

# Maintaining Values: Navigating Diverse Perspectives in Value-Charged Discussions in Open Source Development

JACK JAMIESON, NTT Communication Science Labs, Japan  
 EUREKA FOONG, Tokyo College, the University of Tokyo, Japan  
 NAOMI YAMASHITA, NTT Communication Science Labs, Japan

Communication technologies have significant social impacts, and it is important to consider how designers' and developers' values shape their design. Increasingly, these technologies are released as continually evolving platforms and services, so their development involves ongoing discussions and debates about unforeseen problems and future directions. However, there is a gap in research about how designers, developers, and other stakeholders engage with values during later stages of development. We investigate discussions about values in the context of open source software development, focusing on projects related to the Decentralized Web. We conducted a large-scale analysis of GitHub issues among diverse yet ideologically-related projects. We show that the percentage of discussions about values increases later in development, and we identify features and outcomes of conflicts related to open source participants' values. Finally, we propose suggestions to improve upon existing discussion practices by supporting common ground among collaborators with diverse goals, perspectives, and experiences.

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

Additional Key Words and Phrases: values, open source, conflict, decentralized web

## ACM Reference Format:

Jack Jamieson, Eureka Foong, and Naomi Yamashita. 2022. Maintaining Values: Navigating Diverse Perspectives in Value-Charged Discussions in Open Source Development. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 449 (November 2022), 28 pages. <https://doi.org/10.1145/3555550>

## 1 INTRODUCTION

Communication technologies have profound impacts on our lives, and so it is important to consider how they support or mitigate the achievement of values, such as commitments to human dignity, liberty, and inclusion. Further, it is increasingly important to attend to values across the life-cycles of digital technologies, since they are routinely released as continuously evolving services. In many cases, challenges toward maintaining values in online networking technologies can emerge or be exacerbated by circumstances and decisions long after a technology's first deployment. This challenge has long been a theme in values-oriented HCI research [e.g., 27, 29, 46, 65, 74], such as the recognition of “emergent bias” when designs are introduced to new contexts [29]. In spite of this, researchers have identified that we know relatively little about how designers, developers,

---

Authors' addresses: Jack Jamieson, [jack@jackjamieson.net](mailto:jack@jackjamieson.net), NTT Communication Science Labs, Kyoto, Japan; Eureka Foong, [foong.eureka@tc.u-tokyo.ac.jp](mailto:foong.eureka@tc.u-tokyo.ac.jp), Tokyo College, the University of Tokyo, Tokyo, Japan; Naomi Yamashita, [naomiy@acm.org](mailto:naomiy@acm.org), NTT Communication Science Labs, Kyoto, Japan.

---

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](mailto:permissions@acm.org).

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-0142/2022/11-ART449 \$15.00

<https://doi.org/10.1145/3555550>

and other stakeholders actually engage with values in later stages of design and development, such as software engineering [72], and maintenance and repair [37].

Based on these challenges, we argue that understanding how designers, developers, and other stakeholders engage with values in later stages of design is necessary for long-term maintenance of values in technology. To that end, this study investigates discussions about values in the context of open-source software development on GitHub. Open source development is a continuous process, with traces of discussions and development activity published openly. To highlight values, we focus on the *Decentralized Web*, composed of a variety of software and protocols for building alternatives to the current internet. Specific goals vary across these projects, but they are generally unified around a proposal to minimize concentrations of power by a small number of corporate platforms, and increase agency of individuals and smaller-scale collectives [30, 33]. The Decentralized Web has been proposed as an opportunity for developers to “bake our values into our code” [43], and in 2021 the DWeb community (hosted by the Internet Archive) published a set of principles to guide the Decentralized Web’s development [67].

We investigated patterns of discussion about values in GitHub issues through **RQ1**: “*When do discussions about values occur in GitHub issues?*” We developed a vocabulary of terms related to the principles proposed by the Internet Archive’s DWeb community [67], using the Schwartz theory of basic human values [60] as a foundation. We used this dictionary to identify temporal patterns and other conditions of value-related discussions in GitHub issues for 51 Decentralized Web projects.

Since GitHub issues are used to resolve problems and plan future directions, disagreements are common. Thus, we ask **RQ2**: “*How do people experience disagreements in value-related discussions on GitHub?*” We addressed RQ2 by interviewing 12 individuals who had posted or commented on value-related issues in our dataset. We investigated what kinds of outcomes participants in these discussions expect of them, and how various dimensions of conflict [24] occurred in GitHub issues and related discussions.

We then examined the consequences of these discussions through **RQ3**: “*What are the outcomes of value-related conflicts and current practices to address them?*” We identified both positive and negative consequences, such as extending participants’ understanding of ethical issues related to technological development, and how negative emotional experiences contribute to decreased or abandoned participation.

Across these analyses, we describe how diverse project contributors discuss the DWeb’s foundational values, and we develop insights to support productive and collegial discussions about values in future work. We use the term “diverse” with the same framing as DiSalvo to refer to “varied conditions and experiences” rather than to differences in culture, race, gender, or other such classifications [16, p. 21]. We consider those forms of diversity to be important, but they are not the focus of the present study. Specifically, the contributors we studied have a variety of motivations, stakes, and roles in relation to DWeb technologies—e.g., paid employees, third-party developers, volunteers, end-users, etc.—which contributed to frequent disagreements and conflicts in value-related discussions.

We found that the proportion of discussions in GitHub issues that referenced values increased after the first year of development, reinforcing the need to attend to values during later stages of design and development. We also found the proportion of value-related related comments was positively associated with the number of individuals contributing to a discussion, suggesting that there may be more discussion about values when a greater number of viewpoints is present among stakeholders. Our examination of conflicts showed that value-related disagreements had particularly high stakes since they were routinely viewed as establishing precedent for future normative decisions. Additionally, process conflicts were often exacerbated by mismatched expectations about

project goals and processes for addressing issues, which sometimes contributed to decisions to cease contributing to projects.

This research makes the following main contributions. First, we demonstrate a large-scale analysis of value-related discussions among diverse yet ideologically related projects. This included extending upon prior work about detecting values [54] by tailoring a dictionary of value-related terms to the priorities and linguistic patterns of software contributors in a specific context (projects associated with the decentralized Web). Second, we extend upon past CSCW scholarship about conflict in open source development [23, 24] by identifying features and outcomes of conflicts related to open source participants' values. Through this research, we contribute to knowledge about a complex facet of online collaborative work: The relationship between conflict and values. Based on our findings, we identify opportunities to improve common ground among open source contributors with different motivations and backgrounds, which could contribute to more robust and productive dialogues about values during collaborative development work.

## 2 BACKGROUND

To situate the present study, we review related work about values in collaborative design and development, conflict in open source communities, and the Decentralized Web.

### 2.1 Values in collaborative design and development

A significant body of scholarship has investigated ways in which values are implicated in design decisions [e.g., 27, 29, 65]. Research about values in the context of design focuses on supporting “human well-being, human dignity, justice, welfare, and human rights” [28, p. 1186]. Much of this research has described values as being instantiated in technological artifacts such that their central uses tend to support the realization of specific values [6]. Values and design scholars have asserted that the realization of values is not deterministic, but achieved through interactions between a technology and its users (or other stakeholders) [28, 65]. This renders the work of incorporating values in design as contingent, since outcomes depend on the contexts in which technologies are encountered. Many approaches to values in design emphasize earlier stages of a technology, focusing on anticipating potential future impacts related to designers' intended values [e.g., 7, 42, 63]. However, there is a need for more attention to later stages of technical work such as repair [37] and software engineering [72]. For example, Whittle [72] states that although values-oriented methods are well established in HCI and information systems research, “HCI and information systems do not deal with the business of actually building software, so although they could apply in the early stage of software engineering, they offer little guidance as to how to handle values in the more technical stages of development” (p. 114).

During these “more technical stages,” developers and other stakeholders navigate complex situations with no single, clear solution. To understand this, researchers increasingly interpret values as processes [37, 40, 64]. JafariNaimi et al. [40] asserted that designers do not *apply* pre-defined values to these situations, but instead *employ* values as hypotheses, which leads to richer understandings of both the situations and the nature of values. Accordingly, values should not be considered to be determined in advance of the more technical stages of design, but rather to change and evolve through that work. For example, Iversen et al. [39] described how values emerge and are developed through collaboration in participatory design. They further point out that such fluid and dynamically changing values can lead to conflicts among stakeholders in the design process.

Thus, a significant aspect of engaging with values during design and development is addressing conflicting views among different stakeholders. Researchers have investigated how deliberation and disagreements about values can be navigated and resolved [49, 63] and can be generative for

building consensus [10]. However, the process of resolving value conflicts through collaboration is fragile, as its success depends on the willingness of stakeholders to participate in this process[39].

Based on this background, the present study adopts a process-oriented perspective to investigate how open source developers and other stakeholders navigate discussions and conflicts about values over extended periods of development and maintenance.

## 2.2 Conflict in open source communities

Past work [1, 23, 24, 68] has distinguished between different forms of conflict in open source development and their effects on outcomes such as project performance and identification as part of a team [24]: *task conflict*, *relationship conflict*, *process conflict*, and *normative conflict*. In this section, we briefly summarize the nature and effects of these types of conflict.

*Task conflicts* are disagreements about what goals should be pursued, and what decisions should be made to do so [41]. In moderation, task conflict can be productive since it can contribute to greater understanding about the work being done [68]. However, some studies have failed to identify significant associations between task conflict and team performance [2, 15, 24], and others have observed that task conflicts frequently evolve into or co-occur with other, more disruptive, types of conflicts [14, 23].

*Relationship conflicts* (i.e., affective conflicts) involve interpersonal disagreements and their resulting emotional tensions [24]. Relationship conflicts have been identified with poor quality work and lower productivity in traditional organizations and among Wikipedia editors [2]. However, they were positively associated with team identification and perceptions of team performance in a study of open source developers, with a possible explanation being that relationship conflicts may indicate higher levels of emotional involvement with a project or be indirect indicators of project activity [24].

*Process conflicts* involve disagreements about *how* to perform a task, such as delegation of resources and duties [41]. Past work has struggled to assess the influence of process conflict on work outcomes, since it co-occurs with other conflict-types such as task and relationship conflict [23]. There is partial support for the hypothesis that process conflict is negatively associated with project performance [24].

*Normative conflict* involves disagreements about group values and expected behaviour (prescriptive norms) and actual behaviours by group members (descriptive norms) [55]. They occur at a higher level of abstraction than other conflicts in open source teams, since they involve “disagreements about issues like project policies, governance structures, and project ideology” (p. 1400) [23, p. 1400]. Further, normative conflicts have been identified to have strong negative effects on open source contributors’ identification with their team and perceptions of team performance [24].

We examine conflicts related to values in open source development. In practice, these different types of conflict can overlap, and sometimes one type of conflict evolves into or triggers others [2, 38]. Because we are particularly concerned with projects’ values, normative conflicts are a major consideration for our analysis, and we expect to observe them co-existing with other conflict types. We investigate how discussions and conflicts about values in GitHub issues affect development processes and outcomes, focusing on projects related to the Decentralized Web. In the following section, we described the Decentralized Web and explain why it is a rich site for our analysis.

