

Underspecified Human Decision Experiments Considered Harmful

Jessica Hullman
jhullman@northwestern.edu
Northwestern University
Evanston, Illinois, USA

Alex Kale
kalea@uchicago.edu
University of Chicago
Chicago, Illinois, USA

Jason Hartline
hartline@eecs.northwestern.edu
Northwestern University
Evanston, Illinois, USA

Abstract

Decision-making with information displays is a key focus of research in areas like human-AI collaboration and data visualization. However, what constitutes a decision problem, and what is required for an experiment to conclude that decisions are flawed, remain imprecise. We present a widely applicable definition of a decision problem synthesized from statistical decision theory and information economics. We claim that to attribute loss in human performance to bias, an experiment must provide the information that a rational agent would need to identify the normative decision. We evaluate whether recent empirical research on AI-assisted decisions achieves this standard. We find that only 10 (26%) of 39 studies that claim to identify biased behavior presented participants with sufficient information to make this claim in at least one treatment condition. We motivate the value of studying well-defined decision problems by describing a characterization of performance losses they allow to be conceived.

CCS Concepts

• **Human-centered computing** → **Laboratory experiments; HCI theory, concepts and models.**

ACM Reference Format:

Jessica Hullman, Alex Kale, and Jason Hartline. 2025. Underspecified Human Decision Experiments Considered Harmful. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*, April 26-May 1, 2025, Yokohama, Japan. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3706598.3714063>

1 Introduction

Identifying effective information displays to support data-driven inference and decision-making is a common goal across a number of domains, including research related to human-computer interaction (HCI). For example, visualization researchers emphasize assisting decision-making as a vital goal of data visualization [13]. Similarly, in human-centered artificial intelligence (AI), a growing body of empirical studies on human decision behavior has been called “necessary to *evaluate* the effectiveness of AI technologies in assisting decision making, but also to form a *foundational understanding* of how people interact with AI to make decisions” [28]. Human decisions from predictions and other information displays are often evaluated in studies with the goal of identifying forms of bias: ways

in which study participants’ behavior deviates from expectations of good decision-making. However, within HCI and related fields there is little consensus on what a decision is, and how decisions are a distinct form of human behavior compared to other tasks.

For example, colloquial definitions describe a decision task as a choice between alternatives, or a choice where the stakes are high (such as deciding which of multiple hurricane forecasts is more accurate [34], or whether to screen child welfare cases [43]). Neither isolates a clear set of assumptions that must hold for a judgment to qualify as a decision. Other accounts cast decision-making as a higher-level task akin to sensemaking [13]. For some, the point of studying human decisions is to compare them to a definition of idealized behavior. Sometimes researchers define this standard in a way that is common across studies (e.g., a better forecast comes closer to the true probability of the event [15, 17]). Other times, however, different definitions of the same concept are used: e.g., appropriate reliance as agreeing with an AI when it is correct and disagreeing when it’s not [4, 6], versus providing a higher subjective rating of trust when the model is correct [7], versus choosing the better of the human versus the AI prediction to bet on for a future task [12]. Sometimes the point is to study seemingly “subjective decisions,” like how people react to music recommendations [26], or decide whether a house seems to match its valuation [35], where it can seem unclear whether normative interpretations of task performance can be applied.

Not only is a more coherent definition of decision needed, as recent survey papers recommend [13, 27], but such a definition must clearly define requirements for *evaluating* decisions by identifying forms of bias or performance loss in human decision-making. In particular, we are concerned with the minimum requirements for a controlled study of human behavior to support claims that the human decision-making is flawed or non-ideal in some way. While it is generally recognized that such evaluative studies require ground truth outcomes to be identifiable for decision instances provided to participants, methodological work on experimental design in HCI and beyond lacks much formal guidance on what constitutes a valid research design for drawing normative conclusions about human decisions.

This paper contributes a definition of a decision problem and associated guidelines for experiment design that apply to any study intended to evaluate human decision-making under the provision of information. We show how a widely applicable framework for studying decision-making can be synthesized from several bodies of work outside HCI. First, foundations of statistical decision theory and expected utility theory [38–40] rigorously define what it means to make good decisions under uncertainty. Second, information economics formalizes information structures that capture the relevant information in a set of signals for a decision problem [5]. We argue



that for a decision problem to be well-defined, it must be possible to define a normative (best response) decision to the information provided to the study participants. In other words, whether the results of a study can be used to argue that humans are biased depends on the answer to the question: Could a rational agent who perfectly perceived the information have used it to identify a best response? Importantly, this best response is calculated relative to the information available to the participant. This makes it distinct from scoring responses directly against the ground truth outcome. As an example of the difference, imagine a study where participants are shown weather forecasts, and asked to make decisions about whether to bring an umbrella or not. Rather than concluding that a participant made the wrong decision to bring the umbrella because a simulated outcome was no rain, they should be scored relative to whether the choice they made maximized their expected utility according to the scoring function the researcher uses to tabulate results. Additionally, this scoring function should have been communicated to them if the researchers want to draw conclusions about how good their decisions were.

We motivate this criterion by showing how performance loss can only be conceived of when a study participant could in theory identify how to optimize their behavior from the information they are provided. If, for example, the researchers had not communicated to participants of the weather forecast study sufficient information about how their responses would be judged, then we should expect different participants to make different guesses about what their goals should be. Consequently, their responses are incomparable, and it is left ambiguous what decision problem they thought they were solving. On the other hand, if they were provided sufficient information on the scoring process and provided forecast information, it is in theory possible to distinguish between a limited set of errors that explain observed performance loss, such as not extracting all of the decision-relevant information from the display, or not being able to identify the best response given the information they did extract.

To assess the extent to which recent empirical evaluations in HCI achieve this ideal, we apply our definition to 46 recent studies on AI-assisted decision-making, finding that while 39 of the studies draw conclusions about shortcomings in human decisions with prediction displays, only 10 (26%) of these studies provided participants with a well-defined decision problem in at least one treatment condition.

The benefits of adopting a more rigorous standard based in established theoretical foundations for decision-making [5, 38] for empirical decision research in HCI, human-centered AI/ML, visualization, and related fields are twofold. First, it stands to improve the validity of interpretations of individual studies. The variation induced when a decision problem is not clearly reported nor communicated to participants confounds our ability to interpret the results, in that it cannot be separated from variation induced by true differences in how people do a task under different conditions. Adopting a standard where the normative decision arises from a well-defined process puts researchers in a better position to point to possible sources of performance loss, and raises awareness of the limitations on what can be concluded from study results when a decision problem is under-communicated. Second, it paves the way for the development of theory. Only when a decision problem is well-defined can human responses to that problem be interpreted

relative to the results of other studies (e.g., that explore variations on the problem). This provides HCI researchers an ability to make meaningful contrasts to drive theory development. Otherwise, any apparent differences in the quality of behavior may reflect noise induced by presenting participants with an underspecified problem.

The paper is organized as follows. First, we define a decision problem, normative decision, and sources of performance loss for human decision-makers. Next, we provide study design recommendations for applying this framework in practice. We present two example study designs in detail, and describe how they could be improved to align with the framework we present. We apply the framework to a sample of studies on AI-assisted decisions, illustrating the gap between the ideal it presents and current practice. We conclude with a discussion of limitations and future work.

2 Defining a Decision Problem

We define a decision problem and the normative decision given such a problem and enumerate possible sources of loss. Our definitions are intended for controlled evaluation of human behavior, a.k.a. *normative* decision research. Behavioral data (which may be produced by humans or by simulation) is collected under controlled conditions, with the goal of learning about behavior induced by the provision of information. Such studies are frequently used to describe the quality of human performance in some situation (e.g., how well people make decisions from displays in strategic settings [42]), to rank different assistive elements by human performance (e.g., different visualizations [10] or AI explanation strategies [31]), or to test a hypothesis about how humans make decisions or what will help them do better (e.g., cognitive forcing functions will improve AI-assisted decisions [6]). Critically, such evaluative studies require that ground truth can be identified for any hypothetical states of the world against which study participants' decisions are evaluated.

2.1 Decision Problem

Our definition of a **decision problem** for controlled study starts by assuming a scenario in which there is some state of the world θ which is uncertain at the time of the decision (**uncertain state**). The state is drawn from a **state space** Θ that describes the set of finite, mutually exclusive values that it can take.

The **data-generating model** is a distribution over values of the uncertain state and *signals* that provide information to the agent about the state. The data-generating model associates the state with the information that the agent is presented with (the **signal**) in order to inform their decision; for example a model prediction or visualization of relevant data. Since the signal is correlated with the state, if the participant understands it in the context of the decision problem, they can improve their performance compared to a case where they do not have access to the signal. More precisely, the data-generating model defines a joint distribution (or information structure) $\pi \in \Delta(\mathcal{V} \times \Theta)$ over signals $v \in \mathcal{V}$ and states $\theta \in \Theta$. This joint distribution assigns to each realization $(v, \theta) \in \mathcal{V} \times \Theta$ a probability, denoted $\pi(v, \theta)$.

If we ignore the signal v , π gives rise to a probability distribution over values of the uncertain state, which we denote as the prior distribution $p \in \Delta(\Theta)$. The prior $p(\theta)$ is equivalent to $\sum_{v \in \mathcal{V}} \pi(v, \theta)$.

