



The Many Faces of Fairness: Exploring the Institutional Logics of Multistakeholder Microlending Recommendation

Jessie J. Smith*

Jessie.Smith-1@colorado.edu

Department of Information Science;
University of Colorado, Boulder
Boulder, Colorado, USA

Pradeep Ragothaman

Kiva Microfunds

San Francisco, California, USA
pradeep@kiva.org

Anas Buhayh*

Anas.Buhayh@colorado.edu

Department of Information Science;
University of Colorado, Boulder
Boulder, Colorado, USA

Nicholas Mattei

nsmattei@tulane.edu

Department of Computer Science;
Tulane University
New Orleans, Louisiana, USA

Amy Voids

amy.voids@colorado.edu

Department of Information Science;
University of Colorado, Boulder
Boulder, Colorado, USA

Anushka Kathait

Anushka.Kathait@colorado.edu

Department of Computer Science;
University of Colorado, Boulder
Boulder, Colorado, USA

Robin Burke

robin.burke@colorado.edu

Department of Information Science;
University of Colorado, Boulder
Boulder, Colorado, USA

ABSTRACT

Recommender systems have a variety of stakeholders. Applying concepts of fairness in such systems requires attention to stakeholders' complex and often-conflicting needs. Since fairness is socially constructed, there are numerous definitions, both in the social science and machine learning literatures. Still, it is rare for machine learning researchers to develop their metrics in close consideration of their social context. More often, standard definitions are adopted and assumed to be applicable across contexts and stakeholders. Our research starts with a recommendation context and then seeks to understand the breadth of the fairness considerations of associated stakeholders. In this paper, we report on the results of a semi-structured interview study with 23 employees who work for the Kiva microlending platform. We characterize the many different ways in which they enact and strive toward fairness for microlending recommendations in their own work, uncover the ways in which these different enactments of fairness are in tension with each other, and identify how stakeholders are differentially prioritized. Finally, we reflect on the implications of this study for future research and for the design of multistakeholder recommender systems.

*Both authors contributed equally to this research.



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

FAccT '23, June 12–15, 2023, Chicago, IL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0192-4/23/06.

<https://doi.org/10.1145/3593013.3594106>

CCS CONCEPTS

- Information systems → Recommender systems;
- Computing methodologies → Multi-agent systems;
- Social and professional topics → User characteristics.

KEYWORDS

fairness, institutional logics, international development, microlending, multistakeholder recommendation, recommender systems

ACM Reference Format:

Jessie J. Smith, Anas Buhayh, Anushka Kathait, Pradeep Ragothaman, Nicholas Mattei, Robin Burke, and Amy Voids. 2023. The Many Faces of Fairness: Exploring the Institutional Logics of Multistakeholder Microlending Recommendation. In *2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23), June 12–15, 2023, Chicago, IL, USA*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3593013.3594106>

1 INTRODUCTION

The problem of bias and fairness in algorithmic systems, generally, and in machine learning systems, in particular, is a critical issue for our increasingly data-centric world. Researchers have identified social biases within data sets, algorithms, and the methodologies used in machine learning systems of all kinds. In response, there has been substantial development of algorithmic interventions to enhance the *fairness properties* of these systems [7, 27]. However, application of these ideas has been limited, in part because of an overly-simplistic formulation of fairness [39] and, in part, because the organizational structures in which fairness emerges as a requirement are understudied. With few exceptions [31, 37, 38], published research in machine learning often considers only one protected group of stakeholders around which fairness should be considered, an unrealistic constraint for most applications. In addition, there is a dearth of research that explores the complexity of real-world organizational practices around fairness [33].

As a consequence, results from fairness-aware machine learning research often lack relevance for ML practitioners [20, 51]. This is particularly true of systems that apply fairness to personalized recommendations, even though these systems are critical to how millions of individuals access news, shopping, social connections, and employment opportunities. It is important to situate fairness interventions in specific application contexts. To confront this gap, we report on an interview-based study with 23 employees from the microlending platform, Kiva, to understand how they enact a multiplicity of different *Fairness Logics*—the organizing principles and patterns of practices that guide and shape how individuals carry out fairness in their own work. Specifically, in this study we answer the following research questions:

- **RQ1** What fairness logics are present in Kiva’s sociotechnical ecosystem of loan recommendation? Which stakeholders are prioritized by these different logics?
- **RQ2** How might these fairness logics complement or conflict with one another?
- **RQ3** In what ways do design interventions prioritize different fairness logics and/or different stakeholders?

Synthesizing these accounts of fairness in practice will enable us to understand the broader landscape of fairness within the organization, particularly when there are trade-offs between different logics of fairness that will need to be accommodated through design [3]. We expect that this multiplicity of fairness concepts and logics will be the norm rather than the exception in organizations seeking to implement fairness-aware recommendations. In describing our formative, qualitative study of fairness in this context, we hope to provide a model for how the complexities of fairness can be studied.

Our analysis reveals that interviewees’ diverse considerations of fairness align with four different fairness logics: consequence-based, contract-based, character-based, and duty-based. We explore where tensions arise between these logics. Finally, we conclude with design implications and methodological recommendations to assist practitioners with translating complex relationships among fairness logics into multistakeholder recommendation system design.

1.1 Research Context: Kiva Microfunds

Kiva Microfunds is a nonprofit organization (501(c)(3)) that uses microlending to provide access to capital for individuals, especially in the developing world, who are financially underserved or excluded. Kiva partners with local organizations in countries across the globe and, as of 2022, has lent \$1.82 billion to 4.5 million borrowers in 79 countries. The 2.1 million lenders who have supported these loans have seen a 96.4% repayment rate [46]. Because Kiva’s mission is explicitly aimed at addressing global financial inequity, issues of equity and fairness are paramount within the organization. While the research community grapples publication-by-publication with the complex challenges of algorithmic fairness, Kiva faces these challenges daily. In partnering with Kiva for this research, then, we ally ourselves with a rare breed of organization, one that openly discusses and experiments with strategies for achieving equity and fairness—a complex challenge for any organization, but even more so for an organization operating within the financial constraints of the nonprofit sector.

The key stakeholders at Kiva include the following:

- **Lenders**, the end-users of the Kiva.org online marketplace who fund loan opportunities for Kiva’s borrowers and lending partner organizations;
- **Lending Partners**, which are typically non-governmental organizations or microfinance institutions in borrowers’ local communities who support the borrowers in the loan application and repayment process. Outside of the United States, these lending partners also serve as the point-of-contact in all financial transactions with Kiva, helping to streamline the complexities of international financial services for the borrowers.
- **Borrowers**, the individuals and small groups who ultimately receive loans from Kiva. Borrowers in the United States receive loans directly from Kiva; international borrowers receive their loans via the lending partners. Even in instances when it is the lending partner who initially receives money from Kiva (and technically is the stakeholder borrowing the money), all interviewees refer to the individuals and small groups who ultimately receive the money as the borrower. Throughout the paper, we use the term borrower empirically, in the same way as the interviewees. Notably, it is common practice for many international lending partners to fund borrowers before they actually have money in hand from Kiva; as a result, fairness efforts on the Kiva platform are more likely to directly affect subsequent borrowers and the stakeholder class of borrowers, more generally, rather than the specific borrower whose loan opportunity is posted in the online marketplace.
- **The Organization (i.e., Kiva)**, which has an interest in keeping money flowing through their online marketplace by helping lenders find borrowers to whom they want to lend.

Information about loan opportunities are listed on Kiva’s site, typically for thirty days or until the loan request is fully funded. Loan opportunities are categorized on Kiva.org by their impact area (e.g., Agriculture or Health) and by geographical region. Each individual loan page includes information about the borrower (a photo and a story the borrower writes about themselves); information about any local organization assisting the borrower (their financial track-record in assisting borrowers with repayment); information about the loan (a brief description of its purpose, the total loan amount, the amount already funded, the loan’s expiration date, the length of the loan, and its repayment schedule); as well as the names of other lenders.

