The Impact of Algorithmic Risk Assessments on Human Predictions and its Analysis via Crowdsourcing Studies

RICCARDO FOGLIATO, Carnegie Mellon University, USA

ALEXANDRA CHOULDECHOVA, Carnegie Mellon University, USA

ZACHARY LIPTON, Carnegie Mellon University, USA

As algorithmic risk assessment instruments (RAIs) are increasingly adopted to assist decision makers, their predictive performance and potential to promote inequity have come under scrutiny. However, while most studies examine these tools in isolation, researchers have come to recognize that assessing their impact requires understanding the behavior of their human interactants. In this paper, building off of several recent crowdsourcing works focused on criminal justice, we conduct a vignette study in which laypersons are tasked with predicting future re-arrests. Our key findings are as follows: (1) Participants often predict that an offender will be rearrested even when they deem the likelihood of re-arrest to be well below 50%; (2) Participants do not anchor on the RAI's predictions; (3) The time spent on the survey varies widely across participants and most cases are assessed in less than 10 seconds; (4) Judicial decisions, unlike participants' predictions, depend in part on factors that are orthogonal to the likelihood of re-arrest. These results highlight the influence of several crucial but often overlooked design decisions and concerns around generalizability when constructing crowdsourcing studies to analyze the impacts of RAIs.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; User studies; • **Information systems** → **Decision support systems**.

Additional Key Words and Phrases: Human in-the-Loop, Algorithmic risk assessment instruments, Algorithm-assisted decision-making, User study

ACM Reference Format:

Riccardo Fogliato, Alexandra Chouldechova, and Zachary Lipton. 2021. The Impact of Algorithmic Risk Assessments on Human Predictions and its Analysis via Crowdsourcing Studies. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 428 (October 2021), 24 pages. <https://doi.org/10.1145/3479572>

1 INTRODUCTION

Risk assessment instruments (RAIs) are increasingly deployed in critical domains, including finance, education, criminal justice, child welfare, hiring, and healthcare [8, 21, 63, 68, 71, 74]. Rationales for adopting these tools often center around efficiency: The hope is that RAIs might offer fast, accurate, objective, and cheap decision-making at scale [56]. To assess the potential benefits and disadvantages of RAIs, common practice is to compare their statistical properties against the status quo. However, while most research on algorithmic fairness, accountability, and transparency has focused on RAIs in isolation, a broad recognition is emerging, particularly within the HCI community, that many of these tools must be studied as situated in human-in-the-loop systems.

Authors' addresses: Riccardo Fogliato, rfogliat@andrew.cmu.edu, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, USA, 15213; Alexandra Chouldechova, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, USA, 15213, achoulde@cmu.edu; Zachary Lipton, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, USA, 15213, zlipton@cmu.edu.

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

© 2021 Copyright held by the owner/author(s).

2573-0142/2021/10-ART428. <https://doi.org/10.1145/3479572>

While it can be challenging to assess the predictive performance and fairness properties of RAIs in isolation [17, 33, 46, 47, 51], investigating their impact on systems where the human is the ultimate decision maker introduces a new layer of complexity. Perhaps the most direct evidence bearing on these questions comes from longitudinal studies of real-world systems [16, 20], particularly within the criminal justice setting. Several such studies have argued that these RAIs did not substantially improve the quality of the judicial decisions [7, 77] and in some cases amplified the existing disparities [1]. In studies that reported more positive outcomes, the RAI was often deployed as part of broader criminal justice reforms, making it difficult to isolate the effect of the algorithmic tool from other explanatory factors [5, 37].

In order to overcome the difficulties of characterizing decision-making and the impact of RAIs on systems where the human is in-the-loop, many researchers have turned to lab experiments [6, 65, 85, 86]. Findings from these studies, however, may not generalize to real-world high-stakes contexts because study participants, who are typically recruited on crowdsourcing platforms, are not representative of experts and because, unlike experts, they are not actually making decisions (only predictions). Still, results from these studies can elucidate general phenomena that characterize human-in-the-loop decision-making frameworks. In addition, the predictive performance achieved by study participants can arguably be seen as a benchmark for expert decision makers.

Many studies in this line of research have focused on the criminal justice context [9, 26, 38–40, 43, 53]. These works have mainly examined predictive performance and fairness properties concerning the study participants’ predictions of defendants’ future criminal recidivism, often using the ProPublica COMPAS dataset [3]. However, their survey designs exhibit several key differences, such as the number and type of predictions elicited from participants, the structure of financial rewards, and the requirements participants had to satisfy to take part in the study (longer discussion in §2).

In this paper, we first characterize some of the aforementioned differences across survey designs and then introduce our empirical study to determine how these factors impact crowdworkers’ predictions of future criminal recidivism. More specifically, we identify four research questions (RQs) that we believe are important considerations for assessing the generalizability of results from these studies but, to our knowledge, have not been addressed by prior work:

- RQ1** Do evaluations of participants’ responses depend on the type of predictions that are elicited? More specifically, can we infer participants’ (binary) outcome predictions only based on their (probability) risk estimates? Does this matter for evaluations of participants’ predictive performance and reliance on the RAI?
- RQ2** What is the impact of anchoring effects on participants’ predictions? In other words, do participants respond differently if they are asked to pre-register provisional responses before being shown the RAI’s recommendation?
- RQ3** How much time do participants spend on the assessments?
- RQ4** How do predictions (of re-arrest) and judicial decisions (to incarcerate) differ?

To investigate RQ1–4, we designed a vignette study (the “survey”) where laypersons recruited on Amazon Mechanical Turk were asked to predict future recidivism of a series of criminal offenders with and without the assistance of an RAI (see §3). Participants were shown short descriptions of the offenders and then asked to assess the likelihood and predict whether the given offenders would be rearrested. By employing a between-subjects design, we tested whether participants anchored on the RAI’s recommendations when these were presented at the outset together with the offender’s description.

Our results (§4) indicate the importance of asking participants separately about their *probability* predictions and their (binary) *outcome* predictions: **(1)** The study participants often predicted re-arrest even when they deemed the likelihood of re-arrest to be well below 50%. The distinction

matters for evaluations of participants' predictive performance and reliance on the RAI. For example, we find that participants substantially updated their risk estimates in the direction of the RAI's recommendation, but they rarely revised their binary predictions to match the RAI's (binary) prediction. Interestingly, participants self-reported revising their risk estimates almost twice as often as their binary predictions after seeing the RAI's recommendation. (2) Interestingly, we do not find any evidence of participants anchoring on the RAI's predictions. Instead, surprisingly, the revised risk estimates made by participants in our non-anchoring condition were significantly *closer* to the RAI's. (3) By tracking the time spent on each vignette, we discover that many participants take a long time to complete the survey but actually spend surprisingly little time on each case (average=15 seconds, median=10), taking long breaks. Moreover, because many participants perform similarly and few outperform the RAI, their predictive accuracy does not constitute a reliable proxy of time spent. (4) Finally, we discuss crucial differences between the predictive task and judicial decisions that make it difficult to draw direct conclusions about the latter from crowdsourcing studies such as ours. We present an analysis of real judicial decisions, which suggests that judicial decisions, unlike participants' predictions, depend not only on the offender's likelihood of recidivism but also on the gravity of the crimes committed.

To facilitate the reproducibility of our analysis, we have obtained IRB approval from Carnegie Mellon University and participants' consent to publicly share the data collected as part of our experiment. Our data and code for the analysis are available at www.github.com/ricfog/the-impact-of-algorithmic-rai.

2 BACKGROUND

We focus our treatment of related work primarily on crowdsourcing studies on the prediction of criminal recidivism [9, 26, 38–40, 43, 53, 78]. These studies mainly involve related prediction tasks in which participants are asked to predict the likelihood that a defendant or an offender, if released, will be rearrested within a specified period of time. Factors that would likely influence decisions but not predictions, such as the defendant's ability to pay bail or their culpability, are intentionally excluded. Factors that are not recorded in the data used by the RAI, such as the defendant's and prosecutor's arguments, are also (necessarily) left out. These studies speak to the gains in efficiency that should be expected if judges were to make decisions only based on the defendant's likelihood of re-arrest and the information present in the data. As our analysis will show (see §4.4), these results do not directly inform the impact of future deployments of RAIs; yet, they potentially provide a benchmark for the *predictive* performance of these human-in-the-loop systems.