When we present a person (hereafter agent) with a decision problem, we give them a choice of response (**action**) a from some set of possible responses (**action space**) \mathcal{A} . The action space can be identical to the state space (as in AI-assisted diagnosis (e.g., [2]) or can be defined on a different space. In decision tasks corresponding to prediction tasks, (e.g., “What’s the probability the patient will be readmitted to the hospital?”), the action space is a probabilistic belief over the state space, i.e., $\mathcal{A} = \Delta(\Theta)$. The agent faced with a decision task must consider the probability of different states induced by the signal, and choose the best action under the **scoring rule**, a function that assigns a score representing the quality of the action for the realized state: $S : \Theta \times \mathcal{A} \rightarrow \mathbb{R}$.

Under this definition, the formalization of the joint distribution π enables a definition of normative behavior. Summarizing information about θ , π (as well as p) and communicating it to the participant provides a basis for the experimenter to establish expectations under optimal use of the information.

2.2 Normative Behavior

Given a decision problem defined as above, we calculate the normative (“optimal”) decision by assuming that the agent has coherent preferences under a scoring rule (dictated by a small set of axioms [40]) and uses them to decide between actions under uncertainty about the outcome. Consequently, we can interpret an experiment participant’s performance as an attempt at achieving this standard, and, if sub-optimal, identify sources of error (loss) in performance.

2.2.1 Rational Belief Formation. To define the normative action for a decision problem, we must first characterize how we would ideally expect the agent to form the posterior beliefs q about the state given the signal v . A theory of optimal belief construction starts by assuming that the ideal (i.e., rational) agent perfectly perceives v . Once perceived, how does v inform the agent’s beliefs about θ ? Generally, we assume that the agent has some baseline set of beliefs about π that they could use to choose an action in the absence of v . Encountering v leads the agent to Bayesian update their beliefs.

Specifically, to calculate the optimal decision for a decision task we first define what the agent perceives the probability distribution over the state to be prior to obtaining information from v : $Pr(\theta)$ or $p(\theta)$ as we describe above. Whenever the signaling policy does not reveal $\pi(\theta|v)$ directly via the signal, but does inform on θ , we assume that after seeing the signal, the agent updates their prior belief to a posterior belief over signals and states, using Bayes rule based on what they know about the data-generating model π . When the signaling policy does reveal $\pi(\theta|v)$ directly via the signal (i.e., $\mathcal{V} \subset \Delta(\Theta)$ and $v = q$), this update is trivial. The maximally informative posterior beliefs that any agent can arrive at given the decision problem is:

$$q(\theta) = \pi(\theta|v) = \frac{\pi(v, \theta)}{\sum_{\theta' \in \Theta} \pi(v, \theta')} \quad (1)$$

It is sometimes unreasonable to expect that π can be communicated to any agent, rational or not, without incurring some error. For example, in many studies that involve machine learning predictions based on many features and possibly many states (e.g., multi-label classification), communicating the joint distribution

over signals and states without loss of information is impractical. Under such conditions, we expect the rational agent to arrive at posterior beliefs defined by:

$$q(\theta) = \hat{\pi}(\theta|v) = \frac{\hat{\pi}(v, \theta)}{\sum_{\theta' \in \Theta} \hat{\pi}(v, \theta')} \quad (2)$$

where $\hat{\pi}$ represents the approximation of π that a rational agent arrives at given the information about π employed by the study, such as a sample of labeled instances presented to approximately convey π .

2.2.2 Defining the Normative Action. Given the posterior beliefs as defined in Equation 1 or 2, we use Equation 4 to identify the action that a fully rational agent would choose in order to maximize their expected utility under S . This is the normative decision.

We assume that the agent’s goal is to maximize their expected score under the scoring rule S :

$$S(a, q) = \mathbb{E}_{\theta \sim q}[S(a, \theta)] \quad (3)$$

where q describes the agent’s belief distribution, the probability distribution over the states that the agent believes the states of the world are drawn from, and $S(a, \theta)$ represents the score assigned by S .

The optimal action is then the action that maximizes the agent’s expected score:

$$a_{opt} = \arg \max_{a \in \mathcal{A}} \mathbb{E}_{\theta \sim q}[S(a, \theta)] \quad (4)$$

2.2.3 Eliciting and Scoring Beliefs. A scoring rule induces a “meaning” on actions by demarcating a range of beliefs for which an action is optimal. Consequently, when different rules are applied to incentivize participants versus to evaluate their responses, the *meaning* of a given action differs. One rule will by definition “misinterpret” the actions of the other by implying a different normative standard. This misinterpretation implies that the scoring rule should be consistent across the two purposes of incentivizing participants and evaluating their responses.

An exception to this requirement occurs when beliefs are elicited instead of decisions, and the rule used to incentivize belief reports has certain properties. Among scoring rules for beliefs, we can distinguish those that are *proper scoring rules*: scoring rules with $\mathcal{A} = \Delta(\Theta)$ for which the optimal action is to predict the true distribution, i.e., $p \in \arg \max_{a \in \mathcal{A}} \mathbb{E}_{\theta \sim p} S(a, \theta)$. Under a proper scoring rule, it is not possible to score higher by reporting some alternative beliefs. These beliefs can then be plugged into a different proper scoring rule without the risk of misinterpretation.

For any non-proper scoring rule $S : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ there is an equivalent *proper scoring rule* $\hat{S} : \Delta(\Theta) \times \Theta \rightarrow \mathbb{R}$ defined by playing the optimal action under the reported belief. Formally,

$$\hat{S}(p, \theta) = S(\arg \max_{a \in \mathcal{A}} S(a, p), \theta). \quad (5)$$

For example, applying the construction in Equation 5 to a binary decision gives a threshold rule where for each state, all reported probabilities below a threshold receive one score and those above receive another score.

2.3 Scope of the definitions

The above definitions are applicable to any studies that aim to produce *normative* conclusions about human behavior. In such

Example: AI-assisted flight booking

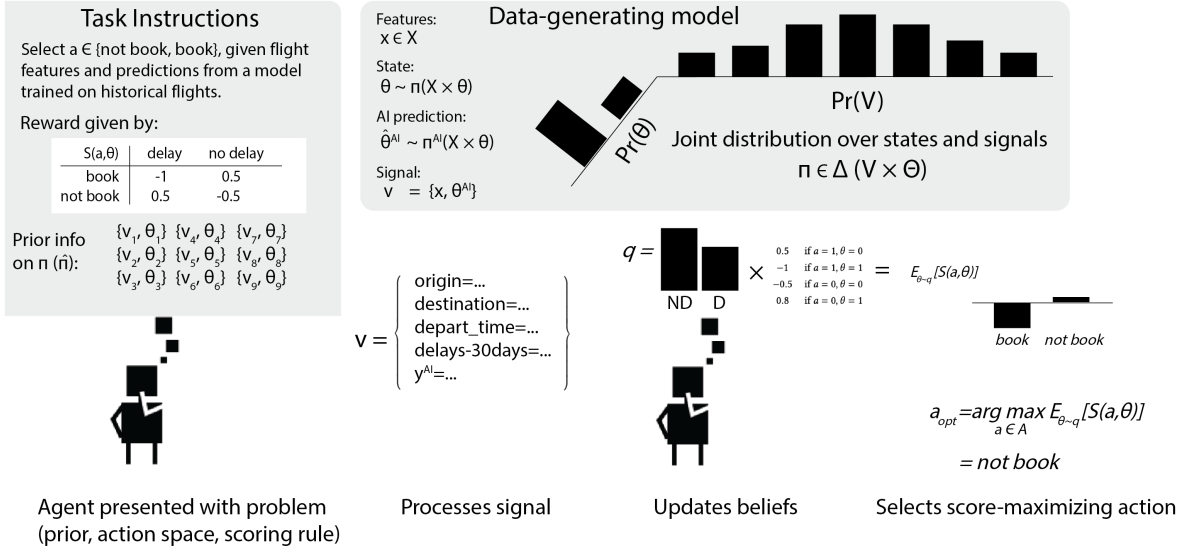


Figure 1: Diagram depicting normative decision for example AI-assisted flight booking scenario. From left to right: The agent is informed of the decision problem, including the action, scoring rule, and prior information about the data-generating model. They next view a signal generated by the data-generating model, which is correlated with the state. The agent updates their beliefs about the state, then chooses the score-maximizing action (in this case, to not book the flight).

studies, behavioral data (produced by humans or simulation) are collected with the goal of learning about how behavior is induced by the provision of information. Such studies are frequently used to describe the quality of human performance in some situation (e.g., how well people make decisions from displays in strategic settings [42]), to rank assistive elements by human performance (e.g., different visualizations [10], AI explanation strategies [31], or prompts [6, 20]), or to test a hypothesis about how humans make decisions [23].

It is important to recognize that even if the researchers do not think about the elements of their experiment in terms of the components identified above (i.e., statespace, data-generating model, signaling policy, action space, scoring rule), the design of their study still induces a decision problem whenever two simple conditions hold [41]. First, there must be some uncertainty about the state (i.e., there is more than one data scenario used to generate stimuli)¹, and second, it must be possible to score decision quality consistently across participants. This requirement eliminates, for example, studies of emotional responses.

2.4 Example

To make these concepts concrete, imagine a study of AI-advised decisions in which people make decisions about whether to book a flight upon seeing features of the flight $\mathbf{x} \in \mathcal{X}$. The state is whether the flight will be delayed beyond some nominal number of minutes, which can take the value 0 for no or 1 for yes ($\theta \in \{0, 1\}$). The action space presented to the agent is whether or not to book the flight ($a \in \{0, 1\}$).

¹A single trial study where all participants view the exact same signal would not qualify.

The data-generating model is the joint distribution over the state (delay or not) and signals comprised of flight feature values and AI prediction information. Assume that participants are presented with a set of feature values describing a flight (origin, destination, time, carrier, prior delays in last 30 days, etc.) from which they must decide on an action. Imagine that the signal also includes the prediction from a model that participants are told was trained on a historical dataset of labeled examples.