Fairness requirements at Kiva are driven by an internal motivation rooted in the organization’s philanthropic mission of improving global financial inclusion, rather than by external demands, such as regulatory requirements that might be found in other financial services institutions. Kiva’s recommendation ecosystem, therefore, is vital to the success of its equity-promoting efforts. Kiva recommends loan opportunities in a variety of ways. While not all of these systems for loan promotion use machine learning, we consider all forms of loan promotion as part of the recommendation ecosystem as it is both plausible and intended that recommender systems will be used more widely in the future. Loans are recommended to users in many ways, including (1) targeted email advertising (particularly to recommend loan opportunities when lenders have

been repaid and can lend again to a new borrower); (2) faceted search; and (3) a personalized loan discovery page with themed horizontal carousels, typically of 10-12 loans each. Carousel themes include “Recommended for [user]”, “Loans for [particular purpose or sector]”, “Trending now”, etc.... The specific carousels and their ordering are chosen (dynamically but globally) by a machine learning system (see [16]). As of this writing, most of the carousels use a personalized ordering that is curated from a content-based recommendation algorithm based on the user’s lending history. There are also several “auto-lending” features by which users can set preferences around their lending and have Kiva automatically fund new loans on their behalf (e.g., when repaid funds are available or via a monthly subscription service).

Since the ongoing pursuit of fairness is “baked in” to Kiva’s mission and work practices, it is a particularly compelling case study for the exploration of recommendation fairness. In addition, as a hybrid organization, embodying the characteristics of a nonprofit along with the characteristics of a financial services institution, research with Kiva is well-situated to have a broader and more transferable impact across genres of institutions and sectors.

2 LITERATURE REVIEW

2.1 Four Classes of Fairness

Four western ethical theories have dominated ML fairness scholarship, resulting in four *classes* of fairness: contract-based, consequence-based, character-based, and duty-based.

Consequence-based fairness has roots in Jeremy Bentham’s theory of *Act Utilitarianism*, and posits that “*an act is right if and only if it results in at least as much overall well-being as any act the agent could have performed*” [22]. For an act utilitarian system, optimizing fairness is synonymous with maximizing positive outcomes and minimizing negative outcomes. This operationalization of fairness is common in computer science research, as it offers two relatively straightforward methods for measuring and optimizing fairness: (1) select a proxy for “positive outcome” and tune the system to maximize that metric; or (2) select a proxy for “negative outcome” and tune the system to minimize that metric (e.g., [36]). However, optimizing an ML system for act utilitarianism might maximize benefits for all users in aggregate, while failing to maximize benefits for subgroups of users that fall within protected demographic categories such as race or gender. An example of this issue is the design of facial recognition technologies that did not correctly identify racial minorities [13].

Contract-based fairness is rooted in *Social Contract Theory* [41], which involves defining an ideal social contract and abiding by its rules. One interpretation of contract-based fairness is Rawls’ theory of *Distributive Justice* [28], which seeks to equally distribute resources within a society. One way to advance distributive justice is through *Equality of Opportunity* [5], which requires that everyone has a fair chance to receive benefits. Equality of opportunity has an extensive history in financial regulation [6], as there are global, systemic inequalities among opportunities to obtain wealth [54]. One way that inequality of opportunity arises in recommendation is through *popularity bias*, when recommendation algorithms exacerbate the difference in exposure between items at different levels of popularity, also known as the long-tail effect. There is a substantial

research literature that seeks to mitigate the effects of popularity bias in recommender systems (see e.g., [2, 19, 50, 66–68]), including in Kiva’s recommendation system [16].

Character-based fairness is rooted in Aristotle’s theory of *Virtue Ethics* [34]. Virtue ethics holds that decisions or outcomes should be based on a person’s *character* rather than their *actions*. However, in recommendation, users’ actions are often used as proxies for one’s “character,” (e.g., when a users’ interests and personality are inferred from their engagement [62]). One example of character-based fairness in recommender systems is the design of trust-based recommender systems, which serve recommendations from individuals that users have chosen to trust, in contrast with collaborative recommendation in which the peers influencing the recommendation are unknown [44, 64].

Duty-based fairness is rooted in Kant’s theory of *Deontology* [40]. Achieving fairness through deontological principles means consistently following a set of rules. Optimizing for deontological fairness in a machine learning system requires selecting a set of rules (e.g., model or algorithm heuristics) and then consistently following those rules. Ferraro et al. [26] provide an example of duty-based fairness as it relates to music recommendation on Spotify. One of their proposed platform “duties” is a promise to musicians (item providers) that less popular or new artists will be recommended at a fair rate in comparison to more popular artists. We note that duty-based fairness is more of a “meta” fairness logic, and is analytically squirrelier, because any one set of rules could align or conflict with any of the other given logics, and there is no specific delineation of what these rules “should” be in practice. Although the concept of duty is more related to moral decision making and less related to fairness explicitly, this notion of what is the “right” thing to do, or the “correct” rules to follow eventually do relate to fairness logics and decision-making processes.

The four classes of fairness characterized above represent a Western, colonial bias within ML scholarship [70]. Other understandings of fairness exist (e.g., Ubuntu fairness [30]), and while these other understandings of fairness were not expressed by the interviewees in this study, we hope that researchers will continue to explore a rich diversity of understandings of fairness, particularly through partnerships with non-Western organizations [18].

2.2 Institutional Logics

Scholars in organizational studies widely concur that broader belief systems shape the way that stakeholders operate and make decisions within organizations; these values-driven organizing principles and patterns of practices are referred to as *institutional logics* [63]. For example, the classes of fairness derived from the research literature serve as institutional logics of fairness when they serve as organizing principles for how individuals enact fairness. Numerous studies have characterized the ways in which multiple institutional logics co-occur within organizations—often in competition and contestation with each other (e.g., [8, 29, 49, 52]). In the nonprofit context, for example, organizations are commonly required to negotiate the competing institutional logics that are part and parcel of being a mission-driven organization (e.g., prioritizing services to clients) and logics that are derived from their

often-public sector funders (e.g., fiscal efficiency and accountability) [10, 24, 47]. As both a nonprofit organization and a financial institution responsible for managing individuals' investments, Kiva sits squarely at the boundaries of multiple institutional traditions with distinct institutional logics. As such, we expect that multiple logics within this organization will manifest in different understandings of how to enact key values such as fairness.

Connecting this body of research to the domain of computing, Voids et al. [65] further found that technologies also embody institutional logics in the ways they instantiate particular values. Even when organizational stakeholders agree on the importance of a particular value (such as fairness), that value may not be operationalized in technology in a way that is harmonious with stakeholders' various assumptions about and orientations towards how to put those values into practice. Such a mismatch creates tensions in practice that can create significant challenges for organizations and the clients they serve. This research, then, suggests that fairness-aware recommendation systems will need to be able to harmonize across multiple, potentially-conflicting logics about how the value of fairness should be operationalized.

2.3 The Multiplicity of Co-Occurring Logics

Research has shown that not everyone thinks of fairness in a similar way [48, 57]. And our literature review has identified at least four different and legitimate logics of fairness (see section 2.1). None of these fairness logics, then, are necessarily either *right* or *wrong*. Indeed, in nearly any context, the co-occurrence of multiple fairness logics is likely the norm and not the exception.

And yet, there has been relatively little recognition of the multiplicity of fairness logics in the ML literature [58]. Most existing research considers the differential impact of a system across only a single protected demographic class (e.g., race or gender). Even in cases where the impact of a system on multiple or intersectional groups is considered (e.g., [32, 37, 69]), the same concept of fairness is applied across all users, a decision that is at times justified through the application of Aristotle's concept of *Justice as Consistency* [55], which requires that similar individuals are treated similarly [21].

In contrast to the abundance of ML research that designs fairness interventions for only one class of stakeholders or applies only one concept of fairness, a subset of ML research has explicitly noted the benefits of combining multiple fairness definitions [9]. In particular, multistakeholder recommendation system research has acknowledged and embraced the multiplicity of stakeholders and fairness concerns that often arise among different groups of stakeholders [1, 14]. This research has explored tradeoffs between consumer (the users who receive recommendations) fairness [23] and provider (those providing the items to be recommended) fairness [60].

3 METHODS

We conducted semi-structured interviews with 23 employees of Kiva's microlending platform. We transcribed interviews and analyzed the transcripts using thematic analysis [11], resulting in a framework of fairness logics and characterizations of tensions among those logics.