## 2.3 Decentralized Web

The Decentralized Web is an emerging set of protocols, software, and approaches for a new internet that moves away from centralized control by a small number of tech giants, and toward peer-to-peer and community driven infrastructures. These goals are pursued through the development of technical, structural interventions, including Blockchain, ActivityPub, and a variety of other system

[3]. Projects in this area share general commitments to openness and decentralization, however they lack consensus on specifics of what decentralization and openness entail [4, 59]. In fact, the indeterminacy of “decentralization” can serve a rhetorical purpose, inviting in “diverse participants and diverse rationales in complex, heterogeneous undertakings” and “obscuring the contradictions at play – at least temporarily” [59, p.11]. One of the hubs for shaping the Decentralized Web is the Internet Archive’s DWeb community, which hosts events for Decentralized Web developers and has published a set of principles that “define the values of a decentralized web based on enabling agency of all people” and building “a more just and equitable world” [67].

There are many challenges for maintaining the values of the Decentralized Web. Raman et al. [56] identified pressures from multiple levels toward recentralized network structures. Halpin [34] described that over a decade of efforts to build decentralized standards through the IETF and W3C failed to constrain centralization of the Internet by large corporations, due significantly to fragmentation and too much complexity among competing standards. Further, in their focus on technical implementations, makers of decentralized systems may fail to integrate their work into surrounding regulatory and economic systems or to ensure their designs are aligned with users’ actual needs [51]. Ultimately, Decentralized Web technologies constitute a productive context for our analysis because their construction is driven by values, shaped by a variety of stakeholders, unfolds over an extended period of time, and is largely open source.

### 3 METHOD

This study used a combination of methods. The first stage of the study, described in Section 3.1 consisted of defining a taxonomy of values related to the DWeb principles, collecting a dataset of issues from Decentralized Web projects, and identifying patterns of discussions related to value keywords. In the second stage, described in Section 3.2, individuals who participated in substantial value-related discussions were identified, and then invited to participate in interviews about those discussions and, more generally, their experiences engaging with values in relation to these projects.

#### 3.1 GitHub issues analysis

*3.1.1 Defining a values dictionary.* We developed a dictionary of keywords related to important values for the decentralized Web. We used a set of principles published by the Internet Archive’s DWeb community as a base [67]. It is important to acknowledge that not every Decentralized Web project adheres to these principles, nor is every contributor familiar with them. Nonetheless, they provide a foundation for framing the basic ethical commitments of these projects, and developers from several Decentralized Web projects have signed the principles to indicate support.

To focus our analysis, we identified the most salient values across the DWeb principles. Since they serve more as a guideline for practitioners than as a taxonomy for researchers, the principles themselves regularly reference the same values across multiple items. Thus, we used a secondary frame of reference to identify where distinct concepts occurred across multiple principles, and then generate a list of the most important values. We adopted the Schwartz theory of basic human values for this purpose [61]. Schwartz’ theory presents a taxonomy of 58 human values in 10 categories, as presented in Table 1. This structure of distinct values has been validated through research in 82 countries with samples that are diverse in terms of geography, culture, language, religion, age, gender, and occupation [60]. Additionally, Schwartz’ framework has been widely adopted across many fields, including research about values in computer science and software engineering [e.g., 22, 47, 54, 71, 73, 75]. Based on its thorough validation and adoption in related research, Schwartz’s theory provided a suitable foundation for our analysis.

Two researchers separately coded each sentence from the DWeb principles to one or more values from the Schwartz theory of human values. Additionally, we assigned a polarity to each value

Table 1. Taxonomy of Schwartz' Theory of Basic Human Values [61].

Note: Privacy is not part of Schwartz's original taxonomy, but was included for its relevance in communication technologies and because it was included in past research using a similar method [53].

Category	Values
Self-direction	Freedom, creativity, independence, choosing own goals, curiosity, self-respect, privacy.*
Stimulation	Excitement in life, a varied life, daring.
Hedonism	Pleasure, enjoying life.
Achievement	Ambitious, influential, capable, successful, intelligent.
Power	Wealth, authority, preserving my public image, recognition, social power.
Security	National security, family security, sense of belonging, social order, healthy, clean, reciprocation of favours.
Conformity	Obedient, self-discipline, honouring of parents and elders, politeness.
Tradition	Respect for tradition, devout, detachment, humble, moderate, accepting my portion in life.
Benevolence	Helpful, responsible, forgiving, honest, loyal, a spiritual life, true friendship, mature love, meaning in life.
Universalism	Equality, wisdom, inner harmony, a world of beauty, social justice, broadminded, a world at peace, unity with nature, protecting the environment.

reference (either supporting or opposing that value). Most references to values were positive, but given that the core motive of decentralization is to avoid centralized concentrations of power, it is unsurprising that references to *social power*, which Schwartz defines as “control over others, dominance” [61], were negative. For example, the statement, “High concentration of organizational control is antithetical to the decentralized web” was coded as opposing *social power*. By contrast, the statement, “We encourage the development of tools and applications in many languages and forms, with a high degree of accessibility” was coded as supporting *broadmindedness* (by encouraging many languages and forms) and *equality* (by emphasizing accessibility). After the first pass, agreement was somewhat low (64% of tags for value categories and 55% for individual values). In most cases, however, the coders were in agreement about at least one value for each statement, which was usually the core value of that point (82% of value categories, 77% of individual values). Disagreements were resolved through discussion until consensus was reached.

We slightly tailored our coding to reflect the language used in the DWeb principles. Specifically, we replaced *politeness*, which Schwartz defines as “courtesy, good manners” [61, p. 61] with *respectfulness*, which has a slightly richer meaning and aligns with the way in which ‘good manners’ are presented in codes of conduct that are common in open source communities. Following Obie et al. [54], we also added a category called *privacy* – while in many instances references to privacy can be mapped to related values (e.g., individual freedom), it is routinely presented as a central concept in technology ethics in its own right. Additionally, during data analysis it became clear that although *equality* (“equal opportunity for all” [61]) and *social justice* (“correcting injustice, caring for the weak” [61]) are distinct concepts, the words used to describe them frequently overlapped in our dataset. We combined them into a single category called *equity and equality*.

We selected the top 6 most frequently occurring values, which each matched with four or more of the 23 statements in the DWeb principles. Both coders agreed that these values represented the main priorities of the document. We validated this interpretation by consulting with one of the

co-stewards of the principles, who agreed and encouraged us to include *protecting the environment* in our analysis. Initially, we had excluded *protecting the environment* because it was only mentioned twice in the principles. However, it was presented prominently (with a dedicated section) so we agreed it was valuable to include.

- 5 mentions each: equality, social justice, broadmindedness.
- 4 mentions each: freedom, respectfulness, (opposing) social power.
- 2 mentions: protecting the environment.

We developed a vocabulary of terms related to these values using the following process: We adapted our vocabulary from one developed by Obie et. al [54], in which two authors created a vocabulary of terms related to 50 Schwartz values, and validated it against a human-coded truthset identifying references to values in Google Play Store user reviews. This was an appropriate base for our vocabulary because it included terms related to a wide range of values relevant to a software context. We removed some terms that had multiple meanings and would likely lead to false matches if mapped to discussions among open source developers. For example, Obie et. al's dictionary included the term "equal" as related to the value of *equality*, but in a software development context "equal" more commonly refers to comparisons of variables with no salience toward social equality. Additionally, we added: synonyms generated using Empath [21] and from Dictionary.com; some terms that are specifically relevant to open source software, such as "code of conduct" as a signifier of *respectfulness*; added variations of verb tenses that were not accommodated by spaCy's lemmatizer; and extended the *equity and equality* category with relevant social justice terms published by Dalhousie University's Human Rights and Equity Services [62].

**3.1.2 Data Collection: GitHub issues.** GitHub is the most popular site for hosting the development of open source software. GitHub repositories contain a detailed history of revisions (called *commits*) to source code and other files. Additionally, they contain *issues*, which are discussion threads for bug reports, feature requests, and related conversations. Often, a single piece of software comprises multiple, related repositories, usually maintained by the same organization or individual [44].

Our dataset was composed of GitHub archives from relevant decentralized Web projects that have (a) been presented at events hosted by the Internet Archive's DWeb community, (b) have published statements indicating ethical commitments that closely align with the DWeb principles, and/or (c) were included in the *Decentralized Tech Ecosystem Map* presented by the Digital Life Collective at the 2018 Decentralized Web Summit [11]. Based on this criteria, we identified 51 projects that constitute a cross-section of decentralized Web projects with related, though not identical, ethical commitments. In most cases, each project was operated by a separate GitHub *organization* (a shared GitHub profile representing a company or other entity), so we included every repository operated by that organization as part of the same "project." Some projects were operated by individuals, who were likely to also maintain irrelevant repositories (e.g., for unrelated software), so we used our judgement to only include relevant repositories.

We downloaded all issue-related events from these repositories over several sessions between June 9, 2021 and June 25, 2021. Although *pull requests* (mechanisms for submitting code changes) are structured as a type of issue in GitHub's API, we excluded them ( $n = 163,633$ ) from our analysis because their practical function is significantly different from other issues. Our dataset consisted of 388,037 comments. Table 2 presents an overview of the objects represented by the dataset.

**3.1.3 Validating and refining the dictionary.** To verify that our values dictionary (developed in Section 3.1.1) could identify comments related to values in the main dataset (collected in Section 3.1.2), we used the following process: First, the values dictionary and the issue comments in the main

Table 2. Explanation of objects represented in our dataset of GitHub issue comments. Hierarchical relationships are indicated by indentation.

Unit	N	Description
- Project	51	We use the term <i>project</i> to refer to the collection of repositories maintained by a single GitHub organization (or rarely, individual maintainer) as part of a single platform, software, or set of closely interrelated technologies.
↪ Repository	1,426	A GitHub repository contains all the files for a single piece of software (or other entity), a record of changes to the code, and all of its issues.
↪ Issue	97,662	An issue is a discussion thread, typically used to report bugs or request features.
↪ Comment	388,037	A comment is a single post to an issue's discussion thread.
- User	21,359	A user is a unique person who posted a comment to an issue in our dataset.

dataset were tokenized and lemmatized using spaCy [66], allowing us to match words across multiple grammatical forms. Then, we identified sentences (and their parent comments) that referenced terms from the values dictionary. We selected a sample of 839 sentences that were matched to the values dictionary. Two authors coded each sentence by hand to validate that it indeed included a reference to that value. We focused on comments posted after January 1, 2020 because we also used this step to identify potential interview participants, as described in Section 3.2.1. Disagreements were resolved iteratively through discussion. We then removed terms from the dictionary where the proportion of correctly identified matches was less than 50%. Among the remaining terms, the mean percentage of correctly identified matches was 78.2%. Terms that scored closer to 50% tended to have multiple possible meanings, such as "user choice," which was used to advocate for enhancing the ability of individuals to choose how to use software (a correct match for *freedom*) or to simply acknowledge that the user may happen to choose a certain option (a false positive for *freedom*).