For a participant to in theory be able to form posterior beliefs about the probability that a flight is delayed after viewing such signals, they must have prior information about the data-generating model. The joint distribution $\pi \in \Delta(\mathcal{V} \times \Theta)$ assigns a probability, for each possible value of $\theta \in \{0, 1\}$, to each possible signal $\mathbf{v} = \{\mathbf{x}, \hat{\theta}\}$ (a combination of feature values and model prediction) used as a trial in the experiment. Rather than attempting to communicate the exact prior to participants, imagine that before beginning the study trials, participants are provided with a finite set of labeled instances $D = \{(\mathbf{x}_1, \hat{\theta}_1, \theta_1), \dots, (\mathbf{x}_m, \hat{\theta}_m, \theta_m)\}$, which defines a joint distribution $\hat{\pi}$ over $\mathcal{V} \times \Theta$. Note that if we did not show the agent any feature values at all (nor model predictions), we would expect them to make their choice based on the prior probability that a randomly drawn flight in the study will be delayed, $\hat{\pi}(\theta)$.

Imagine the researcher uses a binary scoring rule to map the ground truth realization of the state ($\theta \in \{0, 1\}$ for no delay, delay) and the participant's selected action from the action space ($a \in \{0, 1\}$), say, booking the flight ($a = 1$), to an outcome. For example, the agent might be scored by the following rule, which imposes the greatest penalty on false negatives (i.e., booking the flight when it will be delayed):

$$S(a, \theta) = \begin{cases} 0.5 & \text{if } a = 1, \theta = 0 & \text{book, no delay} \\ -1 & \text{if } a = 1, \theta = 1 & \text{book, delay} \\ -0.5 & \text{if } a = 0, \theta = 0 & \text{not book, no delay} \\ 0.8 & \text{if } a = 0, \theta = 1 & \text{not book, delay} \end{cases} \quad (6)$$

where the score could be in units of USD, for example.

The normative response to this decision problem is defined by the action that a rational agent selects after being presented with the problem (including the action space, scoring rule, and approximation of the data-generating model $\hat{\pi}$), viewing the signal, updating their beliefs, and selecting the score-maximizing action. Figure 1 illustrates this example.

2.5 Defining Sources of Loss

Consider the process implied by definitions in Sections 2.1-2.2, and how each step might result in lost performance relative to a rational agent:

- (1) Agent obtains prior \rightarrow prior loss
- (2) Agent processes signal \rightarrow receiver loss
- (3) Agent updates beliefs \rightarrow updating loss
- (4) Agent takes optimal action \rightarrow optimization loss

For these four sources of loss to explain observed performance losses from study participants, we assume the participants understand the decision problem described to them. However, it always possible that they do not, implying a fifth possible source of loss.

Prior loss (1). The agent must obtain prior beliefs that match the prior beliefs the researchers assume in analyzing the problem. Prior loss is the loss in performance due to the difference between the agent's prior beliefs and those used by the researchers to calculate the normative standard.

Receiver loss (2). The agent processes the signal for the information it provides about the value of the uncertain state. Receiver loss is the loss due to the agent not properly extracting the information from the signal, for example, because the human visual system constrains what information is actually perceived, because the process of rendering the signal induces noise, or because participants do not understand how to read the signal.

Updating loss (3). Next the agent updates their prior beliefs with the information obtained. Updating loss is the loss due to the agent not updating their prior beliefs according to Bayes rule with the information they obtained from the signal, for example due to systematic deviations in human belief-updating relative to rational belief updating as prescribed by Bayes rule (e.g., [9]).

Optimization loss (4). The agent decides by choosing a good action under their posterior beliefs. Optimization loss is the loss in performance due to not identifying the optimal action under the scoring rule. Optimization loss may include what perceptual psychologists call "lapse" responses (e.g., from fatigue or inattention), a catch-all term used to denote responses that do not necessarily reflect the underlying function capturing the person's beliefs. Or it may include "elicitation errors," where what participants report on the response scale does not reflect what they had in mind as a result of bias or imprecision induced by the elicitation method.

Misunderstanding the decision problem. A potential contributor to all four sources of loss above is the agent misunderstanding the decision task, including the scoring rule, the statespace, and/or the action space. Even when all necessary components of a decision problem are carefully communicated, the participant might not read the instructions carefully, or might not understand them. Consequently, any tests that the researcher runs to validate hypotheses about biases stemming from particular sources of performance loss should be thought of as tests of the joint hypothesis comprised of the target hypothesis *and* the hypothesis that participants understood the task.

3 Applicability of the Normative Framework to HCAI Studies

To illustrate the applicability of our definition, we present a meta-study in which we analyzed a sample of recent evaluative studies of AI-assisted decisions to estimate how often researchers draw conclusions about flaws in human decisions without a well-defined decision problem. Note that HCAI studies are just one example of the kinds of decision studies to which our work applies (Section 2.3).

Sample. We sampled 46 of the studies surveyed by Lai et al. [28] in their overview of empirical human-subjects studies on human-AI decision-making. According to their inclusion criteria, the sample contains only evaluative human-subject studies targeting a decision task in the context of classification or regression published in ACM or ACL conference venues between 2018 and 2021.

Codes. In coding the studies, we focused on three aspects of the study design and interpretation. First, we identified whether our framework is *applicable*: Does the task given to participants admit a ground truth state *and* induce uncertainty about the state, at least across participants? Next, we identified whether the authors draw *conclusions about flaws or performance loss in human decisions* from their results, i.e., is the study evaluating human decision-making in some way? We identified whether the study targeted a *well-defined problem*; i.e., whether participants in the study were provided with sufficient information for a rational agent to best respond to the task. Finally, we noted whether the researchers *reported conclusions using a different scoring rule* than the one provided to participants, and whether participants were incentivized using a *flat scoring rule* (i.e., a rule that does not distinguish between responses based on their distance to the ground truth). Note that studies that pay participants a flat reward require that researchers report conclusions using a different scoring rule, introducing confounding.

In some cases, coding was challenging because the paper did not provide much detail on what exactly was communicated to participants, nor provide the study interface as a supplement. This demonstrates a lack of sensitivity on the researchers' part when it comes to providing requisite information to interpret their results. We noted such cases, and discuss how specific omissions are problematic.

Two authors participated in the coding. After formulating the coding scheme, the authors divided the sampled studies roughly evenly, each coding independently, then discussing all challenging examples to ensure consensus. The list of coded studies, our coding guide, and full coding results are available as an Appendix.

Results. Of the 46 studies examined, 39 (85%) studied tasks to which the framework was applicable. One study posed questions and drew conclusions about the diagnostic utility of AI recommendations, but elicited only perceived utility [8]. We coded this example as Applicable. The remaining 7 (15%) included studies relegated to perceptions (e.g., of fairness, accuracy) and of subjective appraisals of music or movie recommendations or emotion.

All 39 of the applicable studies drew conclusions about flaws in human decisions, such as by pointing to over- or underreliance on the AI or speculating about why participants failed to get the correct answer. However, when we evaluated whether participants were given sufficient information to identify the best decision response across all conditions, this was true in only 7 of the 39 (18%) cases. In another 3 of 39 (7.7%), participants were given sufficient information to identify the best decision response in some subset of conditions. In the latter cases, informational differences between conditions that the researchers compared could lead to different a priori expectations of decision quality, yet these differences were not accounted for in comparing participants' behavior between conditions². While the framework was clearly applicable, there was not enough information to determine if participants were presented with a well defined problem in 4 (10%) of the remaining studies. In other words, at best, over 70% of the surveyed studies did not give participants an opportunity to do well by the evaluation approach the researchers applied.

Only 7 (18%) of the 39 studies to which the framework is applicable used the same scoring rule to analyze the responses as was used to incentivize participants, illustrating a lack of awareness of the implications of using different rules for interpreting the meaning of responses. 25 (64%) did not, one did in at least one experiment, and another 6 (15%) did not report enough detail for us to know. Finally, 14 (36%) of the 39 used a flat rule for incentives, one did in some experiments, 16 (36%) did not, and 7 (18%) did not report enough detail for us to know.

We observed that even when authors gave participants sufficient information to identify the best response, they rarely compared the observed performance to the best attainable performance. For example, multiple papers defined a scoring rule to incentivize participants to make accurate decisions, but evaluated the responses using a simpler accuracy scale, implying that 100% accuracy was obtainable even when it would not be to a perfectly-perceiving rational agent. In some cases, participants were also given sufficient information to optimize, but it did not accurately reflect the true data-generating model, such that the results interpretation did not hold. The most common scenario was when the researchers gave participants information about the model's held-out accuracy but the participants experienced instances where the model had a different accuracy. Especially when the study design includes feedback, this leads to a contradiction between the instructions and what participants observe which makes it difficult to draw any evaluative claims about behavior.

Altogether, these results suggest that the framework is widely applicable to swaths of studies currently being conducted in HCI and related fields. Moreover, they suggest that researchers stand

to improve the interpretability of their research by adopting it by following recommendations that it implies, as we describe further below.

4 Recommendations

Our primary argument in this work is that to evaluate human decision-making, an experiment must **provide the participant with enough information to in principle identify the normative decision**, which is used to judge their behavior. In other words, does the experiment give the participant enough information to align their understanding of the decision problem with its normative interpretation?