3.1 Interviewees

We recruited 23 Kiva employees, each with a job role that has an investment in the recommendations of loans posted on Kiva's website. The interviewees work from five countries across three continents. They work for nine different teams within Kiva. We recruited our initial interviewees based on the advice of and introductions from the fourth author, who is Head of Data Science at Kiva. We recruited subsequent interviewees through snowball sampling. Due to the relatively small size of some of the teams at Kiva and the number of uniquely identifiable job roles, we are only able to characterize our interviewees in aggregate in order to support their anonymity. In what follows, we refer to our interviewees numerically as P1 through P23. Transcripts from P6 and P10 are not included in this analysis as they were pilot interviews for a future study with a Kiva lender and with an employee of one of Kiva's lending partners.

3.2 Data Collection

The third and last authors conducted semi-structured interviews with all interviewees. We tailored interview questions to each individual and their specific role at Kiva, with a focus on the following themes: their role at Kiva; how they apply the value of fairness in their work; challenges they experience when applying the value of fairness in their work; and how they interact with different technologies at Kiva. We conducted and audio-recorded interviews via Zoom video conferencing. Following each interview, the third author cleaned the transcripts that were generated by Zoom, anonymized them, and posted them to Google Drive for collaborative analysis by the research team. We continued recruiting interviewees until we had interviewed at least one individual in each team that interfaces with loan recommendations and multiple people within the larger teams at Kiva (to ensure breadth in our sample) and then continued recruiting until we had achieved theoretical saturation [12, 53].

3.3 Data Analysis

We conducted collaborative data analysis via the best practices of thematic analysis [11]. Following each interview, the third author drafted a memo capturing initial impressions of the interview, salient themes, and resonances with other interviews. After all interviews were completed, the first, second, and third authors began conducting open coding on interview transcripts; a generative process in which two authors independently tagged low-level themes in the transcript. The first round of open coding resulted in categories that included different definitions of fairness (e.g., fairness is defined as maximizing impact for borrowers), different ways that fairness was operationalized through work (e.g., determining impact scores, auditing for fairness concerns), and different fairness heuristics that focused on considerations for different groups of stakeholders (e.g., prioritizing risk over impact, educating lenders or borrowers). We then turned to explore the relationships among codes. In our next round of coding, we attempted to organize our data by stakeholder, building on related research [42] that has successfully used multi-stakeholder analysis to understand co-occurring fairness concerns. This round of coding ultimately felt insufficient in organizing the data as the relationship among stakeholders and fairness concerns was neither one-to-one nor clear cut.

In our second iteration, we switched to organizing our data by the underlying logic of the fairness consideration. This phase of analysis occurred in parallel with a substantive analysis of the fairness research literature to identify which fairness constructs and preexisting logics were most related to interviewees' descriptions of fairness at Kiva. The results of this coding phase were also used to organize the literature review for this paper. Our next round of coding, then, switched to being more deductive, using the classes of fairness from the literature review to organize the data. The results of this analysis are presented in section 4. Following this analysis, key excerpts of interviews remained that did not fit neatly into any one class of fairness. Instead, they hinged on conflicts between or intersections among classes of fairness. Our final analytic phase, then, focused on classifying the remaining excerpts in terms of these intersections. The tensions that result from these intersections are reported in section 5.

Positionality Statement. All authors live in the United States, and we recognize that our familiarity with Western constructs of fairness and our partnership with a nonprofit organization that is headquartered in the United States likely influenced our analysis and, together, reinforced the prominence of Western fairness logics in this research.

4 RESULTS: THE FOUR FAIRNESS LOGICS ENACTED ACROSS KIVA

What does it mean to be fair? (P21).

Kiva's mission to improve global financial inclusion is, in part, motivated by global, systemic biases that create financial inequity, most often disproportionately for women and girls. The organization seeks to address this inequity by providing capital to women and other underfinanced populations (e.g., historically 80% of Kiva loans have gone to women). Interviewees' diverse accounts of how they enacted fairness in their work aligned with the four classes of fairness outlined in section 2.1: consequence-based, contract-based, character-based, and duty-based. In what follows, we characterize each class of fairness as it is enacted by Kiva employees—each logic of fairness serving as an organizing principle that shapes individual action and decision making. We note that these interviewees speak from the perspective of their individual experiences and roles at Kiva; they do not speak on behalf of the organization, itself. In what follows, then, we characterize each of these fairness logics from the perspectives of our interviewees.

4.1 Consequence-Based Fairness Logic

Considerations of consequence-based fairness focus on achieving the best *outcome*, which can be accomplished in two ways: (1) by maximizing benefit (e.g., by funding a greater number of high-impact loans); or (2) by minimizing cost (e.g., by funding a greater number of low-risk loans). In general, consequence-based fairness logics consider stakeholders such as borrowers and lending partners—when maximizing benefit is equivalent to recommending more high-impact loans over low-impact loans; lenders—when minimizing cost is equivalent to recommending more low-risk loans; and the organization—when both maximizing impact and minimizing risk allows for more flow of capital, which enables the sustainability and growth of the organization's mission.

4.1.1 Benefit-Based Fairness Logic. The first sublogic of consequence-based fairness is benefit-based fairness, where the underlying goal is to maximize benefit.

Maximizing Benefit By Increasing Impact. Interviewees use the word “impact” to describe how a loan might positively change a borrower’s life. In our analysis, we use the word impact and benefit synonymously; while “impact” is used more often by interviewees, “benefit” is the language typically used in fairness literature. P2 noted that Kiva has created an “impact scorecard,” which uses a set of metrics (e.g., the type of loan, the poverty index of the country, and the lending partner’s historical impact) to determine a loan’s likely impact. P5 mentioned that one of Kiva’s goals is to increase funding for the most impactful loans; however, relying solely on an impact score might also be unfair in its own way:

You have to choose one loan that's the most impactful compared to all other loans and it's really hard to do that because everybody needs help to some degree (P5).

Another way to increase impact is to increase funds flowing through Kiva to their borrowers: “we obviously want as much money to go to these borrowers as possible” (P8). The implication here is that in order to maximize benefit, Kiva should recommend loans that lenders are most likely to fund over loans that are less likely to get funded.

Maximizing Benefit Through Purchasing Power. Another method for maximizing benefit is to recommend loans from countries where the value of the dollar is worth more. P4 suggested that some lenders prefer funding loans from low-income countries, believing that the money will go farther and have a greater impact: “there is a perception that the \$25 that you lend can go further for someone in another country which isn’t wrong, but it’s different” (P4). P5 elaborated that the mentality of “why [does Kiva give loans to individuals] in the United States if you can reach 50 more people in Africa for the average cost of the loan” (P5) can bias investment towards low-income countries to increase the number of people served. P16 reaffirmed that people in high-income countries such as the United States also have issues with equality of opportunity: “[there are] plenty of people in the US who also don’t have equal opportunity” (P16) and worried that lenders’ preferences to maximize purchasing power by lending outside of the US could limit the chances of underserved people in high-income countries.

4.1.2 Cost-Based Fairness Logic. A cost-based fairness logic strives to minimize the negative impacts of the lending process (e.g., if a borrower does not receive a loan). Most of the loans that are posted on Kiva’s online marketplace have already been funded through Kiva’s lending partners. Interviewees noted that lending partners depend on being recapitalized through Kiva’s marketplace; if they are not, it will impact their ability to make future loans. While this will not affect the borrowers already featured on the platform (as they are already funded), it could affect the partner’s decisions about which borrowers to finance with future loans.

Minimizing Cost By Minimizing Risk. Interviewees most commonly suggested minimizing cost by reducing risks for lenders—recommending less risky loans. P8 explained that if new lenders are repaid, they will know that they “did good” because “... [the borrowers] were able to repay you” (P8). P3 added that recommending low-risk loans will not only keep lenders coming back to the site

but will also help mobilize more capital for lending, generally—if a lender is paid back, they are more likely to reinvest that money in additional loans and the same dollar becomes a force multiplier for good. However, sometimes global and local issues can affect a lending partner's ability to repay a loan. P8 worried that over-focusing on risk and repayment rates might potentially lead to lower funding for borrowers living in countries that have volatile commodity markets. This interviewee added that focusing only on repayment capabilities might also only help the lending partners who are charging high-interest rates (which can hurt vulnerable borrowers). As a result, and in contrast to a cost-based fairness logic, P8 felt that it is okay to “sacrifice a bit of your risk aversion for the sake of being more impactful” (P8).