The final dictionary is included with this paper as a supplementary document.

**3.1.4 Analysis.** Once we had validated the dictionary, we flagged each comment with whether or not it contained a keyword for each value. Based on this, we generated summary statistics mapping the incidence of value-related keywords per project, per comment author, per issue, and per comment.

We investigated the occurrence of value-related issue comments across two temporal dimensions: absolute time (by calendar date), and project age (time since the first event that was recorded in our dataset for each project). We used visualizations to identify the percentage of each month's comments that contained value-related terms across these dimensions. We then investigated whether the percentage of value-related comments per project was associated with the number of users, since a greater number of people writing comments could contribute to greater diversity of values and normative goals. To assess this, we used logistic regression models for each value (and one overall model indicating if there was a match for *any* value). In these models, each observation represented a single month's activity for one project. The dependent variable for each model was whether or not a comment was posted in relation to that model's value, and the independent variables were the number of comments and number of users for that month's activity per project.

Additionally, we created project-level charts (e.g., Figure 3) showing the frequency of comments that mentioned each value per project over time. We used these to inform some questions during the interview, as described in Section 3.2.2. Additionally, to help understand the relationship between



Table 3. Demographic information of interview participants. Some participants preferred to give less precise responses, indicated by \*.

Age group	N (%)	Gender	N (%)	Location of residence	N (%)
20-29	1 (8.3%)	Female	2 (16.7%)	Australia	1 (8.3%)
30-39	6 (50.0%)	Male	9 (75.0%)	Belgium	1 (8.3%)
40-49	4 (33.3%)	Prefer not to say	1 (8.3%)	Canada	3 (25.0%)
Legal adult*	1 (8.3%)			Czech Republic	1 (8.3%)
				Luxembourg	1 (8.3%)
				Portugal	1 (8.3%)
				United States	3 (25.0%)
				Europe*	1 (8.3%)

value-related discussions in GitHub issues and continued participation, we investigated whether there was a relationship between the proportion of comments containing a value-term and whether the author of the comment had since deleted their GitHub account. To do so, we used Logistic regression with a dependent variable of whether or not the comment contained a value-related term, and an independent variable indicating whether the comment-author had deleted their account. Since the data is strongly skewed toward comments that are not value-related, the Firth procedure was used to reduce the risk of parameter estimation bias [12, 25].

## 3.2 Interviews

**3.2.1 Recruitment.** We reviewed all sentences that contained matched value terms from comments posted after January 1, 2020. We then identified participants who were engaged in debates and discussions about values. This included both participants who were identified directly as having used the value terms, as well as some who did not use a term from our dictionary but participated in relevant discussion with another user who had. In selecting interview participants, we sought to reach people who engaged in discussions that referenced a variety of different values.

We sent an invitation to each prospective participant by email if they had published an email address on their GitHub profile or a linked website, or by Twitter or Mastodon for some who had not listed an email address. Twelve interview participants were recruited (see Table 3 for their demographic information).

**3.2.2 Interview overview.** All of the participants had participated in a value-related discussion in the issues for a project in our dataset. Six of the participants were people who had contributed code to the corresponding project or were members of the organization that maintained that project (*code contributors/organization members*). The other six were people whose only contribution to that project on GitHub was through posting to that project's issues (*issues commenters*). The purpose of this distinction was to have a balance of core and peripheral contributors.

The interviews were structured as follows: First, we asked about participants' overall experience with open source projects, including their motivations. We then asked about their role in relation to the specific Decentralized Web project in which they had written an issue or comment, as well as their understanding about whether that projects' goals had shifted over time. This was followed by asking if they were familiar with the DWeb Principles, and their opinion of how these aligned with project goals. The next questions focused on the specific issue in which participants had commented. We asked about their understanding of the stakes of the discussion, their motivations for posting, and their perception of the outcome. Finally, during interviews with *code contributors/organization*

*members*, we added additional questions about changes in the frequency of value-related comments over time. These questions referenced project-level visualizations as described in Section 3.1.4. We did not include these questions for interview participants whose only participation was commenting on or posting an issue because their participation was much less likely to have led to familiarity with project trends over time.

*Issue commenters* were compensated £26. Because *code contributors/organization members* participated in a longer interview, they were compensated £45. Five participants declined payment. Five interviews were conducted via video call, five by email, one by Twitter direct message, and one was a hybrid, email followed by a video call. The length of the video interviews was 30 minutes for issue commenters and 45 minutes for code contributors. Two code contributor interviews extended into longer, 2-hour conversations.

**3.2.3 Analysis.** Video interviews were transcribed using an automatic transcription tool. Transcripts were then reviewed and corrected by the first author. Two authors used an inductive, initial-coding method [58] to identify themes in the transcripts and logs of text-based interviews. The authors then used affinity diagramming [35] to organize these themes into larger categories. Through a process of iterative refinement, the authors developed a taxonomy that organized codes under the following categories: *Values, motivations, expectations, processes, and outcomes*.

### 3.3 Limitations

The linguistic analysis has limitations due to its reliance on keywords. Specifically, it is unable to identify instances where a comment is implicitly *about* a value but does not contain a matching term. Additionally, our system was unable to disambiguate between homophones. For example, terms such as “pollution” and “environment,” were removed from the dictionary for *protecting the environment* because their most common uses referred to non-ecologically focused meanings. Further, our keyword approach cannot disentangle discussions about values in project outcomes (e.g., enacting ‘respectfulness’ toward end-users) from discussions about values in the open-source development process (e.g., encouraging respectful discussions in GitHub issues).

Most projects in our dataset are based in Western countries, which limits the generalizability of this study to other cultures. Similarly, our straightforward adoption of Schwartz’ theory as a taxonomy of values does not do justice to nuanced conceptions of values in every culture or by every individual. For example, Maori cultural values involve a close relationship among land, family, spirituality, and sovereignty [13]. If our dataset includes comments by Maori people speaking of environmentalism through that lens, those comments would be poorly represented by our simple classification of *protecting the environment* as a value in the *universalism* category. Thus, even though our use of Schwartz’s framework was well-suited to summarize the main priorities of the DWeb principles, the resulting analysis cannot capture the full depth of values across cultures.

Another important consideration is that our research method drew our attention toward issues with lots of comments. Cases where a value-driven proposal is implemented without being challenged are less likely to have been captured in our analysis than those with more back-and-forth discussion, since this increases the likelihood of a keyword-match.

A final point about the linguistic analysis is that our validation of the dictionary accounted for false positives, but did not identify false negatives. We have offered some suggestions for future research to improve upon the linguistic analysis in Section 5.3.1.

Regarding the interviews, all participants had demonstrated a belief that software development should involve a consideration of values insofar as they were identified as having discussed values in GitHub issues and they accepted our invitation to participate in this study. Accordingly, interview participants do not form a representative sample of developers and issues commenters. Due to

Table 4. Proportion of projects, users, issues, and comments with at least one identified reference to values.

Value	Top terms	Freq. of items with at least one reference to values by:			
		Project	User	Issue	Comment
Respectfulness	friendly, cooperative, harassment	47.06% (N = 24)	1.57% (N = 335)	0.41% (N = 404)	0.16% (N = 610)
Broadmindedness	inclusive, diversity, diverse	62.75% (N = 32)	1.47% (N = 313)	0.48% (N = 470)	0.14% (N = 546)
Freedom	freedom, user choice, sovereign	66.67% (N = 34)	1.11% (N = 238)	0.33% (N = 321)	0.10% (N = 389)
Equity and equality	equity, racist, racism	33.33% (N = 17)	0.60% (N = 128)	0.15% (N = 147)	0.05% (N = 210)
(opposing) Social power	central authority, gatekeeping, gatekeeper	43.14% (N = 22)	0.44% (N = 93)	0.10% (N = 97)	0.03% (N = 113)
Protecting the environment	climate change, energy consumption, ecological	21.57% (N = 11)	0.11% (N = 23)	0.03% (N = 27)	0.01% (N = 27)
<b>Any value</b>		<b>76.47%</b> (N = 39)	<b>3.79%</b> (N = 809)	<b>1.36%</b> (N = 1327)	<b>0.47%</b> (N = 1841)
Total N		N = 51	N = 21,359	N = 97,662	N = 388,037

this, our interpretations about how to support value-driven discussions and decision-making in open source emphasize the needs of people who are *currently* engaging with values in software development.

## 4 FINDINGS

When presenting findings, we take several measures to protect participants' anonymity. We have assigned each a random number, which is used to attribute quotations (e.g., *P01*, *P02*, etc.). We use the singular 'they' pronoun to refer to participants rather than identifying them by gender. We have also redacted references to project names. This protects participant anonymity in cases where members of small communities could infer someone's identity from the content of a quotation. Additionally, this helps us focus on the processes through which value are engaged with rather than amplifying criticisms of specific technical decisions and specific developers, which is not the purpose of this research. To add context to participant quotes, we additionally identify participants according to the following three roles in relation to the project about which they had commented:

- *Issue contributor*: People who only engaged with the project by posting and/or commenting on issues, regardless of whether they had coding experience elsewhere (n = 4).
- *Organization member*: People who were members of the GitHub organization that maintains the project (n = 4).
- *Third-party developer*: People who wrote third-party software that interacted with the project (n = 4).

### 4.1 Overall patterns of value-related discussions in GitHub issues

This section addresses **RQ1**: "When do discussions about values occur in GitHub issues?" Table 4 shows each value's three most frequently occurring terms, as well as the percentage of projects, users, issues, and comments with at least one mention for each value. Projects, issues, and comments

have a hierarchical relationship (as illustrated in Table 2), thus the *project* column has a high score because each project contains many issues, and each issue contains many comments.

Overall, value-related comments were relatively uncommon (approximately 1 in each 200 comments). Our readings of value-related discussions showed that a single comment that explicitly referenced one of the value terms in our dictionary was often a good indication that the issue discussion more generally related to that value. The three most commonly identified values were *respectfulness* (n comments = 610), *broadmindedness* (n comments = 565), and *freedom* (n comments = 389). The least referenced by far was *protecting the environment* (n comments = 27). Importantly, this analysis did not distinguish conversations about values in technologies themselves from conversations about values in processes of software development. Notably, we observed many comments urging commenters to be respectful in GitHub issue discussions themselves, regardless of the role of *respectfulness* in the project’s technology.

**4.1.1 Temporal patterns.** The earliest event in our dataset occurred on January 26, 2011, and the most recent occurred on June 23, 2021 (the date of our data collection). Even though the raw data spans approximately 10.5 years, very few of the projects in our dataset existed in 2011. Figure 1a shows that most projects became active during or after 2015. Corresponding to this, Figure 1b shows that most projects have six or fewer years of activity. Thus, visualizations related to calendar time focus on 2015–2021, and visualizations in relation to project age focus on each project’s first six years, starting from its earliest event in our dataset.