The primary motivations for using the framework are *epistemological* in nature: they concern what knowledge can be gained from the results of a study. A lack of consensus around how to define a decision problem and the minimum required components for normative analysis leads to research landscape in which researchers can easily be tempted to overinterpret the results of their studies. Readers may not recognize the overinterpretations, leading to a research literature based on invalid claims about human biases.

In contrast, when the experiment provides the participant sufficient information, we can characterize forms of bias as loss in decision performance according to the four possible sources described in Section 2.5. When an experiment does not provide sufficient information for the participant to solve the task in principle, there is the potential for multiplicity concerning the data-generating model, as we cannot assume that the participant understood the data-generating model as the researcher intended. Consequently, researchers wishing to design a normative decision experiment should communicate all necessary components of the decision problem both to study participants and in publications. Specifically, we recommend that researchers:

Clearly convey action and state spaces and a scoring rule.

The finite action space \mathcal{A} and statespace Θ should be clearly communicated to participants, and the action space reflected in the instruments used to collect participant responses. Because the optimal action can only be determined relative to a scoring rule $S : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ that maps the action and state to a quality or payoff, participants should be given a scoring rule to reduce heterogeneity in their choice of action. This is necessary for interpretation even if the experimenter feels comfortable assuming that the participants will try their best and does not use the scores in assigning monetary rewards.

Provide prior and/or sufficient information about the data-generating model to calculate posterior.

As discussed above, sufficient information about the data-generating model must be provided to arrive at the Bayesian posterior distribution. In designing experiments, researchers should recognize that there is often more than one way for a study to provide sufficient information for a rational decision-maker to identify the normative decision. However, the fact that experiment designs can vary in the amount of information they provide participants about π or how they provide this information does *not* mean that anything goes when it comes to what information is provided. Regardless of the specifics, sufficient

²The calculation of separate performance benchmarks would address such ambiguity [41].

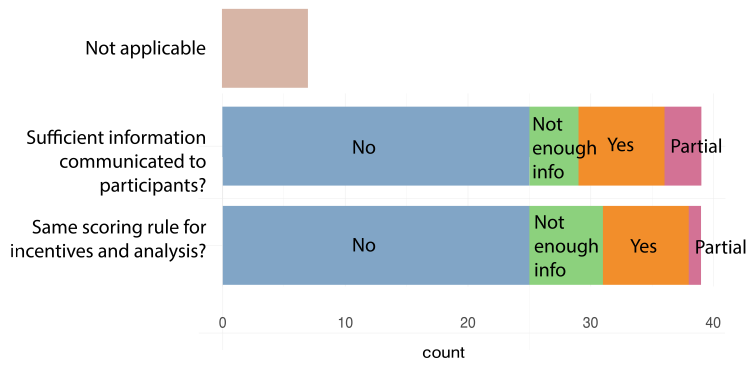


Figure 2: Results from coding 46 studies surveyed by Lai et al. [28]. Seven studies did not evaluate human decisions, instead focusing on capturing perceptions or subjective appraisals. Of the remaining 39, 25 did not communicate sufficient information to participants for them to identify the best response. Additionally, 25 did not use the same scoring rule for incentivizing participants as for analyzing their responses.

information for a rational agent to identify the best response must be provided to interpret the results.

For many experiments, it is critical to convey π or an approximation of π as the prior in order for participants to be able to update these beliefs upon viewing signals. However, for some experiments, it is instead possible for the signal to directly inform of the posterior probability of the state. For example, in a visualization study that asks users to decide whether or not they want to purchase a piece of special equipment expected to improve their score in a game (e.g., [19, 23]), a visualization of the player’s expected score distribution with and without the equipment (from which the realized outcome is drawn depending on their choice) conveys the posterior probability of the state directly.

Whenever the signal on its own does not induce a posterior, the prior should be endowed, and sufficient information about π should be provided to enable estimating $Pr(v|\theta)$ and $Pr(v)$. Exemplar-based training procedures such as have been used to teach humans when to defer to an AI model can be employed for this purpose [32]. More generally, using presentation formats that are more likely to be automatically processed (e.g., conveying distribution information via frequency animations [16, 21]) or validating participants’ priors through elicitation (e.g., [24, 25]) can help researchers understand and minimize prior loss.

Use the same scoring rule for incentivization and evaluation, or elicit beliefs with a proper scoring rule. As described above, using different scoring rules for incentivizing participants versus evaluating their responses means that “good” behavior is defined differently across these two activities, confounding interpretation. Researchers should use the same rule to avoid this confound.

An exception to this rule applies when beliefs are elicited using a proper scoring rule but behavior is analyzed using a different, downstream rule. There are several benefits to eliciting beliefs using a proper scoring rule, in place of or in addition to eliciting decisions. First, when a proper scoring rule is used, the researcher can assume that an agent who understands the rule properly will report their true beliefs. This supports interpretation by eliminating “deviation by design,” referring to how deviation from honest reporting is not technically a loss under a non-proper rule, it should be the

expected behavior under that rule. In other words, in a non-proper scoring rule, we cannot interpret as biased the behavior of a study participant who does not report the appropriate beliefs after examining the signal. Eliciting beliefs using a strictly proper scoring rule enables plugging the elicited beliefs into other scoring rules.³

Second, belief elicitation through proper scoring rules can provide benefits in the form of more information captured in the elicited responses. In general, a coarser action space (i.e., with fewer choices available to participants) will produce less informative responses for the purpose of resolving observed performance relative to a normative standard. For example, if the researcher designed the flight booking study to also elicit the agent’s beliefs, in the form of their prediction of the probability of delay, each trial would provide more information, and they could employ a proper scoring rule such as squared loss (i.e., Brier or quadratic rule) or log score to score the accuracy of the probabilistic beliefs.

When eliciting beliefs, one should also acknowledge that different proper scoring rules imply different expectations about the set of possible decisions [29]. See Li et al. [30] for a discussion of scoring rules that have a high information value for the rational agent.

Compare to best attainable performance. We observed that even when researchers invest effort to design scoring rules to incentivize participants, they do not often compare the observed behavioral performance to the best attainable performance under the rule. Researchers should default to analyzing results in the space of the scoring rule used for incentives, and provide a rationale whenever they deviate.

Account for learning when providing feedback. Some experiment designs support learning about π by providing participants feedback informing them of the realized state after each decision trial. When feedback is provided, we would expect the Bayesian rational decision-maker to update their beliefs according to the information they gain about the joint distribution between signals and states. Imagine that participants in the flight booking study see feedback (i.e., are told the realized state, whether or not the flight

³See, e.g., Wu et al. [41]’s analysis of [23].

was delayed) after each decision they make. Assuming the signal is not revealing of the posterior probability of the state (delayed or not), we would naturally expect the participants to learn about π from this feedback as they complete trials, affecting their posterior beliefs (Equation 2). A challenge in using this approach without providing other information about π (i.e., through signals and the prior) is that how much a participant can learn from any single trial is quite limited. Hence, such designs tend to require many trials if the goal is to study behavior under relatively stable beliefs about the data-generating process.

Elicit measures to estimate specific sources of loss. Assuming the same decision problem has been conveyed to participants as the researchers analyze, it becomes possible to conceive of a set of losses to aid interpretation of results. How to estimate sources of loss in an unbiased way is a rich area for future work. In many cases, we expect separating losses to entail adding carefully constructed questions to the experimental procedure, and/or constructing benchmarks based in rational agent performance to separate sources of loss in observed results. For example, Wu et al. [41] use the notion of a calibrated behavioral agent along with a rational agent benchmark and baseline (expected performance of a rational agent under only the prior) to separate optimization and receiver loss. An important area for future theoretical work is to establish under what conditions the losses we conceptualize can be estimated.

Conduct design analysis. Researchers doing normative decision research can benefit from conducting analyses prior to running the study that simulate the idealized actions of an agent to vet the experiment design (see, e.g., Wu et al. [41]’s pre-study analyses for visualization studies and Guo et al. [18]’s analysis of the potential for complementarity in AI-assisted decisions). Simulation-based design analysis forces the researcher to explicitly define the parameters of their decision problem, which can elucidate what information needs to be communicated to a participant in order to enable an optimal response. It also enables the researcher to test variations on their data-generating model and scoring rule, which can help the researcher realize under what conditions (1) the experiment can have acceptable statistical power and (2) participants can be paid fairly. Design analysis helps the researcher choose conditions (e.g., AI explanations) that provide informative signals for the decision task, narrowing down what kinds of research questions can be answered using a particular decision problem. For example, at a minimum, to be informative of how effectively people use signals, a study design must have sufficient statistical power to separate the rational agent performance with and without the signal.

5 Examples: Using the Framework to Redesign Confounded Studies

We use several examples inspired by the human-centered AI and visualization literature to demonstrate why drawing conclusions about biases in human behavior is misleading in scenarios where sufficient information for a rational agent to optimize is not provided. We discuss how the example study designs could be improved to align with the definitions in Section 2.

5.1 AI-Assisted Flight Booking

We first elaborate the flight booking study mentioned earlier. Imagine the study is designed to investigate whether people over-rely on AI assistance when deciding whether to book a flight or not. The researchers use a standard train-validate-split approach on a large labeled historical data set to train a model that predicts whether or not a flight will be delayed. Each flight in the dataset is represented by a set of feature values, $X = \mathbf{x}$, which comprise part of the signal v shown on each trial. Study participants are randomly assigned to either a control condition where they see the feature values along with only the model’s prediction and what is described as “the model’s confidence in its prediction, ranging between 50% and 100%,” or an AI explanation condition where they are also given additional information about how features in the data contribute to the model’s predictions, e.g., generated by LIME [37]. The model confidence score that is provided is an approximately calibrated probability⁴ of a correct prediction conditional on the feature values, $Pr(correct|X = \mathbf{x})$.