4.2 Contract-Based Fairness Logic

When interviewees' fairness logics aligned with contract-based fairness, it was always focused on enabling equality of opportunity—and nearly always for borrowers. Contract-based fairness aligns with Kiva's mission to improve global financial inclusion.

The mission is to make sure those who are most financially excluded are the ones getting funded [...] Now, something that would be in line with our mission would be making sure those applicants are the first ones seen on the site as far as US loans go (P19).

P2 described this concept as “filling the gaps” in funding: “it's not just that Kiva has less investments, it's that there are less investments in that country so that's where I want to focus. So for me, it's all about the gaps, where are the gaps in funding” (P2). Several interviewees noted specific interventions to address these gaps, including recommending loans for women and borrowers from low-income countries (P1, P2, P3, P4, P9, P12, P13, P22). P11 also mentioned interventions that are aimed at increasing inclusivity among borrowers—for example, providing translation services to non-English speaking borrowers and technical support for elderly borrowers.

However, enacting contract-based fairness is not as straightforward as it may seem. Despite being historically underfunded, on Kiva, female borrowers tend to be much more likely to receive funding when compared to their male counterparts (P3, P9, and P12). Similarly, P4 mentioned that small business owners in the US are also systematically underfunded in the online marketplace, raising the question of whether contract-based fairness applies to equal opportunity globally, or only within the context of Kiva's marketplace.

4.3 Character-Based Fairness Logic

Considerations of character-based fairness most typically explore the decision-making process that lenders go through when deciding who is “worthy” of a loan. Interviewees nearly exclusively discussed this logic while reflecting on the user experience of lenders and the ways in which borrowers are represented in the online marketplace. Interviewees believed that different lenders were more or less comfortable applying a character-based fairness logic and serving as an “arbiter” of which borrowers might be most “worthy.”

There are some lenders who don't like choosing [which loans to fund], because they don't believe in themselves

as an arbiter of who deserves. And there are other lenders who very comfortably sit in this place of evaluating who deserves a loan and who doesn't.... We see both of those mindsets play out in some of these situations where loan worthiness is considered more visible or obvious in certain contexts (P23).

P16 suggested that lenders on Kiva who have more microlending experience tend to prefer to make decisions about who to fund: “*There are some lenders who really care about microfinance and they came to Kiva ten years ago when microfinance was at its peak and they have specific ideas of what type of loan they're looking for*” (P16). However, P16 also shared that other lenders likely felt “*paralyzed by the choice*” when deciding who to fund; that they were not comfortable “*choos[ing] between two faces who [both] need help, and that doesn't feel good*” (P16). Similarly, P1 shared that they know “*friends and folks who've started using Kiva... [and] they struggle with selecting [a borrower] because they feel like they don't have enough criteria to pick like who's worthy*” (P1).

Given interviewees' inferences about lenders' comfort levels with a character-based fairness logic, many reflected on how to best support all lenders in a single online marketplace. If lenders want to make a decision about what borrower to fund using a character-based fairness logic, then Kiva might offer more information to support that decision making—but what information? While the microfinance enthusiasts mentioned earlier tend to want to see lending partners' risk scores, many lenders rely on visual information, especially the photo of the borrower [4, 35]. While P3 explained that the borrower's photo is an important UX component (though not used algorithmically in any way to promote loans), it is also a place where discrimination and bias are very likely to enter the decision-making process: “*The best content really is the best pictures on our website, ... [but] what if [the borrower] didn't feel like smiling that day, what if culturally they're just less likely to smile in a picture, these are all problems... that I can't solve*” (P3).

For lenders who felt “*uneasy*” (P11) making decisions about who to fund (i.e., using a character-based fairness logic), interviewees mentioned the design of auto-lending features that would allow lenders to assign Kiva as a proxy for the decision. Interviewees also mentioned that minimizing the amount of information presented about borrowers might also help these lenders, for example:

A lot of lenders, probably a majority of lenders, do not care how we're doing this thing. They want to come to the website, find somebody whose story appeals to them, do a nice thing, and like move on with their day. They don't want to know the risk rating... they maybe don't even care what country, they just don't want all this information (P3).

4.4 Duty-Based Fairness Logic

Interviewees' considerations of duty-based fairness align with the rules, duties, responsibilities, and accountability that Kiva has towards its various stakeholders. We recognize that duty-based fairness is more of a “meta” logic in that any rules that are chosen could overlap or conflict with any other fairness logic. However, this high-level process of deciding what an organization's values and duties are, and which rules they should follow to align with

these values is an important part of enacting fairness at Kiva. We explore this complexity throughout this section.

P25 shared that “*one of our key core Kiva values is honor and integrity*,” which, for them, meant that Kiva has a duty to “*do the most right thing in the most right way*” (P25). But “*doing the right thing*” means different things for Kiva’s many different stakeholders:

So our algorithm is like determining whether or not people get air time, which, in a high supply environment where there are more loans than dollars, which is what we want... we want to maintain just a little more [loans] than we can fund.... We’re responsible for who gets funded and who doesn’t. We’re creating the environment in which that happens, and there have been at different times like sort of like free market-y ideas at Kiva like “it’s a marketplace, some stuff gets funded, some doesn’t, it’s not our fault,” which is – it’s completely not true. Like we’re curating it to a huge extent, like the number of people involved, the number of rules, systems, processes involved in getting a loan onto the website is super high and the algorithm that presents them to lenders we made so while, yes, it’s a marketplace and stuff, we are responsible for proper rulemaking (P3).

Different interviewees also prioritized different fairness logics when doing the “right thing.” For P1, for example, doing the “right thing” meant maximizing contract-based fairness:

If there’s a way in which our machine is further perpetuating systemic racism or lack of access for certain people based on where they live or how they look or whatever, like those are things that we want to be aware of (P1).

For P2, doing the “right thing” meant protecting lenders’ money, which meant maximizing cost-based fairness:

Their focus is protecting lender money and protecting our repayment rate... trying to ensure that our lenders will get their money back, and that our repayment rate stays high and we’re not losing too much money (P2).

These two operationalizations of duty—while both important and potentially complementary—define a tension or tradeoff in system design, a theme that we turn to next.

5 RESULTS: TENSIONS BETWEEN FAIRNESS LOGICS

When multiple fairness logics co-exist, interviewees often noted tradeoffs or tensions between the logics and between the decisions that would result. P19 characterized these tensions as a “*tough philosophical question*”:

Then comes the question of like, is [intervening on loans] being equitable, and giving people like more opportunity who we feel are deserving of it? Or is that us kind of like playing God and not allowing the market to determine whether or not people should get funded like I don’t know the answer to that question. That’s a tough philosophical question (P19).

More specifically, in this section we identify three key tensions between fairness logics, as visualized in Figure 1¹.

5.1 Tensions Between Benefit-Based and Cost-Based Fairness Logics

The first tension occurs within the consequence-based logic, when a choice may be necessary between maximizing benefit (maximizing impact) and minimizing cost (minimizing risk). Several interviewees explained how, working from a consequence-based fairness logic, it might not be possible to achieve both of these goals at once.

Interviewees reported that Kiva has conducted internal research to better understand which kinds of loans have the highest impact on individuals and their communities. One thing that they learned is that many high-impact loans might also be high-risk, because “*moving [borrowers] out of financially vulnerable to financially stable*” (P12) sometimes requires lending to borrowers who are at a higher risk of not being able to pay the lender back.

However, prioritizing high-risk loans could reduce the flow of money through the online marketplace; if these higher-risk loans are not repaid, there will be less money for lenders to re-loan to other borrowers. Thus, a tension is created: if Kiva prioritizes maximizing impact, borrowers will benefit but lenders might not get repaid; if Kiva prioritizes minimizing risk, lenders will benefit, but certain borrowers might be systematically underserved—which is against Kiva’s mission. Designing for consequence-based fairness, then, requires choosing how to balance these two fairness logics.