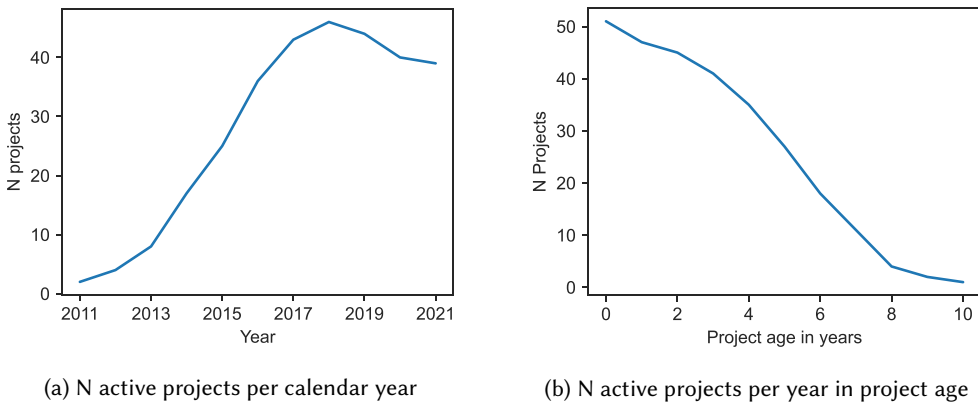


Fig. 1. The majority of projects in our dataset became active during or after 2015, and thus have been active for less than six years.

We then investigated the frequency of comments containing value-related terms over time. Figure 2a shows that the percentage of comments that referenced values was lowest in early 2016. In other periods, the percentage of value-referencing comments was fairly consistent, with a general pattern of increasing into 2018. Figure 2b shows a related pattern, with a lower-than-average percentage of comments that contain value terms in the first year of each project, and an increase in year two. The popularity and visibility of the Decentralized Web was increasing around 2016, which led to an increase in new GitHub projects and more people participating with those projects. Thus, the patterns observed in relation to calendar time and in relation to project age are intertwined.

Interview participants’ reported that project values sometimes fell out of alignment with individual contributors’ goals. This may explain some increase in values-related discussions over time. Some projects changed in overall direction, such as described by P03 (third-party developer): “When

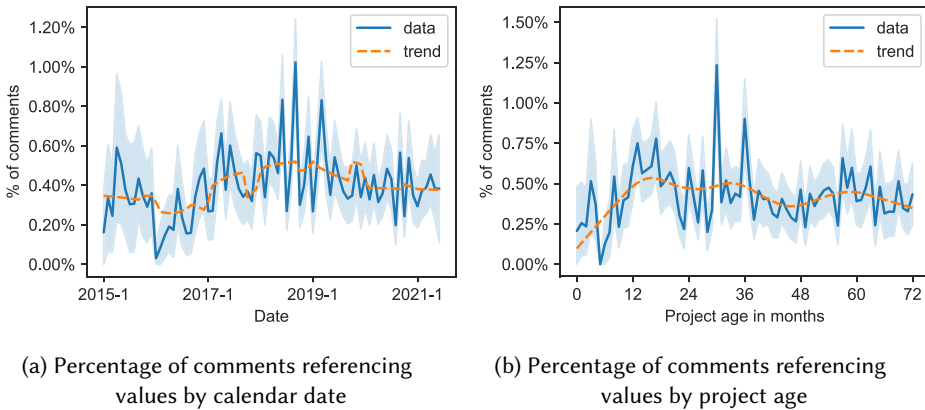


Fig. 2. Left: % of comments referencing any value per calendar month; Right: % of comments referencing any value per month in project age. The percentage of comments that referenced values was lowest in early 2016, and increases in the second year of a project.

we joined, it was a lot more focused on building applications around the concept of keeping ownership of data, etc. [... but later ] they focused a lot more on the currency, on the [smart]-contracts part and less on the technology that deals with the actual data.” Similarly, P08 (organization member) asserted that even when goals remain consistent, practical constraints may guide developers to make compromises: “Goals and intentions and philosophies haven’t changed. But practical realities maybe have reared their heads a little bit more.” By contrast, P04 (issue contributor) remarked that even when projects do *not* change, some people

project their own values and goals onto the project. Then at some point for some people it becomes obvious that their expectations are different to the reality and they react either by deciding they were mistaken or, more often, deciding that they have been betrayed and the project has changed. [P04 - Issue contributor]

All of these possibilities could lead to disagreements about a project’s direction. Value-related terms may be used in these discussions as a rhetorical device to advocate a particular direction. Additionally, disagreements of all sorts could prompt appeals to values of *respectfulness* within GitHub issues, particularly if discussions become heated.

**4.1.2 Diversity of perspectives.** Additionally, we investigated whether the percentage of value-related comments was associated with the number of users, since a greater number of people writing comments could contribute to greater diversity of values and normative goals. Logistic regressions were performed to identify associations between the number of comments per project per month, the number of users who posted those comments, and the likelihood that at least one comment contained a value-term. The number of users and number of comments were converted to z-scores to aid interpretation of effect sizes. Projects’ monthly number of users and number of comments were moderately correlated, with a variance inflation factor (VIF) of 4.35. The results of these regressions are displayed in Table 5. It is unsurprising that the likelihood of value-related comments being posted was proportional to the overall number of comments. However, our results show that even when controlling for the number of comments, a larger number of *people* posting comments was also positively associated with a greater likelihood of any value-related comment being detected (OR = 2.06, 95% CI = [1.57,2.71]). Most notably, mentions of *respectfulness* were much more likely for each increase of one standard deviation in the number of users (OR = 3.53,

Table 5. Logistic Regression models: Likelihood of any value mention per project per month.

\*\*\*  $p < 0.001$ , \*\*  $p < 0.01$ , \*  $p < 0.05$ . Pseudo  $R^2 = \text{McFadden's Pseudo } R^2$ 

Value	N users		N comments		Pseudo $R^2$
	Odds ratio	95% CI	Odds ratio	95% CI	
Respectfulness	<b>3.53***</b>	(2.77, 4.50)	1.02	(0.81, 1.29)	0.30
Broadmindedness	<b>1.33**</b>	(1.08, 1.63)	<b>2.37***</b>	(1.93, 2.90)	0.22
Freedom	1.11	(0.90, 1.37)	<b>2.55***</b>	(2.07, 3.14)	0.21
Equity and equality	<b>1.79***</b>	(1.39, 2.30)	<b>1.56**</b>	(1.21, 2.01)	0.24
(opposing) Social power	0.94	(0.73, 1.21)	<b>2.69***</b>	(2.11, 3.45)	0.23
Protecting the environment	<b>1.85*</b>	(1.13, 3.04)	0.94	(0.52, 1.71)	0.10
<b>Any value</b>	<b>2.06***</b>	(1.57, 2.71)	<b>3.80***</b>	(2.93, 4.94)	0.31

N observations = 2,604

95% CI = [2.77, 4.50]), but were not more likely based on an increase in the number of comments while holding the number of users steady. Only *freedom* and *(opposing) social power* were only associated with the number of comments, and not with the number of users. These results show that, for most values, having more participants was associated with more value-related discussion, independent of the total number of comments.

Our interviews revealed some plausible explanations for why this may have occurred. For example, P08 (organization member) asserted, “The ethical conversations are a heck of a lot more nuanced, the greater diversity of voices that are included in the conversation.” Although this is not guaranteed, a greater *number* of voices likely suggests the presence of commenters from outside a core team, which could introduce more disagreements. For example, P11 recalled a case where a (non-DWeb) software removed a feature:

People using this feature, like me, were angry. But the design team said they ran a survey and nobody was using the feature. Turn[s] out they ran the survey... inside the design team itself. [...] When the decision reach[ed] other users, the mistake [was] immediate. [P11 - Issue contributor]

These comments offer some hints that conversations with many participants may be prone to disagreements about project decisions and more nuanced discussions about project values.

**4.1.3 Patterns in individual projects.** Patterns of discussion among individual projects were more varied. In some cases, specific values became a more significant focus of attention in relation to specific challenges or phenomena. For example, Figure 3 shows that, in one project, comments containing terms associated with *equity and equality* were clustered toward the end of 2019 and into 2020. This corresponds to a period in which project contributors were troubleshooting an funding initiative to incentivize third-party developers to build apps. At that time, there were conversations on GitHub discussing how to maintain equality and fairness when distribute funding. These discussions did not lead to a solution, and the program was paused in February 2020. This example highlights that it could be possible to monitor developer discussions to identify when and where specific values are at stake. Notably, in instances such as this, it is possible to identify bursts of activity, where values that were not relevant or were considered settled in early stages of design became objects of considerable attention in relation to events that occurred later in development.

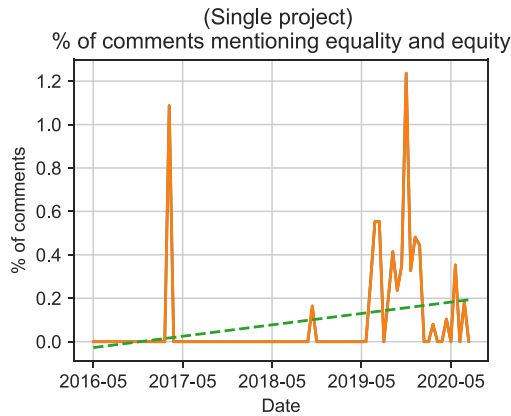


Fig. 3. In one project, there was a burst of comments referencing *equality and equity* were grouped toward the end of 2019 and beginning of 2020.

## 4.2 Participants' experiences of conflicts

This section addresses **RQ2**: “How do people experience disagreements in value-related discussions on GitHub?” Because our analysis focused on issues that had substantial discussions, most of our participants described some level of conflict. Given our focus on values, all of the issues involved some normative claim, proposing decisions with ideological roots. As a result, the stakes of conflicts were high. This section describes how task conflicts about *what* should be done – and process conflicts about *how* decisions should be made each alluded to larger stakes than the immediate task at hand in individual issues.

**4.2.1 Task conflicts.** Many issues were rooted in a task conflict. P12 (issue contributor) posted an issue advocating that one project moved its code repositories away from GitHub and toward a FOSS (Free and Open-Source Software) alternative. The resulting discussion included disagreements about the pros and cons of doing so, and ultimately resulted in no change. P03 (third-party developer) posted a comment to an existing issue about distributing funding for third-party developers, presenting an argument about how to resolve perceived fairness issues in the existing system. P11 (issue contributor) replied to an issue that advocated for removing integration to a third-party software by asserting that the feature should remain. These are all examples of task conflicts in which participants debated about what actions to take by explaining the benefits and demerits of various approaches.

