values into our code" for a new, decentralized web [35]. In 2021, the DWeb community published a list of principles, sourced from community contributors, defining a common set of commitments to justice, equity, and enabling agency of all people [70]. Although there are disagreements about the particulars of what it means to support "decentralization" [62], discussions about this distinct set of values are generally appropriate across DWeb-related projects. This means we can investigate potential

relationships between these discussions and contributor turnover in a context where discussing values is likely to be relevant and on-topic.

We identified projects hosted on GitHub that had either been presented at events hosted by the Internet Archive's DWeb community or had been included in a published list of DWeb-related projects [7]. The resulting set of projects is not exhaustive, but presents a reasonable representation of the DWeb. Most projects were operated by organizations (collective entities representing a company or group). In these cases, we downloaded every repository owned by that organization as part of the same "project." Individuals, by contrast to organizations, often own unrelated repositories, so we used our judgement to only include repositories that were relevant. There were some repositories that contained only automated backups of chat logs or wikis. We omitted these repos from our dataset since they do not involve human contributions.

From May 18-31 2022, we downloaded logs of commits and comments to issues and pull requests from each repository. We only included events by authenticated GitHub users, since this is necessary to consistently distinguish users by login ID. We removed events posted by usernames that included "[bot]," ended in "-bot," and those which we manually verified to be bots based on reading their comments. 60 bot accounts were removed (0.22% of accounts), which had made 1.98% of commits ($n = 15,638$) and 8.66% of comments ($n = 74,255$). The resulting dataset comprised 775,124 commits and 783,151 comments to 277,236 issues and pull requests. These events were distributed across 2,179 repositories in 52 projects and were contributed by 27,783 distinct user IDs.

3.2 Data analysis

3.2.1 Identifying values. Using our log data, we identified comments that contained terms from a dictionary of values-related terms. We used a dictionary that has previously been employed to analyze value-related discussions in DWeb projects [33]. That dictionary itself was adapted from Obie et al. [44] to focus only on values identified as particularly salient to the DWeb, based on an analysis of DWeb's published set of principles [70]. Specifically, DWeb's principles were mapped to Schwartz's theory of human values [65], which presents a classification of 58 values across 10 categories. Schwartz's theory has been validated in many diverse contexts [64] and has been adopted in several prior studies of values in software engineering [e.g., 14, 39, 44, 77, 79, 80]. Jamieson et al. [33] slightly adjusted the Schwartz taxonomy and extending Obie et al.'s [44] dictionary with additional terms to match the language used in the DWeb principles. The most significant changes were to replace Schwartz's 'politeness' category with 'respectfulness' (to match the way "good manners" are enacted in open source codes of conduct and similar policies) and to merge Schwartz's values of 'equality' and 'social justice' into one category called 'equity & equality', since both values included overlapping keywords [33]. This resulted in a list of the following values important to the DWeb: *respectfulness*, *broadmindedness*, *freedom*, *equality and equity*, *(opposing) social power*, and *protecting the environment*.

We used this dictionary to identify value-related terms in issues and pull requests across our dataset. Because both issues and pull

¹https://github.com/jackjamieson2/Jamieson_et_al-ICSE2024-Replication_Package/




Figure 1: Overview of research process.

requests have the same discussion-thread-style structure, we analyzed them together (hereinafter, we use the term “issues” to refer to both issues and pull requests). Hereinafter, when this paper refers to “issues”, this includes both issues and pull requests. To identify value-related discussions in GitHub issues, we tokenized and lemmatized the data using spaCy [68], and then flagged every comment that contained a term matching the dictionary. After reviewing the returned comments, we identified that the terms “inclusive” and “inclusion,” which were part of the *Broadmindedness* value, resulted in a high number of false positives e.g., referring to the “inclusion” of files in a commit, rather than to the concept of being “inclusive” of diverse people. Thus, we removed these terms from our dictionary. The modified dictionary is attached to this submission as a supplementary document.

Some projects in our dataset employed issue templates (forms for submitting a new issue), which sometimes included instructions to read the “code of conduct.” Because this would result in every issue from those repositories being identified as related to the code of conduct, we filtered out matches that were only associated with issue template text.

One feature of this dictionary approach is that there is some overlap between value categories, which are sometimes closely aligned with one another [65]. For example, respectfulness, equity & equality, and broadmindedness all involve terms related to diversity and inclusion, to some extent.

3.2.2 Summarizing values. To answer RQ1, “What are the qualities of value-related discussions in GitHub issues,” we summarized patterns about comments that reference each value. The first author selected a random sample of 100 matches for each value (or all matches if there were fewer than 100), and then categorized them using three questions:

(1) *Is the match a false positive?* False positives were defined as mentions of a keyword that were not related to values (e.g., homophones).

Then, for true positive items, we apply codes based on the answers to (2) and (3):

(2) *Is the value referenced in relation to the discussion on GitHub itself?* This code was added based on prior research [33] showing that value-related discussions, including conflicts, can relate to both project outcomes (e.g., the design of software) and development processes (e.g., people’s conduct during collaboration). This is important because past research has suggested that process-conflicts may be especially disruptive to open source projects [19].

(3) *Does the comment flag a problem related to the value?* This included allegations that a value had been violated (e.g., accusing someone of being rude identifies a violation of *respectfulness*), or more rarely, claims that supporting a given value was undermining some other goal. This question was included because prior work has identified values-violations as a particularly concerning aspect of values in software engineering [44, 45] and because we hypothesized that negative issues might be more strongly related to turnover. References to values-violations outside the project itself, such as criticizing mainstream online platforms, were not included.

Two authors examined and discussed the results for each category, and resolved disagreements through discussion and by slightly refining category definitions, resulting in the three-part scheme described above.

3.2.3 Quantifying value-related issues and turnover. For the purpose of quantifying turnover rates and the proportion of value-related issues, we organized the data into monthly snapshots. We selected monthly units since this is long enough for patterns to accumulate each snapshot and short enough for temporal relationships to remain meaningful.

Quantifying value-related issues: For each project, we calculated the percentage of active issues related to each value, on a monthly basis. “Active issues” includes every issue that received at least one comment or other action during that month. An issue was considered to be related to a value if at least one of its comments contained a value-related term. We use issues as the unit of our analysis because a single value-related comment can contextualize the entire thread as being related to that value.

Quantifying project turnover. For each user, we calculated a running total of months in which they were active in each project. In order to be considered active in a given month, a user must have contributed to at least one commit, issue, or pull-request. To calculate turnover, we define the months of joining and leaving as the first and final months in which a user has any activity, respectively, which were identified using a Python script.

For the purpose of our analysis, we only include data from 2016 (when the DWeb took shape as a recognizable entity) through 2021. By not counting “final” months from January 2022 through our data collection period (late May 2022), we reduce the risk of misidentifying a temporary as a permanent departure, which could occur due to right-censorship, where an event occurs after the period represented by a study’s data.


Table 1: Mean percentage of issues and pull requests that referenced values per project each month, and top-terms by frequency.