One interviewee felt that both of these goals (maximizing benefit and minimizing cost) could be accomplished simultaneously, but leaning on a very different operationalization of impact:

There is this perception, I think a lot of the time, that risk and impact are kind of opposite sides of the spectrum. You can either be impactful or you can be low risk. I think that’s a fallacy. I think they’re kind of two parallel spectrums, spectra. Where you know you can be impactful, and you can be low risk at the same time (P8).

P8 gave as an example that if a lender were to lend \$25 to a low-risk loan, then they are more likely to get paid back, and can consequently lend that same \$25 to more loans in the future—maximizing their overall impact while also minimizing their risk (P8). In this scenario, benefit-based and cost-based fairness logics complement rather than conflict with one another.

5.2 Tensions Between Contract-Based and Character-Based Fairness Logics

The second tension occurs between contract-based fairness and character-based fairness, when prioritizing equality of opportunity might conflict with the funding preferences of lenders. Interviewees described the challenges of improving equality of opportunity (contract-based fairness) while also attending to lenders who have their own preferences about who to fund or preconceptions about

¹Note that duty-based fairness is excluded from the left side of Figure 1 because interviewees were not in consensus about how to prioritize Kiva’s duties in enacting fairness. Depending on which duties interviewees’ emphasized, a duty-based logic could either conflict or complement the others.




Figure 1: Left: key tensions between fairness logics. Right: the four logics and stakeholders that are prioritized by each.

what kind of borrower should be funded (character-based fairness). As P1 summarized:

[Fairness interventions] wouldn't be necessary if we didn't have some people who were deemed more worthy than others, and... those concepts do come from lenders (P1).

And P4 worried about relying solely on lenders to determine who is worthy of funding: *"The average person doesn't understand the impact of one specific loan"* (P4).

Numerous interviewees provided examples of how lenders' preferences have manifested in lending biases across different classes of Kiva borrowers:

- **Types of businesses:** P16, for example, has *"heard [lenders] say like I'm not going to lend to a beauty salon because that doesn't feel important to me."* (P16). But making value judgments about what types of businesses are "worthy" or not imposes Western biases on borrowers. P21 provides the example of lenders not feeling comfortable lending to borrowers whose businesses would be illegal in the United States: *"if it's not illegal in the country of origin and it's not illegal due to international considerations, then what are our rights, from our point of view, to prohibit that type of lending"* (P21).
- **Gender:** Many lenders prefer to fund borrowers who identify as female, systematically underfunding men on Kiva. P12 points out that in certain countries, men may actually need to be funded *more* than women, since they may be the only household member who provides finances for their family because *"gender roles are very different in different countries around the world"* (P12).
- **Geographic regions:** P9 noted that lenders also show regional preferences, for example *"Eastern Europe or certain parts of Asia"* do not fund as quickly as lending partners in Africa.

Each of these lender preferences creates a tension with ensuring equality of economic opportunity for Kiva's borrowers. P22

additionally worried that relying on lenders' preferences could unintentionally amplify lending biases on the online marketplace and result in underserved (long-tail) groups of borrowers:

I think it's not all bad, but it is a cycle that we want to try to not perpetuate too deeply (P22).

And yet, pushing back on lender preferences too much might also have an adverse effect on equality of opportunity. As P11 explained, *"I think we found limits and how much we can like push things on people"* (P11). If lenders feel pushed too far and leave the platform, then equality of opportunity cannot be achieved, since there would be no funds left to give to borrowers.

5.3 Tensions Between Contract-Based and Consequence-Based Fairness Logics

The third key tension occurs between contract-based fairness and consequence-based fairness, when prioritizing equality of opportunity might be at odds with keeping enough money flowing through its online marketplace to keep the nonprofit afloat, pay back its lenders, and pay out to its borrowers and lending partners. P3, P11, and P25 all noted that these pragmatic needs of the organization (consequence-based fairness) all have to co-exist alongside Kiva's mission of increasing financial inclusion (contract-based fairness). Since a contract-based fairness logic prioritizes the "underdog," emphasizing this logic might, for example, entail nudging lenders toward higher risk loans than they might otherwise be drawn towards. P15 believed that it is risky to try to change lenders' behaviors in this way because you might be *"turning people off"* (P15) from lending, and doing so would conflict with the organizational need to keep money flowing through the online marketplace. P3 also worried that if lenders are pushed to fund high-risk loans but are not repaid, they might feel *"betrayed or misled,"* regardless of whether they were lending to a higher-risk borrower for the sake of improving financial equity.

So while working from a consequence-based fairness logic might be best suited to sustaining the organization, it might also reproduce financial inequality—which works against Kiva’s mission. On the other hand, working from a contract-based fairness logic might place Kiva in a position where there are fewer lenders and less funding overall, putting the organization at risk. Balancing these tensions is a critical challenge for the design of microlending recommendation, which we discuss in the next section.

6 DISCUSSION: DESIGN CASES REFLECTING TRADEOFFS, OPPORTUNITIES, AND IMPLICATIONS FOR STAKEHOLDERS

In this research, we have applied logics of fairness as an analytic lens to understand multiple, co-existing fairness considerations. We have characterized four different fairness logics and identified three key tensions that arise at intersections between them. These intersections between logics also represent sites of key design tradeoffs and opportunities to design for multiple co-occurring logics. In this discussion, we explore some of these design opportunities and design tradeoffs. Of particular importance here is the task of understanding how tradeoffs in design impact different stakeholders of Kiva’s recommendation ecosystem.

Different fairness logics prioritize different classes of stakeholders, as shown in Figure 1. Organizing decision making from a contract-based fairness logic centered around Kiva’s mission of financial inclusion prioritizes borrowers. Organizing decision making from a character-based fairness logic privileges lenders’ preferences for lending and their perceptions about what classes of borrowers and loans are “worthy” to be funded. Organizing decision making from a duty-based logic impacts both lenders and borrowers/lending partners, as Kiva has a responsibility to all of these stakeholders. And organizing decision making from a consequence-based fairness logic prioritizes minimizing risks to increase the flow of capital through the online marketplace, benefiting the organization and the lender; it can also benefit borrowers if the emphasis is on maximizing impact; or all of these stakeholders if the loans maximize impact and minimize risk simultaneously.

In what follows, we explore two design cases, where tradeoffs exist between fairness logics and where opportunities for innovation can cater to multiple fairness logics at once. Many of these design options already exist—in varied combinations—across Kiva’s online marketplace; here we detail them discretely in order to make explicit their relationships with fairness logics.

6.0.1 The Presentation of Loans. One tradeoff mentioned in section 4 relates to what kinds of information Kiva might provide to lenders in the presentation of loan opportunities, their borrowers, and lending partners. As previously mentioned, Kiva already provides a wide range of information about each loan on the platform, including:

- Basic information about the loan (e.g., total amount, expiration date, and what the loan will be used for);
- Qualitative and visual information about the borrower (e.g., a photo, their personal story); this design enacts a character-based logic; and

- Quantitative information about the borrower and their lending partner, including repayment rates, risk and impact scores; this design enacts a consequence-based logic.

Providing information like impact and risk scores can be helpful for certain lenders—those motivated by a consequence-based fairness logic—but can also perpetuate biases against lower-impact loans or higher-risk borrowers/lending partners, which undermines the contract-based fairness logic. Providing information about a borrower’s personal life, such as if they are a parent, might help some borrowers and hurt others, leaving this design option open to introducing new biases for those working from a character-based fairness logic. One could combine multiple forms of information in the presentation of a loan in order to help organize decision making from multiple fairness logics. One could also personalize the genres of information on the presentation of loan page for different lenders based on what fairness logics they are inferred to enact when lending (i.e., engaging predominantly with information about a borrower (enacting a character-based logic), engaging predominantly with data about the repayment rates of lenders (enacting a consequence-based logic)).