One way that task conflicts appeared to take on elevated stakes was when participants who posted an issue felt that a project was not living up to the values that had been promised. For example, P05 wrote third-party software that interacted with a large decentralized system. They wrote an issue upon learning about a design feature in this large system, which they believed compromised users' privacy:

By default, [...] everything was configured to go to [A central server operated by the company.] [...] there was also reference about privacy on the website [...] Privacy was one of the values that was put in the spotlight. But it wasn't really privacy focused, was it? [...] The truth is that it's not up to par with the privacy level they claim. It's not decentralized. It's very much centralized because of the default configuration.  
[P05 - Third-party developer]

P05 had additionally contributed to other issues and documents about the system, which involved disagreements that were consistently grounded in commitments to privacy and individual freedom.

P01 demonstrated a related dimension through which task-conflicts could escalate to larger scales. They had expressed support for a proposal to modify a user interface to remind people to include an *alt-text* description for images they upload. During our interview, they asserted that the stakes extended beyond the immediate feature request because the resolution would establish precedence for similar issues in the future:

At stake is blind people's belief in what may amount to one of the only encounters with open communities that they have. Also, the accessibility and openness of the community is at stake, since if this isn't addressed, then by precedent, fewer accessibility issues will be addressed as well. [P01 - Issue contributor]

At face value, the issue was a task conflict about if and how prominently to display an alt-text prompt. Because the purpose of this prompt would be to influence user behaviour, this conflict was perceived to establish future norms about *equity and equality* for people with disabilities.

These examples highlight an important feature of value-related debates in software development. The stakes are rarely as simple as resolving a single bug or adding a feature. Instead, responses to value-related issues influenced the extent to which participants trust project maintainers to live up to the values they believe to have been promised.

**4.2.2 Process conflicts.** The majority of interview participants discussed conflicts connected to the processes through which stakeholders' views were incorporated into decisions. In general, they indicated that substantial discussions about difficult issues were valuable for figuring out projects' directions:

It's not yet clear where we are going. [...] So it's good to have controversial discussions. [P06 - Third-party developer]

Referring to projects centred around small teams of volunteers, P07, a core member of a relatively small project in our study, asserted that "a lot of these participatory projects fail because of internal conflicts", and argued this should be addressed through internal communication styles:

What is at risk? It's not [the specific decision]. It's whether we can come to a conclusion together as a group, and everybody feels they contributed to that and are listened to and we're committing and making compromises in a way that we want [P07 - Organization member]

Almost half of our interview participants expressed that there were cases where decisions were made without sufficient consultation. Several of those individuals indicated this was inconsistent with the values motivating decentralization and open source software in the first place, such as P04, who has some experience maintaining small open source projects and extensive experience contributing bug reports, documentation, and code to larger projects:

I was also really cheese[d] off that there was never any effort made to discuss this change with server admins. [...] There is a lot of talk about freedom and equality in open source and decentralized software [...] yet [this project] is developed under the "Benevolent Dictator for Life model." [P04 - Issue contributor]

There was, however, a significant split in opinions about how much time maintainers should invest in responding to these types of issues. Speaking about a different issue in the same project that P04 had criticized for having a "benevolent dictator" model, P11, who has been a maintainer for several open source projects, applauded much the same behaviour that P04 had criticized:



As you can see, [the lead maintainer] is clearly a smart maintainer and just stopped discussing the feature without even closing it. [P11 - Issue contributor]

Whether or not they were code contributors to these specific projects, this divide generally occurred between participants who did or did not have professional experience writing code for and maintaining open source projects. For example, even though P07, believed in the importance of inclusive discussions, they expressed frustration at perceived expectations that they respond to every request:

If you open a ticket to my repo, am I now obligated to spend three hours reviewing a patch? [“Open source”] doesn’t mean my time is open, you know? But sometimes people expect that. [P07 - Organization member]

This highlights a significant pattern: Various contributors had mismatched expectations of how much time and effort should go into resolving issues. Generally, participants with experience maintaining open source projects were quick to express understanding toward time pressure faced by open source maintainers. And by contrast, those who lacked such experience were more likely to be frustrated by what they perceived as inadequate processes.

### 4.3 Outcomes of value-related conflicts

This section addresses **RQ3**: “*What are the outcomes of value-related conflicts and current practices to address them?*” We describe outcomes across three dimensions: Tangible changes to code or design; effects on participants’ future thinking about values in relation to this work; and influences on participants’ future contributions. We additionally describe how discussions about project directions are sometimes conducted in more private contexts.

**4.3.1 Outcomes in code.** With regard to development itself, the effects of value-related discussions in GitHub issues varied.

The majority of participants indicated that proposals made in the issues about which we interviewed them were not implemented in code, at least not in a way that satisfied the issue-poster. In some cases, participants felt the issue was on its way to being addressed in the future, e.g.,

This issue might not be completely relevant now but may come up in the future when the project is more well-used throughout the decentralized tech industry.  
[P09 - Third-party developer]

However, others felt the issues had been abandoned, such as P01 who expressed frustration that an issue they felt to be important “has mostly been forgotten” [P01 - Issue contributor].

Some participants had more positive experiences, such as when issues led directly to concrete changes. For example, P08 (organization member) proposed renaming the “master” branch of each repository to “main,” since the etymology of the *master-slave* relationships in software is a reference to slavery. P08’s proposed change was implemented in many of the project’s repositories. This was part of a broader movement advocating for moving away from racially charged software terms, and GitHub made changes to the platform itself to support this in mid 2020 [52].

Additionally, when P05, a third-party developer, posted an issue arguing that one project included features that violated individuals’ privacy, the developers of that project engaged in discussion with P05 and made a change to their work based on that discussion. However, even though a change was implemented, P05 argued that it was superficial and did not sufficiently address the concern.

They also used it to their own advantage, saying, “Look, there was this review about privacy. We are addressing all the points. Here is all the work we are doing.” [...] Spoilers: it does not address the issues. [P05 - Third-party developer]

P02, a core member of a different organization, explained that a heated discussion about privacy, censorship, and freedom had been disruptive to implementation, rather than helping to guide it.

The level of heat and debate in this comment thread actually delayed any implementation work here. The conversation was too charged and controversial for the community to reach alignment, and the level of need for the feature by most users wasn't worth having a very disruptive governance battle. Instead [one] team created this feature in a fork, and then published it for others to use without changing the main implementation. [P02 - Organization member]

This highlights that, although discussions rooted in values can lead to changes, this can come with costs to efficiency. Even in the issues among our interview participants that led to concrete technical changes, success was complicated.

*4.3.2 Discussions influenced understanding of values and technology.* Although many value-related discussions in GitHub issues did not lead to immediate changes to software, participants identified that these discussions influenced their understanding of the relationship between values and technology. In one project, there was a rich discussion about how to ensure fairness in an algorithm used to distribute funding to partner developers. The purpose of the program was to encourage third-parties to build an ecosystem of apps, the most popular of which would receive monetary reward. P06, a third-party developer to this project, explained that the program “was dropped because everybody was gaming it,” but that the surrounding discussions shaped his understanding of the value of transparency and accountability. Specifically, people were able to game the system because rules were publicly available, and P06 concluded:

Providing rules publicly [...] is not always a good way. Transparency, it's a good thing. A thing I believe in. But sometimes it's not helpful. [P06 - Third-party developer]

This speaks to a process in which P06's understanding of the value of transparency itself evolved through employing that value to understand and intervene to a complex situation of practice [40].

P02, who was quoted above explaining that a discussion about privacy, censorship, and individual freedom was disruptive to implementation, asserted that this discussion made them consider censorship in a way they hadn't previously:

Some members of the community took it as creating tools for censorship [...] I'm not sure how much of this was just misinterpretation of the issue, or real fundamental disagreement in how this could potentially be abused. [...] I personally hadn't considered that, [for example,] having an optional blacklist feature would be so controversial - it definitely highlighted that it'd be very important for it to be hard to monitor whether a node was using a blacklist. [...] but that's still hypothetical. [P02 - Organization member]

In this case, the most significant impact of the issue is likely its influence on future decision-making, both within this project itself, and potentially in other work.

These patterns indicate that even when discussions about values did not result in tangible changes, they affected participants' understandings of the potential impacts of technical decisions. It is likely that this will shape future decisions, such as by motivating developers to address potential outcomes they had not previously considered.

*4.3.3 Discussions led participants to increase, redirect, or withdraw participation.* Another significant theme was that experiences engaging with values in GitHub issues had an effect on participants' intentions for the future.

In some cases, positive experiences could motivate future participation. For example, P08 experienced “hesitation” and “angst” before creating issues arguing for more inclusive conventions on

their organization's GitHub repos, but explained that it would be easier in the future because they felt encouraged by the discussion and its outcome:

Having attempted to make a change and participated with others in trying to make a change and seeing that change actually come about. No reason to not try it again in the future if I see something else that needs changing. [...] It may be a small thing, but hopefully it's meaningful enough to just make it that much more inclusive of an environment. [P08 - Organization member]

Alongside seeing their proposed change implemented, the feeling of having "participated with others" is likely significant, since P08 could see that other people cared about the same social issues and were willing to take action to support them.

By contrast, negative experiences during value-related conflicts could have a strong emotional effect, which motivated some participants to decrease or redirect their participation. P05 had invested significant time into building third-party software for one of the projects in our study, even making it a significant part of their professional work. However, they stopped development of this third-party software because the main project did not, in P05's view, adequately address important privacy issues. In addition, P05's experience with this conflict resulted in a poor emotional outcome:

You have this frustration that builds over time and over the years. [...] It's sad. I think that's the last remaining feeling I have about all this and open source in general, at the end of the day. It's sad how it's not about the people that you want to build for. It's about the people that build it and how they feel attacked all the time.  
[P05 - Third-party developer]

Similarly, P01 remarked that, after seeing an issue they cared about be dismissed, "I was sufficiently tired of dealing with ableism in the FOSS community to pretty much give up" [P01 - Issue contributor]. Others did not turn away from FOSS altogether, but redirected their attention. P04 expressed frustration that the lead maintainer of one project had developed "a reputation for not listening to anyone else." In response, P04 (issue contributor) shifted their contributions to a different project, whose maintainer "was so welcoming and appreciative that now I'm hooked and intend to make many more contributions."

**Comments by deleted accounts.** Additionally, our quantitative analysis showed evidence suggesting that GitHub users who discussed values were more likely to have deleted their GitHub account altogether. When a GitHub user deletes their account, all code commits, repositories, issue comments, and other data owned by their account are removed, however their contributions to others' repositories are preserved and attributed to a faux-GitHub account named 'Ghost' [17].

In our dataset, 4041 (1.04%) issue comments were attributed to the Ghost user. 0.72% ( $n = 29$ ) of the Ghost user's comments were value-related, compared to 0.47% ( $n = 1,812$ ) of comments made by other users. Logistic regression showed that comments by the Ghost user were significantly more likely to be value-related than comments by other users (odds ratio = 1.55, 95% CI [1.08, 2.33], McFadden's Pseudo  $R^2 = 0.000$ ). 48% ( $n=14$ ) of the Ghost user's value-related comments were associated with *respectfulness*, compared to 33% ( $n=596$ ) of other users' value-related comments. So, we performed a second logistic regression that showed that comments by the Ghost user were significantly more likely to be related to *respectfulness* than comments by other users (odds ratio = 2.31, 95% CI [1.37, 3.90], McFadden's Pseudo  $R^2 = 0.001$ ). The low Pseudo  $R^2$  indicates that whether or not the author deleted their account is, on its own, a weak predictor of whether a comment contains a value-term. Nonetheless, the results suggest there is a positive correlation between deleting one's GitHub account and having posted comments that referenced values (especially *respectfulness*). This is generally commensurate with our interview participants' indications that

value-related conflicts contributed to decisions to stop or change one's participation with the projects where those conflicts occurred.