Existing studies have mainly focused on two key questions: (i) how the predictions of study participants (alone) compare to the RAI's; and (ii) whether and how the RAI's recommendations are taken into account by participants. In terms of the former, the influential study of Dressel and Farid [26] found that, on the ProPublica COMPAS dataset [3], the predictive accuracy of their study participants' (crowdworkers) was comparable to that of the COMPAS RAI. The authors concluded that laypersons performed no worse than COMPAS. Successive studies replicated this finding [40, 43, 53]. However, they also found that participants' predictive performance was considerably lower than the RAI's when outcome feedback was removed, the base rate was decreased, or when the area under the curve (AUC), instead of accuracy, was considered [43]. These results should not be surprising for two reasons. First, humans tend to ignore information about the base rate, a phenomenon known as "base rate neglect" [48, 60]. This bias, however, is mitigated when outcome feedback is provided [36]. Second, even simple RAIs achieve predictive accuracy similar to that of COMPAS [2], but ranking offenders by their risk of recidivism (i.e., regression) represents a "harder" statistical problem. Studies in the second line of work, which investigated human predictions in presence of the RAI, generally found that the RAI led to small or no increments

in the predictive performance of participants' predictions [39, 81]. Two related studies observed differential compliance with the RAI's recommendations across defendants' racial groups: They found that providing participants with the RAI's prediction led to a larger increase in the predicted risk of re-arrest for Black defendants compared to White defendants [38, 39], a phenomenon that the authors named "disparate interactions". In our study, we found that participants, when presented with the RAI's recommendations, performed worse than the RAI across all the metrics that we considered (§??). In addition, we note that we were not able to replicate the disparate-interactions effect (§??).

However, some of the survey designs employed by these experiments present notable differences. In the following paragraphs, we describe some of such differences and discuss, based on both past and our own work, how these design choices can impact the predictions made by study participants.

In eliciting predictions, some studies asked participants for their risk estimates (probabilities) [38, 43], whereas others asked for (binary) outcome predictions of re-arrest [26, 53]. Vaccaro and Waldo [81] collected both and Grgić-Hlača et al. [40] had participants provide confidence ratings. To derive binary predictions from risk estimates, Jung et al. [43] assumed that participants *would have* predicted re-arrest for a defendant whenever their risk estimate was larger than 50%. However, it is well known that many individuals, even expert decision makers, can struggle with poor numeracy skills [10, 52, 72, 82]. In addition, individuals may base their decisions on a distorted version of the estimate of the risk, e.g., by overweighting small probabilities [44, 66]. One recent crowdsourcing study on uncertainty visualization has found that individuals that were poor at probability judgments did well on decision tasks [45]. Consistent with this evidence, our results support the hypothesis that participants' binary predictions cannot be simply obtained by converting risk estimates at a fixed threshold (§4.1).

Past studies also differ in the stage at which the RAI's recommendation was presented to the participants. In one of these studies [38], the defendant's description and the RAI's recommendation were presented to participants at the same time. In another [40], the RAI's information was made available only after the participant had made an initial prediction. However, the psychology and behavioral economics literature suggest that participants might rely on the RAI more heavily in the former setting due to a phenomenon known as the *anchoring effect*, a.k.a. "anchoring and adjustment heuristic" [31, 60, 80]. According to this heuristic, participants that are presented with a novel "anchor" would first try to assess whether the anchor equals the target and then adjust its value. The final answer generally tends to be biased in the direction of the anchor. In the setting of our study, anchor and target are represented by the RAI's recommendation and the defendant's recidivism outcome respectively. Since individuals are often unable to accurately describe their own cognitive mental processes [59, 84], hypotheses around this phenomenon need to be empirically tested. This cognitive bias has been shown to affect judicial decision-making [13, 28–30, 57] and it has been studied in at least three other crowdsourcing experiments [12, 39, 81], two of which focused on recidivism RAIs. Close to our work is Green and Chen [39], which examined the effect of anchoring on participants' risk estimates, finding that participants that had pre-registered their answers achieved higher predictive performance. In our experiment, performance was nearly identical across the two settings. Surprisingly, we found no evidence of anchoring bias: The risk estimates made by the participants that had to pre-register their predictions before seeing the RAI's recommendations were closer to the predictions made by the RAI (§4.2).

The studies also differ considerably in the number of defendants' profiles that were shown, the requirements that workers had to satisfy in order to participate, the structure of the financial compensation, the number and form of attention checks. These factors likely impacted the amount of effort made by participants in the study. But none of the studies has tried to quantify the effort that participants make on each prediction, an aspect that seems especially important for ecological

validity. Problematically, for these tasks, even participants that answer carefully can achieve low predictive accuracy and thus accuracy may not be a reliable proxy for effort. Alternatively, one might consider the time spent by participants on the entire survey, which many past studies have reported [38, 40, 53, 81]. By taking into account the number of defendants shown in each of these surveys, a quick computation suggests that, on average, the prediction for a single defendant may take less than 15 seconds. In our study, we show that, while this actually corresponds to the average time spent by our survey participants on each of the assessments, it varies widely both across participants and across vignettes (median=10 seconds) (§4.3).

Another difference in survey designs that is worth mentioning, but whose consequences we don't explore in this work, is the type of the RAI's recommendation that was displayed in the vignette. In past studies, participants were presented with the raw risk estimate of the RAI [38], the prediction binned into risk scores [81], or just the binary prediction [40]. The type of RAI's prediction that is communicated has been shown to impact participants' judgments in crowdsourcing studies [50, 86] and even decisions made in the criminal justice context [73]. In §4.1, we show that participants often do not revise their binary prediction in the direction of the RAI's. In contrast, Grgić-Hlača et al. [40] found that when study participants revised their (binary) predictions, they did so to match the RAI's in almost the majority of cases. In light of our results around how people convert risk estimates into binary predictions, the seeming contrast between the two results is likely attributable to our design choice of providing participants only with the RAI's risk estimate and not its binary prediction.

We now focus on one last important difference: the target of the studies. As we mentioned before, the goal of some of these experiments was to compare the predictive performance of the RAI with that of laypersons [26, 43]. The goal of other studies was, instead, to highlight potential unintended consequences arising from the use of RAIs in judicial decision-making [38, 40]. However, there is a clear disconnect between *predictions* of re-arrest, which were collected through the surveys, and *decisions* (e.g., of whether to incarcerate), which are made by judges. In particular, the risk of criminal recidivism, on which participants' predictions are based, represents only one of multiple factors that judges (are allowed to) account for in their decision-making process. Our results show that these differences cannot be easily reconciled in the criminal sentencing setting, especially in context of sentencing considered in our study (§4.4).

3 DATA AND METHODS

This section is organized as follows. First, we provide an overview of the dataset and of the statistical modeling used to develop the RAI (§3.1). We then cover task and experimental design (§3.2), procedure (§3.3), and recruitment process (§3.4). Lastly, we present the methodology used for the analysis of the results (§3.5). In Appendix §??, we also describe a small pilot study that we ran while developing the current experiment. The pilot is different in several ways from the main study, such as in the compensation scheme, which was not tied to performance.

3.1 Data and development of RAI

3.1.1 Data. The set of offenders used in our survey comes from a private dataset provided by the Pennsylvania Commission on Sentencing. The data contain information about offenders sentenced in the state's criminal courts. In our analysis, we considered only offenders whose race, as recorded in the data, was White or Black. The final dataset consisted of 117,464 observations, 65% of which corresponded to White offenders. For each of the offenders, we know whether they were rearrested within three years from the time of release from prison or imposition of community supervision. The overall base rate was 41.9%, while the re-arrest rates for Black and White offenders were 51% and 37.1%, respectively. We split the full data into train and test sets (70%-30%) prior to model

	dataset	survey sample
share of male offenders	79.0%	79.1%
share of White offenders	65.3%	65.1%
number of prior arrests (sd)	3.8 (4.9)	3.8 (4.9)
mean age (sd)	31.2 (10.4)	31.4 (10.3)
mean age of first arrest (sd)	23.5 (8.6)	23.5 (8.6)
prevalence of re-arrests	41.9%	42.0%
number of offenders	117464	3523

Table 1. Summary statistics for the offenders contained in the full dataset and in the subset of offenders extracted for the survey. Age is measured in years. Standard deviations (“sd”) are reported in parantheses. The marginal distributions of the offenders’ prior number of charges, type of offense committed, and age of first arrest in dataset and survey sample are also similar, but are omitted from the table.

training. A sample of 3,523 observations were further selected from the test set using stratified random sampling to ensure that the resulting sample reflected the test population on race, sex, age, and re-arrest status. These observations were used in the survey. Summary statistics for the dataset and the survey sample are reported in Table 1.