Value	Mean	Std. Dev	Top terms	False positive	Discussion-focused	Flag problem
Respectfulness	0.19 %	2.33	code of conduct, polite, rude	20.2 %	64.0 %	32.6 %
Freedom	0.21 %	1.26	freedom, user choose, sovereign	27.0 %	0.0 %	11.0 %
Broadmindedness	0.10 %	1.11	diversity, diverse, unconventional	10.4 %	11.5 %	20.7 %
Social Power	0.07 %	1.22	central authority, gatekeeper, monopoly	32.4 %	1.4 %	37.8 %
Equity & Equality	0.03 %	0.27	unfair, fairness, justice	16.7 %	15.3 %	58.9 %
Environment	0.01 %	0.16	climate change, energy consumption, wildlife	15.0 %	0.0 %	17.7 %

We hypothesized that outgoing turnover rates would decrease significantly based on how long one had been contributing. We verified this by calculating the mean turnover rate by month – e.g., 1st-month contributors: 60%, 2nd-month: 39%, 6th-month: 12%, 12th-month: 6%. We concluded that our analysis should account for the significant differences in turnover patterns between short-term and long-term contributors. To that end, we used the Jenks natural breaks [34] clustering algorithm to define groups based on the mean turnover rate by month, setting the number of groups to four, which explained a high degree of variance (goodness of variance fit = .88). Based on this, we calculated four outgoing turnover rates, grouped by each contributor’s rolling count of active months contributing to a given project, in a given month: 1st month, 2-3 months, 4-8 months, 9+ months (e.g., *1st month* turnover is the percentage of *1st month* contributors who left after that month). Additionally, *incoming turnover* was measured as the percentage of users for whom it was their first month, among all active users.

3.2.4 Identifying predictive relationships between value-related discussions and turnover. This section describes how we address RQ2, “To what extent do proportions of value-related discussions predict contributor turnover in open source projects?”

Prior research [76] has demonstrated that Granger analysis has significant potential for predicting changes in open source contributions, particularly given the prevalence of temporal data in open source archives. Given two time series, X and Y, Granger analysis measures whether lagged values of X provide statistically significant information about current values of Y (more than looking at lagged values of Y without X). Figure 2 presents an illustrative example, where X “Granger causes” Y, indicating that past values of X are useful for predicting future values of Y. In our case, we ask whether time series about the proportion of GitHub issues related to a given value can help predict changes in contributor turnover.

**Figure 2: Illustrative example of Granger causality, where time-series X perfectly predicts time-series Y.**

To perform this analysis, we follow the procedure described by Abrigo and Love [2], using the *pvar* package they created for Stata. This involved estimating a panel vector autoregression (PVAR) model for each combination of value and turnover type and then using those models to perform Granger causality analysis. Granger causality does not indicate whether the relationship between two variables is positive or negative so we further evaluate significant Granger results using Impulse Response Functions (IRFs), which are charts measuring the effect of a one-standard-deviation increase of X on Y, over time, including whether Y is expected to increase or decrease in response to X, and the magnitude of that change.

As described by Abrigo and Love [2], several steps are required to validate this analysis. First, Granger analysis requires that the time series are stationary (i.e., the mean values do not consistently change with time). We verified this using Phillips-Perron tests, which are non-parametric and robust to heteroscedasticity. Next, we selected the number of lags for each PVAR model (the number of past months to be included as predictors). To do so, we fit models with lags in the range of 1-6 months, and selected the model that had the best Modified Bayesian Information Criterion (MBIC) score, which balances high model fit with low complexity [3]. We excluded two models which did not satisfy Hansen’s overidentification restriction, which indicates that no matter how many lags were used the model may be misspecified and should not be selected [2]. Finally, we verified that each model was stable, which provides further evidence of stationarity [41, p.25] and that IRFs will provide interpretable results [2].

4 FINDINGS

4.1 RQ1: What are the qualities of value-related discussions?

To address RQ1, we summarize some basic features of discussions related to each value. These summaries serve both to validate the value-detection in our analysis and to describe what aspects of each value were discussed and how. These summaries provide a lens through which we can interpret and propose explanations for predictive patterns identified in RQ2, which we do in Section 5 (“Discussion”).

Table 1 shows the mean percentage (and standard deviation) of issues referencing each value, grouped monthly and by project. Value-related issues were the minority, which is sensible because the majority of discussions on GitHub are technical. Additionally reported are each value’s most frequent terms, and summaries statistics as described in Section 3.2.2. This includes the percentage of false positives per value, then among the remaining true positive

matches, the percentage that was *discussion-focused* (The value referred to the discussion, rather than the product being built) and the percentage that flagged a problem related to that value.

Below, we summarize and provide representative examples of each value. In each example, the matching keywords are indicated using underlines. Even though GitHub issues contain public data, we are aware that people may have mixed opinions regarding the reproduction of their comments in a research article. Further, even if presented without attribution, it could be possible to identify comment authors using a search engine, which has the potential to cause harm [10]. Therefore, we have very slightly rephrased quotations to prevent searchability, taking care to preserve their meaning.

Respectfulness. The majority (64.0%) of matches for *respectfulness* implicated this value in the discussion itself, rather than the product being built. For example, remarking that another contributor's "failure to respond here, while continuing to post in other groups, is rude" or apologizing for "assuming [another commenter's] tone was rude." 35.0% of comments in this category mentioned a project's "code of conduct," often in the context of asking others to behave more respectfully. For example, "This is the second time you opened a new issue specifically to circumvent a prior issue's closure. Continuing to do so is in violation of the code of conduct." Significantly, when project contributors flagged problems related to respectfulness, it was often in the form of calling out a specific individual. Thus, more so than other types of issues, respectfulness problems seemed related to interpersonal conflict. Although the majority of references to respectfulness were discussion-focused, there were also cases of arguing for respectfulness in the technologies being built. For example, "Keeping compatibility for now might be useful (and avoiding incompatibly redefining what motes mean is just polite)."

Freedom. References to freedom were generally about supporting individuals' agency to use the software being built how they wanted. Because many of the projects were targeted at other developers (e.g., infrastructures for third-party developers to build upon), freedom was targeted at both end-users and developers. For example, a proposal might be promoted as allowing "the developer to have more freedom to structure the implementation of the user module instance." There were also many references to freedom as a big-picture goal, such as a claim that "We are building a community-owned infrastructure that gives us [...] autonomy to access information in a free manner."

Additionally, references to freedom were overwhelmingly referred to the product being built rather than concerning the role of freedom in contributor's discussions. And, they were generally positive, rarely flagging a problem or violation of freedom.

Broadmindedness. A strong majority (68.9%) of comments tagged with broadmindedness were references to "diversity" or "diverse". Calls for identity diversity and inclusion were common, such as in a suggestion to "increase the DAO's gender diversity" or noting that an event's "diverse group of attendees" had been helpful for developing "more inclusive descriptions of the project's goals." Additionally, there were many comments in favour of increasing the diversity of perspectives in various aspects of the development

process, such as assertions that "a diversity of implementations will make this system stronger" and that developers are "working really hard to support diversity of contributor's local development environments."

There was also a mix of views about trade-offs between accepting multiple approaches and relying on consistent standards, such as a comment that a particular coding style's "benefits make up for the fact that they are unconventional" as well as critiques, such as a comment that a program's "current behaviour is a bit unorthodox."