**4.3.4 Discussions outside of GitHub.** One significant feature of discussions and conflicts about projects values is that they were not confined to GitHub. Sometimes issue threads made reference to team meetings or other conversations, and there were cases where project maintainers felt it prudent to avoid fully public discussions:

There are internal notes kept in a more private space because as we get into more partnerships and stuff, it doesn't make sense for all that information to be out there.  
[P07 - Organization member]

P10 described how online meetings were important for figuring out processes and governance:

I think we spent like a good six months meeting every other Monday trying to figure out the governance structure. [P10 - Organization member]

These more closed meetings have many advantages, but they also create conditions in which GitHub issues, which were some participants' primary or only way of participating in conversations about project decision-making, may not be regarded as suitable for big-picture decisions that have already been made elsewhere. For example, P02 described that in their organization,

We do make significant decisions in GitHub issues - though not where we tend to define our values per se. [P02 - Organization member]

In sum, external meetings, documents, and conversations can be more controlled than discussions in GitHub issues and other public sites. Of course, this does not mean they are absent of conflict. For example, P10 described many disagreements during their organization's Monday meetings. However, they involve fewer external contributors who may not be on the same page about decision-making processes and they provide greater opportunities for preserving confidentiality when needed. Nonetheless, because contributors whose participation was limited to publicly accessible discussions demonstrated a strong desire to contribute to normative decision-making, closed meetings could contribute to process-conflict in the long-term.

## 5 DISCUSSION

In general, although comments that explicitly referenced value keywords were a minority, our analysis found that the values underlying DWeb's principles were being engaged with across many of these projects. Further, we found that discussions about values increased between the beginning of each project (its earliest event) and through its second year. Insofar as the number of discussions about values increased later in projects' life-cycles, our results support suggestions in prior work to attend to "valuation" as an ongoing process in which values are shaped in relation to specific material and social contexts [37]. Therefore, it is important for maintainers to make sure these discussions are well-managed in order to productively integrate values into development activities.

We also found that the likelihood of values being referenced in an organization's issue comments was positively associated with the number of people posting comments, even when controlling for the number of comments overall. Interview results suggested that the presence of diverse viewpoints in a discussion could increase the potential for disagreements about project decisions generally, and lead to more nuanced conversations about ethical issues. Most dramatically, the likelihood of *respectfulness* being referenced in an organization's comments in any given month was strongly associated with the number of users, and not with the number of comments. One interpretation is that when there are a greater number of people participating in issues, it becomes more important to remind people to be respectful of one another. Another dimension is that conversations that engage with *respectfulness* may attract participation from people who have not

have commented about other issues. These patterns suggest that encouraging discussions that include many contributors with diverse viewpoints may generate more detailed discussions about values, while simultaneously contributing to challenges for maintaining civility.

When speaking with our interview participants, it became clear that conflict was a common feature in value-related issues. Managing conflict in remote and distributed work has been a significant area of research for CSCW [e.g., 18, 24, 70]. Past CSCW work has found that for cognitive conflicts that involve negotiations related to stakeholders' values, team fracture—a loss of team viability that results in team members no longer wishing to work together—can be strongly influenced by early moments during collaboration [70]. Future work may benefit by further investigating value-related discussions in relation to civility and conflict.

Among the cases we discussed with interview participants, extended discussions about values often did not result in what participants felt to be appropriate technical changes. A possible explanation comes from past research finding that maintainers on GitHub utilize implicit mechanisms to prioritize what are perceived to be important issues, and neglect those that are perceived to exceed a project's capacity [45]. Value-related issues often alluded to scopes that extended beyond a single bug-fix or new feature, and thus may be regarded as requiring a high capacity, such as in P02's assertion that a heated debate delayed implementation work related to an issue.

In fact, emotion was a significant feature of participants' descriptions of navigating value-related GitHub issues, especially among those who posted issues or requested changes to projects. Participants expressed frustration and sadness about situations where they felt their concerns hadn't been adequately addressed. These emotions were related to their perceptions that successes or failures to live up to values in the resolution of specific issues signified broader stakes about the future of the project. Emotions were identified as one of eight "grand challenges" for value-sensitive design [26]. Our results suggest that, to meet this challenge, researchers should investigate emotional discord during design processes. Specifically, value-related conflicts appear to have elevated changes of emotional impact, and poor emotional outcomes seem to be associated with decisions to stop contributing to a project.

Finally, among interview participants, the most significant common feature in unproductive or unhealthy value-related discussions was a lack of common ground in relation to process conflict. Misalignment was observed in expectations about how much time maintainers should devote to engaging with specific issues, how open decision-making processes should be, and which values should be prioritized. Identifying and managing differences in expectations is therefore an important task for engaging with values during software development and maintenance. A likely correlate of this mismatch is that project maintainers often relied on sites outside of GitHub for at least some decision-making, which were not always accessible to project outsiders. Thus, one source of conflict was that issues posters, third-party developers and organization members were not always in agreement about *how* and *where* project values should be shaped.

The implications of this study may generally extend to other open source projects in which values are a prominent focus. However, it is important to acknowledge some features that limit this study's generalizability: A high level of participation in GitHub issues by end-users and third-party developers; many contributors have strong, clear beliefs about ethics and technology; and most projects are based in Western countries.

## 5.1 The value of value-related discussions

Given that encouraging and engaging with discussions about project values and future directions requires significant time, some participants identified these discussions as disruptive to development. It is important to consider what sort of advantages could make this worthwhile.

In setting out guidelines to serve the “agency of all people”, DWeb’s principles [67] recognize that the responsibilities of designing communication technology extend beyond providing an enjoyable user experience. This aligns with boy’s argument that “the ideas and concerns of all people need to be a part of the design phase and the auditing of systems, even if this slows down the process” [5]. If values are to be a core design feature, they should be employed not only towards the product itself, but also in the processes of building that product.

Creating a community or organizational culture in which “all people” have an opportunity to contribute is no simple feat. For example, P01 asserted that participation in GitHub issues may be “one of the only encounters” that blind people have with this type of community, and bad experiences may lead to longer-term marginalization if their perspectives are not valued. Nurturing opportunities for “all people” necessitates special consideration toward people who may be marginalized by default.

Beyond specific ethical commitments, there are pragmatic benefits to engaging with users’ normative claims. Increasing decentralization in the *process* of developing the Decentralized Web, such as by better incorporating third-party users into discussions about projects’ future, could be vital for building better products that maintain ideological consistency through situations and contexts about which developers, understandably, lack first-hand experience. One of the reasons that decentralized web projects have struggled in the past is that developers have failed to consider what users really want [34, 51]. While it may be expedient to de-prioritize value-related discussions and requests, this could be a waste of a rare opportunity to learn about varied user experiences that developers may not have considered.

## 5.2 Design implications

In this section, we offer possible directions for making the most of value-related discussions and mitigating process conflict during open source software development. This involves extending upon existing practices as well as considering new tools for discussion.

Our results suggest that process conflicts could be mitigated by improving common ground between project maintainers and people who post and comment upon issues. Open source communities often employ ad-hoc governance, which has been regarded as a strength and source of agility [9, 69]. However, this can lead to “scalar debt” where maintainers must expend more and more energy “putting out fires” as projects grow and scale [31]. This could result in maintainers feeling pressured by time-constraints as complex ethical issues are raised. Further, ad-hoc governance can contribute to a lack of understanding among stakeholders about each other’s roles, which is important because past CSCW scholarship has asserted that a lack of clearly defined roles can be an antecedent to conflict [19]. To improve clarity about roles and contribution processes, many open source projects include a file called “CONTRIBUTING” describing guidelines for contribution. Similarly, we propose that project maintainers could employ GitHub’s *issue templates* to encourage common ground when reporting issues that relate to project values. For example, the default template for *feature requests* asks, “Is your feature request related to a problem? Please describe,” “Describe the solution you’d like,” and “Describe alternatives you’ve considered.” Projects that seek to prioritize values could benefit by including questions about the social or ethical impacts of current designs as well as potential impacts of proposed changes. Additionally, maintainers could benefit by creating standard structures for replying to such issues. This could mitigate the disruptiveness of value-related debates during GitHub issues by establishing clear processes that illuminate developers’ ways of thinking when navigating development decisions.

Discussions about values often involve “wicked problems” with no clear solution. Thus, although GitHub issues are sometimes the clearest path for users to engage with software maintainers, they may lack the required scope for some discussions. This is one of the reasons that open source

projects often use forums, mailing lists, and messaging software (IRC, Slack, Matrix, etc.), since these channels are suited to ongoing discussions that need not achieve closure to be regarded as valuable. However, one of the benefits of GitHub issues compared to those channels is they can put discussions in context with development work (e.g., by linking issues and code commits). This is one of the motivations for GitHub’s launch of a beta feature called *Discussions*, which is intended to help maintainers, contributors, and visitors have different and more largely scoped conversations than are appropriate for issues, without relying on third-party tools [32]. An exploratory analysis of early adopters of GitHub discussions found that they were used for a variety of purposes, including higher level conceptual questions about topics such as design patterns and limitations of current implementations [36]. Additionally, *Discussions* sometimes resulted in the creation of issues and pull-requests to tackle actionable items. The same study, however, found that users were sometimes confused about where to post a given question or conversation, and thus recommended defining guidelines for participating in discussions, which is congruent with our recommendations earlier in this section.

While GitHub Discussions seem promising for encouraging bigger-picture conversations, centralizing all project discussions on a platform owned by Microsoft is generally antithetical to the core principles of decentralization upon which the projects studied in this paper are founded. For projects that prefer to use third-party discussion venues, it would be advantageous to ensure that users are aware of how to participate at these external sites, and to consider integration with GitHub so relevant discussions and code activities can be linked. This sort of linkage could be accomplished using software such as Brid.gy, which supports syndicating between web content feeds and GitHub issues [8].

Additionally, identifying and tagging value-related conversations could help illuminate how values are discussed across multiple issues. This could make it easier to identify like-minded projects and people, which could motivate collaboration and bonding. Further, if DWeb contributors could view how similar challenges are addressed across different projects, they could learn from each other’s successes and failures. This could be partially automated, such as by integrating this study’s linguistic analysis into a GitHub bot<sup>1</sup> or third-party software.

Regardless of the specific tools used, software and other projects that attempt to align ongoing work with value-motivated goals will benefit from nurturing common ground among diverse contributors and facilitating discussions that span scales between specific technical decisions and broader normative directions.