Participants are told the AI model was trained on prior flights whose delay status is known. They are told that the model has an accuracy of 80% (representing the combined false positives and false negatives) on a sample of instances drawn from the test set which comprise the trials in the study.

Each participant makes booking decisions for 20 unique flights. On each trial they are asked to choose whether or not to book the flight ($a \in \{0, 1\}$). Rather than the scoring rule given in Equation 6, imagine a simpler version that assigns different types of errors equivalent costs. Participants are given a base reward of \$10USD, but they are told they will lose \$0.5USD from this initial endowment for every trial where they either book a flight that will be delayed or do not book a flight that will not be delayed. Participants are *not* told how often to expect delays; i.e., the prior p is not endowed, nor is π . Participants are *not* given interim feedback on their decisions.

After collecting data from 500 participants, the researchers model the percentage of trials where the participants incorrectly follow the AI’s advice. They find a statistically significant difference between the AI explanation condition and the control condition, with the rate of incorrectly following the AI’s advice being higher in the former. The researchers interpret this as evidence that AI explanations promote over-reliance.

Interpretation of Results. To assess what can be said about performance loss, we first consider whether the problem is well-defined: could a rational agent use the provided information to select the utility-maximizing decision?

The action space A ($a = 1$ for book, $a = 0$ for not book) is conveyed to participants, as well as the state space Θ ($\theta = 1$ for the flight being delayed, $\theta = 0$ for not). Through the information they are given about the task payouts, they also have knowledge of the scoring rule that assigns a score (reward amount) given the realized state and their decision.

⁴A confidence function $f_B : X \rightarrow [0, 1]$ is α -calibrated with respect to a sample $S \subseteq X$ if there exists some $S' \subseteq S$, with $|S'| \geq (1 - \alpha)|S|$, such that for any $b \in [0, 1]$: $|P(Y = 1 | f_B(X) = b, X \in S') - b| \leq \alpha$ [11]. In other words, the probabilities are approximately calibrated with respect to a sample S when for some subset at least $1 - \alpha$ times the size of S , the difference between the stated and real probability is less than or equal to α .

How would a rational participant form beliefs about the posterior probability of the state? Participants are not provided with information about the prior probability of the state. This would be acceptable if the signal provided them with posterior beliefs about the uncertain state. However, the signal does not provide the required posterior probability: recall that participants receive the estimated probability that the model is correct conditional on the features, $Pr(\text{correct}|X = \mathbf{x})$, but not conditional on the model's prediction. They therefore cannot obtain sufficient information from the data-generating model π to guide their choice of action.

Improving the experimental design. One possibility is to make the signal provide the posterior $q(\theta)$. This could be achieved if the confidence value included in the signal were to give the probability that the model is correct conditional on both the features and the prediction, e.g., $Pr(\text{correct}|X = \mathbf{x}, y_{AI} = 1)$ when the model predicts 1. In this setting, provision of the prior would not be necessary (and may confuse participants).

However, the careful reader may also have noticed another barrier to defining bias for the problem: the model's conditional confidence is described as simply "the model's confidence in its prediction", making it ambiguous from the participant's perspective whether it is conditional on the features and model prediction or only one of the two. The normative response is not well-defined if the probability is not conditional on both, unless the researchers also communicate π or an approximation $\hat{\pi}$ to participants before they begin the decision trials. Hence the model probability description would need to be amended to specify that the confidence is conditional on the predicted label to eliminate ambiguity.

It is worth noting that decision problems do not necessarily have to depend on small, discrete action spaces. Imagine another amendment to the flight booking study where instead of asking participants for a booking decision, the researchers ask them to report the probability that the flight is delayed $Pr(\theta|v)$. Assuming the above changes are made to the design to equip participants with information to form posterior beliefs, participants could in principle report the optimal beliefs. However, if the scoring rule for a decision problem where the action space is over beliefs is not proper, we should not expect the agent to respond with their true beliefs. Optimization loss might be low, but this is not useful to the experimenters' research questions if agents respond with a distortion of their true beliefs.

5.2 Election Forecast Displays

Imagine a study targeting the effect of different ways of visualizing election forecasts. The researchers develop a forecasting model to predict vote share (a percentage) by weighting and aggregating poll results from news sources for ten unique upcoming two-party state-level elections. Participants are randomly assigned to view one of three forecast displays showing the predicted vote share with uncertainty of each party's candidate as a table, an animated graphic, or a static graphic. Assume for simplicity that all three induce informationally equivalent signals to a rational agent⁵.

Upon beginning the experiment, participants are randomly assigned to a display treatment and also assigned a political party preference. They are told they will view election forecasts from a model based on poll results and should imagine they live in the district where the election takes place. Each participant in the study does ten trials corresponding to the ten elections. On each trial, they are asked if they would like to vote for the candidate in the election. Voting carries no cost.

Participants are told they will win a given base reward amount plus an additional \$0.25USD "for each trial where their candidate wins a simulated version of the election that uses the forecast model to predict how people will vote." The researchers generate outcomes by randomly drawing a bivariate vote share outcome from the predicted distribution under the model and generating the uncertain state θ (taking the value 0 or 1 for lose or win) by assuming the higher vote share candidate wins. Participants are *not* given feedback after each trial.

After collecting results for several hundred participants, the researchers fit a linear model that estimates the sensitivity of the vote decision to a variable denoting the closeness of the race, operationalized as the difference between the predicted vote share of the first candidate and 0.5. The researchers postulate that a better display induces more sensitivity. Using statistical significance testing, they reject the null hypothesis of no difference between the estimated slopes for the interaction between display type and race closeness across display types. Using posthoc comparisons, they declare that the animated graphic leads to better voting decisions because it induces the most sensitivity.

Interpretation of Results. To understand why the researchers' conclusions might be misleading as a prediction of real world use of election displays, consider the decision problem from the perspective of the participants. The action space \mathcal{A} ($a = 1$ for vote, $a = 0$ for not vote) is conveyed to participants, and the task description also gives participants the statespace Θ ($\theta = 1$ for their candidate winning, $\theta = 0$ for not). Participants can infer the probability of their candidate winning from the signal showing the model predicted vote share. Because the signal provides the posterior probability of the state, no prior information needs to be communicated.

What do participants know about the realized outcome of the election that the researchers describe simulating? The researchers described the use of the forecast model to simulate voting behavior, and in implementation, do *not* account for the participant's vote in generating the outcome. Consequently, whether a participant votes is irrelevant to the payoff-relevant state under the scoring rule. However, the description provided to participants is ambiguous about whether the participant's vote counts. Some participants might assume, based on their knowledge of real-world elections, that it does. Hence performing well in the experiment requires being able to guess what the researchers have in mind, rather than behaving rationally in the context of the experiment.

Because, in the context of the study, voting carries no value but also no cost, we cannot conceive of the four sources of loss. There is no well-defined decision problem, and consequently, no formal justification for concluding that behavior is more biased or otherwise worse under one display condition.

⁵If any of the displays was not informationally equivalent to the others, the optimal decision may vary by display. See Wu et al. [41] for analysis of such examples using a rational agent framework.

Improving the experimental design. We could create a valid decision problem by changing the scoring rule so that the reward amount depends on whether one votes. For example, we could assign a cost of \$0.25 to voting:

$$S(a, \theta) = \begin{cases} 0 & \text{if } a = 0, \theta = 0 & \text{do not vote, lose election} \\ 0.5 & \text{if } a = 0, \theta = 1 & \text{do not vote, win election} \\ -0.25 & \text{if } a = 1, \theta = 0 & \text{vote, lose election} \\ 0.25 & \text{if } a = 1, \theta = 1 & \text{vote, win election} \end{cases} \quad (7)$$

However, the decision problem does not admit a well-defined normative response until we address the ambiguity in the description of the data-generating model. Only if participants are given a transparent description of how the realized outcome is generated from the model can we conceive of prior loss as well as updating loss, both of which should be zero given the signal design. We can conceive of receiver loss as the loss from participants not being able to extract the information from the signal, and optimization loss as the loss from participants not being able to choose the optimal action under their beliefs.

Of course, if we amend the description to inform participants of how the election outcome is actually generated, we should expect participants who understand the problem to never vote. Rather than informed political engagement, the proportion of participants who vote (e.g., in each treatment condition) could be described as reflecting the level of confusion the participants had about the task. Or, if the researchers did not change the description of how the realized state is generated, but designed the study to give participants feedback on the realized state after each decision, then the participants might in theory learn that their vote does not matter. This would work against the researchers' goals of presenting a situation where people could feel compelled to vote.

This example demonstrates that providing participants sufficient information to determine the utility-maximizing action is necessary but not sufficient for a meaningful decision problem in light of the research questions of interest.

6 Discussion: Goals and Incentives in the Lab versus the World

By arguing that researchers adopt and sufficiently communicate to participants a widely-applicable definition of a decision problem, our goal is to improve the integrity of claims made from human subjects experiments. When participants are not given ample information to “solve” at least in principle the decision problem, we cannot bring well-formed expectations to the results, and should expect heterogeneity in the behaviors they exhibit. Such “heterogeneity by design” renders the results uninterpretable for evaluative purposes like identifying forms of bias.

Imagine that the researchers who designed the voting task intentionally did not include the vote in the scoring rule, and were intentionally ambiguous about the role of the vote in the data-generating model. Because the utility of voting in a real-world election is typically known only implicitly at an individual level and likely to vary between people, the researchers might think that their results would therefore be more representative of real world voting behaviors.