Social power. Comments tagged with "social power" were dominated by the terms "central authority," "gatekeeper," and "monopoly," and generally opposed centralized forms of social power. Many references to these terms were about technical artifacts (e.g., a central authority for authenticating users or managing network connections). When these objects were named without any reference to their social consequences, we marked them as a false positive. However, we marked them as true positives when comments implicitly referenced social power consequences of such artifacts, such as a comment that "what DIDs do is allow one to find cryptographic material proven to be associated with a given identifier, without a central authority involved." Further, some discussions were more explicit, such as an assertion that a proposed system "sets a dangerous precedent since it restricts access to a small number of actors. Who gets to be the gatekeeper?"

Comments in this category almost always referred to the product being built, with only 1.4% referring to social power enacted within contributor's discussions. Further, 37.8% of these comments flagged perceived problems, typically by arguing against centralized power structures, however they often involved productively working through the problem rather than coming into conflict.

Equity & equality. The majority (68.4%) of comments in this category contained the terms "unfair," "fairness," or "unfair." Two projects included significant discussions and debates related to fairness, many of which concerned programs for distributing payment to open source contributors and third-party developers. We carefully considered whether these should be considered "discussion-focused" since they emphasize fairness towards project collaborators, who were often participating in GitHub issues. In general, we determined that they are not "discussion-focused" because the payment methods were part of the novelty of the system being built, discussions about fairness typically referred to the whole system rather than to individuals' conduct in discussions. In a typical case, for example, one contributor alleged that there was an unfair systematic bias in favour of paying developers of new third-party apps, which disadvantaged maintainers of existing apps: "This prioritizing of new apps is undocumented, and could be seen as unfair [to maintainers of existing apps]." On the other hand, a small number of statements related to the fairness of these payment systems concerned individuals' conduct during discussions, such as a disagreement in which one individual was accused of making "unfair demands." In those atypical cases, we considered the comment to be discussion-focused. Ultimately, we have attempted to describe our interpretations with sufficient depth to acknowledge that this was sometimes a grey area.

Other comments in this category included general expressions about overall project goals, such as “We are committed to openness, transparency, and fairness,” and a few references to other keywords, such as working toward a more “equitable design” of specific project features.

Protecting the environment. Issues related to protecting the environment were extremely rare ($n = 20$). To some degree, we believe our count is low because many keywords were omitted due to being commonly used in contexts unrelated to this value (e.g., ‘energy’ and ‘environment’). Most commonly, references to protecting the environment were related to energy consumption, particularly in relation to blockchain technologies. e.g., “Lately, there’s been lots of concern about Bitcoin’s growing energy consumption [...] Has there been a study about whether this software also consumes lots of energy?”

4.2 RQ2: To what extent do proportions of value-related discussions predict contributor turnover among these open source projects?

To answer RQ2, we first provide a brief summary of overall turnover patterns among the projects in our data. Then, Section 4.2.2 reports the results of Granger analysis to identify whether turnover can be predicted by value-related discussions, and Section 4.2.3 reports the magnitudes of predictive changes that were identified in the previous step.

4.2.1 Overall turnover patterns. To provide context for the following analysis, we first present a brief summary of turnover patterns. Figure 3 shows outgoing turnover rates colour-coded to group turnover by contribution length) and the incoming turnover rate (the percentage of new contributors each month). Averaged across all projects, outgoing turnover has a generally increasing pattern, and incoming turnover has a generally decreasing pattern. Overall, the mean number of monthly contributors per project consistently increases, but this growth slows after 2019 (mean monthly contributors by year: 2016: 19.0, 2017: 22.0, 2018: 27.5, 2019: 32.5, 2020: 36.1, 2021: 36.9).




Figure 3: Monthly turnover, averaged across all projects.

4.2.2 Do value-related discussions predict turnover? Table 2 shows the results of Granger causality tests, as well as the number of lags for each PVAR model. To account for potential inflation of Type

Table 2: Granger causality results for hypotheses that the proportion of issues related to each value predicts changes in turnover. Results reported as χ^2 with Bonferroni-adjusted p-values in parentheses. Supported hypotheses indicated in bold. Omitted models indicated with “–” (See Section 3.2.4).

Turnover type:	Out	Out	Out	Out	In
Contrib. length:	1	2-3	4-8	9+	All
Value					
Respectfulness	127.701 (0.000) 1 lag	20.4451 (0.000) 1 lag	28.158 (0.000) 1 lag	28.319 (0.000) 2 lags	57.292 (0.000) 1 lag
Freedom	5.018 (1.00) 1 lag	42.404 (0.000) 1 lag	1.588 (1.00) 1 lag	16.865 (0.012) 2 lags	0.411 (1.00) 1 lag
Broadmindedness	205.009 (.000) 1 lag	0.867 (1.00) 2 lags	4.579 (1.00) 1 lag	–	0.254 (1.00) 2 lags
Social power	2006.876 (.000) 1 lag	2803.359 (.000) 1 lag	144.179 (.000) 1 lag	7.129 (1.00) 2 lags	8.923 (0.672) 2 lags
Equity & equality	1.422 (1.00) 1 lag	0.015 (1.00) 1 lag	3.761 (1.00) 1 lag	7.739 (1.00) 2 lags	0.174 (1.00) 1 lag
Protecting the Environment	7.081 (1.00) 2 lags	50.097 (0.000) 1 lag	7.634 (1.00) 2 lags	–	15.376 (0.026) 2 lags

1 errors due to multiple comparisons, we use Bonferroni-adjusted p-values. Tests where the null hypothesis is rejected at the 5% level of significance are indicated in bold.

For each combination of value and turnover type, the null hypothesis is that a change in the proportion of issues that reference that value does *not* provide information that helps to predict that type of turnover. Thus, a significant result implies that past data about the given *value* variable can be used to predict the given *turnover* variable. The results indicate that:

Respectfulness: The proportion of *respectfulness*-related issues helps predict changes in every turnover category. **Freedom:** The proportion of *freedom*-related issues helps predict outgoing turnover among 2-3 month, and 9+ month contributors. **Broadmindedness:** The proportion of *broadmindedness*-related issues helps predict outgoing turnover among first-month contributors but not among longer-term contributors. **Social power:** The proportion of *social power*-related issues helps predict outgoing turnover among first-month, 2-3 month, and 4-8 month contributors. **Protecting the environment:** The proportion of *protecting the environment*-related issues helps predict outgoing turnover among 2-3 month contributors and also helps predict the rate of incoming turnover (new contributors).

