

Perceptions of the Fairness Impacts of Multiplicity in Machine Learning

Anna P. Meyer*

apmeyer4@wisc.edu

University of Wisconsin - Madison
Madison, Wisconsin, USA

Aws Albarghouthi

aws@cs.wisc.edu

University of Wisconsin - Madison
Madison, Wisconsin, USA

Yea-Seul Kim

yeaseul.kim@gmail.com

Apple

Boulder, Colorado, USA

Loris D'Antoni

ldantoni@ucsd.edu

University of California - San Diego
San Diego, California, USA

Abstract

Machine learning (ML) is increasingly used in high-stakes settings, yet *multiplicity* – the existence of multiple good models – means that some predictions are essentially arbitrary. ML researchers and philosophers posit that multiplicity poses a fairness risk, but no studies have investigated whether stakeholders agree. In this work, we conduct a survey to see how multiplicity impacts lay stakeholders' – i.e., decision subjects' – perceptions of ML fairness, and which approaches to address multiplicity they prefer. We investigate how these perceptions are modulated by task characteristics (e.g., stakes and uncertainty). Survey respondents think that multiplicity threatens the fairness of model outcomes, but not the appropriateness of using the model, even though existing work suggests the opposite. Participants are strongly against resolving multiplicity by using a single model (effectively ignoring multiplicity) or by randomizing the outcomes. Our results indicate that model developers should be intentional about dealing with multiplicity in order to maintain fairness.

CCS Concepts

• **Human-centered computing** → **User studies**; • **Computing methodologies** → *Machine learning*; • **Applied computing** → *Psychology*.

Keywords

Fairness, Fairness Perceptions, Fairness in Machine Learning, Multiplicity, Stakeholder Survey

ACM Reference Format:

Anna P. Meyer, Yea-Seul Kim, Aws Albarghouthi, and Loris D'Antoni. 2025. Perceptions of the Fairness Impacts of Multiplicity in Machine Learning. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26-May 1, 2025, Yokohama, Japan. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3706598.3713524>

*Correspondence to apmeyer4@wisc.edu



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '25, Yokohama, Japan*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1394-1/25/04

<https://doi.org/10.1145/3706598.3713524>

1 Introduction

Machine learning (ML) is ubiquitous in high-stakes decision-making settings, e.g., social services [17], criminal justice [73], healthcare [70], and hiring [59]. Due to the life-altering impacts of these decisions (for instance, whether someone can get hired at a job), it is important that the model making the decision is as accurate, fair, and accountable as possible. The conventional machine learning wisdom is that there exists a single best-performing model and that achieving other desiderata like fairness or robustness requires an accuracy tradeoff [66]. But the literature – both recent [24, 64] and foundational [11] – challenges that paradigm: there are typically many models that achieve similar predictive accuracy by different means (i.e., by making mistakes on different parts of the input distribution). This ambiguity leads to *predictive multiplicity*, which is when two equally good models assign different predictions to a given sample, as illustrated by fig. 1. In socially-salient prediction settings, the set of (approximately) equally good models is often very large [64, 75, 100], and a large fraction of samples receive conflicting outcomes under this multiplicity [64].

Work in computer science and philosophy suggests that multiplicity poses a fairness risk [8, 21, 22]. Fairness has multiple components, including distributional fairness, which considers whether the *outcomes* are fair, and procedural fairness, which considers whether the decision was made *in a fair way*, regardless of outcomes [19]. The existing work points out that multiplicity can violate procedural fairness, because choosing one model over the alternatives results in *arbitrary* decisions [8, 21, 22]. This fairness violation is especially salient when models operate in high-stakes settings where the justifiability of a negative decision is important. For instance, suppose a company uses an ML model m_1 to determine which job applicants to interview, but that there is a different model m_2 that would do equally as well at selecting promising candidates (e.g., m_1 and m_2 have similar accuracy on the available validation data). A candidate x who is not interviewed according to m_1 , but who would have been interviewed if m_2 were used, may feel that procedural fairness is violated if the company did not even consider using m_2 . An additional concern is that when models are deployed at scale, multiplicity can result in systematic arbitrariness that locks people out of opportunities [22]. For instance, if all companies contract with the same software vendor and each uses model m_1 to determine who to interview, x may get no job

interviews even though the existence of m_2 suggests they may be qualified.

While the existing literature makes a strong argument for the fairness risks of multiplicity, no research has interrogated whether lay stakeholders believe that multiplicity impacts the fairness of machine learning models. Likewise, no research has asked how stakeholders think decisions should be made when a machine learning system contains multiplicity. More precisely, consider the same job application scenario where two models m_1 and m_2 have similar performance, yet disagree on whether to interview a job candidate x . The standard machine learning approach is to choose a single well-performing model, for instance, selecting m_1 without even realizing that m_2 exists. Philosophers suggest using randomization – e.g., flipping a coin to determine whether to interview x – as the fairest way to deal with multiplicity [22, 43]. Computer scientists suggest switching to more sophisticated model architectures to reduce multiplicity (e.g., ensemble methods as a way to reduce variance) [7, 21, 74]. Lay stakeholders will not necessarily like those approaches, as research shows that different types of stakeholders (like decision makers and decision subjects) can have different perceptions of what matters for fairness in ML [56, 82, 103]. Critical scholars argue that considering stakeholder preferences is important [10, 85] and this belief underlies frameworks like participatory ML [30] and value-sensitive design [93]. Furthermore, companies deploying ML may care about public opinions of fairness to maintain high levels of trust in their product [98].

We address the gaps in understanding stakeholders' views on multiplicity by conducting a survey of crowdworkers (i.e., laypeople who are decision subjects as ML is increasingly deployed in everyday settings). Our survey aims to determine (1) how the presence of multiplicity impacts the stakeholders' assessment of ML models' fairness and (2) which multiplicity resolution approaches they prefer. We consider the following research questions:

- **RQ1:** Does being alerted to multiplicity in a machine learning system affect participants' perceptions of fairness of the system?
- **RQ2:** What techniques to resolve multiplicity do participants view as most appropriate and fair?
- **RQ3:** How do task characteristics (stakes, uncertainty, and punishment or reward framing) impact **RQ1** and **RQ2**?

To address these questions, we want to evaluate fairness perceptions for a variety of tasks. However, some of the factors we care about for **RQ3** (namely, task stakes and uncertainty) may be subjective, so we did not want to rely solely on our own classification of tasks. Instead, we conducted a preliminary study with $n = 199$ participants, which, to the best of our knowledge, is the first study to get lay stakeholders' opinions on both tasks' perceived stakes and uncertainty (Castelo et al. [14] collect opinions on task uncertainty, but not stakes). Based on the results of this survey and our own assessment of task framing (reward vs. punishment), we selected 8 diverse tasks to use in the main study.

Our main study begins with an educational component describing model multiplicity. After the participant passes a comprehension check, we assign them to a scenario and describe a machine learning model at a high level before asking participants for an initial fairness assessment, consisting of a series of Likert scale

questions. Then, we describe how model multiplicity impacts the situation and ask for updated answers to the same Likert questions (**RQ1**). Finally, we use an open-ended question and choice-based analysis to learn what multiplicity resolution techniques participants view as the most fair (**RQ2**).

Our results indicate that stakeholders do not perceive multiplicity as an overall fairness risk, counter to what the literature suggests. However, this result is nuanced as participants are more likely to express that the model outcomes might be unfair, rather than that the model may be inappropriate to use at all. We also find that participants have clear preferences for how multiplicity is resolved, namely by deferring to a human expert for situations where multiplicity impacts the predictions (**RQ2**). These preferences vary depending on the specific task, with task stakes and task framing significantly impacting resolution preferences (**RQ3**). Despite philosophical arguments for randomization as a multiplicity solution [22, 43], this was the least preferred option in our experiment.

These results have two important implications: first, *the way that multiplicity is discussed in the existing literature does not align with the lay stakeholders' perceptions*. Instead, we hypothesize that other concerns about machine learning (such as recourse and explainability) may dominate participants' fairness assessments. Second, even though overall fairness perceptions do not shift significantly after learning about multiplicity, the strong sentiments that participants have regarding multiplicity resolution suggest that *ML developers must be deliberate in determining how models behave in the presence of multiplicity*. In particular, simply choosing a single well-performing model and not exploring the impacts of multiplicity is viewed as unfair.

2 Background of Fairness and Multiplicity

We overview fairness in the context of machine learning (section 2.1), define multiplicity and its associated fairness risks (section 2.2), and describe various ways to resolve multiplicity (section 2.3).

2.1 Fairness in machine learning

Colquitt and Rodell propose several components of fairness including distributional, procedural, and informational fairness [20]. Distributional fairness asks whether a model's outcomes are fair, irrespective of how the decisions were made. For instance, group fairness metrics such as equalized odds [28, 37] ask whether different demographic groups are treated equally in aggregate. Procedural fairness, by contrast, asks whether a decision was made in a fair way, regardless of what the final outcome is. Procedural fairness has many facets, including counterfactual fairness [52] (how a test sample needs to change to get a different prediction) and the choice of features that are included in a model. For example, in a task such as hiring, procedural fairness may be violated if race is a decisive factor (i.e., changing an applicant's race changes whether they are hired) or even if race is included as a model input at all. Informational fairness requires that details about the decision making process be available, and has been explored in ML through the use of explanations, e.g. [6, 27, 81, 102]. In this paper, we consider perceptions of procedural fairness.

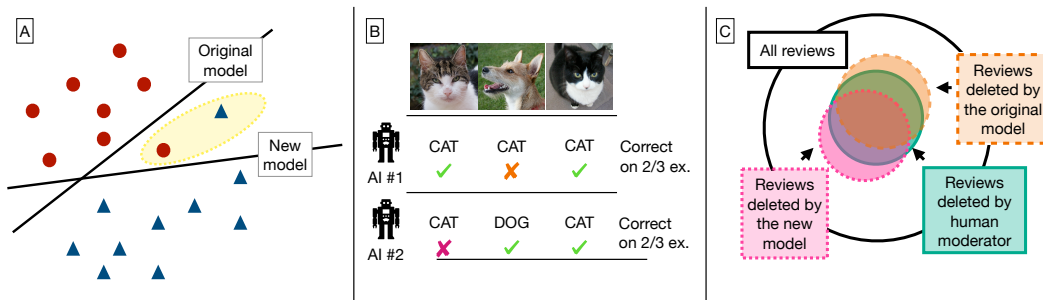


Figure 1: A: Example dataset with two classes (red circle and blue triangle) that, when limited to linear models, exhibits predictive multiplicity because no model achieves perfect accuracy while multiple models are correct on all but one prediction. The predictions in the yellow shaded region get a different prediction depending on whether we choose the “original” or “new” model. B: Example of model multiplicity based on model predictions. C: For the task of detecting fraudulent reviews, shows the overlap between reviews flagged as fraudulent by humans and two models. The two models both have significant overlap with human predictions, but less overlap with each other.

The fairness of a ML model can be measured based on its adherence to some fixed standard (e.g., a group fairness metric), or in a human-centered way through perceived fairness, as in this study. Perceived fairness is a common way to evaluate ML systems [91], perhaps because settling on a formal fairness definition is challenging. Research has established that perceived fairness of a model correlates with trust of that model [87, 88, 98], and suggests that low perceived fairness corresponds to negative consequences (e.g., reputational harm) for the entity deploying ML [91].

2.2 Multiplicity in machine learning and fairness concerns

Predictive multiplicity [64], which we simply call *multiplicity* for brevity, occurs when there are multiple ML models that are equally well-suited for a task, yet give different individual-level predictions. Formally, suppose we have a dataset D , a hypothesis class \mathcal{H} (e.g., all possible decision trees with depth ≤ 3) and some loss function \mathcal{L} (e.g., the 0-1 loss). A standard machine learning pipeline aims to find $h^* \in \mathcal{H}$ that minimizes $\mathcal{L}_h(D)$ (e.g., the total number of misclassifications on D). This solution is typically found via an approximate optimization process such as gradient descent, and yields a model h_1 that hopefully has close-to-optimal loss, i.e., has $\mathcal{L}_{h_1}(D) \approx \mathcal{L}_{h^*}(D)$. But sometimes, the optimization problem is underconstrained and there is another model h_2 such that $\mathcal{L}_{h_2}(D) \approx \mathcal{L}_{h_1}(D)$. (Intuitively, there may be multiple global minima.) If there is some input $\mathbf{x} \in D$ with $h_1(\mathbf{x}) \neq h_2(\mathbf{x})$ we say that the models exhibit (*predictive*) *multiplicity*.

Prevalence of multiplicity. Computing the fraction of inputs whose predictions are affected by multiplicity is generally intractable, so existing approaches quantify multiplicity by limiting their analysis to specific model types, approximating the impacts of multiplicity, and/or using brute force. Estimates on the COMPAS dataset, which is used to predict risk of recidivism in defendants, indicate that multiplicity affects 44% of test inputs when using linear models [64] and 69.4% of test inputs when using decision trees [21]. For probabilistic models that compute risk scores, model outputs vary by $> 10\%$ on 82-100% of samples across three tabular datasets [96]. These levels

of multiplicity indicate that many predictions are a function of the specific model that is chosen, rather than purely a function of the data.