Such beliefs stem from a convention among many researchers toward being optimistic about how their experimental observations relate to people's behavior in the “real world.” Most controlled decision studies are designed to help us understand how people behave in some realistic setting that the experimental task is meant to proxy for. However, in the real world, people have goals and prior beliefs. We might not be able to perceive what utility function each individual person is using, but we can assume that behavior is goal-directed in one way or another. Savage's axioms [39] and the derivation of expected utility theory tell us that for behavior to be rationalizable, a person's choices should be consistent with their beliefs about the state and the payoffs they expect under different outcomes.

Consequently, when people are in an experiment, we should not expect the analogous real world goals and beliefs for that kind of task to apply. For example, people might take actions in the real world for intrinsic value—e.g., voting in order to feel like a good citizen, or paying close attention to features of a flight they are deciding whether to book because the decision will impact whether or not they arrive in time for an important event. It is difficult to motivate people to take actions based on intrinsic value in an experiment, unless the experiment is designed to elicit and describe those intrinsic values, in which case the normative framework here no longer applies. Hence, the idea that as experimenters we can skip fully describing the decision problem to our participants while still retaining the ability to identify their behavior as faulty or biased arises from a failure to recognize our fundamental uncertainty about how the experimental context relates to the real world.

6.1 Implications for Research Ethics

The requirement of communicating to study participants sufficient information for a rational agent to optimize may seem unnecessarily strict if one assumes that participants' domain knowledge alone should enable them to evaluate the information about the state captured in signals. For example, in the flight booking example, a researcher may be tempted to assume that even without being given any information about the probability that the model prediction is correct, a participant will be able to use their prior knowledge about features of delayed flights to infer the probability that the model prediction is correct. The problem is that once a researcher makes this assumption, they are no longer evaluating participants on only the information they provided. They are also implicitly evaluating them on how well their prior beliefs about the state and the validity of the model's learned information align with the researcher's choice of data-generating model. There are various implications of this misalignment that we may find problematic upon reflection. For example, under such conditions, participants must bring identical beliefs to the researchers' in order to have a chance of scoring well. A related point has been made in research on crowdsourcing, where it has been observed that there are often many plausible ways to interpret a crowdsourced task based on the information workers are provided, but work requesters often use quality control algorithms that recognize and pay only one of these [1, 22].

Fortunately, this does not mean that human decisions can only be evaluated by assuming that participants bring no additional information. Instead, the data-generating model and decision task should account for the presence of human prior information. For example, for AI-assisted decision tasks like the flight booking example (2.4), a human prediction of the delay status of a flight, governed by a distinct generating process from the AI prediction, can be included in the data-generating process. The decision problem for which the normative decision is sought can then be cast as a decision between the human and AI prediction. If independent human predictions are not available, they can be elicited from the subjects before providing the signal. Recent work by Guo et al. [18] demonstrates this approach.

6.2 Limitations & Future Work

Our argument is that the interpretation of results is restricted when a well-defined decision problem is not conveyed to participants, not that researchers cannot conduct comparisons at all without a well-defined decision. Researchers can conduct and report on studies that do not conform to this approach, but they cannot conclude that people made bad decisions if they were not given sufficient information to know what a good decision is. For example, imagine a study comparing the decision performance of users with and without the use of recommendations from an AI, where no prior is given to participants, and no scoring rule is given. Imagine that performance in the study is measured in terms of accuracy, and average accuracy is higher among participants who had access to the AI. Given an appropriately designed randomized experiment, the researcher could conclude that there is a causal effect of the AI on decisions. However, they could not claim that people are making worse or biased decisions without the AI, because they are scoring their decisions using information they never provided to participants. When we say that one condition did better, we should acknowledge that we did not give them a chance to fully understand and optimize their behavior for the problem that we analyzed.

A common misinterpretation of the use of a rational agent model is that we are implying that human behavior is expected to be perfectly rational, or approximately rational. This is not the case. While a rational agent approach formally justifies what it means to best use information to inform the selection of an action using statistical decision theory, participants in our experiment may use different strategies. The point is that without a definition of what it means to best respond to the information we have provided, we are hard-pressed to evaluate participant behavior. Even if we do not expect this procedure to characterize what human agents do, it provides a basis for interpreting the results in terms of deviations from a well-understood standard.

Several challenges have been raised against expected utility theory as put forth by Savage [38]; we recommend [39] for an overview. Another critique of a normative approach dubbed “value collapse” [33] is that presenting a clear set of values as a target for behavior limits variance in ways that can be harmful. Our argument is that we should not mistake the variation that emerges from an under-communicated decision problem as representative of behavior under a particular set of values or a particular real-world

decision task. If a researcher wishes to evaluate decisions but contests Bayesian decision theory as a normative standard, they should be prepared to motivate whatever standard they adopt instead with equivalent rigor (e.g., derivation from interpretable assumptions combined with mathematical laws). We note that alternative models have been well-explored, but rarely achieve provable grounding relative to mathematical laws as statistical decision theory does.

Sometimes researchers are explicitly interested in understanding the range of values at play in a naturalistic decision scenario rather than providing a well-defined decision problem. In this case, their research is better characterized as descriptive, and they should refrain from claims that hinge on evaluating participant decisions. While our framework is intended for normative decision research, the definitions may also be useful if one is doing such descriptive or exploratory decision research, by making clear which aspects of the decision problem are underspecified and therefore presumably the object of study. For example, a researcher might aim to characterize challenges in model-assisted decisions in a particular domain [43], or conduct an exploratory study to elicit information about the utility function (i.e., scoring rule) that operates in a setting and then communicate it back to new participants as part of the decision problem in an evaluative experiment (e.g., [14]).

Some have voiced pessimistic views on the value of monetary incentives for experiment participants (e.g., [36]). As Wu et al. [41] show, this view is to be expected under a rational agent approach whenever the value of information to the decision problem (i.e., the difference in the rational agent’s expected score given only the prior versus with access to the signals [41]) is small. This motivates analyzing whether a decision study is likely to provide agents with a sufficient incentive to motivate them to consult the signals [41].

Of course, scoring rules can sometimes be challenging for study participants to interpret. Theoretic incentive compatibility does not guarantee understanding on the part of subjects. There is a rich literature discussing possible descriptions that can be used with common rules like the quadratic scoring rule (e.g., [3]). Interfaces researchers may be in a good position to extend this literature through careful observational studies. More broadly, however, the insensitivity of human participants to monetary incentives are orthogonal to the crux of our argument, which is philosophical and concerned with what is knowable more so than the pragmatics of how much behavior changes when a problem is well-defined.

7 Conclusion

We present definitions of a decision problem and normative response synthesized from statistical decision theory and information economics for the purpose of improving the interpretability of decision studies in HCI, visualization, human-centered AI, and related fields. We demonstrate the applicability of the framework to recent HCAI studies, and provide specific recommendations to researchers and examples of how confounded study designs can be made interpretable or well-formed within the framework. Ultimately, the goal of our work is to give researchers the tools to recognize that when we do experiments, we are restricted to learning about behavior in the artificial worlds we create. As much as we might want to equate our results with some real world setting used as motivation, extrapolating from the world of the controlled experiment to the real

world will always be a leap of faith. It is therefore the responsibility of the experimenter to fully understand their experimental world, and, for the purposes of offering the results as a form of potentially generalizable knowledge, to give their participants a chance to do so as well.