4.2.3 What is the magnitude of predicted change in turnover? To evaluate effect sizes and distinguish positive vs. negative effects,

we created orthogonalized Impulse Response Functions (IRF). IRFs show that a sudden one standard deviation increase to the proportion of issues related to a specific value (the impulse) will be followed by a change in turnover in the following months (the response), also measured in units of Standard Deviations (SD). 95% confidence intervals are displayed in grey, calculated using 200 Monte Carlo draws as described in Abrigo and Love [2]. The most significant results were related to discussions about *respectfulness*, and are presented in Figure 4. Figure 5 presents IRFs related to other values. In most cases, a peak (or valley) occurs one month after the impulse, showing that the largest predicted effect of value-related discussions on turnover is within the following month, after which the predicted effect size returns to zero. These figures indicate the following results:

- **Respectfulness:** A shock to the proportion of *respectfulness* issues increases outgoing turnover among all contributors. For 1st-month contributors, the effect size peaks at about 0.6 SD. For longer-term contributors, the effect size is smaller (0.2, 0.4, and 0.1 SD for 2-3 month, 4-8 month, and 9+ month contributors, respectively). Incoming turnover (the proportion of new contributors) is forecasted to decrease by about 0.2 standard deviations.
- **Broadmindedness:** A shock to the proportion of *broadmindedness* issues predicts a 1 SD increase in outgoing turnover among first-month contributors.
- **Freedom:** A shock to the proportion of *freedom* issues predicts a 0.4 SD decrease in turnover among 2-3 month contributors. It also predicts a drop in turnover among long-term contributors (9+ months), but the magnitude is tiny (0.01 SD).
- **Social power:** A shock to the proportion of *social power* issues predicts a significant decrease in outgoing turnover in the following month: 1.5 SD for first-month contributors, 0.8 SD for 2-3 month contributors, and about 0.2 SD among 4-8 month contributors.
- **Protecting the environment:** A shock to the proportion of *protecting the environment* issues predicts decreased outgoing turnover among 2-3 month contributors (about .15 SD) and an increase in new contributors (0.5 SD). However, the extreme rarity of these issues overall discourages us from interpreting this result with too much credence.

4.2.4 Checking for mutual Granger causality. We also tested the hypotheses about the opposite direction of prediction (does turnover predict the proportion of value-related issues?). Although there were two significant Granger results (the proportion of *protecting the environment* issues was predicted by (i) incoming turnover and (ii) outgoing turnover among first-month contributors predicts), the effect size indicated by IRFs was tiny (<0.001 standard deviations). No other significant results were observed. This adds veracity to our findings by showing that the relationship between value-related issues and turnover is not merely correlational, but is in fact a unidirectional predictive relationship. For brevity, these insignificant results are omitted from this paper but are included in the linked replication package.




Figure 4: Orthogonalized IRFs with 95% confidence intervals. Patterns related to Respectfulness.

5 DISCUSSION

Our analysis in Section 4.2 identified that the proportions of issues related to several values predict changes in turnover. *Respectfulness*-related discussions predicted an increase in outgoing turnover and a decrease in newcomers. By contrast, discussions about *social power* predicted contributor retention, with fewer contributors leaving. There were also some notable effects related to issues that discuss *broadmindedness* (increased turnover among 1st-month contributors) and *freedom* (decreased turnover among 2-3 month contributors). These relationships were unidirectional, with no evidence that turnover significantly predicts changes in the proportion of value-related discussions.

Interpreting predictive causality. While prior research has identified links between value-related discussions and contributor turnover through self-report surveys and interviews [19, 33], our temporal analysis of GitHub logs provides empirical evidence about the statistical significance and effect size of such relationships. These quantitative results demonstrate predictive causality but are insufficient to determine “true causality” explaining *why* such patterns exist [76]. Therefore, we reflect on these results through our observations about value-related discussions (Section 4.1) and by drawing from prior research, which allows us to develop empirically grounded theories to explain the predictive patterns we observed. In doing so, we focus on clear and significant patterns. Thus, we do not further discuss patterns related to *protecting the environment* because the number of issues in this category was extremely small.

Predictive causality as an early warning sign. The most striking set of results was related to discussions about respectfulness, which




Figure 5: Orthogonalized IRFs with 95% confidence intervals. Patterns related to freedom, social power, broadmindedness, equity & equality, and protecting the environment.

were associated with an increase in outgoing turnover among all contributors and a decrease in the proportion of new contributors. In Section 4.1, we noted that many discussions referencing *respectfulness* called attention to an already-existing problem by asserting that another contributor’s conduct had been disrespectful. Thus, respectfulness issues may draw attention to incivility, which could discourage both new and existing contributors. For example, prior research found that 48% consider a welcoming community to be very important when deciding whether to join an open source project [23] and may use prior issues discussions to evaluate a project’s welcomeness [51]. Additionally, incivility has been identified as a cause of workplace turnover in general [55].

Therefore, comments referring to past instances of disrespect may *cause* turnover by raising awareness of incivility. Still, it is probably more appropriate to assert that both turnover and *respectfulness* issues are influenced by inappropriate conduct. And, given the temporal ordering implied by our results, issues calling out past instances of disrespect appear to provide an early warning sign of impending turnover.

- Discussions referencing respectfulness may draw attention to incivility, providing early warning signs about turnover.

Contentious and uncontentious values. In addition, respectfulness issues may be contentious to some contributors. Specifically, over a third of respectfulness issues referenced the term “code of conduct.” Codes of conduct are used both proactively and reactively for community governance, and controversial acts of moderation can incite community backlash [40]. More generally, codes of conduct are related to diversity and inclusion initiatives, which some open source contributors view as politicized [27] or detracting from perceived “meritocracy” [63]. To some extent, *broadmindedness* issues may be similar since they predicted increases in outgoing turnover (albeit only among first-month contributors) and often referred to gender and cultural diversity. Further, 64% of the time, the topic of respectfulness referred to the discussion on GitHub rather than to the product being built, which aligns with the observation that some regard code of conduct discussions as a distraction from the work of software development itself [40].

- Frequent discussions about topics such as codes of conduct may cause increased turnover among those who criticize those discussions as drawing attention away from actual software development.

By contrast, the topics that predicted improved contributor retention—*freedom* and *social power*—were almost entirely focused on the product being built and are widely agreed to be core ideological motivations of the decentralized web (and of open source software more broadly [6, 36]). Social power, specifically, is embodied in the technical architectures being built, such that terms like “central authority” and “gatekeeper” served both as markers for discussing social aspects of the decentralized web and as identifiers of technical components.

- Discussions that demonstrate commitments to core, widely agreed-upon values by integrating them with technical work may motivate sustained contributions.

Moderating turnover-predicting discussions. Based on this study’s results, it would be prudent for open source maintainers to monitor for discussions that could predict increased turnover. Particularly in large projects, existing monitoring tools could be adapted to identify discussions using keywords or other automated methods.

- Open source community monitoring dashboards [26, 52] could be adapted to identify discussions that predict turnover (e.g., about respectfulness) and prompt maintainers to review them manually.

Once identified, there may be instances where straightforward content moderation is appropriate, such as removing² extremely rude comments that could make a project appear unwelcoming [51]. However, maintainers should be cautious about removing comments since this can cause a backlash if the removal is perceived to be unjust [40, 72]. For example, inappropriate comments sometimes result from miscommunication rather than malice (e.g., due

²For example, GitHub’s documentation suggests hiding or removing “disruptive” comments (<https://docs.github.com/en/communities/moderating-comments-and-conversations/managing-disruptive-comments>)

to working in a second language) [40], and the values-related discussions we observed were generally not abusive or disruptive to an extent that typically warrants removal.