### 5.3 Future research directions

Based on our analysis, we suggest two directions for future research to build upon this work.

**5.3.1 Improved linguistic analysis.** The computational model we built for detecting references to values is a promising early step, albeit one with several limitations. Most notably, our keyword-based approach is unable to distinguish between multiple meanings of homonyms. Although we tuned our keyword selection using the method described in Section 3.1.3, more sophisticated models could have better performance. For example, topic modelling could be used to identify the contexts in which certain words are used, which could probably assist in unpacking the difference between, for example, references to a software “environment” and discussions about protecting earth’s “environment” through sustainable practices.

Another fruitful direction could be to include measures of sentiment in future work. Past work has found that developers do express emotions in software issue-tracking systems, however even human coders have struggled to evaluate the nuance of developers’ emotions, posing a challenge

<sup>1</sup>e.g., using <https://probot.github.io>

for automated sentiment analysis in this context [50]. Nonetheless, sufficiently validated sentiment models could help unpack positive and negative emotional experiences on a large scale. Further, it may be helpful to combine sentiment analyses with measures of behaviour (such as leaving a project, or increasing participation) and the patterns such as the timing of messages.

While Schwartz’s theory of basic human values provided a suitable foundation for this study, other theories may be better-suited to study values in different contexts. Even if using a “universal” theory of values, we recommend customizing the values dictionary to accommodate local communication norms. Further, large-scale analyses cannot tell the whole story, and closer forms of analysis (such as ethnographies) could lead to more grounded insights about specific collaboration sites.

These suggestions may help future work address one of this study’s limitations by disentangling references to values in technologies from references to values as features of collaboration processes. However, we advise that technological products can not be fully disentangled from the processes of their development. This is especially evident with much open-source software, since, for example, the ability to view code and interact with developers shapes user-experience.

**5.3.2 Absence and migration.** Our analysis has shown that negative experiences in GitHub issues, such as feeling that one’s contributions were not valued or taken seriously, can motivate people to leave specific projects. In some cases, this involves migration to projects that seem more like-minded, in others it could entail simply reducing one’s participation in open-source in general. We posit that, although contributor turnover causes knowledge loss for specific projects [57], there could be advantages for ecosystems of related projects (such as the multitude of projects oriented around the Decentralized Web). Past work has tended to focus on contributor disengagement resulting from career changes and time pressure [48]. By contrast, we identified cases where participants stopped contributing to a project they perceived to have become poorly fitted to their priorities, and turned their attention toward another open source project. Aligning with work about non-use and abandonment of technology more generally [20], we suggest that future work should study absent and departing users. This could involve mapping contributors’ movements between projects, potentially identifying bridging ties [76] who disseminate knowledge and approaches among different projects in the course of their transition.

Related to this, our analysis of comments by the *Ghost* user suggests that value-related discussions may be slightly associated with deleting one’s GitHub account. We interpret this with a number of caveats since it was not possible to distinguish between different people represented by the Ghost user and the overall explained variance was low. Thus, rather than drawing strong conclusions from this observation, we believe this identifies a thread for future research. Specifically, although the anonymity afforded by the use of the “Ghost” pseudonym poses a methodological challenge, future work could further leverage the Ghost user to investigate when and why people withdraw participation from online discussions, while respecting those people’s right to privacy.

## 6 CONCLUSION

We have investigated how contributors to open source projects related to the decentralized web discuss values in GitHub issues over time. Our analysis included identifying patterns of discussions that referenced values across a large dataset, including that value-related discussions become more common later in development. Through interviews with issues contributors, third-party developers and project organization members, we examined features of task and process conflicts related to project values. Specifically, we described that the perceived stakes of these conflicts are high due to their association with large normative goals. We also identified a lack of consensus among participants with regard to the processes of navigating value-related issues. We then presented



recommendations for addressing establishing common ground to mitigate conflict and supporting productive engagement with values by diverse stakeholders during long development cycles.

## REFERENCES

- [1] Ofer Arazy, Hila Lifshitz-Assaf, and Adam Balila. 2019. Neither a Bazaar nor a Cathedral: The Interplay between Structure and Agency in Wikipedia’s Role System. *Journal of the Association for Information Science and Technology* 70, 1 (2019), 3–15. <https://doi.org/10.1002/asi.24076>
- [2] Ofer Arazy, Lisa Yeo, and Oded Nov. 2013. Stay on the Wikipedia Task: When Task-related Disagreements Slip into Personal and Procedural Conflicts. *Journal of the American Society for Information Science and Technology* 64, 8 (2013), 1634–1648. <https://ideas.repec.org/a/bla/jamist/v64y2013i8p1634-1648.html>
- [3] Chelsea Barabas, Neha Narula, and Ethan Zuckerman. 2017. *Defending Internet Freedom through Decentralization: Back to the Future?* Technical Report. The Center for Civic Media & The Digital Currency Initiative, MIT Media Lab.
- [4] Irina Bolychevsky. 2018. There’s More to Decentralisation than Blockchains and Bitcoin. Retrieved {2019-09-15} from <https://medium.com/@shevski/how-decentralised-are-you-a6539eeb27ff>
- [5] danah boyd. 2019. Facing the Great Reckoning Head-On. Retrieved {2019-09-16} from <https://medium.com/@zephoria/facing-the-great-reckoning-head-on-8fe434e10630>
- [6] Philip Brey. 2000. Disclosive Computer Ethics. *SIGCAS Comput. Soc.* 30, 4 (Dec. 2000), 10–16. <https://doi.org/10.1145/572260.572264>
- [7] Philip Brey. 2012. Anticipating Ethical Issues in Emerging IT. *Ethics and Information Technology* 14, 4 (Dec. 2012), 305–317. <https://doi.org/10.1007/s10676-012-9293-y>
- [8] Bridgy. 2020. About - Bridgy. Retrieved {2018-05-01} from <https://bridgy/about>
- [9] Eugenio Capra, Chiara Francalanci, and Francesco Merlo. 2008. An Empirical Study on the Relationship Between Software Design Quality, Development Effort and Governance in Open Source Projects. *IEEE Transactions on Software Engineering* 34, 6 (Nov. 2008), 765–782. <https://doi.org/10.1109/TSE.2008.68>
- [10] EunJeong Cheon, Stephen Tsung-Han Sher, Šelma Sabanovic, and Norman Makoto Su. 2019. I Beg to Differ: Soft Conflicts in Collaborative Design Using Design Fictions. *San Diego* (2019), 14.
- [11] Digital Life Collective. 2018. Decentralized Tech. Retrieved {2022-01-13} from <https://kumu.io/DigLife/decentralized-tech>
- [12] Joseph Coveney. 2021. FIRTHLOGIT: Stata Module to Calculate Bias Reduction in Logistic Regression. Boston College Department of Economics. <https://ideas.repec.org/c/boc/bocode/s456948.html>
- [13] Lucy J. Cowie, Lara M. Greaves, Taciano L. Milfont, Carla A. Houkamau, and Chris G. Sibley. 2016. Indigenous Identity and Environmental Values: Do Spirituality and Political Consciousness Predict Environmental Regard Among Māori? *International Perspectives in Psychology* 5, 4 (Oct. 2016), 228–244. <https://doi.org/10.1037/ipp0000059>
- [14] Carsten K. W. De Dreu and Laurie R. Weingart. 2003. Task versus Relationship Conflict, Team Performance, and Team Member Satisfaction: A Meta-Analysis. *Journal of Applied Psychology* 88, 4 (2003), 741–749. <https://doi.org/10.1037/0021-9010.88.4.741>
- [15] Frank R. C. de Wit, Lindred L. Greer, and Karen A. Jehn. 2012. The Paradox of Intragroup Conflict: A Meta-Analysis. *Journal of Applied Psychology* 97, 2 (2012), 360–390. <https://doi.org/10.1037/a0024844>
- [16] Carl DiSalvo. 2022. *Design as Democratic Inquiry: Putting Experimental Civics into Practice*. The MIT Press, Cambridge, Massachusetts.
- [17] GitHub Docs. 2022. Deleting Your User Account. Retrieved {2022-01-14} from <https://docs.github.com/en/account-and-profile/setting-up-and-managing-your-github-user-account/managing-user-account-settings/deleting-your-user-account>
- [18] S. M. Easterbrook. 1993. *CSCW: Cooperation or Conflict?* Springer-Verlag. <http://books.google.com/books?id=oJJXAAAAAAAJ>
- [19] S. M. Easterbrook, E. E. Beck, J. S. Goodlet, L. Plowman, M. Sharples, and C. C. Wood. 1993. A Survey of Empirical Studies of Conflict. In *CSCW: Cooperation or Conflict?*, S. M. Easterbrook (Ed.). Springer-Verlag, London, 1–68.
- [20] Daniel A. Epstein, Monica Caraway, Chuck Johnston, An Ping, James Fogarty, and Sean A. Munson. 2016. Beyond Abandonment to Next Steps: Understanding and Designing for Life after Personal Informatics Tool Use. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1109–1113. <http://doi.org/10.1145/2858036.2858045>
- [21] Ethan Fast, Binbin Chen, and Michael S. Bernstein. 2016. Empath: Understanding Topic Signals in Large-Scale Text. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, San Jose California USA, 4647–4657. <https://doi.org/10.1145/2858036.2858535>
- [22] 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>
- [23] 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>
- [24] 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>
- [25] David Firth. 1993. Bias Reduction of Maximum Likelihood Estimates. *Biometrika* 80, 1 (March 1993), 27–38. <https://doi.org/10.1093/biomet/80.1.27>
- [26] Batya Friedman, Maaik Harbers, David G. Hendry, Jeroen van den Hoven, Catholijn Jonker, and Nick Logler. 2021. Eight Grand Challenges for Value Sensitive Design from the 2016 Lorentz Workshop. *Ethics and Information Technology* (April 2021). <https://doi.org/10.1007/s10676-021-09586-y>
- [27] Batya Friedman and David Hendry. 2019. *Value Sensitive Design: Shaping Technology with Moral Imagination*. The MIT Press, Cambridge, Massachusetts.
- [28] Batya Friedman and Peter H. Kahn. 2003. Human Values, Ethics, and Design. In *Handbook of Human-Computer Interaction*, J Jacko and A Sears (Eds.). Lawrence Erlbaum Associates, Mahwah, NJ, 1177–1201.
- [29] Batya Friedman and Helen Nissenbaum. 1997. Bias in Computer Systems. In *Human Values and the Design of Computer Technology*, Batya Friedman (Ed.). Number no. 72 in CSLI Lecture Notes. Cambridge University Press, New York, 21–40.
- [30] Robert W. Gehl. 2015. The Case for Alternative Social Media. *Social Media + Society* 1, 2 (2015), 1–12. <https://doi.org/10.1177/2056305115604338>
- [31] 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>
- [32] GitHub.com. [n.d.]. About Discussions. Retrieved {2022-01-09} from [https://docs.github.com/\\_next/data/a5AXZ1cRu1KS7LbQZV33G/en/free-pro-team%40latest/discussions/collaborating-with-your-community-using-discussions/about-discussions.json?versionId=free-pro-team%40latest&productId=discussions&restPage=collaborating-with-your-community-using-discussions&restPage=about-discussions](https://docs.github.com/_next/data/a5AXZ1cRu1KS7LbQZV33G/en/free-pro-team%40latest/discussions/collaborating-with-your-community-using-discussions/about-discussions.json?versionId=free-pro-team%40latest&productId=discussions&restPage=collaborating-with-your-community-using-discussions&restPage=about-discussions)
- [33] Jay Graber. 2021. Ecosystem Review. [https://matrix.org/\\_matrix/media/r0/download/twitter.modular.im/981b258141aa0b197804127cd2f7d298757bad20](https://matrix.org/_matrix/media/r0/download/twitter.modular.im/981b258141aa0b197804127cd2f7d298757bad20)
- [34] Harry Halpin. 2018. Decentralizing the Social Web: Can Blockchains Solve Ten Years of Standardization Failure of the Social Web?. In *INSCF 2018- 5th International Conference 'Internet Science'*. St. Petersburg, Russia, 16.
- [35] Gunnar Harboe and Elaine M. Huang. 2015. Real-World Affinity Diagramming Practices: Bridging the Paper-Digital Gap. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 95–104. <https://doi.org/10.1145/2702123.2702561>
- [36] Hideaki Hata, Nicole Novielli, Sebastian Baltes, Raula Gaikovina Kula, and Christoph Treude. 2021. GitHub Discussions: An Exploratory Study of Early Adoption. *Empirical Software Engineering* 27, 1 (Oct. 2021), 3. <https://doi.org/10.1007/s10664-021-10058-6>
- [37] Lara Houston, Steven J. Jackson, Daniela K. Rosner, Syed Ishtiaque Ahmed, Meg Young, and Laewoo Kang. 2016. Values in Repair. In *CHI '16 Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM Press, 1403–1414. <https://doi.org/10.1145/2858036.2858470>
- [38] Haiyan Huang and Rosalie Ocker. 2006. Preliminary Insights into the In-Group/out-Group Effect in Partially Distributed Teams: An Analysis of Participant Reflections. In *Proceedings of the 2006 ACM SIGMIS CPR Conference on Computer Personnel Research Forty Years of Computer Personnel Research: Achievements, Challenges & the Future - SIGMIS CPR '06*. ACM Press, Claremont, California, USA, 264. <https://doi.org/10.1145/1125170.1125232>
- [39] Ole Sejer Iversen, Kim Halskov, and Tuck W. Leong. 2012. Values-Led Participatory Design. *CoDesign* 8, 2-3 (June 2012), 87–103. <https://doi.org/10.1080/15710882.2012.672575>
- [40] Nassim JafariNaimi, Lisa Nathan, and Ian Hargraves. 2015. Values as Hypotheses: Design, Inquiry, and the Service of Values. *Design Issues* 31, 4 (Oct. 2015), 91–104. [https://doi.org/10.1162/DESI\\_a\\_00354](https://doi.org/10.1162/DESI_a_00354)
- [41] Karen A. Jehn and Elizabeth A. Mannix. 2001. The Dynamic Nature of Conflict: A Longitudinal Study of Intragroup Conflict and Group Performance. *Academy of Management Journal* 44, 2 (April 2001), 238–251. <https://doi.org/10.5465/3069453>
- [42] Deborah G. Johnson. 2011. Software Agents, Anticipatory Ethics, and Accountability. In *The Growing Gap between Emerging Technologies and Legal-Ethical Oversight*, Gary E. Marchant, Braden R. Allenby, and Joseph R. Herkert (Eds.).