Fairness risks of multiplicity. Fairness risks of multiplicity include harms to an individual, society or a group of individuals, and the entity who is deploying the model. Fairness risks to individuals include a lack of justifiability for a negative decision [8] and systemic loss of opportunity when models containing multiplicity are deployed at scale [22].

Example 2.1 (Lack of justifiability). Black et al [8] point out that US credit laws like the Equal Credit Opportunity Act aim to reduce arbitrariness and increase equity in credit decisions, such as by requiring lenders to provide justification for denied loan requests. However, if another multiple equally-accurate model approves a loan request, it may be more challenging to justify why the bank needed to use the model that denies the loan.

Example 2.2 (Systemic loss of opportunity). Creel and Hellman [22] discuss how individual employers may have idiosyncratic hiring criteria, but that the harm from this to job applicants is minimal, as they can apply to many jobs. Since each employer has their own potentially arbitrary criteria (in other words, multiplicity exists in the process of choosing who to hire), the job applicant will eventually be hired by someone. However, if every company uses the same hiring system, then certain job applicants will not be able to be hired anywhere, causing them significant harm. This setup can happen, for instance, if every company contracts with the same third-party software to screen resumes.

Society or groups of individuals may be adversely impacted by multiplicity if a sub-optimal model, with regards to fairness, is selected from the set of equivalently-accurate models. The entity using the model is also at risk, first for potential lawsuits or reputational harm if their decisions are not justifiable [8]. Second, if a single model is deployed at scale when multiplicity exists, this standardization can provide a worse set of outcomes.

Example 2.3 (Worse outcomes for the decision-maker). Kleinberg et al. [51] show that companies who all use the same recruiting

Table 1: The multiplicity resolution techniques that we consider. We explain to participants that the proposed method is only used if the two AI systems disagree. We describe the techniques in more detail tailored to the specific task (e.g., “Let a doctor decide” for human in a medical setting); full text is in the appendix.

Method name	When the two models disagree...
default-bad	Default to the undesirable outcome
default-good	Default to the desirable outcome
fancy	Train a more complex model and use its output
human	Let a human expert decide
ignore	Use the original model’s decision
random	Choose the outcome randomly

software may end up with lower-quality new hires than if each company used its own model, an outcome that is harmful for the company itself and not just potential job applicants.

However, no work that we are aware of has investigated the fairness risks of multiplicity *from the perspective of stakeholders*. Previous work shows that formal fairness definitions often do not line up with stakeholders’ fairness perceptions [90], so, it seems likely that experts and lay stakeholders may have distinct insights about whether multiplicity is (un)fair.

2.3 “Solving” multiplicity

A natural question, after finding evidence of multiplicity, is how to go about making decisions. In this paper, we consider the fairness of six different multiplicity-resolution techniques (see table 1). The standard approach in machine learning is to ignore¹ multiplicity; that is, to find a single well-performing model and use it, without considering whether multiplicity exists. Of work that focuses on multiplicity, three schools of thought exist on how to resolve it: reducing multiplicity through algorithmic interventions, standardizing the procedural response to multiplicity, and using multiplicity as an opportunity to optimize for second-order criteria like fairness or robustness.

Reducing multiplicity. Refining the modeling technique can reduce multiplicity, such as by using cascading classifiers [74], bootstrapping [21], or selective ensembles that abstain if the confidence is not high enough [7]. Similar ensembling approaches have long been used to improve accuracy in ML applications [76]. The difference when using ensembling to combat multiplicity is that multiplicity is often discussed in the context of social prediction tasks where an objective ground truth is elusive or even does not exist. Thus, using ensembling to get rid of multiplicity often does not yield significantly higher accuracy than using a non-ensembled model. It does, however, mask that multiplicity exists, making it harder to challenge the validity of an algorithm. We refer to ensembling and other approaches that aim to reduce multiplicity through

¹Here, and throughout the rest of the paper, we use this font to denote the names of the multiplicity-resolution techniques.

modeling choices as fancy, i.e., training a fancier (more complex) model.

Dealing with multiplicity procedurally. ML researchers have suggested abstaining, and, presumably, deferring to a human expert (denoted in this paper as human), when a decision is subject to multiplicity [7]. However, the design of “human-in-the-loop systems” is not straightforward [99], and may interact with other criteria we are trying to measure, e.g., algorithm aversion [32]. The idea of abstaining and deferring to a human decision-maker approach parallels a long line of work in robust machine learning that similarly suggests abstaining on decisions when robustness cannot be verified [3, 18, 39, 54].² However, these works rarely present a plan for what abstaining looks like in practice: often, abstaining on a decision results in not taking action, thereby defaulting to the negative label (for instance, abstaining on a loan application is effectively denying it). We consider abstaining and deferring to a default option in our survey. However, attitudes may vary depending on whether the outcome is desirable or not, so we use default-good and default-bad to represent the outcomes that are good and bad, respectively, from the perspective of the decision subject.

An alternate procedural approach is random, i.e., randomizing over the set of equally-good models [22, 43]. ML in social prediction tasks often serves to distribute a limited resource, e.g., admissions to a college with a fixed enrollment limit. Lotteries are typically viewed as fair system for dividing limited resources [12, 86]. However, there are technical challenges in implementing randomization effectively [43].

Optimizing for second-order criteria. Traditional ML practice assumes that there is a trade-off between model accuracy and fairness, as defined by group-level metrics such as equalized odds or statistical parity [46, 66]. Recent research to explore the whole set of equally-good models challenges this paradigm, showing that accuracy and other criteria are not in tension [75, 100]. But even if we can use multiplicity to find fairer (or more robust) models, there are still multiple mutually incompatible fairness (or robustness) definitions. We do not include optimizing for fairness as one of the possible multiplicity resolution techniques because the task of choosing the fairest model is typically not well-defined [16, 50].

3 Related Work and Hypotheses

Fairness perceptions of machine learning algorithms have been studied extensively, as evidenced by multiple survey papers on the topic [67, 91]. Here, we will focus on dimensions most related to our research questions.

Research in this area investigates distributional fairness (e.g., matching lay stakeholder’s fairness notions to formal definitions [80, 90]), procedural fairness (e.g., considering what features to include [34]), and informational fairness (e.g., the impact of explanations [81, 102]). In this work, we focus on perceptions of procedural fairness, presenting multiplicity as a hurdle that must be dealt with

²Note that non-robustness constitutes a form of multiplicity, as it implies there is another reasonable model that outputs an alternate prediction for the test sample. Within the ML literature, robustness is typically discussed for deep networks and/or text and image data, while multiplicity research largely focuses on social prediction tasks using tabular data.

procedurally. The fairness perceptions of ML models are often considered relative to a baseline of human decision making, and can either increase from algorithmic appreciation [60], or decrease from algorithmic aversion [26]. Existing work establishes that fairness perceptions of ML models vary based on task characteristics, including stakes [36, 79, 102], uncertainty [78], complexity [77], and framing (reward vs. punishment) [36].

In the following subsections, we describe our hypotheses about how multiplicity impacts perceived fairness (H_{1a} - H_{1d}) and which multiplicity resolution techniques stakeholders prefer (H_{2a} - H_{2d}). See table 3 for a summary of all hypotheses.

3.1 Algorithm aversion

Algorithm aversion is when people prefer human decision-makers, even when the available AI system generally outperforms humans [26]. It may be triggered when people see a model make a mistake, even if they would have been forgiving of a human making the same mistake [26], or when the model's performance is worse than stated [101]. In settings where a correct answer can take many forms (e.g., generative AI), seeing multiple conflicting output lowers trust in a model [57]. Algorithm aversion also tends to be higher for subjective tasks, a characterization that is malleable depending on how the task is described [14]. We expect that witnessing model multiplicity will increase algorithm aversion for three reasons: first, the existence of two good (but different) models means that the models make mistakes, similar to what was observed for conflicting generative AI responses [57]. Second, the existence of multiple good models indicates that the decisions these models make are subjective. In particular, multiplicity means that choosing a model is not a purely objective task (i.e., optimizing accuracy), but rather a subjective choice of choosing between multiple models with near-equivalent accuracy. Finally, existing work assumes (without consulting stakeholders) that multiplicity threatens fairness (section 2.2).

H_{1a} : Stakeholders will view machine learning systems as less fair once they are made aware of multiplicity.

Between algorithm aversion and prior work showing preference for human over AI arbitrators [84], we think participants will prefer human over other resolution techniques (recall that human and the other approaches are defined in table 1). We think that – counter to what philosophical work suggests – random will be viewed as the least fair approach. Explanations are important to fairness [81, 102], so we think the sense of unpredictability from randomized decisions dominate stakeholders' reactions to random.

H_{2a} : Stakeholders will have significant preferences between different multiplicity solutions. In particular, they will view human as the fairest multiplicity solution and random as the least fair solution.

3.2 Task stakes

Competing with algorithm aversion is *algorithm appreciation*, when people over-rely on algorithmic decisions [60]. There is evidence that algorithm appreciation occurs in high-stakes tasks [1, 79], and stakeholders rate accuracy as more important in high-stakes settings [90]. Thus, we think that fancy (i.e., building a more complex model) may be preferred for high-stakes tasks because it may be

perceived as being more accurate. However, the importance of accuracy may also trigger algorithm aversion and increase preference for human. Fairness is more important to stakeholders in high-stakes settings than low-stake ones [44]. We suspect that higher-stakes tasks are more emotionally charged. People view emotional decision making as a strength of humans [38], and can have negative emotional responses to AI decision-making [55].

H_{1b} : The presence of multiplicity will have a larger impact on fairness perceptions in high-stakes settings.

H_{2b} : Participants' preferred multiplicity resolution technique will be moderated by task stakes. In particular, high-stakes tasks will correspond to stronger preferences for human and fancy, and low-stakes tasks will correspond to stronger preferences for ignore and random.

3.3 Task framing

Harm that is caused by action, rather than omission, causes harsher fairness judgments [23]. Tasks that are punitive commit active harm in the event of a mistake, whereas reward-based tasks cause harm when they make an omission. Thus, we think that punitive tasks will have greater decreases in fairness perceptions in the presence of multiplicity. Indeed, others have found that when people adopt critical mindsets, they are more likely to find issues with a model [36]. Algorithms are seen as more rational than humans [60], and we suspect that rationality will be more important in punitive settings where mistakes are judged more harshly. Conversely, there are positive connotations towards human accuracy [55], so being primed to think about reward settings may influence participants to select human.

H_{1c} : The presence of multiplicity will have a larger impact on fairness perceptions in punitive settings.

H_{2c} : Preferences about how multiplicity is resolved will be moderated by task framing. In particular, for punitive tasks, fancy will be more preferred and for reward-based tasks, human and default-good will be more preferred.

3.4 Task uncertainty

High-uncertainty tasks typically do not have a single objectively-right answer, such as predicting future human behavior (e.g., job performance) or matters of taste (e.g., movie recommendation). Low-uncertainty tasks typically have an objectively-correct answer (e.g., detecting spam emails). Algorithm appreciation is known to be stronger in objective (i.e., low-uncertainty) tasks [2, 60], aligning with beliefs that algorithms are not good at subjective (i.e., high-uncertainty) tasks [14]. There is more algorithm aversion and preference for human decision makers in high-uncertainty settings [25]. When a task has lower uncertainty, it may be possible to have overall higher accuracy, which increases algorithm appreciation [79]. Intuitively, people may find it easier to rationalize that there are multiple good solutions to a high-uncertainty task, resulting in lower impact to the perceived fairness.

H_{1d} : The presence of multiplicity will have a larger impact on fairness perceptions for low-uncertainty tasks.

H_{2d} : Preferences of how to resolve multiplicity will be moderated by task uncertainty. In particular, for low-uncertainty tasks

stakeholders will see fancy as more fair, while for high-uncertainty tasks, they will see random as more fair.

4 Method

Figure 2 shows an overview of our study design, which was approved by our institution’s IRB board. In this section, we first describe our preliminary study (section 4.1), which informed the task selection for our main study. Then, we describe our pre-registered study³ design (section 4.2) and data analysis plan (section 4.4).

4.1 Task selection (Preliminary study)

Our main study asks participants about tasks that differ on three between-subjects axes: stakes, uncertainty, and framing. Framing (i.e., whether a positive model output is a reward or a punishment) is objective; however, individuals may differ on how high-stakes or uncertain they think a task is. We conduct a preliminary study ($n = 199$) to solicit perceptions of task stakes and uncertainty. Our goal is to select one task for each of the eight categories [punishment, reward] \times [low uncertainty, high uncertainty] \times [low stakes, high stakes].