3.1.2 Development of the RAI. We used the available data to train RAIs that predict 3-year post-release re-arrest using demographic features (age, sex, and race),¹ information about the current charge (type of offense and whether it is a misdemeanor or a felony), and several features reflecting prior criminal history (e.g., past number of arrest and charges for several offense categories). We trained logistic regression, Lasso [79], random forest [11], and extreme gradient boosting (XGBoost) [15] models on the training data, tuning the hyperparameters via cross-validation. The four models showed similar performance on the holdout set: Prediction accuracy was around 66% and the area under the curve (AUC) was approximately 70%. These AUCs are comparable to those of other recidivism RAIs that are deployed in jurisdictions across the country, such as COMPAS [25] and the Public Safety Assessment (PSA) [22], among others [24]. All models were fairly well calibrated overall and, at a classification threshold of 0.5, they all predicted re-arrest for approximately 30% of the offenders in the holdout set. We also assessed the models for racial predictive bias, finding that the Lasso and random forest models were both racially well-calibrated. More details on our assessment of predictive bias in the predictions are provided in §???. Given the similarity in overall performance across the four models, we decided to adopt the Lasso for our experiments, since it was the simplest model that we found to be racially well-calibrated. For this part of the analysis, we used the tidyverse [83] and tidymodels [49] sets of packages in R [67].

¹We included race among the predictors because all modeling approaches on our data that excluded this feature resulted in re-arrest risk being overestimated for White offenders relatively to Black offenders. This over-estimation phenomenon has been documented for other RAIs, including some presently in use or under consideration for use [22, 62]. The use of race as an input to improve the predictive bias properties of the model has recently been discussed in Skeem and Lowenkamp [75]. Here, we acknowledge one strong objection to the use of risk assessment in criminal justice decision-making: Re-arrest may be a racially biased measure of offending—due to disparities in how different groups are policed. Thus an unbiased predictor of re-arrest may nevertheless be a biased predictor of offending [33]. However, the participants in our study are explicitly asked to assess the likelihood of re-arrest, so the RAIs predictive bias as a predictor of re-arrest is precisely the property at question.

3.2 Task and experimental design

In the survey, participants were tasked with predicting the likelihood and the occurrence of a re-arrest for a series of 40 offenders. The profile of each offender was presented through a descriptive paragraph based on a subset of the features that had been used to train the RAI.² For each offender, participants were asked the two following questions: (Q^p) "What is the likelihood of this offender being rearrested in the three years following release?" and (Q^b) "Do you think that the offender was rearrested in the three years following release?".³ The possible answers to Q^p were probabilities on a scale 0%-100% in bins of 1% which could be selected using a slider scale. The initial value of such slider was randomly set at either 0% or 100% when the participant entered the survey. Instead, Q^b required a negative or positive answer which could be selected using radio buttons. Two examples of vignettes are shown in Figure 1 (see also the structure of the vignette in §??). The core task in the survey consisted of three consecutive parts, which consisted of 14, 25, and 1 offenders' assessments respectively (see Figure 2). We now describe each of these in turn.

In the first part of the survey (offenders #1-14 in Figure 2), participants were provided with outcome feedback, but the RAI's predictions were not made available to them. Participants were asked to record their responses (Q_p^p and Q_p^b) and given feedback after each prediction. The feedback was as follows: "The offender was/was not rearrested in the three years following release", which was shown in green if the participant's prediction for Q^b , the binary prediction question, was accurate and in red otherwise.⁴ Offenders' profiles were randomly drawn from the survey sample.

In the second part of the survey (offenders #15-39), outcome feedback was removed but participants were shown the RAI's predictions. We employed a between-subjects design to study the effect of anchoring, assigning participants to one of two conditions, which we call *anchoring* and *non-anchoring*. The participants that were assigned to the anchoring condition were shown the offender's profile and the prediction of the RAI in the same vignette (Q_{P+RAI}^p and Q_{P+RAI}^b). Here, the RAI serves as the anchor. Instead, if assigned to the non-anchoring setting, participants were first asked to estimate the likelihood (Q_p^p) and predict the occurrence (Q_p^b) of a re-arrest based on the offender's profile alone, as they had done during the first part of the study. After submitting a response, they were presented with the RAI's prediction for the given case and were allowed to revise their own predictions (Q_{P+RAI}^p and Q_{P+RAI}^b). The randomization scheme was based on Efron's biased coin design [27].⁵ In this part of the survey, participants were shown the descriptions of a set of offenders that were drawn from the survey sample either randomly or controlling for the RAI's predictive accuracy to be around 67%. However, due to a glitch, some of the participants were shown all cases for which the RAI's binary predictions were accurate first. One past work found that the order in which the offenders' profiles were ordered did not affect participants' reliance on the RAI, even when outcome feedback was given [40]. Since the study participants affected by the glitch were equally split across the anchoring and non-anchoring conditions (see Table ?? in §??)

²The accuracy of the models trained on this subset of features were all around 65%-66%, which is nearly identical to the 66% accuracy achieved by the Lasso model selected as our final RAI. The additional features relied upon by the Lasso model were helpful in achieving model calibration, but were excluded from the vignettes to make the offenders' descriptions sufficiently brief for participants to process. Those features are counts of the offender's past arrests for specific crime types.

³As shown in Figure 1, in the survey we used the term "defendant" rather than "offender".

⁴In these 14 trials, we provided outcome feedback to make participants more comfortable with the task and, at the same time, to mitigate the possibility of base rate neglect. However, it is likely participants learnt even after the end of this part of the survey. For example, Jung et al. [43] reported using only the last 10 (out of 50) responses given by participants because of learning effects.

⁵More specifically, the probability of assignment was 1/2 for the first participant of the survey, and was adjusted according to Efron's biased coin design [27] for successive participants. According to this biased coin design, participants were assigned to the anchoring setting with probability 2/3 if the majority of past participants had been assigned to the non-anchoring setting, 1/3 if the majority of past participants had been assigned to the anchoring setting, and 1/2 otherwise.

Defendant #15 of 40.

Demographics: The defendant is black and male. He is 42 years old.

Current charge: The defendant has been charged with a property offense (felony).

Criminal history: The defendant has previously been arrested 3 times and has been charged 9 times. Of these arrests, all occurred when he was an adult. He was arrested for the first time when he was 22. The previous charges were for the following categories of offenses: sex, public administration.

The algorithmic tool estimates that the defendant's likelihood of re-arrest is 38%.

You have previously estimated that the defendant's likelihood of re-arrest is 54%.

You have previously estimated that the defendant WAS going to be rearrested.

What is the likelihood of this defendant being rearrested in the three years following release?

0 100

0 10 20 30 40 50 60 70 80 90 100

Do you think that the defendant was rearrested in the three years following release?

☐ No ☐ Yes

Defendant #15 of 40.

Demographics: The defendant is black and male. He is 36 years old.

Current charge: The defendant has been charged with a drug offense (felony).

Criminal history: The defendant has previously been arrested 12 times and has been charged 34 times. Of these arrests, all occurred when he was an adult. He was arrested for the first time when he was 24. The previous charges were for the following categories of offenses: drugs, sex, public administration, property, weapons. Of these, 4 charges were for violent offenses.

The algorithmic tool estimates that the defendant's likelihood of re-arrest is 69%.

What is the likelihood of this defendant being rearrested in the three years following release?

0 100

0 10 20 30 40 50 60 70 80 90 100

Do you think that the defendant was rearrested in the three years following release?

☐ No ☐ Yes

Fig. 1. The top and bottom panels provide examples of vignettes that were presented to participants in the non-anchoring and anchoring settings respectively. Participants assigned to the former setting were first asked to pre-register provisional responses that they were allowed to revise after being shown the RAI's recommendation (example here). Participants assigned to the anchoring setting were shown directly the RAI's recommendation. Both the likelihood estimate and the binary prediction were required to proceed to the following vignette. The initial value of the slider was randomly set at either 0% or 100%.

and we could not detect any interaction of the two effects in some preliminary analysis of the data, we decided to consider their responses in the analysis.