- Instead of content removal, which may itself repel contributors, it may be more appropriate to focus on predicting and de-escalating disruptive discussions that could drive contributors away [see for e.g., 4, 61]

Balancing turnover with other project goals. Although attracting and retaining contributors is important for sustainability, this must be balanced against other concerns. For example, code of conduct discussions could drive turnover among some who disapprove of “politicizing” software development. However, codes of conduct typically aim to (and succeed at) increasing contributions from underrepresented groups [58, 67]. Achieving this may require accepting some degree of turnover.

- Future work should examine if and how *respectfulness* and related discussion predict different turnover patterns according to gender and other contributor characteristics.

Further, the relationships we found between value-related discussions and turnover were generally strongest for relatively new contributors, with small or no effects on those who had been contributing for nine months or more. This is unsurprising because social capital, which is developed over time, predicts longevity in open source contributions [53], so long-term contributors may be less easily swayed to leave a project. Although increases in new contributors can make an online community appear more active and attract increased participation [8], attracting new contributors is not always a top priority, since an influx of newcomers can lead to reduced code quality [20], and rapid, unplanned growth can lead to “catastrophic success” where maintainers are burdened with unsustainable governance labour [24].

- Monitoring value-related discussions that predict turnover is more important for projects in a growth stage than for projects that are less focused on attracting and retaining new contributors.

Lessons for toxicity detection. Finally, the results regarding respectfulness include lessons for another area of research – systems to identify “toxic” online comments [e.g., 1, 29]. General purpose detection systems perform poorly on software engineering discussions [59], so researchers have built software engineering domain-specific automated detection systems [e.g., 17, 54, 56, 60] and conducted qualitative analyses to understand different forms of toxicity and incivility [e.g., 15, 16, 42]. However, there is a need for further improvement.

Our observation that *respectfulness*-related issues, which often included allegations of disrespect, predicted increased turnover aligns with the Google Perspective API’s definition of toxic comments as “rude, disrespectful, or unreasonable comment[s] that [are] likely to make people leave a discussion” [1]. Our results have demonstrated that comments calling out other people for violating respectfulness could be used to identify subtly toxic utterances and toxic non-verbal behaviours that are beyond the capabilities of current systems.

- Toxicity detection systems may be improved by using replies to help evaluate potentially toxic comments and behaviours.

6 THREATS TO VALIDITY

Internal validity. Statistical methods can provide strong evidence of predictive causality, but “true causality” is a philosophically complex concept that requires careful application of theory [76]. We have proposed explanations based on our observations about the properties of value-related discussions and theories from prior work. Nonetheless, it is impossible to eliminate the possibility that the links we identified in Section 4.2 are related to unobserved confounding factors.

Our definition of value-related issues does not differentiate between instances where a value was the central unit of discussion in an issue compared to cases where a value was mentioned in passing. If anything, this may introduce random noise into our analysis, so we do not believe it undermines this paper’s core results, given their high degree of statistical significance.

Contributors were distinguished by user name. Prior research has identified that some open source contributors use multiple user names. Most of this research has referred to cases where a single person uses distinct user names across multiple sites. Therefore, we do not believe this would significantly affect our analysis, but if our data includes individuals who use multiple user names within a single GitHub project, they may be incorrectly counted as multiple people.

Finally, we evaluated all contributors together. Since prior research has identified different discussion styles and turnover patterns among different types of contributors (e.g., developers and users) [47] and based on characteristics like gender, future research may investigate the extent to which relationships between turnover and issues vary across people with different contribution patterns.

Construct validity. Our keyword-based approach cannot capture value-related discussions where (a) values are only discussed implicitly, and (b) where we have omitted a relevant term. For example, the term “inclusive” was omitted from the *broadmindedness* category because it returned many false positives, even though it also returned some relevant results. Ultimately, the definition of each value category is shaped by the values dictionary. Thus, we have attached the dictionary as a supplementary document and have summarized the value categories in Section 4.1, including false positive rates. Unfortunately, this method is not suitable for estimating a false negative rate. Future work may refine this sort of analysis by employing more sophisticated language models or using interviews and surveys to gather richer details.

External validity. We focused on GitHub projects related to the decentralized web because they are loosely aligned by a common set of values. Therefore, investigating other open source projects may identify different patterns. Additionally, there are many other sites where values and other topics are discussed in relation to software engineering, such as mailing lists, forums, and chat rooms. Future work could use similar methods to examine those sites, potentially linking data across multiple sites to provide a more holistic view of how value-related and other human-centric discussions affect turnover among open source contributors.

7 CONCLUSION

Values-related discussions are essential for building software that meets high ethical standards and positively impacts society. These

discussions can sometimes motivate people to work toward a common goal, but they can also be disruptive and demotivating. Therefore, we investigated the extent to which value-related discussions predict turnover among open source contributors. Using logs of GitHub activity, we identified and summarized issue comments that reference core values among a community of 52 related projects. We then used Granger causality analysis and Impulse Response Functions to identify how the proportion of issues related to those values each month can predict several types of turnover. We found several patterns, the strongest of which is that issues referencing *respectfulness* predict an increase in people leaving a project and a decrease in newcomers. We theorized reasons for significant patterns based on prior research and our observations about the content of value-related discussions. We discussed implications about how open source maintainers may manage discussions to improve turnover and how the results can inform future research about open source contributors. These results extend prior work by finding empirical evidence that value-related discussions can contribute to significant turnover, especially regarding respectful conduct among collaborating peers. This work contributes to knowledge about how contributors' discussions can lead to unexpected, tangible outcomes in open source collaborations. This is an essential step toward improving how values such as respectfulness are incorporated into collaborative processes.