Study design. We collect 18 scenarios through brainstorming and borrowing from previous studies [6, 14, 47]. We randomly assign each participant to see 4 scenarios to limit participant fatigue when completing the study. For each scenario, participants answer a series of 5-point Likert-scale questions relating to task clarity (2 questions), task stakes (3 questions), and task uncertainty (2 questions). Then, participants rank the four tasks according to stakes and uncertainty levels. We use Prolific to find respondents (US-based adults with a 97%+ approval rate). We compensate participants at Prolific’s recommended rate (\$12/hour) based on the estimated survey completion time.

Data analysis. Our survey yields four pieces of data: stakes rating, uncertainty rating, stakes ranking, and uncertainty ranking. We convert all Likert ratings to a scale from -2 (strongly disagree) to 2 (strongly agree). For the rating data, we average the three questions relating to task stakes into a single stakes rating quantity. We elect to use one of the task uncertainty questions alone, due to better data quality. To recover the underlying stakes and uncertainty rankings given the partial ranking information from each participant, we use assume a Plackett-Luce distribution [62, 71] and use the I-LSR (iterative Luce Spectral Rating) algorithm [65]. The basic assumption behind Plackett-Luce is that choice rankings are transitive (i.e., if $A > B$ and $B > C$ then $A > C$), and the I-LSR algorithm uncovers the full underlying ranking. We combine the rating and ranking data into a single metric by scaling each set of data to have mean 0 and variance 1, and then taking their arithmetic mean.

Preliminary study results. We obtain responses from 201 participants and, after excluding responses that failed more than one attention check, use 199 responses in the analysis.

Task stakes and task uncertainty are not meaningfully correlated ($r=0.19$, $p=0.42$). We exclude three tasks that scored poorly on clarity metrics, leaving 15 tasks to choose from. There are trade-offs in terms of achieving either a wide range of scores on each axis

or relative parity between tasks in different (stakes \times uncertainty) quadrants. That is, we can either select tasks to maximize the range of stakes and uncertainty scores, or we can select tasks with a narrower range of scores but with a better one-to-one comparison between tasks with different framing. We decide to maximize the range stakes and uncertainty scores since our analysis plan involves using numeric stakes and uncertainty labels rather than solely binary “low” and “high” categories. The tasks we select are described in table 2.

4.2 Study design

The survey flow is depicted in Figure 2. The survey begins with education about multiplicity (with a comprehension question) and then has two main sections (“Part 1” and “Part 2”). Participants are randomly assigned to one task at the beginning of Part 1. Then we describe a scenario and solicit the participant’s fairness perceptions of the model on a Likert scale consisting of four seven-point questions. In Part 2, we present a series of binary choices between different ways to resolve multiplicity and the user selects their preferred option. We conclude by collecting demographic information from the participant. We provide further details on all parts of the survey below.

Details of multiplicity education. We describe to participants, at a high level, how AI models are built (e.g., “AI developers [...] choose the model with the best performance”). Then, we reveal that there are sometimes multiple models that can perform equally well, but “make different mistakes”. We illustrate model multiplicity using a graphic similar to fig. 1B. We tell participants that this phenomenon, model multiplicity, is common. Then, participants have two chances to correctly answer the multiple-choice question “What does model multiplicity mean in the context of AI?” The correct answer is written using similar phrasing as in the descriptive text. We include full text of the multiplicity education section in the appendix.

Details of Part 1. In Part 1, we aim to address **RQ1** and **RQ3** by seeing how participants’ fairness perceptions of ML systems change when multiplicity is present. We first describe the scenario including what decision is being made, what factors the AI has access to, and how decisions will be operationalized (e.g., whether users will have recourse). We also say that the AI is validated and “agrees with human experts most of the time”. This description of accuracy is vague to avoid anchoring participants to specific numbers and because reporting an “accuracy” suggests a ground truth that does not exist for all scenarios we consider.

Next, we ask users to respond to four seven-point Likert scale questions to elicit fairness perceptions. Two questions focus on the model outcomes, e.g., whether the proposed model is likely to be able to make “good” decisions, and two questions focus on whether using the model is appropriate. Then, we describe how multiplicity impacts the situation, explaining that there is a second model that agrees with human experts equally often as the first model, but disagrees less often with the original model. We accompany the description with a graphic such as the one shown in Figure 1C. At this point, participants answer two critical-thinking questions:

³<https://osf.io/fv3bt>

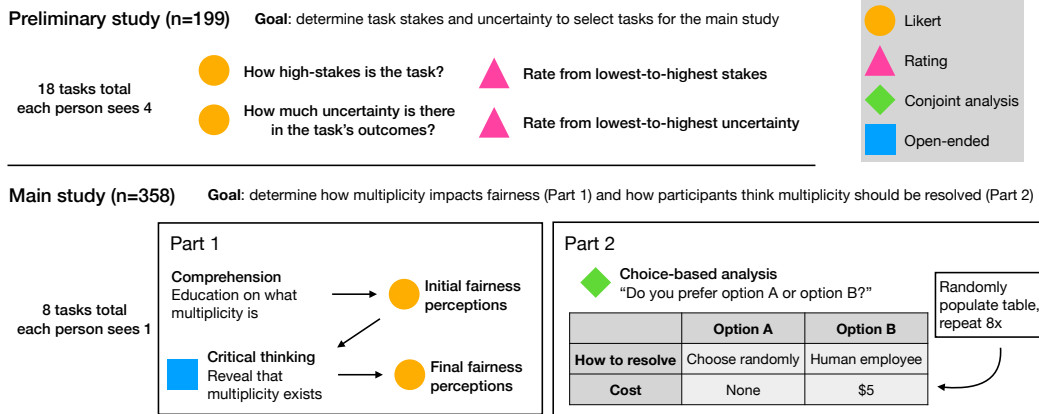


Figure 2: Overview of our survey design, including the preliminary study (top row) and main study (bottom row).

Table 2: Summary of tasks and their stakes, uncertainty, and framing. Starred columns (*) are based on the results of our preliminary study (section 4.1). (The raw stakes and uncertainty scores are binarized based on whether they are above or below the medians across all 18 tasks from that study.) A full description of tasks and their associated stakes and uncertainty numeric values are available in the appendix.

Domain	Task	Stakes*	Uncertainty*	Framing
Healthcare	Access to a preventative health care plan	High	High	Reward
Business	Job probation based on employee monitoring	High	High	Punishment
Healthcare	Tumor detection from medical imaging	High	Low	Reward
Government	Tax audit	High	Low	Punishment
Travel	Airline seat upgrade	Low	High	Reward
Travel	Rebooking passengers on an over-full flight	Low	High	Punishment
Business	Customers get beta access to an app with discounts	Low	Low	Reward
Business	Removing online reviews that are spam or fake	Low	Low	Punishment

one is multiple-choice and comprehension based; the other is open-ended and asks how decisions should be made when the two models disagree.

The next page of the survey asks the participant to respond again to the same set of Likert questions. We refer to the first set of responses as the “pre” timestamp, and this set of responses as the “post” timestamp, because the latter are collected after the participant learns that multiplicity exists for the task. The participant does not have access to their previous responses at the “post” timestamp, because we do not want them to be anchored by their original responses (i.e., we want to allow for a change in their frame of reference [42]).

Details of Part 2. In part 2, we investigate RQ2 and RQ3 by seeing which multiplicity resolution techniques participants prefer. We provide a description of 5-6 ways (depending on the task) that multiplicity can be resolved, i.e., how to make decisions when the two models disagree (see table 1). For some tasks, default-good is not realistic due to logical constraints, e.g., airline seat upgrades likely cannot be given to every passenger selected by either model because there is a firm limit in the number of available first-class seats. In cases like this, we exclude that option. We also explain to

participants that some options may be more expensive to implement, so might require a trade-off such as a monetary cost. The cost that participants see for each option is either “no extra cost” or a small scenario-specific cost that is meant to be an annoyance, but not a barrier for most people (e.g., airline tickets will cost \$5 extra).

The participant then sees 8 randomly constructed binary choices between two randomly chosen multiplicity resolution techniques and associated costs, such as the example in fig. 2.

Demographic information. We collected demographic information from participants including gender, age, education level, and experience with AI. Experience with AI was collected through the question “Do you interact with artificial intelligence or automated decision-making systems as part of your work or school?”, with the choices “yes,” “no,” and “not currently, but in the past”.

4.3 Data collection

We collect data using Prolific. Our participation criteria is US-based adults (18+) who have at least a 95% approval rate. We plan to collect data from 360 participants, i.e., 45 per task. We estimated that 35-40 participants per task would be sufficient based on trends in pilot data; we rounded up to 45/task to have additional buffer in case we need to reject some samples. We follow Prolific’s guidance [72]

and exclude survey respondents who fail the initial comprehension check question twice, fail two or more attention check questions, or are statistical outliers in survey completion time (in the fast direction).

4.4 Statistical analyses

4.4.1 Part 1 analysis.

Data transformation. We recode the Likert scale data from -3 (strongly disagree) to 3 (strongly agree) and compute the average across the four Likert scale questions.

Testing H_{1a} - H_{1d} . To test H_{1a} - H_{1d} , we fit a linear mixed model (LMM), which generalizes a linear model by accounting for both between- and within-subject factors. The formula we fit to the LMM predicts the transformed Likert data as a function of time, task stakes, task framing, task uncertainty, respondent ID (to account for individual differences), and the first-order interaction terms of time with each task characteristic. We use the `lme4` package [4] in R to construct the LMM model and then use the `car` package [33] to conduct a type-III ANOVA. Using ANOVA procedures on Likert-scale data is controversial because Likert scales are not continuous and may not produce normally distributed data. Scholars have argued both for [9, 13, 69] and against [45, 53] applying ANOVA when its assumptions are violated. We decided that LMMs fit our needs best, so to minimize concerns about using ANOVA, we included a large population in each group (as required by the central limit theorem for normality of the data sample) and used the average of several Likert items rather than a single item [13]. The main alternative technique we considered is Aligned-Rank-Transform (ART) ANOVA [97], a procedure that transforms ordinal (e.g., Likert-scale) data to fit the ANOVA assumptions. Ultimately we chose LMMs because ART-ANOVA does not allow for continuous factors. However, for readers who are interested, we include two additional analyses in the appendix: ART-ANOVA and a non-parametric test, the Mann-Whitney U-Test [63].⁴

To test H_{1a} , we check whether the ANOVA p-value is significant, and if so, run a post-hoc contrast test to gauge directionality. To test H_{1b} and H_{1d} , we find the estimated marginal trends for how task stakes (H_{1b}) and uncertainty (H_{1d}) vary with time using `emtrends` in the `emmeans` package in R [83]. Then, we conduct a pairwise contrast test on the result. To test H_{1c} , we perform a pairwise contrast test on the estimated marginal means of the interaction term of time and task framing, again using `emmeans`.

Multiple hypothesis correction. In total, since we perform four post-hoc contrast tests on the Likert scale data, we will use the Holm-Bonferroni correction [41] with $\alpha = 0.05$ and $m = 4$ to evaluate the hypotheses.

4.4.2 Part 2 analysis.

Conjoint analysis. We test H_{2a} - H_{2d} using choice-based conjoint analysis [35]. Choice-based analysis is a technique that solicits user preferences over two options with various factors set to different attribute levels. The analysis generates the marginal means of each attribute level. Our analysis uses the `cregg` package in R [58]. For

⁴To clarify, we preregistered the analysis with LMMs. Readers should view the results of the other tests as a partial *multiverse analysis* [92].

Table 3: Summary of the hypotheses described in section 3. The final column indicates whether we found statistically significant support for the hypothesis (section 5.2)

	Description	Sig?
H_{1a}	Multiplicity reduces fairness perceptions	No
H_{1b}	H_{1a} is modulated by task stakes	No
H_{1c}	H_{1a} is modulated by task framing	No
H_{1d}	H_{1a} is modulated by task uncertainty	No
H_{2a}	Participants care how multiplicity is resolved	Yes
H_{2b}	H_{2a} is modulated by task stakes	Yes
H_{2c}	H_{2a} is modulated by task framing	Yes
H_{2d}	H_{2a} is modulated by task uncertainty	No

H_{2a} , we look directly at the marginal mean effect of each of the multiplicity resolution techniques and costs. For each condition, the marginal mean procedure outputs a point estimate and a confidence interval of how preferred an option is relative to random chance. For instance, an output of 0.8 indicates that participants presented with that option chose it 80% of the time. For H_{2b} , H_{2c} , and H_{2d} , we compute the marginal means stratified by task characteristics. To the best of our knowledge, there is no tool to compute marginal means for a continuous factor, so we use the binary encodings (low vs. high) for task stakes and task uncertainty.

Testing H_{2a} . We will consider H_{2a} to be weakly supported if the confidence interval for human is strictly greater than 0.5 and the confidence interval for random is strictly less than 0.5. Then, we will test whether the estimate for human is strictly higher than the estimates for the other methods. To do this, we will perform five two-sided t-tests⁵ (one against each other multiplicity resolution technique). Then, we will perform four t-tests (random vs. each other multiplicity resolution technique other than human) to see whether random is strictly less preferred than each other technique.