Vol. 7. Springer Netherlands, Dordrecht, 61–76.

- [43] 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/>
- [44] Eirini Kalliamvakou, Georgios Gousios, Kelly Blincoe, Leif Singer, Daniel M. German, and Daniela Damian. 2016. An In-Depth Study of the Promises and Perils of Mining GitHub. *Empirical Software Engineering* 21, 5 (Oct. 2016), 2035–2071. <https://doi.org/10.1007/s10664-015-9393-5>
- [45] Riivo Kikas, Marlon Dumas, and Dietmar Pfahl. 2015. Issue Dynamics in Github Projects. In *Product-Focused Software Process Improvement (Lecture Notes in Computer Science)*, Pekka Abrahamsson, Luis Corral, Markku Oivo, and Barbara Russo (Eds.). Springer International Publishing, Cham, 295–310. [https://doi.org/10.1007/978-3-319-26844-6\\_22](https://doi.org/10.1007/978-3-319-26844-6_22)
- [46] Bruno Latour. 2008. Where Are the Missing Masses? Sociology of a Few Mundane Artifacts. In *Technology and Society, Building Our Sociotechnical Future*, Deborah J. Johnson and M Wetmore Jameson (Eds.). MIT Press, Cambridge, Mass, 151–180.
- [47] 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) <http://arxiv.org/abs/2109.14359>
- [48] Courtney Miller, David Gray Widder, Christian Kästner, and Bogdan Vasilescu. 2019. Why Do People Give Up FLOSSing? A Study of Contributor Disengagement in Open Source. In *Open Source Systems (IFIP Advances in Information and Communication Technology)*, Francis Bordeleau, Alberto Sillitti, Paulo Meirelles, and Valentina Lenarduzzi (Eds.). Springer International Publishing, Cham, 116–129. [https://doi.org/10.1007/978-3-030-20883-7\\_11](https://doi.org/10.1007/978-3-030-20883-7_11)
- [49] Jessica K. Miller, Batya Friedman, Gavin Jancke, and Brian Gill. 2007. Value Tensions in Design: The Value Sensitive Design, Development, and Appropriation of a Corporation’s Groupware System. In *Proceedings of the 2007 International ACM Conference on Supporting Group Work (GROUP ’07)*. ACM, New York, NY, USA, 281–290. <https://doi.org/10.1145/1316624.1316668>
- [50] Alessandro Murgia, Parastou Tourani, Bram Adams, and Marco Ortu. 2014. Do Developers Feel Emotions? An Exploratory Analysis of Emotions in Software Artifacts. In *Proceedings of the 11th Working Conference on Mining Software Repositories (MSR 2014)*. Association for Computing Machinery, New York, NY, USA, 262–271. <https://doi.org/10.1145/2597073.2597086>
- [51] Arvind Narayanan, Solon Barocas, Vincent Toubiana, Helen Nissenbaum, and Dan Boneh. 2012. A Critical Look at Decentralized Personal Data Architectures. *arXiv:1202.4503 [cs]* (Feb. 2012). [arXiv:1202.4503 \[cs\]](https://arxiv.org/abs/1202.4503) <http://arxiv.org/abs/1202.4503>
- [52] BBC News. 2020. GitHub Abandons ‘master’ Term to Avoid Slavery Row. Retrieved {2021-12-30} from <https://www.bbc.com/news/technology-53050955>
- [53] Humphrey O. Obie, Waqar Hussain, Xin Xia, John Grundy, Li Li, Burak Turhan, Jon Whittle, and Mojtaba Shahin. 2020. A First Look at Human Values-Violation in App Reviews. *arXiv:2012.10095 [cs]* (Dec. 2020). [arXiv:2012.10095 \[cs\]](https://arxiv.org/abs/2012.10095) <http://arxiv.org/abs/2012.10095>
- [54] 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>
- [55] Dominic J. Packer and Alison L. Chasteen. 2010. Loyal Deviance: Testing the Normative Conflict Model of Dissent in Social Groups. *Personality and Social Psychology Bulletin* 36, 1 (Jan. 2010), 5–18. <https://doi.org/10.1177/0146167209350628>
- [56] Aravindh Raman, Sagar Joglekar, Emiliano De Cristofaro, Nishanth Sastry, and Gareth Tyson. 2019. Challenges in the Decentralised Web: The Mastodon Case. In *Proceedings of the Internet Measurement Conference (IMC ’19)*. Association for Computing Machinery, New York, NY, USA, 217–229. <https://doi.org/10.1145/3355369.3355572>
- [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.ijinfomgt.2018.11.015>
- [58] Johnny Saldaña. 2013. *The Coding Manual for Qualitative Researchers* (2nd ed ed.). SAGE, Los Angeles.
- [59] 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>
- [60] 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>
- [61] 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.
- [62] Dalhousie University Human Rights & Equity Services. [n.d.]. Social Justice Terms. Retrieved {2021-12-01} from <https://www.dal.ca/dept/hres/education-campaigns/educational-resources/definitions.html>
- [63] Katie Shilton. 2015. Anticipatory Ethics for a Future Internet: Analyzing Values During the Design of an Internet Infrastructure. *Science and Engineering Ethics* 21, 1 (Feb. 2015), 1–18. <https://doi.org/10.1007/s11948-013-9510-z>

- [64] Katie Shilton. 2018. Engaging Values Despite Neutrality: Challenges and Approaches to Values Reflection during the Design of Internet Infrastructure. *Science, Technology, & Human Values* 43, 2 (March 2018), 247–269. <https://doi.org/10.1177/0162243917714869>
- [65] Katie Shilton, Jes A. Koepfler, and Kenneth R. Fleischmann. 2013. Charting Sociotechnical Dimensions of Values for Design Research. *The Information Society* 29, 5 (2013), 259–271. <http://www.tandfonline.com/doi/abs/10.1080/01972243.2013.825357>
- [66] spaCy. 2020. spaCy · Industrial-strength Natural Language Processing in Python. Retrieved {2020-03-03} from <https://spacy.io/>
- [67] Mai Ishikawa Sutton and John Ryan. 2021. DWeb Principles. Retrieved {2021-05-06} from <https://getdweb.net/principles>
- [68] 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>
- [69] Steven Weber. 2004. *The Success of Open Source*. Harvard Univ. Press, Cambridge, Mass.
- [70] Mark E. Whiting, Allie Blaising, Chloe Barreau, Laura Fiuza, Nik Marda, Melissa Valentine, and Michael S. Bernstein. 2019. Did It Have To End This Way? Understanding The Consistency of Team Fracture. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (Nov. 2019), 209:1–209:23. <https://doi.org/10.1145/3359311>
- [71] 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>
- [72] Jon Whittle. 2019. Is Your Software Valueless? *IEEE Software* 36, 3 (May 2019), 112–115. <https://doi.org/10.1109/MS.2019.2897397>
- [73] 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>
- [74] Langdon Winner. 1980. Do Artifacts Have Politics? In *Daedalus*. <http://www.jstor.org/stable/10.2307/20024652>
- [75] 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>
- [76] Y. Connie Yuan and Geri Gay. 2006. Homophily of Network Ties and Bonding and Bridging Social Capital in Computer-Mediated Distributed Teams. *Journal of Computer-Mediated Communication* 11, 4 (July 2006), 1062–1084. <https://doi.org/10.1111/j.1083-6101.2006.00308.x>

Received January 2022; revised April 2022; accepted August 2022