REFERENCES

- [1] 2021. Perspective API - How It Works. <https://www.perspectiveapi.com/how-it-works/>. Retrieved {2023-01-31} from <https://www.perspectiveapi.com/how-it-works/>
- [2] Michael R. M. Abrigo and Inessa Love. 2016. Estimation of Panel Vector Autoregression in Stata. *The Stata Journal: Promoting communications on statistics and Stata* 16, 3 (Sept. 2016), 778–804. <https://doi.org/10.1177/1536867X1601600314>
- [3] Donald W. K. Andrews and Biao Lu. 2001. Consistent Model and Moment Selection Procedures for GMM Estimation with Application to Dynamic Panel Data Models. *Journal of Econometrics* 101, 1 (March 2001), 123–164. [https://doi.org/10.1016/S0304-4076\(00\)00077-4](https://doi.org/10.1016/S0304-4076(00)00077-4)
- [4] Jonathan P. Chang, Charlotte Schluger, and Cristian Danescu-Niculescu-Mizil. 2022. Thread With Caution: Proactively Helping Users Assess and Deescalate Tension in Their Online Discussions. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 545:1–545:37. <https://doi.org/10.1145/3555603>
- [5] An-Shou Cheng and Kenneth R. Fleischmann. 2010. Developing a Meta-Inventory of Human Values. *Proceedings of the American Society for Information Science and Technology* 47, 1 (2010), 1–10. <https://doi.org/10.1002/meet.14504701232>
- [6] Gabriella Coleman. 2004. The Political Agnosticism of Free and Open Source Software and the Inadvertent Politics of Contrast. *Anthropological Quarterly* 77, 3 (2004), 507–519. <https://doi.org/10.1353/anq.2004.0035>
- [7] Digital Life Collective. 2018. Decentralized Tech. <https://kumu.io/DigLife/decentralized-tech>. Retrieved {2022-01-13} from <https://kumu.io/DigLife/decentralized-tech>
- [8] Laura Dabbish, Rosta Farzan, Robert Kraut, and Tom Postmes. 2012. Fresh Faces in the Crowd: Turnover, Identity, and Commitment in Online Groups. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12)*. Association for Computing Machinery, New York, NY, USA, 245–248. <https://doi.org/10.1145/2145204.2145243>
- [9] Laura Dabbish, Colleen Stuart, Jason Tsay, and Jim Herbsleb. 2012. Social Coding in GitHub: Transparency and Collaboration in an Open Software Repository. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (CSCW '12)*. Association for Computing Machinery, New York, NY, USA, 1277–1286. <https://doi.org/10.1145/2145204.2145396>
- [10] Phillip Dawson. 2014. Our Anonymous Online Research Participants Are Not Always Anonymous: Is This a Problem? *British Journal of Educational Technology* 45, 3 (2014), 428–437. <https://doi.org/10.1111/bjet.12144>
- [11] Françoise Détienne, Michael Baker, and Chloé Le Bail. 2019. Ideologically-Embedded Design: Community, Collaboration and Artefact. *International Journal of Human-Computer Studies* 131 (June 2019). <https://doi.org/10.1016/j.ijhcs.2019.06.003>
- [12] Christina Dunbar-Hester. 2020. *Hacking Diversity: The Politics of Inclusion in Open Technology Cultures*. Princeton University Press, Princeton, New Jersey.
- [13] Maria Angela Ferrario, Will Simm, Stephen Forshaw, Adrian Gradinar, Marcia Tavares Smith, and Ian Smith. 2016. Values-First SE: Research Principles in Practice. In *2016 IEEE/ACM 38th International Conference on Software Engineering Companion (ICSE-C)*. 553–562.
- [14] Maria Angela Ferrario, Will Simm, Jon Whittle, Christopher Frauenberger, Geraldine Fitzpatrick, and Peter Purgathofer. 2017. Values in Computing. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '17)*. Association for Computing Machinery, New York, NY, USA, 660–667. <https://doi.org/10.1145/3027063.3027067>
- [15] Isabella Ferreira, Bram Adams, and Jinghui Cheng. 2022. How Heated Is It? Understanding GitHub Locked Issues. In *Proceedings of the 19th International Conference on Mining Software Repositories (MSR '22)*. Association for Computing Machinery, New York, NY, USA, 309–320. <https://doi.org/10.1145/3524842.3527957>
- [16] Isabella Ferreira, Jinghui Cheng, and Bram Adams. 2021. The “Shut the F**k up” Phenomenon: Characterizing Incivility in Open Source Code Review Discussions. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (Oct. 2021), 353:1–353:35. <https://doi.org/10.1145/3479497>
- [17] Isabella Ferreira, Ahlaam Rafiq, and Jinghui Cheng. 2022. Incivility Detection in Open Source Code Review and Issue Discussions. arXiv:2206.13429 [cs]
- [18] Anna Filippova and Hichang Cho. 2015. Mudslinging and Manners: Unpacking Conflict in Free and Open Source Software. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '15)*. Association for Computing Machinery, Vancouver, BC, Canada, 1393–1403. <https://doi.org/10.1145/2675133.2675254>
- [19] Anna Filippova and Hichang Cho. 2016. The Effects and Antecedents of Conflict in Free and Open Source Software Development. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. Association for Computing Machinery, San Francisco, California, USA, 705–716. <https://doi.org/10.1145/2818048.2820018>
- [20] Matthieu Foucault, Marc Palyart, Xavier Blanc, Gail C. Murphy, and Jean-Rémy Falleri. 2015. Impact of Developer Turnover on Quality in Open-Source Software. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2015)*. Association for Computing Machinery, New York, NY, USA, 829–841. <https://doi.org/10.1145/2786805.2786870>
- [21] Batya Friedman and Helen Nissenbaum. 1996. Bias in Computer Systems. *ACM Transactions on Information Systems* 14, 3 (July 1996), 330–347. <https://doi.org/10.1145/230538.230561>
- [22] Hana Frulkaj, Laura Dabbish, David Gray Widder, Huilian Sophie Qiu, and James D. Herbsleb. 2022. Gender and Participation in Open Source Software Development. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 299:1–299:31. <https://doi.org/10.1145/3555190>
- [23] R. Stuart Geiger. 2017. Summary Analysis of the 2017 GitHub Open Source Survey. arXiv:1706.02777 [cs] (2017). <https://doi.org/10.17605/OSF.IO/ENRQ5>
- [24] R. Stuart Geiger, Dorothy Howard, and Lilly Irani. 2021. The Labor of Maintaining and Scaling Free and Open-Source Software Projects. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 175:1–175:28. <https://doi.org/10.1145/3449249>
- [25] Marco Gerosa, Igor Wiese, Bianca Trinkenreich, Georg Link, Gregorio Robles, Christoph Treude, Igor Steinmacher, and Anita Sarma. 2021. The Shifting Sands of Motivation: Revisiting What Drives Contributors in Open Source. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*. 1046–1058. <https://doi.org/10.1109/ICSE43902.2021.00098>
- [26] Sean P. Goggins, Matt Germonprez, and Kevin Lumbard. 2021. Making Open Source Project Health Transparent. *Computer* 54, 8 (2021), 104–111.
- [27] Mariam Guizani, Bianca Trinkenreich, Aileen Abril Castro-Guzman, Igor Steinmacher, Marco Gerosa, and Anita Sarma. 2022. Perceptions of the State of D&I and D&I Initiative in the ASF. In *Proceedings of the 2022 ACM/IEEE 44th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS '22)*. Association for Computing Machinery, New York, NY, USA, 130–142. <https://doi.org/10.1145/3510458.3513008>
- [28] Anja Guzzi, Alberto Bacchelli, Michele Lanza, Martin Pinzger, and Arie van Deursen. 2013. Communication in Open Source Software Development Mailing Lists. In *2013 10th Working Conference on Mining Software Repositories (MSR)*. 277–286. <https://doi.org/10.1109/MSR.2013.6624039>
- [29] Laura Hanu and Unitary Team. 2020. Detoxify. Unitary.
- [30] Guido Hertel, Sven Niedner, and Stefanie Herrmann. 2003. Motivation of Software Developers in Open Source Projects: An Internet-based Survey of Contributors to the Linux Kernel. *Research Policy* 32, 7 (July 2003), 1159–1177. [https://doi.org/10.1016/S0048-7333\(03\)00047-7](https://doi.org/10.1016/S0048-7333(03)00047-7)
- [31] Giuseppe Iaffaldano, Igor Steinmacher, Fabio Calefato, Marco Gerosa, and Filippo Lanubile. 2019. Why Do Developers Take Breaks from Contributing to OSS Projects? A Preliminary Analysis. In *Proceedings of the 2nd International Workshop on Software Health (SoHeal '19)*. IEEE Press, Montreal, Quebec, Canada, 9–16. <https://doi.org/10.1109/SoHeal.2019.00009>
- [32] Mai Ishikawa Sutton. 2022. The Debate Over DWeb vs. Web3 & The Decentralized Elephant in the Room. Retrieved {2023-06-27} from <https://hypha.coop/dripline/debate-over-dweb-vs-web3/>