Testing H_{2b} - H_{2d} . We compute the marginal means model for each task characteristic. We then perform post-hoc t-tests on the estimates. For instance, to test H_{2b} , we first check whether the estimate for human for low-stake tasks is different from its estimate for high-stake tasks. In total, we perform four post-hoc tests for H_{2b} , three for H_{2c} , and two for H_{2d} .

Multiple hypothesis correction. Since there are six multiplicity resolution techniques, we use the Bonferroni⁶ correction to get a p-value of 0.0083 and use this to compute the confidence intervals for each marginal means estimate. We perform a total of 18 post-hoc t-tests on the data, so will evaluate their significance using the Holm-Bonferroni correction with $\alpha = 0.05$ and $m = 18$.

5 Data Analysis

5.1 Summary of responses

We receive 367 survey responses and after excluding ones that fail the initial comprehension check twice or fail multiple attention checks, we use 358 in our analysis. 58.7% of respondents are women, 2.0% are non-binary, and 39.4% are men. The average respondent age is 36.3 (range: 18-78, stdev: 12.3). For the “pre” timestamp, Cronbach’s Alpha between the four Likert items is 0.88 and for the “post” timestamp it is 0.94.

5.2 Evaluation of hypotheses

We find no evidence for H_{1a} ($\chi^2(1) = 0.146, p = 0.703$), H_{1b} ($\chi^2(1) = 1.262, p = 0.261$), H_{1c} ($\chi^2(1) = 0.308, p = 0.579$), or H_{1d} ($\chi^2(1) = 1.405, p = 0.236$).

H_{2a} is significant since different multiplicity resolution techniques are selected at statistically significantly different rates (see fig. 3). For instance, human is selected in a pairwise comparison at a rate of 0.795 (SE = 0.015, $p < 0.0001$) and random is selected at a rate of 0.279 (SE = 0.017, $p < 0.0001$). A post-hoc test (detailed results in appendix) reveals that the preference for human is statistically stronger than all other options, while random is significantly less preferred than all options other than default t-bad.

H_{2b} is significant as well, and one of our four specific predictions holds: human is more strongly preferred when a task is high stakes. H_{2c} is significant, as is one of our four predictions: when the task is reward-based, default t-good is more strongly preferred than when the task is punishment-based. H_{2d} is not significant. Marginal mean estimates and details on t-test results are in the appendix.

5.3 Exploratory analyses

We describe five exploratory results (labeled E1-E5).

E1 (Differences between different dimensions of fairness). While the set of 4 Likert items has high internal agreement for both the “pre” and “post” timestamps, the changes that occur with time are not distributed randomly across the four items. Overall, there is not a significant difference between “pre” and “post” fairness scores (H_{1a}). But the first two Likert items, which ask about model outcomes (whether a system will be able to make good/fair decisions), see the scores decrease by an average of 0.122 points per item, while the second two Likert items, which ask about the appropriateness of using the AI system, increase with time by an average of 0.131 points per item. This difference is statistically significant according to the Mann-Whitney U-test ($W=74118, p=0.0001$).

E2 (Survey duration’s impact on change in fairness scores). We compute Spearman’s rank correlation between survey duration and average change in fairness score. There a very weak positive correlation ($\rho = 0.137, p=0.0002$).

E3 (Trends in multiplicity solution preferences). We observe several trends that our hypotheses did not predict. (All results in the




Figure 3: Marginal means for each multiplicity resolution technique. A score of 0.5 indicates the frequency the option would have been chosen if the selection is purely random, while a score greater (less) than 0.5 indicates an option is chosen more (less) frequently than would be expected by random chance.

following paragraph are for t-tests based on marginal means estimates.) When stratifying tasks by stakes, default t-bad is more preferred in low-stakes settings than in high-stakes settings (difference in marginal means -0.166, SE=0.0386, $t=-4.307, p<10^{-4}$). We hypothesized that random would be more preferred in low-stakes solutions (H_{2b}); this was not significant with the multiple hypothesis correction but shows a trend in the right direction (difference -0.079, SE=0.033, $t=-2.411, p=0.016$). The default t-bad option is also more preferred in punishment settings than in reward settings (difference -0.183, SE=0.039, $t=-4.692, p<10^{-5}$), while random is more preferred in reward settings (difference 0.108, SE=0.033, $t=3.31, p=0.0009$). In high-uncertainty settings, default t-good and default t-bad are both more preferred than in low-uncertainty settings (difference 0.108, SE=0.044, $t=2.441, p=0.015$ and difference 0.089, SE=0.0387, $t=2.310, p=0.021$, respectively).

E4 (Role of cost). A marginal-means analysis shows that no extra cost is usually significantly preferred over extra cost (t-test result: difference 0.225, SE=0.016, $t=13.91, p<10^{-5}$). However, in three settings (tumor detection, care plan access, and job probation), the marginal means for no extra cost and extra cost are not statistically different (see the appendix for complete data).

E5 (Qualitative analysis of open-ended question). We analyze the responses to the open-ended question as follows: First, we exclude responses that are unclear or do not suggest a specific decision-making strategy. Then, we categorize responses according to the resolution techniques from table 1 and an “other” category. We then review the responses in the “other” category and group similar responses into new categories. If a response mentions multiple approaches, we count it in each category. Note that at this stage

⁵The choice-based analysis data may not be normally distributed; however, we have over 800 observations for each resolution technique so the central limit theorem makes the use of this test acceptable.

⁶Note, we do not use Holm-Bonferroni here because it is not supported by cregg.

of the survey, participants had not seen the options in table 1. We report general trends here; full statistics are available in the appendix.

The majority of participants across all tasks mention human, including over 80% of participants for job probation, care plan access, and tumor detection. `default-good` is mentioned most frequently for tumor detection, tax audit, and beta app access, while `default-bad` is suggested for beta app access and flight upgrade. `random` is suggested frequently for flight upgrade and occasionally for tax audit and flight rebooking. Participants also frequently suggested specific rules they would implement or indicated that a more complex model could be built to resolve disagreements. This type of answer was especially common for flight upgrades and rebooking, review moderation, and beta app access. Finally, a few participants suggested involving the decision subject directly, e.g., asking a flight passenger how they feel about being rebooked or upgraded. This option was mentioned by at least one respondent for flight rebooking, flight upgrade, and care plan access.

6 Discussion

6.1 The standard ML approach is unfair (ignore is disliked)

Our most impactful finding is that `ignore` is generally not preferred.

Implication: Standard ML workflows need to adapt to consider multiplicity, such as by increasing transparency when multiplicity impacts a decision.

Depending on the target audience, interventions to increase transparency could take different forms. For instance, common machine-learning packages could alert developers when multiplicity exists (e.g., by automatically performing a bootstrapping analysis in the background). Model documentation should include an analysis of what proportion of test samples are subject to receive a conflicting prediction under multiplicity, so that anyone who decides to use or adapt the model in the future is aware. In public communications about a model, the fraction of inputs (based on validation data) that are subject to receive a different prediction under multiplicity could be reported.

Implication: Model developers should be able to justify how multiplicity is handled, whether this is by deferring to a human decision-maker, employing randomization, using ensemble methods, or something else.

It is possible that decision subjects would be okay with using a single model, and knowingly ignoring multiplicity, if this decision is justified. But, just as model developers need to justify decisions such as feature selection and statistical fairness trade-offs, they should justify why the existence of multiplicity does not pose an existential threat to their project. In particular, they should consider potential harms to individuals, e.g., whether the model's use will be widespread and lead to systemic exclusion of some people.

6.2 Is multiplicity an inherent fairness risk?

Our results show that participants' fairness perceptions did not decrease when multiplicity was present (i.e., H_{1a} was not supported). This contradicts work in ML and philosophy that assumes and argues that multiplicity poses a fairness risk (see section 2.2). Since H_{1a} was not supported by our data, either (1) the experts are incorrect about the fairness risks of multiplicity, (2) lay stakeholders do not see fairness risks of multiplicity due to low technical literacy, or (3) our study design was insufficient to capture stakeholders' true views on multiplicity. We discuss (3) when describing our work's limitations (section 6.6), so we will overview (1) and (2) here.

If (1), the fairness risks of multiplicity summarized in section 2.2 must paint an incomplete picture. Perhaps the philosophical view that only systematic arbitrariness is a fairness risk [22] is correct. Alternatively, maybe multiplicity poses a *theoretical* fairness risk, but one that is either too small to matter in practice, or that is overpowered by the benefits of using ML. Participants may feel as if the ML system is a closed box, and thus that its decisions are impenetrable. This perception would, understandably, lead participants to feel disenfranchised in their interactions with the ML model. From this perspective, multiplicity is not the primary reason why a model's usage could be unfair.

Implication: Theoretical researchers should articulate how the fairness risks of multiplicity relate to other fairness issues in ML – such as accountability and recourse – to help clarify when, exactly, multiplicity poses a fairness risk. Future research about the perceived fairness of multiplicity should do more to access the “why:” what are the driving forces in fairness perceptions?

If (2), then H_{1a} was not supported because the stakeholders' technical literacy was not high enough to realize that multiplicity poses the procedural fairness risks suggested by computer scientists and philosophers. As ML technology powers more and more decisions that used to be made by humans, it is important that stakeholders are able to recognize its limitations to avoid inappropriate reliance on and trust in ML systems [61]. One way that reliance and trust can be calibrated is through improving technical/AI literacy [15]. We discuss how our survey design could be modified to promote better understanding in section 6.6. But also, technical literacy can be built gradually over time by exposure to news stories and public discourse [68].

Implication: We encourage people working at the intersection of technology and communication to discuss multiplicity and to take a critical lens, more broadly, when describing new ML systems in socially-impactful domains.

6.3 Randomization is disliked

Scholars in ML and philosophy typically suggest using a more complex model (e.g., via ensembles) [21, 74] or randomization [22, 43]. Participants in the survey expressed comfort with `fancy` – choosing this option more often than would be expected by random chance – but strongly disliked `random` (H_{2a}). This contradicts the established philosophical view that lotteries are fair in situations where a scarce resource must be allocated [12, 86]. A possible discrepancy between

our setup and a run-of-the-mill lottery is the inclusion of machine learning in the decision-making process. Perhaps randomization via a lottery is familiar, but the addition of machine learning suggests that there must be some way to “learn” the best allocation, thereby making randomization too crude of a tool.

Implication: More research should be conducted into why random is disliked, even though lotteries are generally viewed as fair. Researchers can investigate whether providing additional education about machine learning (e.g., how models can make mistakes and that there is often not a single, optimal solution) improves comfort with random.

6.4 Participants prefer human intervention, but is that scalable?

Participants generally prefer human over other multiplicity-resolution techniques (H_{2a}). However, when large proportions of predictions can vary under multiplicity, deferring to a human decision-maker is difficult to operationalize and negates cost and time savings associated with using ML. In these cases, one option is to simply forgo using ML at all to avoid the appearance of automation and the expense of developing a system that only handles a small number of cases. (Indeed, it may not be scientifically valid to use machine learning when many predictions vary under multiplicity, as has been argued for “predictive optimization” tasks [94].) Instead, decision-makers could rely on an interpretable, publicly-available set of guidelines. Alternatively, model developers could use a strategy that combines fancy and human. Recall that fancy (e.g., using ensemble methods) is favored by ML researchers as a technique to reduce multiplicity. Model developers could use ensembles or more complex model architectures to reduce multiplicity, and then defer to a human decision maker only for the smaller number of cases whose predictions remain ambiguous. In our experiments, participants were generally comfortable with fancy, so while further research is needed to confirm that a combination fancy and human approach would be viewed as fair, we are optimistic about that possibility.

Implication: Model developers may not be able to align with stakeholders’ preferences for human in all cases, but should consider whether other, more interpretable decision-making rules are more appropriate when multiplicity is present. Alternatively, using ensemble methods to reduce multiplicity, and then deferring to a human on the smaller number of ambiguous cases, could strike a good balance.

6.5 Impact of task choice

Preferences for multiplicity resolution techniques varied across tasks. E3 shows that random is more preferred in general for low-stakes scenarios, but also for one high-stakes scenario (tax audit). We hypothesize that this is because decisions about whom to audit are seen as more impersonal than the other high-stakes scenarios. In medical and employment domains, decision subjects typically have a personal relationship with their doctor or manager, and thus trust that they will have the ability to appeal to human emotion

and get an explanation of the decision. By contrast, decisions by the IRS (the US tax agency) about whom to audit likely already feel random to many individuals. Conversely, human is more strongly preferred in high-stakes settings where a personal connection is the norm (H_{2b} & E3). We also want to call attention to how ignore is usually preferred to random (H_{2a}). From the perspective of a decision subject who interacts with a model a single time, there is not a mathematical difference between ignore and random because the personal outcome is unknown and (relatively) unintelligent in each. (By contrast, default-good and default-bad give a known outcome, and people may assume that they can appeal to a human decision maker to get the correct/favorable outcome with human and that fancy will be smart enough to give the correct decision.) So, individuals must either be imagining repeated encounters with a system or just view having a (presumably) deterministic model as providing benefits such as accountability and predictability.