The third and last part of the survey consisted of only one assessment (offender #40), which the participant made right before exiting the survey and was analogous to those in the anchoring setting. Through this assessment, we wished to test any meaningful differences in the predictions of the participants that had been assigned to the two different treatments. We did not find any substantial difference in their responses and consequently we omitted the discussion of these results from the paper.

The survey also included three questionnaires, which we call *perception questions*, in which participants were asked to reflect and elaborate on the predictions that they had made. The questionnaires were located at the end of the first part of the survey (where feedback was given, after offender #14), in the middle (after offender #27) and at the end (after offender #39) of the second part of the survey. Perception questions, which include the participant's self-reported level of confidence, accuracy, trust, and use of the algorithmic tool, are described in §???. For each set of questions, participants were asked to refer only to those predictions that they had made in

Fig. 2. Structure of the survey. $Q_{P+RAI}i$ and Q_Pi indicate the answers given by the participant with and without the assistance of the RAI, respectively, for the i^{th} offender. At the initialization of the web app, the offender's profiles and the other parameters were chosen. Then the participant went through the instructions (see also §??). In the first part of the survey, consisting of 14 offenders plus one attention check, outcome feedback was provided. At the end, the participant completed the "perception questions" and then proceeded to the second part of the survey. Here, the participant did not receive feedback and was assigned to either the anchoring or to the non-anchoring setting. This part consisted of 25 offenders (#15-39) plus two attention checks and contained two sets of perception questions. Demographics were collected right before the third and last part of the survey, which consisted of only one question (#40).

the corresponding part of the survey.⁶ Participants' demographics were collected at the end of the survey (see §??). We deployed the survey through an interactive web application, using the shiny [14] and shinyjs [4] packages in R.

3.2.1 Compensation. The payment scheme consisted of a base amount of \$1.5 awarded at the completion of the survey and of a bonus up to \$5 proportional to the predictive performance achieved by the participant in the prediction task. The bonus was based on an incentive-compatible payment mechanism based on the answers provided in the second part of the survey. For the non-anchoring setting, only the predictions made by the participant with the assistance of the RAI were taken into account. The computation of the bonus worked as follows. The total reward was split evenly across the 25 assessments, i.e., the highest possible reward for each assessment was \$0.20. Then, for 12 out of the 25 total assessments, performance was measured as the accuracy of the binary predictions (Q^b). For the remaining cases, the bonus was computed according to the Brier score, a proper score function which has been used by prior work [38, 43].

3.2.2 Attention checks and exclusion criteria. We designed three "attention checks" to assess whether participants were reading carefully the offenders' descriptions. The attention checks looked like other vignettes, except that participants were explicitly told what answer to provide for Q^b in the text. Specifically, the statement "The offender was/was not rearrested in the three years after release" was inserted in the offender's description in some random position between two other sentences, with the inclusion of 'not' chosen randomly. Participants who answered incorrectly were not allowed to proceed with the successive assessments or to retry the survey. We placed the first attention check in the first part of the survey and the other two in the second part.⁷

To further ensure that participants included in the final analysis had spent sufficient time on each question, we recorded the time taken to complete each offender's assessment and the entire survey

⁶Throughout the paper, we will eventually omit the answers given to the questions in the second questionnaire because a technical issue affected the answers given by approximately the first 200 participants to the second set of perception questions. We provide more details about the issue in §??.

⁷The second attention check was added only after 30 participants had already completed the survey. This additional check decreased the likelihood that participants would pass all attention checks by random chance from $1/4 (=1/2^2)$ to $1/8 (=1/2^3)$.

as well. Participants were not informed that time was recorded independently of the Mechanical Turk system. To calculate the time elapsed per question, we created a new timestamp when a new offender's profile appeared on the participant's screen and another once both predictions were made. The difference between the two timestamps represents a measure of the time that the participant spends reading the offender's description and making the two predictions. Since participants could take breaks, this measure likely represents an overestimation of the time that they were fully engaged with the task. Participants that spent less than one second on one or more of the prediction tasks were excluded from our analysis.

3.3 Procedure

We now provide a brief overview of how a hypothetical participant navigated our survey (see §?? for more details). Upon accepting the task on Mechanical Turk and logging into the initial web page of our survey, the participant was shown a consent form. In case of consent, their identifier was collected and we checked whether the participant had not attempted or completed the survey before. If that was not the case, a short description of the task with a sample question was then displayed. Once an answer was provided, the participant was shown the full set of instructions. To ensure that participants paid sufficient attention, information regarding the content of the task (e.g., sample of offenders, RAI) and the requirements (e.g., presence of attention checks) were presented in two consecutive web pages. These two pages could still be accessed throughout the rest of the survey. Participants were told that the RAI had been trained to predict re-arrest on the same population of offenders using a superset of the information available to them and that it was well calibrated, a property that was explained in detail and illustrated in a plot. Participants were informed that their compensation would be based on the accuracy of both their likelihood estimates and binary predictions (Q^p and Q^b). Those participants that were assigned to the non-anchoring setting were not told that only the predictions they had made after having access to the RAI's recommendation would be taken into account. This choice was made to ensure that participants invested an equal amount of effort in all answers and, at the same time, to guarantee that payments would be similar across settings if having access to the RAI's predictions led to a large increase in performance. At the end of the instructions, the participant was shown one offender's description at a time and was required to make both likelihood estimates (Q^p) and binary (Q^b) predictions before proceeding to the following offender. Navigating backward to revise submitted answers was not possible.

3.4 Recruitment and participants

The study was advertised on Amazon Mechanical Turk as follows: "This is a research study about the prediction of criminal behavior. Participants can receive up to \$6.5 based upon their own performance. Average total reward is expected to be \$5". We limited our pool of participants to workers that had an historical HIT approval rate higher than 90%, over 500 Human Intelligence Tasks (HITs) approved, were located in the US, and had not completed or attempted the survey before. The maximum time allowed to submit the HIT on Mechanical Turk was set at 90 minutes. A total of 1438 Mechanical Turk workers tried the survey. Approximately 900 of them failed one of the attention checks and were not allowed to retry the survey. Less than 10 workers were rejected after the completion of the survey because they completed one or more prediction tasks in under one second. We were left with valid data from 531 participants. On average, these participants completed the survey in 28 minutes (standard deviation=13.5, median=25.4) and earned a bonus of \$3.37 (sd=\$0.47, median=\$3.37). The average reward was \$12.7 per hour (sd=\$5.9, median=\$11.4).

Participants were mostly male (proportion=62.5%, total=332), White (76.1%, 404), and had a college degree (80.2%, 426). The average age was 38.7 years (sd=11.9, median=36). College graduates and males were overrepresented in our sample compared to their prevalence in the US population.

Additional details and comparison with demographics from the Census are presented in Table ?? of §??.

3.5 Data analysis

Throughout the paper, we use the following methods and notational shortcuts (inside parentheses). Correlations on both ordinal and cardinal data are computed using Spearman's rank correlation coefficient (S). We use Kruskal-Wallis test (KW) as the omnibus test of association between a numeric outcome and a categorical factor variable. The Kruskal-Wallis test is a rank-based nonparametric analog to one-way ANOVA. We use Mann-Whitney U test (MW) for post-hoc comparisons or comparisons between only two groups. To test the statistical significance of pairs of differences in means between independent samples, we use Welch's t-test (TT). Standard deviation (sd) and standard error (se) are sometimes reported together with (or in place of) the results of these tests. The reported confidence intervals generally rely on the asymptotic normality of the distribution of the corresponding test statistic. Confidence intervals for the AUC, however, are obtained using the bootstrap. Before conducting the experiment, we chose the significance level of $\alpha = 0.001$ which we use for testing all the hypotheses. Accordingly, we report only whether the p-values (p), which we adjust via Bonferroni correction in case of multiple testing, are lower or higher than the chosen significance level (α). Throughout the analysis, we make the simplifying assumption that all pairs of predictions (i.e., likelihood estimates and binary predictions (Q^p, Q^b)), both between and within participants, are independent. We relax this assumption only in case of the pre-registered and corresponding revised predictions in the non-anchoring setting.