- [33] Jack Jamieson, Eureka Foong, and Naomi Yamashita. 2022. Maintaining Values: Navigating Diverse Perspectives in Value-Charged Discussions in Open Source Development. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 481:1–481:28. <https://doi.org/10.1145/3555550>
- [34] GF Jenks. 1967. The Data Model Concept in Statistical Mapping. *Statistical Mapping. International Yearbook of Cartography* 7, 186–190 (1967).
- [35] Brewster Kahle. 2015. Locking the Web Open: A Call for a Decentralized Web. Retrieved {2019-01-17} from <http://brewster.kahle.org/2015/08/11/locking-the-web-open-a-call-for-a-distributed-web-2/>
- [36] Christopher M. Kelty. 2008. *Two Bits: The Cultural Significance of Free Software*. Duke University Press.
- [37] Hourieh Khalajzadeh, Mojtaba Shahin, Humphrey O. Obie, and John Grundy. 2022. How Are Diverse End-User Human-Centric Issues Discussed on Github?. In *Proceedings of the 2022 ACM/IEEE 44th International Conference on Software Engineering: Software Engineering in Society*. 79–89.
- [38] Karim Lakhani and Robert G. Wolf. 2003. Why Hackers Do What They Do: Understanding Motivation and Effort in Free/Open Source Software Projects. *SSRN Electronic Journal* (2003). <https://doi.org/10.2139/ssrn.443040>
- [39] Conghui Li, Humphrey O. Obie, and Hourieh Khalajzadeh. 2021. A First Step Towards Detecting Values-violating Defects in Android APIs. *arXiv:2109.14359 [cs]* (Sept. 2021). [arXiv:2109.14359 \[cs\]](https://arxiv.org/abs/2109.14359)
- [40] Renee Li, Pavithra Pandurangan, Hana Fruckaj, and Laura Dabbish. 2021. Code of Conduct Conversations in Open Source Software Projects on Github. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (April 2021), 19:1–19:31. <https://doi.org/10.1145/3449093>
- [41] Helmut Lütkepohl. 2005. *New Introduction to Multiple Time Series Analysis*. New York : Springer, Berlin.
- [42] Courtney Miller, Sophie Cohen, Daniel Klug, Bogdan Vasilescu, and Christian Kästner. 2022. "Did You Miss My Comment or What?": Understanding Toxicity in Open Source Discussions. In *Proceedings of the 44th International Conference on Software Engineering (ICSE '22)*. Association for Computing Machinery, New York, NY, USA, 710–722. <https://doi.org/10.1145/3510003.3510111>
- [43] Arif Nurwidyantoro, Mojtaba Shahin, Michel R.V. Chaudron, Waqar Hussain, Rifat Shams, Harsha Perera, Gillian Oliver, and Jon Whittle. 2022. Human Values in Software Development Artefacts: A Case Study on Issue Discussions in Three Android Applications. *Information and Software Technology* 141 (Jan. 2022), 106731. <https://doi.org/10.1016/j.infsof.2021.106731>
- [44] Humphrey O. Obie, Waqar Hussain, Xin Xia, John Grundy, Li Li, Burak Turhan, Jon Whittle, and Mojtaba Shahin. 2021. A First Look at Human Values-Violation in App Reviews. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. IEEE Computer Society, 29–38. <https://doi.org/10.1109/ICSE-SEIS52602.2021.00012>
- [45] Humphrey O. Obie, Idowu Ileku, Hung Du, Mojtaba Shahin, John Grundy, Li Li, Jon Whittle, and Burak Turhan. 2022. On the Violation of Honesty in Mobile Apps: Automated Detection and Categories. <https://doi.org/10.48550/arXiv.2203.07547> *arXiv:2203.07547 [cs]*
- [46] Shaul Oreg and Oded Nov. 2008. Exploring Motivations for Contributing to Open Source Initiatives: The Roles of Contribution Context and Personal Values. *Computers in Human Behavior* 24, 5 (Sept. 2008), 2055–2073. <https://doi.org/10.1016/j.chb.2007.09.007>
- [47] Marco Ortu, Tracy Hall, Michele Marchesi, Roberto Tonelli, David Bowes, and Giuseppe Destefanis. 2018. Mining Communication Patterns in Software Development: A GitHub Analysis. In *Proceedings of the 14th International Conference on Predictive Models and Data Analytics in Software Engineering (PROMISE '18)*. Association for Computing Machinery, New York, NY, USA, 70–79. <https://doi.org/10.1145/3273934.3273943>
- [48] Harsha Perera, Waqar Hussain, Jon Whittle, Arif Nurwidyantoro, Davoud Mougouei, Rifat Ara Shams, and Gillian Oliver. 2020. A Study on the Prevalence of Human Values in Software Engineering Publications, 2015 – 2018. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (ICSE '20)*. Association for Computing Machinery, New York, NY, USA, 409–420. <https://doi.org/10.1145/3377811.3380393>
- [49] Harsha Perera, Gunter Mussbacher, Waqar Hussain, Rifat Ara Shams, Arif Nurwidyantoro, and Jon Whittle. 2020. Continual Human Value Analysis in Software Development: A Goal Model Based Approach. In *2020 IEEE 28th International Requirements Engineering Conference (RE)*. IEEE, Zurich, Switzerland, 192–203. <https://doi.org/10.1109/RE48521.2020.00030>
- [50] Huilian Sophie Qiu. 2022. *Understanding and Designing Mechanisms for Attracting and Retaining Open-Source Software Contributors*. Ph. D. Dissertation. Carnegie Mellon University, Pittsburgh, PA, USA.
- [51] Huilian Sophie Qiu, Yucen Lily Li, Susmita Padala, Anita Sarma, and Bogdan Vasilescu. 2019. The Signals That Potential Contributors Look for When Choosing Open-source Projects. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 122:1–122:29. <https://doi.org/10.1145/3359224>
- [52] Huilian Sophie Qiu, Anna Lieb, Jennifer Chou, Megan Carneal, Jasmine Mok, Emily Amspoker, Bogdan Vasilescu, and Laura Dabbish. 2023. Climate Coach: A Dashboard for Open-Source Maintainers to Overview Community Dynamics. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (CHI '23). Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3544548.3581317>
- [53] Huilian Sophie Qiu, Alexander Nolte, Anita Brown, Alexander Serebrenik, and Bogdan Vasilescu. 2019. Going Farther Together: The Impact of Social Capital on Sustained Participation in Open Source. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. 688–699. <https://doi.org/10.1109/ICSE.2019.00078>
- [54] Huilian Sophie Qiu, Bogdan Vasilescu, Christian Kastner, Carolyn Egelman, Ciera Jaspan, and Emerson Murphy-Hill. 2022. Detecting Interpersonal Conflict in Issues and Code Review: Cross Pollinating Open- and Closed-Source Approaches. In *2022 IEEE/ACM 44th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS)*. IEEE, Pittsburgh, PA, USA, 41–55. <https://doi.org/10.1109/ICSE-SEIS55304.2022.9793879>
- [55] Afzalur Rahim and Dana M. Cosby. 2016. A Model of Workplace Incivility, Job Burnout, Turnover Intentions, and Job Performance. *Journal of Management Development* 35, 10 (Jan. 2016), 1255–1265. <https://doi.org/10.1108/JMD-09-2015-0138>
- [56] Naveen Raman, Minxuan Cao, Yulia Tsvetkov, Christian Kästner, and Bogdan Vasilescu. 2020. Stress and Burnout in Open Source: Toward Finding, Understanding, and Mitigating Unhealthy Interactions. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER '20)*. Association for Computing Machinery, New York, NY, USA, 57–60. <https://doi.org/10.1145/3377816.3381732>
- [57] Mehvish Rashid, Paul M. Clarke, and Rory V. O'Connor. 2019. A Systematic Examination of Knowledge Loss in Open Source Software Projects. *International Journal of Information Management* 46 (June 2019), 104–123. <https://doi.org/10.1016/j.jinfomgt.2018.11.015>
- [58] Neill Robson. 2018. 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 (ESEC/FSE 2018)*. Association for Computing Machinery, New York, NY, USA, 986–987. <https://doi.org/10.1145/3236024.3275441>
- [59] Jaydeb Sarker, Asif Kamal Turzo, and Amiangshu Bosu. 2020. A Benchmark Study of the Contemporary Toxicity Detectors on Software Engineering Interactions. In *2020 27th Asia-Pacific Software Engineering Conference (APSEC)*. IEEE Computer Society, 218–227. <https://doi.org/10.1109/APSEC51365.2020.00030>
- [60] Jaydeb Sarker, Asif Kamal Turzo, Ming Dong, and Amiangshu Bosu. 2023. Automated Identification of Toxic Code Reviews Using ToxiCR. *ACM Transactions on Software Engineering and Methodology* (Feb. 2023). <https://doi.org/10.1145/3583562>
- [61] Charlotte Schluger, Jonathan P. Chang, Cristian Danescu-Niculescu-Mizil, and Karen Levy. 2022. Proactive Moderation of Online Discussions: Existing Practices and the Potential for Algorithmic Support. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (Nov. 2022), 370:1–370:27. <https://doi.org/10.1145/3555509>
- [62] Nathan Schneider. 2019. Decentralization: An Incomplete Ambition. *Journal of Cultural Economy* 0, 0 (April 2019), 1–21. <https://doi.org/10.1080/17530350.2019.1589553>
- [63] Nathan Schneider. 2021. The Tyranny of Openness: What Happened to Peer Production? *Feminist Media Studies* (Feb. 2021), 1–18. <https://doi.org/10.1080/14680777.2021.1890183>
- [64] Shalom Schwartz. 2012. An Overview of the Schwartz Theory of Basic Values. *Online Readings in Psychology and Culture* 2, 1 (Dec. 2012). <https://doi.org/10.9707/2307-0919.1116>
- [65] Shalom H Schwartz. 1992. Universals in the Content and Structure of Values: Theory and Empirical Tests in 20 Countries. *Advances in experimental social psychology* 25 (1992), 1–65.
- [66] Rifat Ara Shams, Waqar Hussain, Gillian Oliver, Arif Nurwidyantoro, Harsha Perera, and Jon Whittle. 2020. Society-Oriented Applications Development: Investigating Users' Values from Bangladeshi Agriculture Mobile Applications. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS '20)*. Association for Computing Machinery, New York, NY, USA, 53–62. <https://doi.org/10.1145/3377815.3381382>
- [67] Vandana Singh, Brice Bongiovanni, and William Brandon. 2022. Codes of Conduct in Open Source Software—for Warm and Fuzzy Feelings or Equality in Community? *Software Quality Journal* 30, 2 (June 2022), 581–620. <https://doi.org/10.1007/s11219-020-09543-w>
- [68] spaCy. 2020. spaCy - Industrial-strength Natural Language Processing in Python. <https://spacy.io/>. Retrieved {2020-03-03} from <https://spacy.io/>
- [69] M. Stevenson. 2021. Cultures of Peer Production. In *The Handbook of Peer Production*, Mathieu O'Neil, Christian Pentzold, and Sophie Toupin (Eds.). John Wiley & Sons, Inc, Hoboken, NJ, 56–69.
- [70] Mai Ishikawa Sutton and John Ryan. 2021. DWeb Principles. <https://getdweb.net/principles>. Retrieved {2021-05-06} from <https://getdweb.net/principles>
- [71] Kaitlyn Tiffany. 2022. The Battle for the Soul of the Web. Retrieved {2023-06-27} from <https://www.theatlantic.com/technology/archive/2022/10/internet-archive-decentralized-web-web3-brewster-kahle/671647/>