Yet another design option would be to back away from interventions at the level of the individual loan and provide more generalizable information about *thematic loan categories*. Kiva already provides options for lenders to filter loans categorically (e.g., “agriculture” or “education”). At this categorical level, one could present information about the importance of lending to each of these categories. For example, P9 pointed out that borrowers from certain regions could be systematically underfunded on Kiva. If a certain loan category is found to be underfunded, one could design a thematic cover page for all loan opportunities within that theme, with information about why this class of borrowers is important. This is similar to the “Spotlighted by Kiva” carousel that is already available on the platform, which highlights loans that are likely to not receive full funding. Another similar, yet alternative design could involve allowing lenders to choose to place a loan in this category, leaving the decision of what specific loan should be funded to Kiva (for lenders who are not comfortable enacting a character-based fairness logic) or providing additional context for lenders who want to select their own loan opportunities. This design would enact a combination of character-based, consequence-based, and contract-based fairness logics.

6.0.2 The Loan Discovery Page. As described above, Kiva currently uses a loan discovery page to help lenders find loan opportunities via a set of themed horizontal carousels, typically with 10–12 loans each. One set of design options revolves around choosing themes for these carousels, for example:

- Loans from categories, geographic regions, or borrower demographics that are the same as those the lender has lent to previously; this theme enacts a character-based fairness logic;
- Loans from categories, geographic regions, or borrower demographics that are more likely to go un-funded, regardless of lenders’ funding preferences; this theme enacts a contract-based fairness logic; and
- Loans with high impact scores; this theme enacts a consequence-based fairness logic.

While all of these carousels could easily be included on the loan discovery page, the order in which carousels appear can also be an object of design. Their ordering could prioritize different fairness logics or could be personalized based on inferences about what fairness logics the lender has enacted through their previous loans.

Another set of design options revolves around choosing the subset of loan opportunities that are promoted on each carousel and the order in which they are displayed—their ranking—for example:

- Ranking well-funded (“popular”) loans higher; this option enacts a character-based fairness logic, though it tends to reinforce lender biases;
- Ranking under-funded loans higher; this option enacts a contract-based fairness logic and mitigates the ‘popularity bias’ effect that can fall out of enacting a character-based logic;
- Ranking high-risk loans (less likely to get repaid) higher; this option also enacts a contract-based fairness logic but, instead, rebalances against some of the effects that can fall out of enacting a consequence-based logic; and
- Ranking low-risk loans (more likely to get repaid) higher; this option enacts a consequence-based fairness logic.

However, as multiple interviewees in our study shared, if one goes too far in “*push[ing] things*” (P11) on lenders, the sustainability of the online marketplace is put at risk. As a result, one might also want to explore ways of designing the loan discovery page—both the carousels and the loan rankings within each carousel—to prioritize multiple of these co-occurring fairness logics at once and experiment with weighting them in different ways in different contexts [45, 60]. Social choice theory, for example, has been successfully used as a framework for algorithm design that combines multiple fairness logics [15, 17, 59]. One might also personalize the design approach by inferring which fairness logics lenders have previously enacted in their lending decisions and weighting their recommendations based on a similar balance of fairness logics (e.g., [25, 43, 45, 60]).

Finally, one might consider providing additional transparency about the fairness logics enacted through design. Sonboli and Smith et al. [61], for example, found that a ‘tool tip’ providing a fairness explanation was received favorably by lenders, as it allows them to have the *option* to act on explicit fairness considerations rather than feeling nudged or manipulated into doing so. Selecting *which* logics to prioritize in *what contexts* on Kiva’s online marketplace and *how* to operationalize these logics is a compelling challenge for future work.

7 LIMITATIONS & FUTURE WORK

Our research has a number of limitations. First, our interviewees were all employees of Kiva. While their wide-ranging roles across the organization enabled them to offer a diversity of insights into organizational experiences and perspectives of fairness, they could only provide anecdotal or subjective opinions about how fairness considerations might impact lenders, borrowers, or lending partners. Future research should explore fairness considerations of an even broader sample of stakeholders.

In this research, we offer a case study of fairness considerations in one exemplar organization that has been incredibly transparent

about its fairness mission and work in pursuit of fairness. Although generalizability is not a benchmark for qualitative research, it is appropriate to look to future research to ask how transferable the findings of this study might be [56]—for example, whether some of our findings might transfer to similar organizations from either the financial or nonprofit sectors (near transfer) or whether some of the findings might transfer to organizations that differ in more substantive ways (far transfer). Studying the diversity of fairness considerations in other cross-sector organizations would be a particularly compelling direction for future research and to assess the far transferability of these findings.

The *methods* that we introduce in this research—conducting semi-structured interviews across a breadth of stakeholders within an organization to understand how they enact key values in practice, identifying the fairness logics that are revealed through those interviews, and analyzing the tensions and trade-offs among fairness logics—are a generalizable contribution that should be useful for future researchers who seek to explore multistakeholder fairness in practice. In general, we believe that fairness concerns in recommender systems will always be complex, multiple in kind, and in tension with each other. Any organization seeking to enhance fairness in their recommendation platform will need to conduct a detailed study of how fairness is understood from the perspective of multiple stakeholders. We have provided an example of such a study here.

Finally, the qualitative research we have carried out here also serves as a form of formative requirements gathering for recommender system design. Future research will be needed to design algorithms that can balance multiple co-occurring fairness considerations (e.g., [17]) as well as to explore methods of transparency with stakeholders about who is being prioritized in fairness-aware recommendations.

8 CONCLUSION

In this research, we have explored the rich and diverse landscape of fairness considerations related to recommending loans on Kiva.org. Drawing from semi-structured interviews with 23 employees at Kiva, we identified and characterized (a) four fairness logics that are enacted through their work; and (b) three fairness tensions that are situated at the intersections of these logics. We unpack how each fairness logic prioritizes certain stakeholders over others and explore several design cases in order to demonstrate how different design decisions might differentially prioritize different logics or stakeholders. We hope that this research provides methodological inspiration for how others might explore the complex contexts of fairness—exploring and modeling the multiple co-occurring logics of fairness before metrics are decided upon or interventions are designed.

ACKNOWLEDGMENTS

We are indebted to our Kiva collaborators and the leadership team there for making this collaboration possible. Many thanks, as well, to all of the Kiva interviewees, whose passionate pursuit of fairness has brought so much joy to this work. This publication is based upon research supported by the National Science Foundation under grants #2107577 and #2107505.