Implication: Model developers should involve participants in the design process to figure out how they prefer that multiplicity be handled for the specific task.

6.6 Limitations and opportunities for future work

We conclude our discussion with an overview of limitations and how they could be addressed in future work.

Cognitive load for participants. A limitation of this work is that multiplicity is not an easy concept to grasp, especially for people without machine learning experience. We attempted to explain it clearly and in simple terms and iterated on the survey text through multiple pilot iterations. Our results do not show significant differences in the responses for participants with or without ML experience. We attempted to check that participants understood multiplicity through an initial comprehension check, but some participants may have gotten lucky and passed this question by chance, only looked for the correct answer rather than engaging critically with the reading, or may have forgotten details when answering the remaining survey questions.

Unknown role of other task-related factors. To counteract the cognitive load required to reason about multiplicity, we tried to make the other details about the tasks and models as simple as possible, including by providing only high-level details of the AI system’s functionality and context. Although we tried to select tasks that are broadly relevant, some participants may not have seen the relevance to their own lives or imagined themselves as the decision subjects. Some participants may have felt like they did not have enough information to evaluate the system’s fairness, and the presence of multiplicity may have seemed insignificant relative to other unknown factors. In particular, we did not include any explanations of how the models made their decisions to keep the setup simpler. Informational transparency is a component of overall fairness [19], though, so future work could explore how different explanations of both model behavior and of multiplicity itself impact fairness perceptions. Likewise, we do not include personal impact – i.e.,

whether the prediction a model makes on a specific decision subject varies under multiplicity – as a factor due to the experiment’s complexity. We hypothesize – and future work should check – that if participants are told that the decision the model makes for them varies under multiplicity, fairness perceptions would decrease, mirroring work on outcome favorability [95]. To combat issues with the large number of unknowns inherent to hypothetical situations, future research could be situated in a real-world task, where explaining more details about the model would not necessitate a large cognitive load due to baseline familiarity with an existing algorithm. Existing research also shows that people behave differently in hypothetical settings [31], again highlighting the need for additional research situated in a real use case.

Unknown role of other personal factors. We did not find significant trends between demographic groups for any of the factors we recorded. However, our analysis was not powerful enough to detect how any intersectional identities (e.g., gender and education level) may impact perceptions of multiplicity. Furthermore, our collection of some demographic attributes, especially experience with AI, may have been too coarse to see meaningful trends. Future work could consider contrasting specific participant groups (e.g., ML engineers versus adults with no ML experience) to gain a better sense of how fairness perceptions vary with ML experience.

Addressing the “why” behind multiplicity preferences. Our results indicate that stakeholders have strong preferences for how multiplicity is handled; however, our experiment does not determine why these preferences exist. In particular, the discrepancy between philosophical arguments for the fairness of random and the strong dislike that survey participants expressed towards random merits more attention. Is there truly a mismatch between what is theoretically fair based on philosophical principles and what people perceive as fair? Or is there a breakdown in the communication, e.g., would participants actually be comfortable with randomization if more details about the model and its operationalization were provided? Future studies could vary the amount of detail provided to see whether attitudes towards randomization change.

Task choice and characterization. We consider tasks that differ along three axes, but there may be alternate ways to characterize tasks. It is also not necessarily straightforward to introduce new tasks into our classification system, particularly for the continuous scale of task stakes and task uncertainty. Our analysis also lost detail about the task characteristics in the conjoint analysis, as this framework does not support continuous factors. Future work could continue to study how to systematically characterize tasks to promote reproducible and generalizable research.

US-centric perspective. We conducted our survey with US-based crowdworkers. It is known that fairness perceptions vary cross-culturally [5, 49], as do attitudes towards uncertainty [40]. It therefore seems likely that attitudes towards the fairness of multiplicity in ML would vary across cultures, too, which future work could explore.

7 Conclusions

We conducted a study (n=357) of crowdworkers to determine how lay stakeholders (1) perceive the fairness of ML models when multiplicity impacts decisions, and (2) prefer that this multiplicity be resolved. Our results indicate that stakeholders have strong preferences about how multiplicity should be resolved. In particular, they express a preference for human decision makers, and a dispreference for both ignoring multiplicity (arbitrarily using a single model) and randomizing between different models’ outputs. These findings counter not only the established ML practice of looking for a single well-performing model, but also philosophical arguments that randomization is the fairest way to deal with multiplicity. Despite these strong opinions about resolving multiplicity, the presence of multiplicity did not impact how fair the participants thought the models were, which is a discrepancy that future work will need to address. One promising avenue is to consider how perceptions of informational fairness relating to multiplicity, e.g., the use of explanations, intersects with the procedural components of fairness that we investigated. Ultimately, a greater understanding of when ML design choices – including how multiplicity is resolved – are perceived as fair can help guide model developers to create technology that aligns with stakeholders’ values.

Acknowledgments

The authors thank members of the MadPL group for serving as pilot testers for initial survey iterations, and Michael Coblentz for his comments on an earlier draft. This work was conducted with the support of NSF grants 1918211 and 2446711.

References

- [1] Theo Araujo, Natali Helberger, Sanne Kruikeimeier, and Claes H De Vreese. 2020. In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI & society* 35, 3 (2020), 611–623.
- [2] Sarah Bankins, Paul Formosa, Yannick Griep, and Deborah Richards. 2022. AI decision making with dignity? Contrasting workers' justice perceptions of human and AI decision making in a human resource management context. *Information Systems Frontiers* 24, 3 (2022), 857–875.
- [3] Peter L Bartlett and Marten H Wegkamp. 2008. Classification with a Reject Option using a Hinge Loss. *Journal of Machine Learning Research* 9, 8 (2008), 1823–1840.
- [4] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- [5] John J Berman, Virginia Murphy-Berman, and Purnima Singh. 1985. Cross-cultural similarities and differences in perceptions of fairness. *Journal of Cross-Cultural Psychology* 16, 1 (1985), 55–67.
- [6] Reuben Binns, Max Van Kleek, Michael Veale, Ulrik Lyngs, Jun Zhao, and Nigel Shadbolt. 2018. 'It's Reducing a Human Being to a Percentage' Perceptions of Justice in Algorithmic Decisions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3173951>
- [7] Emily Black, Klas Leino, and Matt Fredrikson. 2022. Selective Ensembles for Consistent Predictions. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=HfUyCRBeQc>
- [8] Emily Black, Manish Raghavan, and Solon Barocas. 2022. Model multiplicity: Opportunities, concerns, and solutions. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 850–863. <https://doi.org/10.1145/3531146.3533149>
- [9] M José Blanca Mena, Rafael Alarcón Postigo, Jaume Arnau Gras, Roser Bono Cabré, and Rebecca Bendayan. 2017. Non-normal data: Is ANOVA still a valid option? *Psicothema* 29, 4 (2017), 552–557.
- [10] Elizabeth Bondi, Lily Xu, Diana Acosta-Navas, and Jackson A Killian. 2021. Envisioning communities: a participatory approach towards AI for social good. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 425–436. <https://doi.org/10.1145/3461702.3462612>
- [11] Leo Breiman. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science* 16, 3 (2001), 199–231.
- [12] John Broome. 1990. Fairness. In *Proceedings of the Aristotelian society*, Vol. 91. JSTOR, [Aristotelian Society, Wiley], 87–101.
- [13] James Carifio and Rocco Perla. 2008. Resolving the 50-year debate around using and misusing Likert scales. , 1150–1152 pages.
- [14] Noah Castelo, Maarten W Bos, and Donald R Lehmann. 2019. Task-dependent algorithm aversion. *Journal of Marketing Research* 56, 5 (2019), 809–825.
- [15] Chun-Wei Chiang and Ming Yin. 2022. Exploring the Effects of Machine Learning Literacy Interventions on Laypeople's Reliance on Machine Learning Models. In *Proceedings of the 27th International Conference on Intelligent User Interfaces (Helsinki, Finland) (IUI '22)*. Association for Computing Machinery, New York, NY, USA, 148–161. <https://doi.org/10.1145/3490099.3511121>
- [16] Alexandra Chouldechova. 2017. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data* 5, 2 (2017), 153–163.
- [17] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, 134–148. <https://proceedings.mlr.press/v81/chouldechova18a.html>
- [18] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. 2019. Certified Adversarial Robustness via Randomized Smoothing. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 1310–1320. <https://proceedings.mlr.press/v97/cohen19c.html>
- [19] Jason A Colquitt. 2001. On the dimensionality of organizational justice: a construct validation of a measure. *Journal of applied psychology* 86, 3 (2001), 386.
- [20] Jason A Colquitt and Jessica B Rodell. 2015. Measuring justice and fairness. *The Oxford Handbook of Justice in the Workplace* 1 (2015), 187–202.
- [21] A. Feder Cooper, Katherine Lee, Madiha Zahrah Choksi, Solon Barocas, Christopher De Sa, James Grimmelman, Jon Kleinberg, Siddhartha Sen, and Baobao Zhang. 2024. Arbitrariness and Social Prediction: The Confounding Role of Variance in Fair Classification. *Proceedings of the AAAI Conference on Artificial Intelligence* 38, 20 (Mar. 2024), 22004–22012. <https://doi.org/10.1609/aaai.v38i20.30203>
- [22] Kathleen Creel and Deborah Hellman. 2022. The Algorithmic Leviathan: Arbitrariness, Fairness, and Opportunity in Algorithmic Decision-Making Systems. *Canadian Journal of Philosophy* 52, 1 (2022), 26–43. <https://doi.org/10.1017/can.2022.3>
- [23] Fiery Cushman, Liane Young, and Marc Hauser. 2006. The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science* 17, 12 (2006), 1082–1089.
- [24] Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. 2022. Underspecification Presents Challenges for Credibility in Modern Machine Learning. *Journal of Machine Learning Research* 23, 226 (2022), 1–61.
- [25] Berkeley J Dietvorst and Soham Bharti. 2020. People reject algorithms in uncertain decision domains because they have diminishing sensitivity to forecasting error. *Psychological Science* 31, 10 (2020), 1302–1314.
- [26] Berkeley J Dietvorst, Joseph P Simmons, and Cade Massey. 2015. Algorithm aversion: people erroneously avoid algorithms after seeing them err. *Journal of experimental psychology: General* 144, 1 (2015), 114.
- [27] Jonathan Dodge, Q Vera Liao, Yunfeng Zhang, Rachel KE Bellamy, and Casey Dugan. 2019. Explaining models: an empirical study of how explanations impact fairness judgment. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. Association for Computing Machinery, New York, NY, USA, 275–285. <https://doi.org/10.1145/3301275.3302310>
- [28] 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 (Cambridge, Massachusetts) (ITCS '12)*. Association for Computing Machinery, 214–226. <https://doi.org/10.1145/2090236.2090255>
- [29] Lisa A Elkin, Matthew Kay, James J Higgins, and Jacob O Wobbrock. 2021. An aligned rank transform procedure for multifactor contrast tests. In *The 34th annual ACM symposium on user interface software and technology*, 754–768.
- [30] Michael Feffer, Michael Skirpan, Zachary Lipton, and Hoda Heidari. 2023. From preference elicitation to participatory ML: A critical survey & guidelines for future research. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 38–48. <https://doi.org/10.1145/3600211.3604661>
- [31] Oriel FeldmanHall, Dean Mobbs, Davy Evans, Lucy Hiscox, Lauren Navrady, and Tim Dalgleish. 2012. What we say and what we do: The relationship between real and hypothetical moral choices. *Cognition* 123, 3 (2012), 434–441.
- [32] Riccardo Fogliato, Shreya Chappidi, Matthew Lungren, Paul Fisher, Diane Wilson, Michael Fitzke, Mark Parkinson, Eric Horvitz, Kori Inkpen, and Basmira Nushi. 2022. Who goes first? Influences of human-AI workflow on decision making in clinical imaging. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 1362–1374. <https://doi.org/10.1145/3531146.3533193>
- [33] John Fox and Sanford Weisberg. 2019. *An R Companion to Applied Regression* (third ed.). Sage, Thousand Oaks CA. <https://socialsciences.mcmaster.ca/jfox/Books/Companion/>
- [34] Nina Grgic-Hlaca, Elissa M. Redmiles, Krishna P. Gummadi, and Adrian Weller. 2018. Human Perceptions of Fairness in Algorithmic Decision Making: A Case Study of Criminal Risk Prediction. In *Proceedings of the 2018 World Wide Web Conference (Lyon, France) (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 903–912. <https://doi.org/10.1145/3178876.3186138>
- [35] Jens Hainmueller, Daniel J Hopkins, and Teppei Yamamoto. 2014. Causal inference in conjoint analysis: Understanding multidimensional choices via stated preference experiments. *Political analysis* 22, 1 (2014), 1–30.
- [36] Jacqueline Hannan, Huei-Yen Winnie Chen, and Kenneth Joseph. 2021. Who gets what, according to whom? An analysis of fairness perceptions in service allocation. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 555–565. <https://doi.org/10.1145/3461702.3462568>
- [37] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems* 29 (2016).
- [38] Natali Helberger, Theo Araujo, and Claes H de Vreese. 2020. Who is the fairest of them all? Public attitudes and expectations regarding automated decision-making. *Computer Law & Security Review* 39 (2020), 105456.
- [39] Radu Herbei and Marten H Wegkamp. 2006. Classification with reject option. *The Canadian Journal of Statistics/La Revue Canadienne de Statistique* 34, 4 (2006), 709–721.
- [40] Geert Hofstede. 1984. *Culture's consequences: International differences in work-related values*. Vol. 5. sage.
- [41] Sture Holm. 1979. A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics* 6, 2 (1979), 65–70.
- [42] George S Howard, Kenneth M Ralph, Nancy A Gulanic, Scott E Maxwell, Don W Nance, and Sterling K Gerber. 1979. Internal invalidity in pretest-posttest self-report evaluations and a re-evaluation of retrospective pretests. *Applied Psychological Measurement* 3, 1 (1979), 1–23.
- [43] Shomik Jain, Kathleen Creel, and Ashia Camage Wilson. 2024. Position: Scarce Resource Allocations That Rely On Machine Learning Should Be Randomized. In *Forty-first International Conference on Machine Learning*. <https://openreview>.