References

- [1] Ali Alkhatib and Michael Bernstein. 2019. Street-level algorithms: A theory at the gaps between policy and decisions. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13.
- [2] Rohan Alur, Manish Raghavan, and Devavrat Shah. 2024. Distinguishing the Indistinguishable: Human Expertise in Algorithmic Prediction. *arXiv preprint arXiv:2402.00793* (2024).
- [3] Florian Artinger, Filippas Exadaktylos, Hannes Koppel, and Lauri Sääksvuori. 2010. *Applying quadratic scoring rule transparently in multiple choice settings: A note*. Technical Report. Jena Economic Research Papers.
- [4] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the whole exceed its parts? the effect of ai explanations on complementary team performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [5] Dirk Bergemann and Stephen Morris. 2019. Information design: A unified perspective. *Journal of Economic Literature* 57, 1 (2019), 44–95.
- [6] Zana Bućina, Maja Barbara Malaya, and Krzysztof Z Gajos. 2021. To trust or to think: cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–21.
- [7] Adrian Bussone, Simone Stumpf, and Dymprna O'Sullivan. 2015. The role of explanations on trust and reliance in clinical decision support systems. In *2015 international conference on healthcare informatics*. IEEE, 160–169.
- [8] Carrie J Cai, Emily Reif, Narayan Hegde, Jason Hipp, Been Kim, Daniel Smilkov, Martin Wattenberg, Fernanda Viegas, Greg S Corrado, Martin C Stumpe, et al. 2019. Human-centered tools for coping with imperfect algorithms during medical decision-making. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–14.
- [9] Colin Camerer. 1995. Individual decision making. *Handbook of experimental economics* (1995).
- [10] Michael Correll, Dominik Moritz, and Jeffrey Heer. 2018. Value-suppressing uncertainty palettes. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [11] Nina Corvelo Benz and Manuel Rodriguez. 2024. Human-aligned calibration for ai-assisted decision making. In *Advances in Neural Information Processing Systems*, Vol. 36.
- [12] 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.
- [13] Evanthis Dimara and John Stasko. 2021. A critical reflection on visualization research: Where do decision making tasks hide? *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 1128–1138.
- [14] Michael Fernandes, Logan Walls, Sean Munson, Jessica Hullman, and Matthew Kay. 2018. Uncertainty displays using quantile dotplots or cdfs improve transit decision-making. In *Proceedings of the 2018 CHI conference on human factors in computing systems*. 1–12.
- [15] Riccardo Fogliato, Alexandra Chouldechova, and Zachary Lipton. 2021. The impact of algorithmic risk assessments on human predictions and its analysis via crowdsourcing studies. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–24.
- [16] Daniel G Goldstein and David Rothschild. 2014. Lay understanding of probability distributions. *Judgment and Decision making* 9, 1 (2014), 1–14.
- [17] Ben Green and Yiling Chen. 2019. Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. In *Proceedings of the conference on fairness, accountability, and transparency*. 90–99.
- [18] Ziyang Guo, Yifan Wu, Jason D Hartline, and Jessica Hullman. 2024. A Decision Theoretic Framework for Measuring AI Reliance. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*. 221–236.
- [19] Jake M Hofman, Daniel G Goldstein, and Jessica Hullman. 2020. How visualizing inferential uncertainty can mislead readers about treatment effects in scientific results. In *Proceedings of the 2020 chi conference on human factors in computing systems*. 1–12.
- [20] Jessica Hullman, Matthew Kay, Yea-Seul Kim, and Samana Shrestha. 2017. Imagining replications: Graphical prediction & discrete visualizations improve recall & estimation of effect uncertainty. *IEEE transactions on visualization and computer graphics* 24, 1 (2017), 446–456.
- [21] Jessica Hullman, Paul Resnick, and Eytan Adar. 2015. Hypothetical outcome plots outperform error bars and violin plots for inferences about reliability of variable ordering. *PLoS one* 10, 11 (2015), e0142444.
- [22] Sanjay Kairam and Jeffrey Heer. 2016. Parting crowds: Characterizing divergent interpretations in crowdsourced annotation tasks. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 1637–1648.
- [23] Alex Kale, Matthew Kay, and Jessica Hullman. 2020. Visual reasoning strategies for effect size judgments and decisions. *IEEE transactions on visualization and computer graphics* 27, 2 (2020), 272–282.
- [24] Yea-Seul Kim, Paula Kayongo, Madeleine Grunde-McLaughlin, and Jessica Hullman. 2020. Bayesian-assisted inference from visualized data. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 989–999.
- [25] Yea-Seul Kim, Logan A Walls, Peter Krafft, and Jessica Hullman. 2019. A bayesian cognition approach to improve data visualization. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–14.
- [26] Todd Kulesza, Simone Stumpf, Margaret Burnett, and Irwin Kwan. 2012. Tell me more? The effects of mental model soundness on personalizing an intelligent agent. In *Proceedings of the sigchi conference on human factors in computing systems*. 1–10.
- [27] Vivian Lai, Chacha Chen, Q Vera Liao, Alison Smith-Renner, and Chenhao Tan. 2021. Towards a science of human-ai decision making: a survey of empirical studies. *arXiv preprint arXiv:2112.11471* (2021).
- [28] Vivian Lai, Chacha Chen, Alison Smith-Renner, Q Vera Liao, and Chenhao Tan. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. 1369–1385.
- [29] Benjamin Anders Levinstein. 2017. A pragmatist's guide to epistemic utility. *Philosophy of Science* 84, 4 (2017), 613–638.
- [30] Yingkai Li, Jason D Hartline, Liren Shan, and Yifan Wu. 2022. Optimization of scoring rules. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 988–989.
- [31] Han Liu, Vivian Lai, and Chenhao Tan. 2021. Understanding the effect of out-of-distribution examples and interactive explanations on human-ai decision making. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021), 1–45.
- [32] Hussein Mozannar, Arvind Satyanarayan, and David Sontag. 2022. Teaching humans when to defer to a classifier via exemplars. In *Proceedings of the aaai conference on artificial intelligence*, Vol. 36. 5323–5331.
- [33] C. Thi Nguyen. 2023. Value Collapse. Presented as the Annual Cardiff Lecture of the Royal Institute of Philosophy. <https://royalinstitutephilosophy.org/event/value-collapse/>
- [34] Lace M Padilla, Grace Hansen, Ian T Ruginski, Heidi S Kramer, William B Thompson, and Sarah H Creem-Regehr. 2015. The influence of different graphical displays on nonexpert decision making under uncertainty. *Journal of Experimental Psychology: Applied* 21, 1 (2015), 37.
- [35] Snehal Prabhudesai, Leyao Yang, Sumit Asthana, Xun Huan, Q Vera Liao, and Nikola Banovic. 2023. Understanding Uncertainty: How Lay Decision-makers Perceive and Interpret Uncertainty in Human-AI Decision Making. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 379–396.
- [36] Daniel Read. 2005. Monetary incentives, what are they good for? *Journal of Economic Methodology* 12, 2 (2005), 265–276.
- [37] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*. 1135–1144.
- [38] Leonard J Savage. 1972. *The foundations of statistics*. Courier Corporation.
- [39] Katie Steele and H Orri Stefánsson. 2015. Decision theory. (2015).
- [40] John Von Neumann and Oskar Morgenstern. 2007. *Theory of games and economic behavior (60th Anniversary Commemorative Edition)*. Princeton university press.
- [41] Yifan Wu, Ziyang Guo, Michails Mamakos, Jason Hartline, and Jessica Hullman. 2023. The Rational Agent Benchmark for Data Visualization. *IEEE Transactions on Visualization and Computer Graphics (forthcoming)* (2023).
- [42] Dongping Zhang, Jason Hartline, and Jessica Hullman. 2023. Designing Shared Information Displays for Agents of Varying Strategic Sophistication. *arXiv preprint arXiv:2310.10858* (2023).
- [43] Alexandra Zytek, Dongyu Liu, Rhema Vaithianathan, and Kalyan Veeramachaneni. 2021. Sibyl: Understanding and addressing the usability challenges of machine learning in high-stakes decision making. *IEEE Transactions on Visualization and Computer Graphics* 28, 1 (2021), 1161–1171.

8 Appendix

8.1 List of coded studies

We randomly sampled and coded the following studies from Lai et al. [27]:

1. A. Abdul, C. von der Weth, M. Kankanhalli, B. Lim. (2020). COGAM: Measuring and Moderating Cognitive Load in Machine Learning Model Explanations.
2. A. Alqaraawi, M. Schuessler, P. Weiß, E. Costanza, N. Berthouze. (2020). Evaluating Saliency Map Explanations for Convolutional Neural Networks: A User Study.
3. G. Bansal, T. Wu, J. Zhou, R. Fok, B. Nushi, E. Kamar, M. T. Ribeiro, and D. Weld. (2021). Does the whole exceed its parts? the effect of AI explanations on complementary team performance.
4. G. Bansal, B. Nushi, E. Kamar, W.S. Lasecki, D.S. Weld, E. Horvitz. (2019). Updates in Human-AI Teams: Understanding and Addressing the Performance/Compatibility Tradeoff.
5. G. Bansal, B. Nushi, E. Kamar, W.S. Lasecki, D.S. Weld, E. Horvitz. (2019). Beyond Accuracy: The Role of Mental Models in Human-AI Team Performance.
6. R. Binns, M. Van Kleek, M. Veale, Ulrik Lyngs, J. Zhao, N. Shadbolt. (2018). 'It's Reducing a Human Being to a Percentage': Perceptions of Justice in Algorithmic Decisions.
7. Z. Buçinca, M.B. Malaya, K.Z. Gajos. (2021). To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making.
8. Z. Buçinca, P. Lin, K.Z. Gajos, E.L. Glassman. (2020). Proxy Tasks and Subjective Measures Can Be Misleading in Evaluating Explainable AI Systems.
9. A. Bussone, S. Stumpf, D. O'Sullivan. (2015). The Role of Explanations on Trust and Reliance in Clinical Decision Support Systems.
10. C. Cai, E. Reif, N. Hegde, J. Hipp, B. Kim, D. Smilkov, M. Wattenberg, F. Viegas, G.S. Corrado, M.C. Stumpe, M. Terry. (2019). Human-Centered Tools for Coping with Imperfect Algorithms during Medical Decision-Making.
11. S. Carton, Q. Mei, P. Resnick. (2020). Feature-Based Explanations Don't Help People Detect Misclassifications of Online Toxicity.
12. A. Chandrasekaran, V. Prabhu, D. Yadav, P. Chattopadhyay, D. Parikh. (2018). Do Explanations make VQA Models more Predictable to a Human?
13. H-F. Cheng, R. Wang, Z. Zhang, F. O'Connell, T. Gray, F.M. Harper, H. Zhu. (2019). Explaining Decision-Making Algorithms through UI: Strategies to Help Non-Expert Stakeholders.
14. M. Chromik, M. Eiband, F. Buchner, A. Krüger, A. Butz. (2021). I Think I Get Your Point, AI! The Illusion of Explanatory Depth in Explainable AI.
15. D. Das and S. Chernova. (2020). Leveraging Rationales to Improve Human Task Performance.
16. B.J. Dietvorst, J.P. Simmons, C. Massey. (2018). Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them.
17. J. Dodge, Q.V. Liao, Y. Zhang, R. Bellamy, C. Dugan. (2019). Explaining Models: An Empirical Study of How Explanations Impact Fairness Judgment.
18. J. Dressel, H. Farid. (2018). The accuracy, fairness, and limits of predicting recidivism.
19. K. Gero, Z. Ashktorab, C. Dugan, Q. Pan, J. Johnson, M. Ruiz, S. Miller, D. Millen and W. Geyer. (2020). Mental Models of AI Agents in a Cooperative Game Setting.
20. B. Ghai, Q.V. Liao, Y. Zhang, R. Bellamy, K. Mueller. (2021). Explainable Active Learning (XAL): An Empirical Study of How Local Explanations Impact Annotator Experience.
21. B. Green, Y. Chen. (2019). Disparate Interactions: An Algorithm-in-the-loop Analysis of Fairness in Risk Assessments.
22. S. Guo, F. Du, S. Malik, E. Koh, S. Kim, Z. Liu, D. Kim, H. Zha, and N. Cao. (2019). Visualizing Uncertainty and Alternatives in Event Sequence Predictions.
23. P. Hase, M. Bansal. (2020). Evaluating Explainable AI: Which Algorithmic Explanations Help Users Predict Model Behavior?
24. R. Kocielnik, S. Amershi, P.N. Bennet. (2019). Will you Accept an Imperfect AI? Exploring Designs for Adjusting End-user Expectations of AI Systems.
25. T. Kulesza, S. Stumpf, M. Burnett, S. Yang, I. Kwan, W-K. Wong. (2013). Too Much, Too Little, or Just Right? Ways Explanations Impact End Users' Mental Models.
26. J. Kunkel, T. Donkers, L. Michael, C-M. Barbu, J. Ziegler. (2019). Let Me Explain: Impact of Personal and Impersonal Explanations on Trust in Recommender Systems.
27. I. Lage, E. Chen, J. He, M. Narayanan, B. Kim, S. Gershman, F. Doshi-Velez. (2019). An Evaluation of the Human-Interpretability of Explanation.
28. A. Levy, M. Agrawal, A. Satyanarayan, D. Sontag. (2021). Assessing the Impact of Automated Suggestions on Decision Making: Domain Experts Mediate Model Errors but Take Less Initiative.
29. V. Lai, C. Tan. (2019). On Human Predictions with Explanations and Predictions of Machine Learning Models: A Case Study on Deception Detection.
30. Z. Lin, J. Jung, S. Goel, J. Skeem. (2020). The limits of human predictions of recidivism.
31. J.M. Logg, J. Minson, D.A. Moore. (2019). Algorithm Appreciation: People Prefer Algorithmic To Human Judgment.
32. Lu, Z., M. Yin. (2021). Human reliance on machine learning models when performance feedback is limited: Heuristics and risks.
33. K. Mallari, K. Inkpen, P. Johns, S. Tan, D. Ramesh, E. Kamar. (2020). Do I Look Like a Criminal? Examining how Race Presentation Impacts Human Judgement of Recidivism.
34. S. McGrath, P. Mehta, A. Zytek, I. Lage, H. Lakkaraju. (2020). When Does Uncertainty Matter?: Understanding the Impact of Predictive Uncertainty in ML Assisted Decision Making.
35. D. Nguyen. (2018). Comparing Automatic and Human Evaluation of Local Explanations for Text Classification.
36. J.S. Park, R. Barber, A. Kirlik, K. Karahalios. (2019). A Slow Algorithm Improves Users' Assessments of the Algorithm's Accuracy.
37. F. Poursabzi-Sangdeh, D.G. Goldstein, J.M. Hofman, J. Wortman Vaughan, H. Wallach. (2021). Manipulating and Measuring Model Interpretability.
38. E. Rader, K. Cotter, J. Cho. (2018). Explanations as Mechanisms for Supporting Algorithmic Transparency.
39. A. Smith-Renner, R. Fan, M. Birchfield, T. Wu, J. Boyd-Graber, D. Weld, L. Findlater. (2020). No Explainability without Accountability: An Empirical Study of Explanations and Feedback in Interactive ML.

40. A. Springer, S. Whittaker. (2019). Progressive Disclosure: Designing for Effective Transparency.
41. K. Stowers, N. Kasdaglis, M. Rupp, J. Chen, D. Barber, and M. Barnes. (2017). Insights into Human-Agent Teaming: Intelligent Agent Transparency and Uncertainty.
42. M. Szymanski, M. Millecamp, K. Verbert. (2021). Visual, textual or hybrid: the effect of user expertise on different explanations.
43. N. van Berkel, J. Goncalves, D. Russo, S. Hosio, M.B. Skov. (2021). Effect of Information Presentation on Fairness Perceptions of Machine Learning Predictors.
44. X. Wang, M. Yin. (2021). Are Explanations Helpful? A Comparative Study of the Effects of Explanations in AI-Assisted Decision-Making.
45. M. Yin, J. Wortman Vaughan, H. Wallach. (2019). Understanding the Effect of Accuracy on Trust in Machine Learning Models.
46. K. Yu, S. Berkovsky, R. Taib, J. Zhou, F. Chen. (2019). Do I trust my machine teammate?: an investigation from perception to decision.

8.2 Coding guide

Our coding approach is designed to separate the following questions:

- Does the study include a task for which a ground truth response can be defined?
- Does the study give the participant a decision problem that a rational agent could best respond to?
- Do the researchers interpret the results assuming the same decision problem given to participants, or do they assume a different problem (e.g., by changing the scoring rule when they evaluate responses)?

Through iterative collaborative coding, the first two authors developed five codes:

- (1) **Applicability of framework:** Do the research questions concern a type of task for which a ground truth state can be defined?
- (2) **Evaluative conclusions:** Do the authors draw *conclusions about flaws or performance loss in human decisions* from their results, such as ranking different conditions, or implying that participants could have done better in any condition?
- (3) **Well-defined problem:** Does the task presented to participants represent *a well-defined decision problem*? In other words, are participants in the study were provided with sufficient information for a rational agent to best respond to the task?
- (4) **Same rule (problem) reported?** Do the authors interpret the results assuming the same scoring rule given to participants, such that the results bear on the same decision problem given to participants?
- (5) **Flat rule?** Does the rule used to incentivize participants assign the same score to all responses regardless of the ground truth?

When the framework was not applicable, we did not proceed to code the remaining items. In cases where not enough information was provided in the paper text or the supplemental material to determine the code, we noted this. When papers reported multiple

studies or multiple conditions in the same study that differed in their codes, we used reported Sometimes and Partially, respectively.

We coded the third item above—Is the problem given to participants well-defined?—by requiring that the following three criteria were fulfilled:

- Does the task given to participants admit a ground truth state and induce uncertainty about the state, at least across participants?
- Are participants provided with ample information about the data-generating model for a rational agent to understand how to best respond (i.e., optimize)?
- Are participants provided with a scoring rule that unambiguously assigns a score to each possible response in the action space?

Note that the first criteria is distinct from Applicability of the framework code described above: the former asks about the task given to participants, while the latter asks more broadly whether the researchers were interested in behaviors for which a ground truth could be defined. In other words, the Applicability criteria asks, regardless of how the researchers scored the responses, would it have been possible to score against a ground truth? For the second criteria, we analyzed whether the problem put to participants was one which a rational agent could best respond to in principle, regardless of whether the authors analyzed the results assuming the same problem definition. For the third criteria, we considered flat scoring rules (i.e., paying every participant the same amount regardless of response) unambiguous, since they clearly assign the same score to every response. In cases where other rules were verbally described only (e.g., "Do your best to submit accurate responses") but a flat rule was used, we considered the flat reward to be the scoring rule for incentives for that problem since it is the rule that a rational agent would respond to.

In coding the fourth item—Was a flat scoring rule used to incentivize participants?—we encountered some cases where the researchers paid participants a given amount regardless of quality but filtered responses in analysis to only those that correctly answered one or more “check” questions. We coded such cases as using a flat rule unless researchers specifically noted that they did not pay participants who failed these questions.

8.3 Coding results

Study ID	Applicable?	Evaluative?	Well defined?	Same rule?	Flat rule?
1	Y	Y	N	N	Y
2	Y	Y	Y	Y	N
3	Y	Y	Y	N	N
4	Y	Y	Y	Y	N
5	Y	Y	Y	Y	N
6	N	NA	NA	NA	NA
7	Y	Y	N	N	Y
8	Y	Y	N	N	Y
9	Y	Y	N	N	Y
10	Y	Y	N	N	NA
11	Y	Y	N	Y	N
12	Y	Y	Not enough info	N	N
13	Y	Y	N	N	N
14	Y	Y	N	N	Y
15	Y	Y	Not enough info	N	Y
16	Y	Y	N	N	N
17	N	NA	NA	NA	NA
18	N	NA	NA	NA	NA
19	Y	Y	N	Not enough info	Not enough info
20	Y	Y	Y	N	N
21	Y	Y	Y	N	N
22	Y	Y	N	Not enough info	Not enough info
23	Y	Y	Y	Y	Y
24	Y	Y	Not enough info	N	Y
25	Y	Y	N	N	Not enough info
26	N	NA	NA	NA	NA
27	Y	Y	N	Not enough info	Not enough info
28	Y	Y	N	N	Y
29	Y	Y	Partially	N	N
30	Y	Y	Not enough info	N	N
31	Y	Y	N	Sometimes	Sometimes
32	Partially	Y	N	Not enough info	Not enough info
33	Y	Y	Partially	N	N
34	Y	Y	N	Y	N
35	Y	Y	N	N	Y
36	Y	Y	N	N	Y
37	Y	Y	Partially	N	Y
38	N	NA	NA	NA	NA
39	N	NA	NA	NA	NA
40	N	NA	NA	NA	NA
41	Y	Y	N	Not enough info	Not enough info
42	Y	Y	N	Not enough info	Not enough info
43	Y	Y	N	N	Y
44	Y	Y	N	N	N
45	Y	Y	N	N	Y (most conditions)
46	Y	Y	N	Y	N