REFERENCES

- [1] Himan Abdollahpouri, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction* 30 (2020), 127–158. <https://doi.org/10.1007/s11257-019-09256-1>
- [2] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. 2019. Managing popularity bias in recommender systems with personalized re-ranking. In *Proceedings of the Thirty-Second International FLAIRS Conference*. AAAI Press, Washington, DC, USA, 413–418.
- [3] Amanda Aird, Paresha Farastu, Joshua Sun, Amy Voda, Nicholas Mattei, and Robin Burke. 2023. Dynamic fairness-aware recommendation through multi-agent social choice. *CoRR* abs/2303.00968 (2023). <https://doi.org/10.48550/arXiv.2303.00968> arXiv:2303.00968
- [4] Kenton Anderson and Gregory D Saxton. 2016. Smiles, babies, and status symbols: The persuasive effects of images in small-entrepreneur crowdfunding requests. *International Journal of Communication* 10 (2016), 1764–1785.
- [5] Richard Arneson. 2002. Equality of opportunity. In *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, Palo Alto, CA, USA. <https://plato.stanford.edu/archives/sum2015/entries/equal-opportunity/>
- [6] Solon Barocas, Elizabeth Bradley, Vasant Honavar, and Foster Provost. 2017. Big data, data science, and civil rights. arXiv preprint arXiv:1706.03102.
- [7] Solon Barocas, Moritz Hardt, and Arvind Narayanan. 2017. Fairness in machine learning. <https://fairmlbook.org/pdf/fairmlbook.pdf>.
- [8] Marya L Besharov and Wendy K Smith. 2014. Multiple institutional logics in organizations: Explaining their varied nature and implications. *Academy of Management Review* 39, 3 (July 2014), 364–381. <https://doi.org/10.5465/amr.2011.0431>
- [9] Alex Beutel, Jilin Chen, Tulsee Doshi, Hai Qian, Li Wei, Yi Wu, Lukasz Heldt, Zhe Zhao, Lichan Hong, and Ed H Chi. 2019. Fairness in recommendation ranking through pairwise comparisons. arXiv:1903.00780 [cs.CY]
- [10] Amy Binder. 2007. For love and money: Organizations' creative responses to multiple environmental logics. *Theory and Society* 36, 6 (Dec. 2007), 547–571. <https://doi.org/10.1007/s11186-007-9045-x>
- [11] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. <https://doi.org/10.1191/147808706qp063oa>
- [12] Virginia Braun and Victoria Clarke. 2016. (Mis)conceptualising themes, thematic analysis, and other problems with Fugard and Potts' (2015) sample-size tool for thematic analysis. *International Journal of Social Research Methodology* 19, 6 (2016), 739–743. <https://doi.org/10.1080/13645579.2016.1195588>
- [13] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*. Proceedings of Machine Learning Research, Vol. 81. PMLR, New York, NY, USA, 77–91. <https://proceedings.mlr.press/v81/buolamwini18a.html>
- [14] Robin Burke. 2017. Multisided fairness for recommendation. 2017 Workshop on Fairness, Accountability and Transparency in Machine Learning (FAT/ML 2017). arXiv:1707.00093 [cs.CY]
- [15] Robin Burke, Nicholas Mattei, Vladislav Grozin, Amy Voda, and Nasim Sonboli. 2022. Multi-agent social choice for dynamic fairness-aware recommendation. In *Adjunct Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '22)*. ACM Press, New York, NY, USA, 234–244. <https://doi.org/10.1145/3511047.3538032>
- [16] Robin Burke, Pradeep Raghavan, Nicholas Mattei, Brian Kimmig, Amy Voda, Nasim Sonboli, Anushka Kathait, and Melissa Fabros. 2022. A performance-preserving fairness intervention for adaptive microfinance recommendation. In *Proceedings of the KDD Workshop on Online and Adapting Recommender Systems (OARS)*. ACM Press, New York, NY, USA, 6 pages.
- [17] Robin Burke, Amy Voda, Nicholas Mattei, and Nasim Sonboli. 2020. Algorithmic fairness, institutional logics, and social choice. Harvard CRCS Workshop: AI for Social Good. https://projects.iq.harvard.edu/files/crcs/files/ai4sg_2020_paper_29.pdf
- [18] Emanuelle Burton, Judy Goldsmith, Nicholas Mattei, Cory Siler, and Sara-Jo Swiatek. 2023. *Computing and Technology Ethics: Engaging through Science Fiction*. MIT Press.
- [19] Sushma Channamsetty and Michael D Ekstrand. 2017. Recommender response to diversity and popularity bias in user profiles. In *Proceedings of the Thirtieth International FLAIRS Conference*. AAAI Press, Washington, DC, USA, 657–660.
- [20] Henriette Cramer, Jeni Wortman Vaughan, Ken Holstein, Hanna Wallach, Jean Garcia-Gathright, Hal Daumé III, Miroslav Dudík, and Sravana Reddy. 2019. Challenges of incorporating algorithmic fairness into industry practice. FAT* 2019 Tutorial. https://www.microsoft.com/en-us/research/uploads/prod/2020/10/FAT_2019-tutorial_-algorithmic-bias-in-practice.pdf
- [21] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*. ACM Press, New York, NY, USA, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [22] Ben Eggleston. 2014. Act utilitarianism. In *The Cambridge Companion to Utilitarianism*, Ben Eggleston & Dale E. Miller (Ed.). Cambridge University Press, Cambridge, UK, 125–145.
- [23] Michael D. Ekstrand, Mucun Tian, Ion Madrazo Azpiroz, Jennifer D. Ekstrand, Oghenemaro Anuyah, David McNeill, and Maria Soledad Pera. 2018. All the cool kids, how do they fit in?: Popularity and demographic biases in recommender evaluation and effectiveness. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*. PMLR, New York, NY, USA, 172–186. <https://proceedings.mlr.press/v81/ekstrand18a.html>
- [24] Adalbert Evers. 2005. Mixed welfare systems and hybrid organizations: Changes in the governance and provision of social services. *International Journal of Public Administration* 28, 9–10 (2005), 737–748. <https://doi.org/10.1081/PAD-200067318>
- [25] Paresha Farastu, Nicholas Mattei, and Robin Burke. 2022. Who Pays? Personalization, Bossiness and the Cost of Fairness. <https://doi.org/10.48550/ARXIV.2209.04043>
- [26] Andres Ferraro, Xavier Serra, and Christine Bauer. 2021. What is fair? Exploring the artists' perspective on the fairness of music streaming platforms. In *Human-Computer Interaction–INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30–September 3, 2021, Proceedings, Part II*. Springer, Cham, Switzerland, 562–584. https://doi.org/10.1007/978-3-030-85616-8_33
- [27] Jessica Fjeld, Nele Achten, Hannah Hilligoss, Adam Nagy, and Madhulika Srikrumar. 2020. *Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI*. Technical Report 2020-1. Berkman Klein Center. <https://doi.org/10.2139/ssrn.3518482>
- [28] Jason Gabriel. 2022. Toward a theory of justice for artificial intelligence. *Daedalus* 151, 2 (2022), 218–231. https://doi.org/10.1162/daed_a_01911
- [29] Royston Greenwood, Amalia Magán Díaz, Stan Xiao Li, and José Céspedes Lorente. 2010. The multiplicity of institutional logics and the heterogeneity of organizational responses. *Organization Science* 21, 2 (March–April 2010), 521–539. <https://doi.org/10.1287/orsc.1090.0453>
- [30] Arthur Gwagwa, Emre Kazim, and Airlie Hilliard. 2022. The role of the African value of Ubuntu in global AI inclusion discourse: A normative ethics perspective. *Patterns* 3, 4 (8 April 2022), 100462. <https://doi.org/10.1016/j.patter.2022.100462>
- [31] Tatsunori Hashimoto, Megha Srivastava, Hongseok Namkoong, and Percy Liang. 2018. Fairness without demographics in repeated loss minimization. In *Proceedings of the 35th International Conference on Machine Learning*. Proceedings of Machine Learning Research, Vol. 80. PMLR, New York, NY, USA, 1929–1938. <https://proceedings.mlr.press/v80/hashimoto18a.html>
- [32] Ursula Hébert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. 2018. Multicalibration: Calibration for the (computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning*. Proceedings of Machine Learning Research, Vol. 80. PMLR, New York, NY, USA, 1939–1948. <https://proceedings.mlr.press/v80/hebert-johnson18a.html>
- [33] Kenneth Holstein, Jennifer W. Vaughan, Hal Daumé III, Miroslav Dudík, and Hanna Wallach. 2019. Improving fairness in machine learning systems: What do industry practitioners need? In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. ACM Press, New York, NY, USA, Article 600, 16 pages. <https://doi.org/10.1145/3290605.3300830>
- [34] Rosalind Hursthouse. 2007. Virtue ethics. In *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (Ed.). Metaphysics Research Lab, Stanford University, Palo Alto, CA, USA. <https://plato.stanford.edu/archives/fall2007/entries/ethics-virtue/>
- [35] Christina Jenq, Jessica Pan, and Walter Theseira. 2015. Beauty, weight, and skin color in charitable giving. *Journal of Economic Behavior & Organization* 119 (Nov. 2015), 234–253. <https://doi.org/10.1016/j.jebo.2015.06.004>
- [36] Stamatia Karnouskos. 2021. The role of utilitarianism, self-safety, and technology in the acceptance of self-driving cars. *Cognition, Technology & Work* 23, 4 (Nov. 2021), 659–667. <https://doi.org/10.1007/s10111-020-00649-6>
- [37] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2018. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. arXiv:1711.05144 [cs.LG]
- [38] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. 2019. An empirical study of rich subgroup fairness for machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19)*. ACM Press, New York, NY, USA, 100–109. <https://doi.org/10.1145/3287560.3287592>
- [39] Os Keyes, Jevan Hutson, and Meredith Durbin. 2019. A mulching proposal: Analysing and improving an algorithmic system for turning the elderly into high-nutrient slurry. In *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (CHI EA '19)*. ACM Press, New York, NY, USA, Article alt06, 11 pages. <https://doi.org/10.1145/3290607.3310433>
- [40] Alexander Larry and Michael Moore. 2016. Deontological ethics. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Stanford, CA, USA. <https://plato.stanford.edu/archives/win2016/entries/ethics-deontological/>
- [41] Manzoor Laskar. 2013. Summary of social contract theory by Hobbes, Locke and Rousseau. SSRN. <https://doi.org/10.2139/ssrn.2410525>

[42] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, and Alexandros Psomas. 2019. WeBuildAI: Participatory framework for algorithmic governance. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW, Article 181 (Nov. 2019), 35 pages. <https://doi.org/10.1145/3359283>

[43] Yunqi Li, Hanxiong Chen, Shuyuan Xu, Yingqiang Ge, and Yongfeng Zhang. 2021. Towards personalized fairness based on causal notion. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, NY, 1054–1063.