- net/forum?id=44qxX6Ty6F
- [44] Maurice Jakesch, Zana Bućinca, Saleema Amershi, and Alexandra Olteanu. 2022. How different groups prioritize ethical values for responsible AI. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 310–323. <https://doi.org/10.1145/3531146.3533097>
- [45] Susan Jamieson. 2004. Likert scales: How to (ab) use them? *Medical education* 38, 12 (2004), 1217–1218.
- [46] Faisal Kamiran and Toon Calders. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems* 33, 1 (2012), 1–33.
- [47] Maria Kasinidou, Styliani Kleanthous, Pinar Barlas, and Jahna Otterbacher. 2021. I agree with the decision, but they didn't deserve this: Future Developers' Perception of Fairness in Algorithmic Decisions. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (Virtual Event, Canada) (FAccT '21)*. Association for Computing Machinery, New York, NY, USA, 690–700. <https://doi.org/10.1145/3442188.3445931>
- [48] Matthew Kay, Lisa A. Elkin, James J. Higgins, and Jacob O. Wobbrock. 2021. *mjskay/ARTool: ARTool 0.11.0*. <https://doi.org/10.5281/zenodo.4721941>
- [49] Tae-Yeol Kim and Kwok Leung. 2007. Forming and reacting to overall fairness: A cross-cultural comparison. *Organizational Behavior and Human Decision Processes* 104, 1 (2007), 83–95.
- [50] Jon Kleinberg, Sendhil Mullainathan, and Manish Raghavan. 2017. Inherent Trade-Offs in the Fair Determination of Risk Scores. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany. <https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2017.43>
- [51] Jon Kleinberg and Manish Raghavan. 2021. Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118, 22 (2021), e2018340118.
- [52] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. *Advances in Neural Information Processing Systems* 30 (2017).
- [53] William Kuzon, Melanie Urbanchek, and Steven McCabe. 1996. The seven deadly sins of statistical analysis. *Annals of plastic surgery* 37 (1996), 265–272.
- [54] Cassidy Laidlaw and Soheil Feizi. 2019. Playing it safe: Adversarial robustness with an abstain option. *arXiv preprint arXiv:1911.11253* (2019).
- [55] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (2018), 2053951718756684.
- [56] Min Kyung Lee and Katherine Rich. 2021. Who Is Included in Human Perceptions of AI?: Trust and Perceived Fairness around Healthcare AI and Cultural Mistrust. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 138, 14 pages. <https://doi.org/10.1145/3411764.3445570>
- [57] Yoonjoo Lee, Kihoon Son, Tae Soo Kim, Jisu Kim, John Joon Young Chung, Eytan Adar, and Juho Kim. 2024. One vs. Many: Comprehending Accurate Information from Multiple Erroneous and Inconsistent AI Generations. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 2518–2531. <https://doi.org/10.1145/3630106.3662681>
- [58] Thomas J. Leeper. 2020. *cregg: Simple Conjoint Analyses and Visualization*. R package version 0.4.0.
- [59] Lan Li, Tina Lassiter, Joohye Oh, and Min Kyung Lee. 2021. Algorithmic hiring in practice: Recruiter and HR Professional's perspectives on AI use in hiring. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. Association for Computing Machinery, New York, NY, USA, 166–176. <https://doi.org/10.1145/3461702.3462531>
- [60] Jennifer M Logg, Julia A Minson, and Don A Moore. 2019. Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes* 151 (2019), 90–103.
- [61] Duri Long and Brian Magerko. 2020. What is AI Literacy? Competencies and Design Considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–16. <https://doi.org/10.1145/3313831.3376727>
- [62] R.Duncan Luce. 1977. The choice axiom after twenty years. *Journal of Mathematical Psychology* 15, 3 (1977), 215–233. [https://doi.org/10.1016/0022-2496\(77\)90032-3](https://doi.org/10.1016/0022-2496(77)90032-3)
- [63] Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics* 18, 1 (1947), 50–60.
- [64] Charles Marx, Flavio Calmon, and Berk Ustun. 2020. Predictive Multiplicity in Classification. In *Proceedings of the 37th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 119)*. Hal Daumé III and Aarti Singh (Eds.). PMLR, 6765–6774. <https://proceedings.mlr.press/v119/marx20a.html>
- [65] Lucas Maystre and Matthias Grossglauser. 2015. Fast and accurate inference of Plackett–Luce models. *Advances in Neural Information Processing Systems* 28 (2015).
- [66] Aditya Krishna Menon and Robert C Williamson. 2018. The cost of fairness in binary classification. In *Conference on Fairness, accountability and transparency (Proceedings of Machine Learning Research, Vol. 81)*, Sorelle A. Friedler and Christo Wilson (Eds.). PMLR, PMLR, 107–118. <https://proceedings.mlr.press/v81/menon18a.html>
- [67] Devesh Narayanan, Mahak Nagpal, Jack McGuire, Shane Schweitzer, and David De Cremer. 2024. Fairness perceptions of artificial intelligence: A review and path forward. *International Journal of Human–Computer Interaction* 40, 1 (2024), 4–23.
- [68] Dennis Nguyen and Erik Hekman. 2024. The news framing of artificial intelligence: a critical exploration of how media discourses make sense of automation. *AI & Society* 39, 2 (2024), 437–451.
- [69] Geoff Norman. 2010. Likert scales, levels of measurement and the “laws” of statistics. *Advances in health sciences education* 15 (2010), 625–632.
- [70] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [71] R. L. Plackett. 1975. The Analysis of Permutations. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 24, 2 (1975), 193–202. <http://www.jstor.org/stable/2346567>
- [72] Prolific. 2024. Prolific's Attention and Comprehension Check Policy. <https://researcher-help.prolific.com/en/article/fb63bb>
- [73] Propublica. 2016. Propublica Compas Analysis. <https://github.com/propublica/compas-analysis>
- [74] Aaron Roth, Alexander Tolbert, and Scott Weinstein. 2023. Reconciling Individual Probability Forecasts*. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (<conf-loc>, <city>Chicago</city>, <state>IL</state>, <country>USA</country>, </conf-loc>)* (FAccT '23). Association for Computing Machinery, New York, NY, USA, 101–110. <https://doi.org/10.1145/3593013.3593980>
- [75] Cynthia Rudin, Chudi Zhong, Lesia Semenova, Margo Seltzer, Ronald Parr, Jiachang Liu, Srikar Katta, Jon Donnelly, Harry Chen, and Zachery Boner. 2024. Position: Amazing Things Come From Having Many Good Models. In *Forty-first International Conference on Machine Learning*. <https://openreview.net/forum?id=oFDFGd9Age>
- [76] Omer Sagi and Lior Rokach. 2018. Ensemble learning: A survey. *Wiley interdisciplinary reviews: data mining and knowledge discovery* 8, 4 (2018), e1249.
- [77] Sara Salimzadeh, Gaole He, and Ujwal Gadiraju. 2023. A missing piece in the puzzle: Considering the role of task complexity in human-ai decision making. In *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization*. Association for Computing Machinery, New York, NY, USA, 215–227. <https://doi.org/10.1145/3565472.3592959>
- [78] Sara Salimzadeh, Gaole He, and Ujwal Gadiraju. 2024. Dealing with Uncertainty: Understanding the Impact of Prognostic Versus Diagnostic Tasks on Trust and Reliance in Human-AI Decision-Making. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3613904.3641905>
- [79] Melissa Saragih and Ben W Morrison. 2022. The effect of past algorithmic performance and decision significance on algorithmic advice acceptance. *International Journal of Human–Computer Interaction* 38, 13 (2022), 1228–1237.
- [80] Nripsuta Ani Saxena, Karen Huang, Evan DeFilippis, Goran Radanovic, David C. Parkes, and Yang Liu. 2019. How Do Fairness Definitions Fare? Examining Public Attitudes Towards Algorithmic Definitions of Fairness. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society (Honolulu, HI, USA) (AI/ES '19)*. Association for Computing Machinery, New York, NY, USA, 99–106. <https://doi.org/10.1145/3306618.3314248>
- [81] Jakob Schoeffer, Niklas Kuehl, and Yvette Machowski. 2022. “There is not enough information”: On the effects of explanations on perceptions of informational fairness and trustworthiness in automated decision-making. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, New York, NY, USA, 1616–1628. <https://doi.org/10.1145/3531146.3533218>
- [82] Ian A Scott, Stacy M Carter, and Enrico Coiera. 2021. Exploring stakeholder attitudes towards AI in clinical practice. *BMJ Health & Care Informatics* 28, 1 (2021).
- [83] Shayle R Searle, F Michael Speed, and George A Milliken. 1980. Population marginal means in the linear model: an alternative to least squares means. *The American Statistician* 34, 4 (1980), 216–221.
- [84] Ayelet Sela. 2018. Can computers be fair: How automated and human-powered online dispute resolution affect procedural justice in mediation and arbitration. *Ohio St. J. on Disp. Resol.* 33 (2018), 91.
- [85] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*.

- Association for Computing Machinery, New York, NY, USA, 59–68. <https://doi.org/10.1145/3287560.3287598>
- [86] George Sher. 1980. What makes a lottery fair? *Noûs* 14, 2 (1980), 203–216.
- [87] Donghee Shin and Yong Jin Park. 2019. Role of fairness, accountability, and transparency in algorithmic affordance. *Computers in Human Behavior* 98 (2019), 277–284.
- [88] Donghee Shin, Bu Zhong, and Frank A Biocca. 2020. Beyond user experience: What constitutes algorithmic experiences? *International Journal of Information Management* 52 (2020), 102061.
- [89] Joseph P Simmons, Leif D Nelson, and Uri Simonsohn. 2011. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological science* 22, 11 (2011), 1359–1366.
- [90] Megha Srivastava, Hoda Heidari, and Andreas Krause. 2019. Mathematical Notions vs. Human Perception of Fairness: A Descriptive Approach to Fairness for Machine Learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Anchorage, AK, USA) (KDD '19)*. Association for Computing Machinery, New York, NY, USA, 2459–2468. <https://doi.org/10.1145/3292500.3330664>
- [91] Christopher Starke, Janine Baleis, Birte Keller, and Frank Marcinkowski. 2022. Fairness perceptions of algorithmic decision-making: A systematic review of the empirical literature. *Big Data & Society* 9, 2 (2022), 20539517221115189.
- [92] Sara Steegen, Francis Tuerlinckx, Andrew Gelman, and Wolf Vanpaemel. 2016. Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science* 11, 5 (2016), 702–712.
- [93] Steven Umbrello and Ibo Van de Poel. 2021. Mapping value sensitive design onto AI for social good principles. *AI and Ethics* 1, 3 (2021), 283–296.
- [94] Angelina Wang, Sayash Kapoor, Solon Barocas, and Arvind Narayanan. 2024. Against predictive optimization: On the legitimacy of decision-making algorithms that optimize predictive accuracy. *ACM Journal on Responsible Computing* 1, 1 (2024), 1–45.
- [95] Ruotong Wang, F Maxwell Harper, and Haiyi Zhu. 2020. Factors influencing perceived fairness in algorithmic decision-making: Algorithm outcomes, development procedures, and individual differences. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3313831.3376813>
- [96] Janelle Watson-Daniels, David C. Parkes, and Berk Ustun. 2023. Predictive Multiplicity in Probabilistic Classification. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 9 (Jun. 2023), 10306–10314. <https://doi.org/10.1609/aaai.v37i9.26227>
- [97] Jacob O Wobbrock, Leah Findlater, Darren Gergle, and James J Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 143–146. <https://doi.org/10.1145/1978942.1978963>
- [98] Allison Woodruff, Sarah E Fox, Steven Rousso-Schindler, and Jeffrey Warshaw. 2018. A qualitative exploration of perceptions of algorithmic fairness. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–14. <https://doi.org/10.1145/3173574.3174230>
- [99] Xingjiao Wu, Luwei Xiao, Yixuan Sun, Junhang Zhang, Tianlong Ma, and Liang He. 2022. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems* 135 (2022), 364–381.
- [100] Rui Xin, Chudi Zhong, Zhi Chen, Takuya Takagi, Margo Seltzer, and Cynthia Rudin. 2022. Exploring the Whole Rashomon Set of Sparse Decision Trees. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (Eds.). Curran Associates, Inc., 14071–14084. <https://openreview.net/forum?id=WHqVVk3UHR>
- [101] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the effect of accuracy on trust in machine learning models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300509>
- [102] Mireia Yurrita, Tim Draws, Agathe Balayn, Dave Murray-Rust, Nava Tintarev, and Alessandro Bozzon. 2023. Disentangling Fairness Perceptions in Algorithmic Decision-Making: the Effects of Explanations, Human Oversight, and Contestability. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 134, 21 pages. <https://doi.org/10.1145/3544548.3581161>
- [103] Angie Zhang, Olympia Walker, Kaci Nguyen, Jiajun Dai, Anqing Chen, and Min Kyung Lee. 2023. Deliberating with AI: Improving Decision-Making for the Future through Participatory AI Design and Stakeholder Deliberation. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW1, Article 125 (apr 2023), 32 pages. <https://doi.org/10.1145/3579601>