- [72] Kristen Vaccaro, Christian Sandvig, and Karrie Karahalios. 2020. "At the End of the Day Facebook Does What It Wants": How Users Experience Contesting Algorithmic Content Moderation. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (Oct. 2020), 167:1–167:22. <https://doi.org/10.1145/3415238>
- [73] Ruben van Wendel de Joode. 2004. Managing Conflicts in Open Source Communities. *Electronic Markets* 14, 2 (June 2004), 104–113. <https://doi.org/10.1080/10196780410001675059>
- [74] Bogdan Vasilescu, Alexander Serebrenik, Prem Devanbu, and Vladimir Filkov. 2014. How Social Q&A Sites Are Changing Knowledge Sharing in Open Source Software Communities. In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing (CSCW '14)*. Association for Computing Machinery, New York, NY, USA, 342–354. <https://doi.org/10.1145/2531602.2531659>
- [75] Georg von Krogh, Stefan Haefliger, Sebastian Spaeth, and Martin W. Wallin. 2012. Carrots and Rainbows: Motivation and Social Practice in Open Source Software Development. *MIS Quarterly* 36, 2 (2012), 649–676. <https://doi.org/10.2307/41703471> jstor:41703471
- [76] Zhendong Wang, Yi Wang, and David Redmiles. 2018. Competence-Confidence Gap: A Threat to Female Developers' Contribution on Github. In *Proceedings of the 40th International Conference on Software Engineering: Software Engineering in Society (ICSE-SEIS '18)*. Association for Computing Machinery, New York, NY, USA, 81–90. <https://doi.org/10.1145/3183428.3183437>
- [77] Jon Whittle. 2019. Human Values in Software: A New Paradigm for Requirements Engineering?. In *2019 IEEE 27th International Requirements Engineering Conference (RE)*. 4–4. <https://doi.org/10.1109/RE.2019.00011>
- [78] Jon Whittle. 2019. Is Your Software Valueless? *IEEE Software* 36, 3 (May 2019), 112–115. <https://doi.org/10.1109/MS.2019.2897397>
- [79] Jon Whittle, Maria Angela Ferrario, Will Simm, and Waqar Hussain. 2021. A Case for Human Values in Software Engineering. *IEEE Software* 38, 1 (Jan. 2021), 106–113. <https://doi.org/10.1109/MS.2019.2956701>
- [80] Emily Winter, Steve Forshaw, and Maria Angela Ferrario. 2018. Measuring Human Values in Software Engineering. In *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM '18)*. Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3239235.3267427>