To quantify participants' reliance on the RAI in the second part of the survey (i.e., offenders #15-39), we employ metrics based on Q_P (when available), Q_{P+RAI} , and Q_{RAI} for both Q^p and Q^b . These measures include analogs to measures that were adopted by prior work, such as *deviation* [65], *influence* [38, 39, 65], *agreement fraction*, and *switch fraction* [85]. Here, we formally introduce only influence. This metric, which is also known as "weight of advice" [35], quantifies the magnitude of the revision of the participant's risk estimate relatively to the difference between the RAI's recommendation and the participant's initial prediction. In mathematical terms, for each offender's assessment made by the participant, influence is defined as $(Q_{P+RAI}^p - Q_P^p) / (Q_{RAI}^p - Q_P^p)$. We excluded cases for which $|Q_{RAI}^p - Q_P^p| \leq 5\%$. According to this definition, influence can only be measured for the predictions made by the participants that were assigned to the non-anchoring setting. It is 1 if the participant's revised prediction matches the RAI's, 0 in case of no revision.

4 RESULTS

Together, the 531 participants made 28015 pairs of predictions (Q^p, Q^b) on 3521 different offenders. Approximately half of the participants were assigned to the anchoring setting (proportion=49%, total=260). The average time spent on each offender's assessment was 15 seconds ($sd=34.2$, median=10). The total time taken to complete the evaluation of the 40 offenders was 14.9 minutes ($sd=8.3$, median=12.9) for the participants assigned to the non-anchoring setting and 11.4 minutes ($sd=6.6$, median=10.0) for those in the anchoring setting. The time gap is due to the 25 additional pairs of questions asked in the non-anchoring setting.

In the sample shown to the participants, 42% of the offenders were rearrested. The RAI produced an average probability of re-arrest of 42% (median=40%) and predicted re-arrest for 29.5% of the offenders. The risk estimates made by participants were slightly higher but of similar magnitude to the RAI's (mean=50.4%, median=50%). But participants predicted that 59.2% of the offenders would be rearrested, twice as many as the RAI had predicted.

Fig. 3. Comparison of risk estimates Q_{P+RAI}^p and corresponding binary predictions Q_{P+RAI}^b made by participants in presence of the RAI. (left) Share of assessments with predictions of re-arrest as a function of the (binned) risk estimates. Error bars represent 95% confidence intervals. For example, participants predicted that the offender would be rearrested in more than 40% of the assessments in which they estimated the probability of re-arrest to be between 40% and 49%. (right) Share of assessments with predictions of re-arrest as a function of the risk estimate and the participant. Each row corresponds to a different participant. A large fraction of the participants often—but not always—predicted re-arrests while assigning likelihoods well below 50%. Note that participants rarely applied a consistent probability threshold in predicting re-arrest.

We now turn to the presentation of our key results. The structure of the rest of the section is as follows. In §4.1, we discuss how participants' risk estimates (Q^p) did not uniformly map onto binary predictions (Q^b). In §4.2, we show that participants did not anchor on the RAI's predictions. In §4.3, we show that performance was not related to time spent on the survey and total time spent on the survey is only a poor proxy for time truly spent on each question. Last, in §4.4 we compare predictions of re-arrest and judicial decisions regarding incarceration.⁸

4.1 Divergence between risk estimates and binary predictions

In §3.2, we described how the payment scheme in the second part of the survey was designed to incentivize participants to report their most accurate likelihood estimates and binary predictions. If participants adopted the profit-maximizing strategy, they would convert their likelihood estimates (Q_{P+RAI}^p) into binary predictions (Q_{P+RAI}^b) by predicting re-arrest for all those offenders for whom the value of Q_{P+RAI}^p was 50% or greater. Only a small fraction of the participants showed this behavior.⁹ Figure 3 shows that participants assisted by the RAI predicted re-arrest even when their risk estimates were well below the 50% threshold. Approximately one fourth of the offenders (26%, $se=0.5\%$) whose estimated probability of re-arrest was below 50% had predictions of re-arrest. Conversely, one tenth of the offenders (9.1%, $se=0.4\%$) whose estimated risk was larger than 50% had predictions of no re-arrest. The left panel of Figure 3 describes this phenomenon by displaying the share of predicted re-arrests (i.e., mean of Q_{P+RAI}^b) as a function of the corresponding estimated

⁸For the interested reader, in the supplementary material we include the following four analyses and related findings. In §??, we show that participants' predictive performance, both with and without the assistance of the RAI, was lower than the RAI's with respect to all metrics but the false negative rate. In §??, we show that, as in Green and Chen [38], participants' self-reported predictions accuracy was correlated with their real accuracy only when outcome feedback was provided. In §??, we show that, in contrast with past work [38, 39], we did not find evidence of disparate-interactions in the uptake of RAI's predictions. In §?? we show that the ordering of the offenders' profiles did not affect self-reported trust and reliance on the tool. Lastly, in §??, we examine how participants assigned to the non-anchoring setting updated their risk estimates and binary predictions. This subsection extends the discussion in §4.1.

⁹Here, we consider only the predictions made by participants in the second part of the survey and in the presence of the RAI. We found very similar patterns also for the answers given by participants in absence of the RAI (i.e., Q_P^p and Q_P^b) throughout the entire survey. This indicates that the presence of the RAI did not have any substantial impact on how participants translate risk estimates into binary predictions.

Fig. 4. Analysis of predictions made by participants assigned to the non-anchoring setting in the second part of the survey. The two figures show the distribution of the pre-registered Q^P and revised Q^P_{P+RAI} predictions made by participants (vertical axis), grouped by the RAI's risk estimate Q^P_{RAI} binned (horizontal axis). (left) Density estimates of the participants' risk estimates Q^P . Compared to the pre-registered predictions, the densities of the revised risk estimates put more mass on the interval where also the RAI's predictions lie. This indicates that, unsurprisingly, the participants' revised estimates were closer to the RAI's. (right) Share of assessments with predictions of re-arrest in participants' pre-registered Q^P and revised Q^P_{P+RAI} binary predictions. Error bars represent 95% confidence intervals. The agreement between the RAI's and the participants' predictions increased after the RAI's prediction was shown, but the disagreement remained high in some of the buckets, even for low values of the RAI's risk estimates.

likelihood of re-arrest (Q^P_{P+RAI}), appropriately binned. We observe a gradual increment in the share of predicted re-arrests as the risk estimates increase. As one could expect, the share of predicted re-arrests drastically increases around 50%: A binary prediction is twice as likely to correspond to re-arrest if the likelihood is just above the optimal threshold (mean $Q^P_{P+RAI} = 44.1\%$ for $Q^P_{P+RAI} \in [40\%, 49\%]$, 76.4% for $Q^P_{P+RAI} = 50\%$, 90.8% for $Q^P_{P+RAI} \in [51\%, 60\%]$; all $p_{TT} < \alpha/3$).

The right panel of Figure 3 shows the significant heterogeneity in the strategies adopted by the participants. Approximately one third of them (32.4%, 172 out of 531) always predicted re-arrest when their probability estimate was greater than 50% and no re-arrest when their estimate was less than 50%. Approximately half of the participants (56.1%, 298) followed this strategy in at least 90% of their predictions. Interestingly, Figure 3 shows that participants did not even use fixed thresholds, i.e., for most participants there was no given threshold t for which $Q^b = \text{rearrest}$ whenever $Q^P \geq t$.

Consequences for evaluations of participants' predictive performance. As discussed in §2, the results of Jung et al. [43] were obtained by soliciting only likelihood estimates and converting them into binary predictions by applying a uniform threshold (50%) across all participants. Our findings indicate that this assumed correspondence is not borne out in practice, and analyses based on this assumption can lead to incorrect conclusions. For example, consider the question of assessing the predictive performance of our study participants' binary predictions with converted binary predictions (obtained by thresholding Q^P at 50%) compared to actual binary predictions Q^b .¹⁰ While the overall accuracy of the converted predictions is similar to that of the actual binary predictions (59.2% and 57.6% resp., $p_{TT} > \alpha$), other performance metrics are quite different. The false positive rate is 8.4% lower for the converted binary predictions compared to the actual binary predictions (40.8% vs. 49.2%, $p_{TT} < \alpha$), and, instead, the false negative rate is 8% higher (40.4% vs. 32%, $p_{TT} < \alpha$).