A Preliminary study

A.1 Survey materials

We ask the following questions (on 5-point Likert scales from strongly disagree to strongly agree) to gauge task stakes (S), task uncertainty (C), and task clarity and realistic-ness (R). Questions are flagged with (*) to indicate that the phrasing changes with the scenario to provide more detail (e.g., for the task of loan approval “in the positive direction” would become “approves a loan for someone who will not be able to repay it”, since getting a loan is the positive outcome from the perspective of the decision subject).

- (S) This scenario describes a high-stakes situation.
- (S) It is a big problem if the system makes a mistake in the positive direction.*
- (S) It is a big problem if the system makes a mistake in the negative direction.*
- (C) For each decision the system makes, it is objectively right or wrong (i.e., there is not much ambiguity in decisions).
- (C) If the AI system is designed well, I think it would be able to make better decisions than a human expert.
- (R) Is this scenario described clearly?
- (R) Is this scenario realistic?

A.2 Tasks and results

Table 4 shows all tasks we considered in the preliminary study, along with their final stakes and uncertainty scores.

B Complete survey materials

B.1 Tasks

This section contains the full text used to describe each task at the beginning of Part 1.

Care plan. A hospital plans to use an AI system to identify patients who qualify for a preventative care plan, which includes access to educational materials and more frequent doctor check-ins. The AI will base its decisions on the patient’s health and past disease progression patterns. Doctors can override the AI’s decisions if needed. The hospital checks that the AI works well by seeing how often it agrees with doctors’ choices. The AI and the doctors agree most of the time.

Job probation. A company is planning to use an AI system to identify underperforming employees who need job probation. The AI will make its decisions based on performance metrics and work history. All flagged employees will be placed on probation, but they will not be demoted or fired unless a human review of their probationary period is negative. The company checks that the AI works well by seeing how often it agrees with human managers’ choices. The AI and the managers agree most of the time.

Tumor detection. A hospital plans to use AI system to look at CT scans to find potential cancerous tumors. If the AI detects signs of a tumor, the patient will receive additional tests. However, if no tumor is detected, additional tests will not be conducted. Doctors can override the AI’s decision if needed. The hospital checks that the AI works well by seeing how often it agrees with doctors’ choices. The AI and the doctors agree most of the time.

Table 4: Summary of tasks and their stakes, uncertainty, and framing as determined by the preliminary study. Tasks marked with (*) were excluded from the main study because participants indicated they were unclear and/or unrealistic. The final column “Use” indicates whether we use the task in the main study.

Domain	Task	Stakes	Uncertainty	Framing	Use
Healthcare	Access to a preventative health care plan	0.42	0.91	Reward	Yes
Business	Job probation based on employee monitoring	0.72	1.12	Punishment	Yes
Healthcare	Tumor detection from medical imaging	1.61	-1.37	Reward	Yes
Government	Tax audit	0.44	-1.34	Punishment	Yes
Travel	Airline seat upgrade	-1.94	1.18	Reward	Yes
Travel	Rebooking passengers on an over-full flight	-0.27	0.50	Punishment	Yes
Business	Customers get beta access to an app with discounts	-1.58	-0.21	Reward	Yes
Business	Removing online reviews that are spam or fake	-1.26	-0.16	Punishment	Yes
Education	College admissions decisions*	0.40	1.13	Reward	No
Education	Automated plagiarism detection	-0.16	-1.47	Punishment	No
Healthcare	Discharge from hospital*	1.41	0.50	Punishment	No
Business	Hiring based on resume screening	-0.18	1.04	Reward	No
Business	Approving online reviews for publication	-1.15	-0.06	Reward	No
Banking	Detecting fraudulent credit card transactions	0.44	-1.19	Reward	No
Banking	Approving loan applications	0.49	-0.01	Reward	No
Government	Determining benefits eligibility*	0.92	-0.31	Reward	No
Insurance	Adding an insurance surcharge if high-risk	0.11	-0.12	Punishment	No
Insurance	Given an insurance discount if low-risk	-0.40	-0.15	Reward	No

Tax audit. A government tax agency is planning to use an AI system to detect potential tax fraud. Everyone flagged by the AI will be audited, and audits will occur solely based on AI flags. However, taxpayers will only face penalties if a human investigation confirms the fraud. The tax agency checks that the AI works well by seeing how often it agrees with human experts’ choices of who to audit. The AI and the human experts agree most of the time.

Upgrade. An airline is planning to use an AI system to decide which passengers get upgrades to empty first-class seats. The AI will make its decisions based on factors like passenger loyalty status, ticket price, and check-in time. The AI’s decisions are final, but passengers can decline the upgrade and keep their original seat if they prefer. The airline checks that the AI works well by seeing how often it agrees with trained airline employees’ choices. The AI and the airline employees agree most of the time.

Rebook. An airline wants to use an AI system to manage over-booked flights. The AI will re-route passengers based on their destination, personal details (like age and group size), and booking details. While the AI’s decisions are final, passengers can opt for a refund or select an alternative re-routing plan if they are unhappy with the AI’s choice. The airline checks that the AI works well by seeing how often it agrees with trained airline employees’ choices. The AI and the employees agree most of the time.

App. A company is planning use an AI system to decide which customers get beta access to their new app, which offers exclusive deals and coupons. The AI will make its decisions based on expected financial gain for the company based on the customers’ past spending habits. The decisions about which customers get early access are final until the app is available to the public. The company checks that the AI works well by seeing how often it agrees

with human marketing experts’ choices. The AI and the marketing experts agree most of the time.

Reviews. A website where people post reviews of businesses is planning to use an AI system to detect and remove fake reviews, e.g., posts by bots or by people who have not actually visited a business. Once a review is deleted, it cannot be recovered, but the person who wrote it can re-post a similar review if they want. The review website checks that the AI works well by seeing how often it agrees with trained human moderators’ choices. The AI and the moderators agree most of the time.

B.2 Multiplicity education

We provide the following text to participants at the beginning of the survey:

When building Artificial Intelligence (AI), choosing the best model is crucial. AI developers consider a range of AI models and choose the one with the best performance. For example, if the goal is to distinguish between cats and dogs, developers will test different AIs using a set of images and choose the AI that correctly identifies the most pictures.

In the example below, AI #1 correctly identifies 4 out of 5 images. Suppose this is the highest-performing AI model available.

We display the top half of fig. 4 (i.e., images and AI #1 only).

Sometimes, however, multiple AI models perform equally well but differ in their specific errors. This is

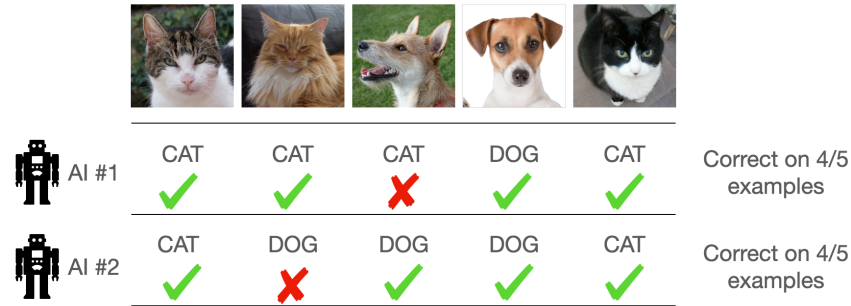


Figure 4: Image that we include with the multiplicity education.

called model multiplicity. For example, AI #2 also correctly identifies 4 out of 5 images but makes a different mistake than AI #1.

We display fig. 4.

AI #1 and AI #2 are both well-suited for classifying cat and dog images. If no other AI performs better, developers must choose between AI #1 and AI #2.

Model multiplicity is common in AI: multiple models can perform equally well overall but make different mistakes. In some settings, like predicting cats and dogs, model multiplicity is inevitable but choosing between models is low-stakes. In other settings, choosing between models is more challenging, such as when the model predictions directly affect people.

Then, we ask the comprehension question “What does model multiplicity refer to?” with these choices: (a) It refers to a model making multiple correct predictions, (b) It means there are several models that perform equally well but makes different mistakes, (c) It describes an AI model that performs better than all other models, and (d) It refers to an AI model that can perform multiple task simultaneously.

B.3 Part 1: Fairness questions

We use two questions relating to distributive fairness, and two relating to procedural fairness. Each question is on a 7-point Likert scale (“Strongly disagree” to “Strongly agree”). The exact phrasing of each question is tailed to each scenario to be more descriptive (e.g., replacing the word “entity” with the context-appropriate word).

- (Distributive) I think the AI will make good decisions.
- (Distributive) I think the AI will make fair decisions.
- (Procedural) I think it is appropriate for the entity to use AI in this context.
- (Procedural) I would feel comfortable if the entity used AI to make a decision about me.

B.4 Part 1: Critical-thinking questions

We ask two critical-thinking questions midway through Part 1 of the survey to encourage participants to slow down and better absorb the material. The first question is multiple-choice, asking which of four statements is true: (a) The original AI model is the only good way to do [task], (b) The new AI and human experts always choose the same people, (c) Multiple models are equally good at [task], or

(d) The original AI and new AI never agree on [task]. The correct answer is (c) and this is visible to participants as soon as they make any selection. The open-ended question asks “If the original AI and the new AI disagree about whether or not to [predict + or -], how should the [decision-making entity] decide what to do?” Respondents much enter at least 8 characters in their answer, but apart from that requirement there is no in-survey validation of the answer.

B.5 Part 2: Costs

Table 5 shows the (non-zero) cost that we use for each task.

C Alternate statistical analyses

We describe two alternate statistical analyses for the Part 1 data (change in fairness perceptions after learning about multiplicity). As stated in section 4.4, we preregistered our study with a plan to use linear mixed models. Those results are in the main text. We include two alternate analyses here for two reasons: first, there is disagreement in the statistical community about what type of analysis is appropriate for Likert-scale data [9, 13, 45, 53, 69]. Including multiple options serves to “future-proof” our analysis in the case of future statistical consensus. Second, while preregistration is one tool to guard against “researcher degrees of freedom” [89], another option is a multiverse analysis [92]. In a true multiverse analysis, researchers present the outcomes of all feasible statistical analyses on their data. Then, if the same conclusions hold for all or most of the analyses, they can be considered more robust. We present a partial multiverse analysis.

C.1 ART-ANOVA

Aligned-Rank-Transform (ART) ANOVA [97] is a procedure that transforms ordinal (e.g., Likert-scale) data to fit the ANOVA assumptions. We fit our ART-ANOVA model and perform post-hoc contrast tests [29] using the ARTTool package in R [48]. Our formula includes the factors time, task stakes (binary), task uncertainty (binary), and task framing and all possible interactions between factors. We also include the participant ID to account for within-subjects factor. We test H_{1a} using a post-hoc contrast test with “time” as the factor of interest. We test H_{1b} , H_{1c} , and H_{1d} with a post-hoc contrast test of the interaction term between “time” and the relevant task characteristic (with the “Interaction” flag set to true).

Table 5: The non-zero cost for each task.