[44] Yung-Ming Li, Jyh-Hwa Liou, and Yi-Wen Li. 2020. A social recommendation approach for reward-based crowdfunding campaigns. *Information & Management* 57, 7 (Nov. 2020), 103246. <https://doi.org/10.1016/j.im.2019.103246>

[45] Weiwen Liu, Jun Guo, Nasim Sonboli, Robin Burke, and Shengyu Zhang. 2019. Personalized fairness-aware re-ranking for microlending. In *Proceedings of the 13th ACM Conference on Recommender Systems (RecSys '19)*. ACM Press, New York, NY, USA, 467–471. <https://doi.org/10.1145/3298689.3347016>

[46] Kiva Microfunds. 2023. Make a loan, change a life. <https://kiva.org/>

[47] David Mullins. 2006. Competing institutional logics? Local accountability and scale and efficiency in an expanding non-profit housing sector. *Public Policy and Administration* 21, 3 (Sept. 2006), 6–24. <https://doi.org/10.1177/0952076706021003>

[48] Arvind Narayanan. 2018. 21 fairness definitions and their politics. Tutorial presented at the Conference on Fairness, Accountability, and Transparency (FAT* Conference 2018). <https://fairmlbook.org/tutorial2.html>

[49] Anne-Claire Pache and Filipe Santos. 2013. Embedded in hybrid contexts: How individuals in organizations respond to competing institutional logics. In *Institutional logics in action, part B*, Michael Lounsbury and Eva Boxenbaum (Eds.). Emerald Group Publishing Limited, Bingley, UK, 3–35.

[50] Yoon-Joo Park and Alexander Tuzhilin. 2008. The long tail of recommender systems and how to leverage it. In *Proceedings of the 2008 ACM Conference on Recommender Systems (RecSys '08)*. ACM Press, New York, NY, USA, 11–18. <https://doi.org/10.1145/1454008.1454012>

[51] Bogdana Rakova, Jingying Yang, Henriette Cramer, and Rumman Chowdhury. 2021. Where responsible AI meets reality: Practitioner perspectives on enablers for shifting organizational practices. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1, Article 7 (April 2021), 23 pages. <https://doi.org/10.1145/3449081>

[52] Trish Reay and C Robert Hinings. 2009. Managing the rivalry of competing institutional logics. *Organization Studies* 30, 6 (June 2009), 629–652. <https://doi.org/10.1177/0170840609104>

[53] Benjamin Saunders, Julius Sim, Tom Kingstone, Shula Baker, Jackie Waterfield, Bernadette Bartlam, Heather Burroughs, and Clare Jinks. 2018. Saturation in qualitative research: Exploring its conceptualization and operationalization. *Quality & Quantity* 52, 4 (July 2018), 1893–1907. <https://doi.org/10.1007/s11135-017-0574-8>

[54] Steve Sawyer and Mohammad Hossein Jarrahi. 2014. Sociotechnical approaches to the study of information systems. In *Computing Handbook, Third Edition: Information Systems and Information Technology*, Heikki Topi (Ed.). CRC Press, Boca Raton, FL, 5–1.

[55] Frederick Schauer. 2018. On treating unlike cases alike. SSRN. <https://ssrn.com/abstract=3183939>

[56] Andrew K. Shenton. 2004. Strategies for ensuring trustworthiness in qualitative research projects. *Education for Information* 22, 2 (2004), 63–75. <https://doi.org/10.3233/EFI-2004-22201>

[57] Jessie Smith, Nasim Sonboli, Casey Fiesler, and Robin Burke. 2020. Exploring user opinions of fairness in recommender systems. Position paper presented at the CHI 2020 Workshop on Human-Centered Approach to Fair & Responsible AI. arXiv:2003.06461 [cs.IR]

[58] Jessie J Smith, Lex Beattie, and Henriette Cramer. 2023. Scoping Fairness Objectives and Identifying Fairness Metrics for Recommender Systems: The Practitioners' Perspective. In *Proceedings of the ACM Web Conference 2023*, 3648–3659.

[59] Nasim Sonboli, Robin Burke, Nicholas Mattei, Farzad Eskandanian, and Tian Gao. 2020. "And the winner is...": Dynamic lotteries for multi-group fairness-aware recommendation. arXiv:2009.02590 [cs.IR]

[60] Nasim Sonboli, Farzad Eskandanian, Robin Burke, Weiwen Liu, and Bamshad Mobasher. 2020. Opportunistic multi-aspect fairness through personalized re-ranking. In *Proceedings of the 28th ACM Conference on User Modeling, Adaptation and Personalization (Genoa, Italy) (UMAP '20)*. ACM Press, New York, NY, USA, 239–247. <https://doi.org/10.1145/3340631.3394846>

[61] Nasim Sonboli, Jessie J Smith, Florencia Cabral Berenfus, Robin Burke, and Casey Fiesler. 2021. Fairness and transparency in recommendation: The users' perspective. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21)*. ACM Press, New York, NY, USA, 274–279. <https://doi.org/10.1145/3450613.3456835>

[62] Harald Steck. 2018. Calibrated recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18)*. ACM Press, New York, NY, USA, 154–162. <https://doi.org/10.1145/3240323.3240372>

[63] Patricia H. Thornton, William Ocasio, and Michael Lounsbury. 2012. *The Institutional Logics Perspective: A New Approach to Culture, Structure, and Process*. Oxford University Press, Oxford, UK. <https://doi.org/10.1093/acprof:oso/9780199601936.001.0001>

[64] Patricia Victor, Martine De Cock, and Chris Cornelis. 2011. Trust and recommendations. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer, Boston, MA, USA, 645–675. https://doi.org/10.1007/978-0-387-85820-3_20

[65] Amy Voida, Lynn Dombrowski, Gillian R. Hayes, and Melissa Mazmanian. 2014. Shared values/conflicting logics: Working around e-government systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3583–3592. <https://doi.org/10.1145/2556288.2556971>

[66] Tianxin Wei, Fuli Feng, Jiawei Chen, Ziwei Wu, Jinfeng Yi, and Xiangnan He. 2021. Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD '21)*. ACM Press, New York, NY, USA, 1791–1800. <https://doi.org/10.1145/3447548.3467289>

[67] Emre Yalcin and Alper Bilge. 2022. Evaluating unfairness of popularity bias in recommender systems: A comprehensive user-centric analysis. *Information Processing & Management* 59, 6 (Nov. 2022), 103100. <https://doi.org/10.1016/j.ipm.2022.103100>

[68] Ziwei Zhu, Yun He, Xing Zhao, and James Caverlee. 2021. Popularity bias in dynamic recommendation. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD '21)*. ACM Press, New York, NY, USA, 2439–2449. <https://doi.org/10.1145/3447548.3467376>

[69] Ziwei Zhu, Xia Hu, and James Caverlee. 2018. Fairness-aware tensor-based recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*. ACM Press, New York, NY, USA, 1153–1162. <https://doi.org/10.1145/3269206.3271795>

[70] John Zoshak and Kristin Dew. 2021. Beyond Kant and Bentham: How ethical theories are being used in artificial moral agents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. ACM Press, New York, NY, USA, Article 590, 15 pages. <https://doi.org/10.1145/3411764.3445102>