Consequences for evaluations of participants' reliance on the RAI. We now examine how participants assigned to the non-anchoring setting updated their pre-registered probability and binary

¹⁰For this analysis, we excluded all risk estimates exactly equal to 50%. As before, we used only the risk estimates and predictions made by participants in presence of the RAI.

predictions after they were shown the RAI. When it came to risk estimates, Figure 4 shows that participants tended to update their estimates in the direction of the RAI's recommendation. These revisions were often substantial (mean influence=37%, see also Figure ??; in addition mean of $|Q_{RAI}^p - Q_P^p| = 18.8\%$ vs. $|Q_{RAI}^p - Q_{P+RAI}^p| = 13.3\%$). Participants changed their pre-registered binary answers in 10.2% (se=0.3%) of all predictions, switching in slightly more than half of these cases from a prediction of re-arrest to a prediction of no re-arrest (see breakdown by RAI's score in Figure 4). However, the overall agreement with the RAI's (binary) predictions only increased by less than 4% from the pre-registered to the revised answers (mean $Q_P^b = Q_{RAI}^b : 61.7\%$, mean $Q_{P+RAI}^b = Q_{RAI}^b : 65.4\%$, $p_{TT} < \alpha$) because participants' final prediction often did not match the RAI's even if their pre-registered prediction did. It is certainly possible that, had participants been provided also with the RAI's binary predictions, we would not have observed such a phenomenon (c.f. "result 2" in Grgić-Hlača et al. [40]). We also asked participants to self report for what share of the assessments they had revised their own predictions after taking into account the RAI's recommendation. Interestingly, they indicated that they had revised their risk estimates and binary predictions in 59% and 34% of the assessments respectively (medians=60% and 20%). Thus, evaluations based solely on one of the two types of predictions could lead to drastically different conclusions regarding participants' (especially self-reported) reliance on the RAI.

4.2 Absence of anchoring effects

In the introduction, we hypothesized that the participants could anchor on the RAI's predictions. If this had happened, then the risk estimates of the participants assigned to the anchoring setting would have been closer to the RAI's than those made by the participants assigned to the non-anchoring setting. As we might expect, we found that the pre-registered risk estimates of the participants in the non-anchoring setting were 16% further from the RAI's than those of the participants in the anchoring setting. Yet, very surprisingly, the revised risk estimates made in the non-anchoring setting were 17% closer to the RAI's than those made in the anchoring setting (mean and median absolute diff. $|Q_{P+RAI}^p - Q_{RAI}^p|$ in the non-anchoring setting: 13.3% and 8% resp., in the anchoring setting=16.1% and 11% resp.; $p_{MW} < \alpha$). In contrast, participants' binary predictions matched the RAI's at exactly the same rate across the two settings (mean $Q_{RAI}^b \neq Q_{P+RAI}^b : 34.6\%$ in both). We also found that performance did not differ across the two settings: accuracy (acc. in anchoring=57.4% and in non-anchoring=57.6%), false positive rate, false negative rate, positive predicted values, and the AUC were not significantly different ($p_{TT} > \alpha$ for all comparisons).¹¹

In the perception questions, participants self reported similar levels of trust in the tool across the two settings (76.5% and 79.3% of participants in anchoring and non-anchoring resp.; $p_{TT} > \alpha$), and also of confidence and accuracy of their own predictions. Participants reported revising their binary predictions after accounting for the RAI's at approximately the same rate in the two settings (mean reported revision=31.8% in anchoring and 33.7% in non-anchoring; $p_{MW} > \alpha$). However, the story is different in case of risk estimates. Participants assigned the non-anchoring setting self reported that they had revised their answers on average 65% more often than those in the anchoring setting after looking at the RAI's recommendation (mean reported revision=58.9% in non anchoring, 35.7% in anchoring; $p_{MW} < \alpha$). Despite the risk estimates of the participants assigned to the non-anchoring setting were only slightly closer to the RAI's, their perceived use of the tool to adjust risk was substantially higher. It seems possible that, by pre-registering their risk estimates, participants became more aware of the influence of the RAI on their risk estimates.

¹¹We conducted analogous analyses after dropping the predictions that were made in less than 3, 5, and 10 seconds. Again, we found no evidence of anchoring effects.

Fig. 5. Analysis of time. (left) Time taken to complete the HIT on Mechanical Turk (vertical axis) as a function of the time spent on the survey (horizontal axis) for each participant. The blue dotted line corresponds to 90 minutes, which was the maximum time allowed on Mechanical Turk for the submission of the HIT. Each black dot corresponds to an individual participant. The observations below the diagonal line correspond to participants that started the survey before accepting the HIT. If the time on Mechanical Turk and on the survey were similar for each of the participants, the observations would lie close to the dashed red line at 45 degrees. We observe that the time taken to complete the HIT was generally substantially larger than the time spent on the survey. (right) Boxplot of the time spent on each offender's assessment by every participant. The whiskers extend from the hinges to the smallest or largest values at most 1.5·IQR of the hinge. For visualization purposes we randomly sampled 100 participants from those that had been assigned to the anchoring setting and limited the length of the vertical axis.

To summarise, we did not find evidence of anchoring effects. Instead, we found the opposite: The predictions made by participants in the non-anchoring setting were closer (in absolute distance) to the RAI's. These participants also self-reported that the tool had influenced more heavily their risk estimates.

4.3 Time spent on the survey as a poor proxy for time spent on the assessment

Participants completed all the steps in the survey, from login to submission, in an average time of 28 minutes ($sd=13.5$). In comparison, the average time recorded on Mechanical Turk, from the acceptance to the submission of the HIT, was 68% greater, averaging 46.7 minutes ($sd=20.2$). The left panel of Figure 5 shows the time spent on the survey and on Mechanical Turk for each participant. The time spent on the survey was on average 89% greater than the time recorded on the platform. For 27.9% of the participants (139 out of 499)¹² the submission of the HIT took 5 minutes less than the completion of the survey; for 43.5% of them (217) it took 10 minutes less, and for 63.1% of them (315) it took 20 minutes less. Figure 5 also reveals that the time spent on the survey varied considerably across participants, from less than 10 minutes up to more than one hour. An analysis at a more granular level reveals that the assessment of an individual vignette was carried out on average in 15 seconds ($sd=34.2$, median=10). Once we excluded the pre-registered and revised predictions made by the participants assigned to the non-anchoring setting in the second part of the survey,¹³ the average time was 17.5 seconds ($sd=39.2$, median=12.2).

The substantial difference between the mean and the median raises the following question: Is the time taken to complete the survey (divided by the number of questions) a reliable proxy for the time spent on each question? Equivalently, do participants spend the same amount of time on

¹²We could not match all survey participants to their Mechanical Turk identifiers.

¹³These participants made two pairs of predictions for each offender and the revised prediction was made more quickly than the pre-registered one (mean time of revised=9.4 seconds, $sd=21.7$).

all question?¹⁴ The right panel of Figure 5 shows that participants generally did not allocate their time equally across questions. On average, they spent more than one tenth of their time on the assessment of only one offender (out of 40, mean proportion of time per participant=12.1% and median=7.7%) and one fifth of their time on four assessments (20.3%, median=16.6%). Once the four assessments that took longest were removed for each participant, the time spent that each of them spent on the assessments decreased on average by 19.4%. We also found, as might be expected, that participants tended to spend more time on the first question of each part of the survey. For example, in the first part of the survey, participants spent on average almost 9 seconds more on the first question than in the others (average time first question=26.8 and median=17.6 seconds, others mean=17.9 and median=12.4; $p_{MW} < \alpha$). Similarly, the first question in the second part of the survey took twice as much time as the others (average time first question=30.6 seconds, others=17.2; $p_{MW} < \alpha$). We could not identify any other clear reasons for why some predictions took longer than others. It is likely that participants simply took breaks from the survey. Nonetheless, they allocated the time proportionally to the number of questions in each of the two parts of the survey (i.e., on average 37.8% of their time on the first part of the survey that contained 35%=14/40 of the offenders).

We also examined whether time was related to predictive performance. For the participants that had been assigned to the anchoring setting, the time spent on the entire survey was weakly correlated with the accuracy of the binary predictions ($\rho_S = 0.3$, $p < \alpha$). However, we observe that this correlation is mainly due to the participants that completed the surveys very quickly and also achieved very low accuracy. Assessments on which predictions were accurate did not take longer than inaccurate predictions on average ($p_{TT} > \alpha$) and only the median time was slightly higher (median time accurate=12.4 seconds and inaccurate=11.9; $p_{MW} < \alpha$). Given the large variance in the time spent on the assessments across participants, we also tested a similar hypothesis. For each participant, we ranked their 40 predictions in increasing order of time. The mean ranking of the predictions that were accurate was, perhaps unexpectedly, *lower* than that of the those that were inaccurate. One possible explanation is that some vignettes are easy to get right, and thus do not require much time to complete.