Task	Cost
Care plan	15 minute longer wait time at appointments to see the doctor
Job probation	Employees will spend 15 extra minutes per week filling out paperwork about their work
Tumor detection	15 minute longer wait time at appointments to see the doctor
Tax audit	There will be an additional \$5 fee to file taxes
Upgrade	Airline tickets will cost \$5 more
Rebook	Airline tickets will cost \$5 more
App	The shipping time on all products will be 2 days longer
Reviews	The website will have to host more ads, including pop-up ads

Results. Our ART-ANOVA model indicates that the following factors have significance: task stakes ($F=35.2$, $Df=1$, $p < 10^{-8}$), task framing ($F=5.0$, $Df=1$, $p=0.026$), and the interaction term between task stakes and task uncertainty ($F=7.9$, $Df=1$, $p=0.005$). As was the case when using LMMs in the analysis, none of the hypotheses H_{1a} - H_{1d} was significant.

C.2 Non-parametric tests

The Mann-Whitney U-test [63] is another way we can evaluate H_{1a} - H_{1c} . This test is non-parametric, so does not require any assumptions about the data’s normality or distribution. We perform the tests using the `wilcox.test` function in R. For H_{1a} , we test whether the average fairness perception is different as stratified by “time”. For H_{1b} - H_{1d} , we compute the change in score over time for each sample, and then conduct a Mann-Whitney U test on the data stratified by the relevant task characteristic.

Results. As is the case with using LMMs or ART-ANOVA, none of the hypotheses H_{1a} - H_{1d} is significant.

D Additional results (Fairness perceptions)

D.1 Task choice and baseline fairness perceptions

Table 6 shows the fairness perceptions for each task at the “pre” timestamp. These values should be interpreted as a baseline for how fair participants think it is to ML in each scenario. We observe that high-stakes tasks have lower fairness scores, as do punishment-framed tasks.

D.2 Demographic factors

We consider binary gender, integer age, binary experience with AI (defined as whether the participant currently uses AI in their work or school), and binary education (whether the participant has a bachelor’s degree or higher). There is no significant difference in initial fairness perceptions based on gender (t-test $p=0.822$), age (linear regression $p=0.184$), experience with AI (t-test $p=0.088$), or education (t-test $p=0.328$). There is no significant different in change in fairness scores after learning about multiplicity for gender (t-test $p=0.921$), age (linear regression $p=0.944$), AI experience (t-test $p=0.460$), or education (t-test $p=0.413$).

For high-stakes tasks, the mean fairness score is 0.05 on a 7-pt scale (IQR -0.81 to 1.0) and for low-stakes tasks it is 0.79 (IQR 0.25 to

Table 6: The average fairness score (on the questions listed in appendix B.3) at the “pre” timestamp for each task. Scores are from 7-point Likert scales; re-coded as -3 (Strongly disagree that the model is fair) to 3 (Strongly agree that the model is fair). We report the mean with the interquartile range in parentheses. We also indicates stakes (binary low/high), uncertainty (binary low/high), and framing for reference.

Task	Score	Stakes	Unc.	Framing
Care plan	0.26 (-0.75, 1.13)	High	High	Rew.
Job probation	-0.38 (-1.50, 0.88)	High	High	Pun.
Tumor detection	0.12 (-0.75, 1.00)	High	Low	Rew.
Tax audit	0.19 (-0.25, 1.00)	High	Low	Pun.
Upgrade	1.01 (0.75, 1.75)	Low	High	Rew.
Rebook	0.83 (0.50, 1.75)	Low	High	Pun.
App	0.78 (0.00, 1.50)	Low	Low	Rew.
Reviews	0.55 (-0.25, 1.50)	Low	Low	Pun.

1.5). This difference is statistically significant according to a Mann-Whitney U-test ($W = 43476$, $p < 10^{-13}$). Punishment-framed tasks have a median fairness score of 0.30 (IQR -0.5 to 1.25), while for reward-framed tasks it is 0.55 (IQR 0.0 to 1.5). This difference is also statistically significant according to a Mann-Whitney U-test ($W = 71016$, $p = 0.012$). Task uncertainty does not have an impact on fairness score according to a Mann-Whitney U-test ($W = 64358$, $p = 0.918$).

E Additional results (Multiplicity resolution)

E.1 Details on testing H_{2a} - H_{2d}

Table 7 shows the effect sizes and significant for the post-hoc analyses conducted for H_{2a} - H_{2d} (section 5.2). We see that all effects except two are in the expected direction, and, using the Holm-Bonferroni correction, 10 of the 18 tests are significant.

E.2 Qualitative results

Table 8 shows the results of the open-ended question, where participants were asked how to resolve multiplicity prior to seeing the techniques that we proposed. The table contains 10 response categories and tallies the number of participants whose answer referenced that category for each of the 8 tasks. We did our best to interpret the responses, however, in some cases, we could not

determine what the participant meant so we marked it as “unclear”. Most of these responses were on-topic and were a clear attempt to answer the question.

E.3 Task-by-task preferences

Figure 5 shows the relative preferences for each multiplicity resolution technique for tasks stratified by stakes, uncertainty, and framing.

E.4 Costs

As noted in finding E4 in section 5.3, no extra cost is significantly preferred over extra cost in aggregate. However, for three high-stakes settings (tumor detection, care plan access, and job probation), this distinction is not significant. The point estimates for marginal mean preferences for extra cost are 0.436 (SE 0.032, $p < 0.0001$) for tumor detection, 0.474 (SE 0.038, $p < 0.0001$) for care plan access, and 0.445 (SE 0.035, $p < 0.0001$) for job probation. (A

point estimate less than 0.50 indicates that no extra cost is more preferred, and a point estimate of 0.50 indicates that there is no preferences between the cost options.)

E.5 Demographic factors

There is no significant difference in preference for multiplicity resolution technique based on education level or AI experience. That is, a t-test comparing the results for each resolution technique has a p-value greater than 0.05. When we consider age as a binary factor (divided at 35, roughly the median age and a common cutoff for “young adult”), fancy is preferred by people younger than 35 (diff=0.081, SE=0.039, $t = 2.054$, $p = 0.040$) and default_good is preferred by people older than 35 (diff=-0.101, SE=0.040, $t = -2.515$, $p = 0.012$). Looking at gender, fancy is preferred by men (diff=0.0783, SE=0.039, $t = 2.009$, $p = 0.045$) and random is preferred by women (diff=-0.070, SE=0.035, $t = -2.027$, $p = 0.043$). We encourage future work to explore whether these trends hold as the experiment setup is varied.

Table 7: Results of post-hoc t-tests for evaluating H_{2a} - H_{2d} . The “diff” column is the difference in marginal means between the two listed factors (first 9 rows) or between the two test conditions for a single factor (remaining rows). (*) indicates that an effect is in the opposite direction as expected.

Hypothesis	Factor	Diff	SE	t-stat	p-val	Sig?
H_{2a}	human v. fancy	0.188	0.025	7.563	<0.0001	Yes
H_{2a}	human v. default-good	0.214	0.026	8.382	<0.0001	Yes
H_{2a}	human v. default-bad	0.462	0.025	18.340	<0.0001	Yes
H_{2a}	human v. ignore	0.375	0.024	15.594	<0.0001	Yes
H_{2a}	human v. random	0.511	0.023	22.407	<0.0001	Yes
H_{2a}	random v. fancy	-0.323	0.026	-12.592	<0.0001	Yes
H_{2a}	random v. default-good	-0.297	0.026	-11.304	<0.0001	Yes
H_{2a}	random v. default-bad	-0.049	0.026	-1.887	0.0593	No
H_{2a}	random v. ignore	-0.136	0.025	-5.477	<0.0001	Yes
H_{2b}	human	0.097	0.030	3.211	0.0013	Yes
H_{2b}	fancy	0.043	0.038	1.128	0.2596	No
H_{2b}	ignore	-0.089	0.036	-2.468	0.0137	No
H_{2b}	random	-0.079	0.033	-2.411	0.0160	No
H_{2c}	human*	-0.004	0.031	-0.133	0.8943	No
H_{2c}	default-good	0.016	0.040	3.965	< 0.0001	Yes
H_{2c}	fancy	-0.059	0.039	-1.510	0.1313	No
H_{2d}	random	0.010	0.033	0.300	0.7641	No
H_{2d}	fancy*	0.015	0.039	0.389	0.6970	No

Table 8: Results of our qualitative analysis. The number in the header indicates how many participants were assigned to that group. If a response suggests multiple techniques, we count it in all categories. (*) Unclear indicates we could not interpret the response, while Not sure means that the participant wrote that they did not know how to resolve the multiplicity.

	Care plan <i>n</i> =43	Job probation <i>n</i> =43	Tumor detection <i>n</i> =45	Tax audit <i>n</i> =45	Upgrade <i>n</i> =46	Rebook <i>n</i> =45	App <i>n</i> =46	Reviews <i>n</i> =45	Total <i>n</i> =358
human	35	35	38	23	22	28	28	35	244
default-good	1	1	7	7			3	2	21
ignore	1			1	5	5	5	2	19
fancy	1	2			5	1	3	2	14
random				2	5	2	1		10
default-bad	1	1		3	1		3		9
Specific rule suggestion	1	1		1	4	5	3	2	17
Ask the decision subject	1				2	3			6
Unclear*	2	3	2	7	4	6	2	5	31
Not sure*		1		2	1				4

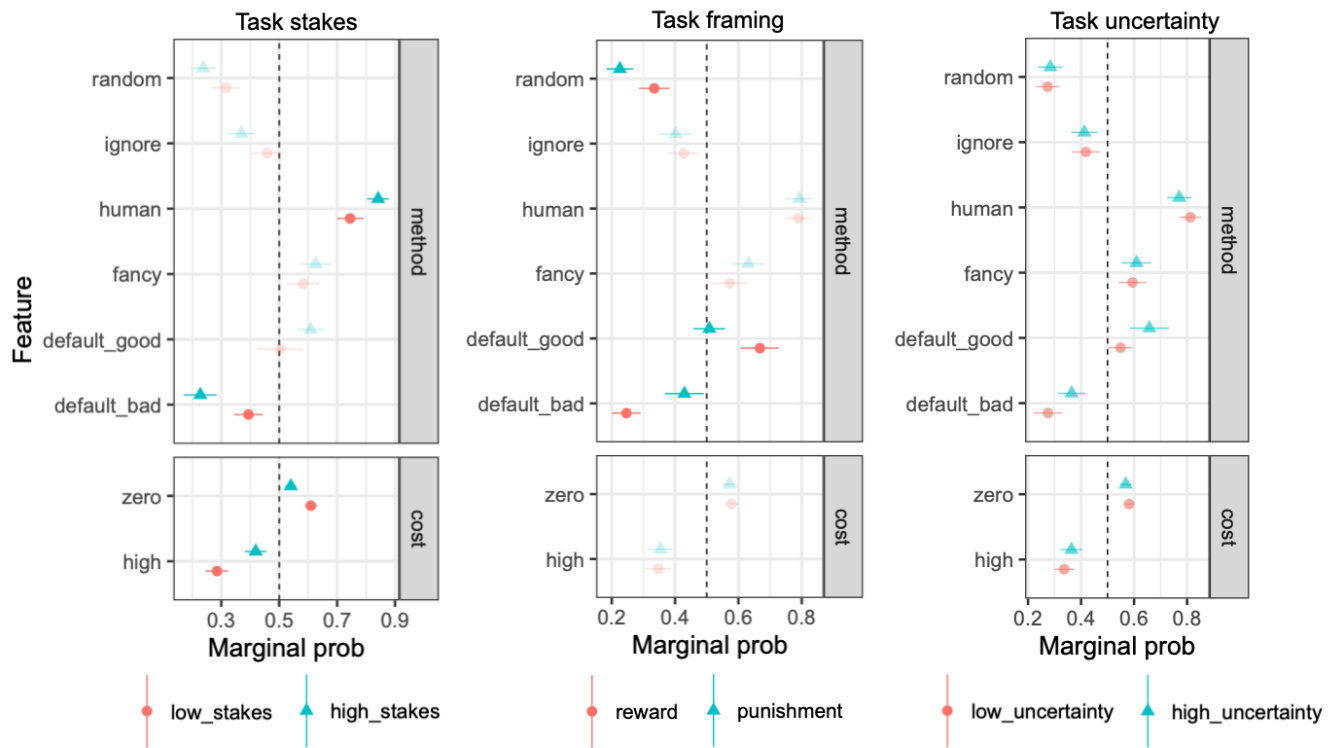


Figure 5: Marginal probabilities stratified by task stakes, task framing, and task uncertainty. For each “feature” (multiplicity resolution technique or cost level), the marginal probability is the chance of choosing that attribute, relative to a baseline random choice of 0.5. Results are further broken down by task characteristic. Results are darker when the confidence intervals for the two task levels do not overlap as a marker of greater significance. Confidences intervals were computed using $\alpha = 0.05$ to ease readability, see the discussion in section 5.2 for which differences are statistically significant with a multiple hypothesis correction.