In summary, we found that the time taken to complete the HIT on Mechanical Turk was a severe overestimation of the time that participants spent on the survey. The time spent of the survey was also a poor proxy for the time spent on each vignette, given that participants typically spent a substantial share of their time on only a few assessments. Lastly, we found little evidence of any association between time spent by the participants on the assessment and accuracy of their predictions.

4.4 Predictions are not decisions

In our study, participants were asked to provide likelihood estimates and binary predictions of offenders' re-arrest outcomes. As we discuss in this section, our findings concerning how humans update their predictions when presented with the RAI's recommendations—and findings from these types of studies more generally—do not readily translate into implications for decision-making. Firstly, there is the clear issue that our study participants are not judges, and are not being presented with the full set of information that judges would have access to in real world decision-making settings. However, we would not be surprised if conducting our experiment on a population of judges produced largely the same qualitative findings regarding risk *predictions*. The bigger issue

¹⁴If this was the case, then we could infer—with some certainty—bounds for the time taken to complete an individual assessment based on the total time spent on the survey. In this analysis we only considered all 40 predictions made by the participants that were assigned to the anchoring setting.

Fig. 6. Comparison of judicial decisions of incarceration, binary predictions of the RAI, and binary predictions of participants (both Q_P and Q_{P+RAI}). The threshold for converting the RAI’s risk estimates into binary predictions for this analysis was set such that the overall share of predicted re-arrests was equal to the incarceration rate. The color in each cell is proportional to the share of offenders that were incarcerated (judge) or with predictions of re-arrest (RAI and participants), for specific values of the offense gravity score (OGS) and prior record score (PRS). In this figure, the heatmaps for judicial decisions and RAI’s predictions were generated using all offenders in the test set.

that we wish to draw attention to here is that risk is just one of many factors that judges are asked to consider in their *decisions*.

In the context of sentencing, considerations of risk are often secondary to factors such as the severity of the offense and restitution to victims. The offender’s risk of re-arrest often enters into decisions indirectly through considerations of prior criminal history, which is a leading predictor of future recidivism. For instance, the Pennsylvania Commission on Sentencing has adopted a set of sentencing guidelines that “provide sanctions proportionate to the severity of the crime and the severity of the offender’s prior conviction record” [64]. These guidelines are primarily a function of two factors: the offense gravity score (OGS), which is a measure of the gravity of the most serious crime among the offender’s current charges; and the prior record score (PRS), which is a composite measure summarizing prior criminal records. Higher values of OGS and PRS correspond to more serious offenses and more severe prior criminal histories [42]. While the recommendations are advisory, judges need to provide a written justification when their sentencing decisions deviate from the recommended range.

Our data contain information not only on re-arrest outcomes, but also on judicial sentencing decisions. This allows us to compare judicial decisions of whether to sentence the offender to a period of incarceration to the predictions of re-arrest made by participants and RAI. Figure 6 shows judicial decisions along with participants’ and RAI’s predictions broken down by levels of the OGS and the PRS. For judges, we observe that the likelihood of incarceration increases with both the PRS and OGS. But the pattern is more complex in case of participants’ and RAI’s predictions of re-arrest: The likelihood of a predicted re-arrest appears to increase with the PRS, but it does not seem to be associated with the OGS.

We investigated how the likelihood of incarceration and the predictions of re-arrest made by participants and RAI depended on the PRS and the OGS by fitting three separate logistic regression models of the form $outcome \sim PRS + OGS$. We treated both scores as numeric (repeated felony offender, “RFEL”, was converted into a score of 6) and set the threshold for the RAI such that the share of predicted positives was equal to the share of offenders that were incarcerated. The coefficients of the OGS were 0.05, 0.08, and 0.49 for the models targeting predicted re-arrest by

participants, by the RAI, and incarceration respectively (all $p < \alpha$). These results would indicate that, *ceteris paribus*, an increase in the OGS by one level was associated to a 63% increase in the odds ratio of the likelihood of incarceration but only to a 5% increase in that of the participants' predictions of re-arrest (see the full results in Table ??).

To take a concrete example, consider the subgroup of offenders charged with drug offenses. Among these offenders, 68% of those charged with felonies were incarcerated, compared to only 21% in case of misdemeanors. Yet, the rate of re-arrest was virtually identical across the two types of charges (47% and 45% respectively).¹⁵ In comparison, participants and the RAI, once equalized for overall rates, predicted re-arrest for only a slightly higher share of the felony offenders (share of predictions of re-arrest for felonies: RAI=43% and participants=66%; misdemeanors: RAI=52% and participants=59%). This indicates that there are cases where differences in decisions are unrelated to the risk of recidivism (here, the decision of whether to incarcerate).

Because risk predictions are not aligned with the OGS, and the OGS is one of the primary determinants of criminal sentences, we should not expect findings of how judges' predictions change when provided with RAI information to be directly informative of likely changes in the resulting decisions. Even if our findings that participants revised their risk predictions when presented with the RAI's recommendation generalized to judges, it is unclear what effect, if any, those revised assessments would have on sentencing outcomes.

5 GENERAL DISCUSSION

5.1 Limitations

There are several limitations related to the generalizability of our results. First, despite our strict exclusion criteria, the analysis in §4.3 revealed that a substantial share of the assessments were carried out very quickly, often in less than 5 seconds. This indicates that participants often did not fully read the offenders' descriptions or did not think carefully about their answers. The post-hoc analyses, however, have revealed that our conclusions were largely unaffected by the presence of these predictions in the data. It is also unclear whether our findings would generalize to other experimental setups (such as those with different incentive structures), populations of participants, domains, or to real-world decision-making settings. For example, we might expect judges to exhibit many of the cognitive biases exhibited by our survey participants [41] and suspect they might also overestimate their own predictive abilities [23], which may result in lesser reliance on the RAI. Consequently, the impact of the RAI on their judgment may be different than what has been observed in our experiment. Another potential limitation is that our results could have been affected by our explanations of the incentive structure, information regarding the RAI, and framing of the questions. We also note that our pool was composed of Amazon Mechanical Turk workers that were 20–40 years old, a demographic group that is likely tech-savvy and more open to the introduction of technologies than other populations [69]. As such, it is not a representative sample of the broader US population, and is not demographically representative of the US judiciary. Lastly, it is possible that, due to the negative experiences that workers often have on Mechanical Turk [34, 55], running the experiment on an alternative platform could lead to different results, especially those discussed in §4.3.

5.2 Discussion

In our study, participants often predicted re-arrest even when assigning a low numeric score to the likelihood of that event. We offer four hypotheses that could explain this behavior. The first is

¹⁵Even if we assumed that imprisonment helped reduce re-arrest rates (and evidence suggests that this is unlikely to be the case [58]), the large gap in incarceration rates could still not be explained by differences in the likelihood of re-arrest.

poor numeracy skills: Even highly educated individuals tend to struggle with simple numeracy questions [10, 52, 72, 82]. Some of our participants may have not realized that the expected profit-maximizing strategy is to predict that the given offender would be re-arrested if and only if they believed the risk of rearrest is greater than 50%. A second explanation might be that participants' subjective likelihood judgments did not simply reflect the estimated frequencies of events [32, 54], e.g., due to an asymmetry in the perceived cost of errors. This hypothesis seems plausible, even though participants were incentivized through the payment scheme to report calibrated probability estimates. A third explanation might be lack of care: We found that, while the observed phenomenon held in general, it was more common among the predictions that took the least time. Consequently, it is possible that some of the participants made the predictions quickly and without paying adequate attention, e.g., satisficing. The fourth explanation might be that participants did not understand the slider instrument, despite being provided with a clear set of instructions (see §??). However, given that this instrument is prevalent in current practice [70] and it generally has high concurrent validity with respect to other types of instruments such as Likert scales and radio buttons [18], this hypothesis seems unlikely to explain our results. Future work might explore why this phenomenon occurred.

It is of paramount importance to employ measures that are comprehensible to the human and can also be effectively and consistently estimated by the human. It is also equally important that the RAI's recommendations be easily understood by the human. Different ways to communicate the RAI's prediction can help decision makers calibrate their trust in the RAI [50, 86] and, potentially, also reduce variability in their decisions. The comparison between our results around how people often switch their (binary) predictions not in the direction of the RAI's and the findings in Grgić-Hlača et al. [40] likely represents an example of the importance of such design choices. Another notable implication of this result is that the type of prediction on which participants' predictive performance and reliance on the RAI's recommendations are evaluated can affect the conclusions that researchers draw from the study. In our experiment, for example, an analysis that focused only on binary predictions would have neglected the role of the RAI in influencing participants' risk estimates. While the choice of the type of prediction might be context-dependent, the potential limitations of the assessment should be recognized and acknowledged.

One notable and unexpected finding in our work is the absence of anchoring effects. We designed the second part of the survey to determine how much (if at all) participants anchored on the RAI's predictions. For binary predictions, we found no difference in the agreement between the participants and the RAI across the two settings. For risk estimates, not only did our participants not anchor on the RAI—participants who were shown the RAI directly gave predictions further from the RAI's than participants that had pre-registered their estimates before seeing the RAI's outputs. One possible explanation of this behavior is that participants felt that they had to—or wanted to—demonstrate a sense of agency. According to this hypothesis, participants assigned to the non-anchoring setting, who had already delivered risk estimates without the assistance of the RAI, would have been more comfortable with accepting the RAI's predictions, having already demonstrated their agency. At the same time, participants assigned to the anchoring setting might have felt compelled to offer predictions that differed from those of the RAI to demonstrate that they were performing the task earnestly and not simply copying the RAI. Another possible explanation for this phenomenon might be that the users who pre-registered their predictions also had the opportunity to notice how similar the RAI's predictions were to theirs and might consequently have trusted the RAI more. However, these users did not report higher levels of trust in the tool. Yet they reported making a heavier use of the tool's predictions. This finding opens an interesting new direction for future research: If these results held in other experiments, could we improve

experts' perceived and actual reliance on algorithmic tools by eliciting predictions from the experts both before and after revealing the RAI's recommendation?

Like past works, we found that even when participants were shown the RAI's predictions, they continued to underperform the RAI in terms of predictive accuracy. It seems possible that this observation may characterize many forecasting settings where both the RAI and humans have low accuracy. Tan et al. [78] demonstrated that even if the predictions made by the RAI and the human alone were combined, the resulting unavoidable error may still be very large. We note that in our experimental setting we would not expect participants' predictions to outperform those of the RAI. This is because the RAI has access to all of the features that participants are able to consider, and is trained to optimize for predictive accuracy. In practice, judges have access to information that the tool does not. Further studies of human-in-the-loop systems are needed to examine the influence of the overlap in information sets between humans and the RAI on predictive performance, participants' trust in the tool, and decision-making. It is likely that even in the case of complementary information sets, or where participants have access to more information than is available to the tool, it will be difficult for humans to match the RAI's performance without carefully crafted feedback.

Unsurprisingly, the time that participants actually spent on the survey was substantially lower than the time taken to submit the HIT and did not even represent a reliable proxy for the time spent on the assessment of a single vignette. In addition, participants who spent more time on the survey did not, in general, achieve higher predictive accuracy. Ideally, especially in tasks characterized by high degrees of uncertainty such as the prediction of criminal recidivism, researchers might want to design compensation structures in which the assigned reward is proportional to the effort made by the participant. While time itself is not a perfect measure of effort, future work could employ exclusion criteria based on the time that participants spend on each assessment. Time spent is also worth taking into account when assessing the generalizability of study findings, especially with respect to real-world decision-making in high-stakes settings. For example, researchers could run post-hoc analyses to assess whether their conclusions still hold once the predictions that participants made quickly are dropped. In our experiment, as we already mentioned, the main results still held even when these predictions were not considered in the analysis.

Lastly, in §4.4 we showed that, unlike judicial decisions, the predictions made by participants were weakly associated to the seriousness of the offender's current charge. For drug felonies and misdemeanors, predictions were nearly orthogonal to the gravity of the offense. This finding sheds light on the limits of the use of human predictions of criminal recidivism as proxies for judicial decision-making: Even if studies found large effects of the impact of RAI's recommendation on human predictions (which they have not), further investigation would be required to understand how or whether those would translate to decisions. Such an analysis would not only require understanding for which offenders decisions and predictions would diverge, but also how the introduction of the RAI would reshape the existing decision-making framework. For example, in the Pennsylvania sentencing system, the newly adopted RAI is used to determine when a pre-sentencing investigation report is to be generated, which would provide judges with additional information on the offender [61]. In the New Jersey pretrial system, the RAI represents a building block of the decision-making framework, but it is not its only element [19]. There, the RAI is also used to decide whether a summons or a warrant should be issued. Thus, the RAI potentially affects not one but many sequential decisions made at several stages of the pipeline by interacting with other elements in a larger framework.

6 ACKNOWLEDGMENTS

We are grateful to PwC USA for funding this research through the Digital Transformation and Innovation Center sponsored by PwC. We also deeply thank Shamindra Shrotriya and Pratik Patil for providing valuable feedback on the survey.

REFERENCES

- [1] Alex Albright. 2019. *If You Give a Judge a Risk Score: Evidence from Kentucky Bail Decisions*. Technical Report. Working paper.
- [2] Elaine Angelino, Nicholas Larus-Stone, Daniel Alabi, Margo Seltzer, and Cynthia Rudin. 2017. Learning certifiably optimal rule lists for categorical data. *The Journal of Machine Learning Research* 18, 1 (2017), 8753–8830.
- [3] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. ProPublica. See <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (2016).
- [4] Dean Attali. 2020. shinyjs: Easily Improve the User Experience of Your Shiny Apps in Seconds. <https://CRAN.R-project.org/package=shinyjs> R package version 1.1.
- [5] James Austin and Wendy Ware. 2020. Why Bail Reform is Safe and Effective: The Case of Cook County. Available at SSRN 3599410 (2020).
- [6] 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.
- [7] Richard Berk. 2017. An impact assessment of machine learning risk forecasts on parole board decisions and recidivism. *Journal of Experimental Criminology* 13, 2 (2017), 193–216.
- [8] Richard Berk, Hoda Heidari, Shahin Jabbari, Michael Kearns, and Aaron Roth. 2018. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research* (2018), 0049124118782533.
- [9] Arpita Biswas, Marta Kolczynska, Saana Rantanen, and Polina Rozenshtein. 2020. The Role of In-Group Bias and Balanced Data: A Comparison of Human and Machine Recidivism Risk Predictions. In *Proceedings of the 3rd ACM SIGCAS Conference on Computing and Sustainable Societies*. 97–104.
- [10] William C Black, Robert F Nease Jr, and Anna NA Tosteson. 1995. Perceptions of breast cancer risk and screening effectiveness in women younger than 50 years of age. *JNCI: Journal of the National Cancer Institute* 87, 10 (1995), 720–731.
- [11] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [12] Zana Bućinca, 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.
- [13] Shawn D Bushway, Emily G Owens, and Anne Morrison Piehl. 2012. Sentencing guidelines and judicial discretion: Quasi-experimental evidence from human calculation errors. *Journal of Empirical Legal Studies* 9, 2 (2012), 291–319.
- [14] Winston Chang, Joe Cheng, JJ Allaire, Yihui Xie, and Jonathan McPherson. 2020. shiny: Web Application Framework for R. <https://CRAN.R-project.org/package=shiny> R package version 1.4.0.2.
- [15] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [16] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. 134–148.
- [17] Amanda Coston, Ashesh Rambachan, and Alexandra Chouldechova. 2021. Characterizing fairness over the set of good models under selective labels. *arXiv preprint arXiv:2101.00352* (2021).
- [18] Mick P Couper. 2008. *Designing effective web surveys*. Cambridge University Press.
- [19] New Jersey Courts. 2018. Pretrial Release Recommendation Decision Making Framework (DMF). <https://www.njcourts.gov/courts/assets/criminal/decmakframework.pdf?cacheID=nDrtJn4>
- [20] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 2020. A Case for Humans-in-the-Loop: Decisions in the Presence of Erroneous Algorithmic Scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.
- [21] Dursun Delen. 2010. A comparative analysis of machine learning techniques for student retention management. *Decision Support Systems* 49, 4 (2010), 498–506.
- [22] Matthew DeMichele, Peter Baumgartner, Michael Wenger, Kelle Barrick, Megan Comfort, and Shilpi Misra. 2018. The public safety assessment: A re-validation and assessment of predictive utility and differential prediction by race and gender in kentucky. Available at SSRN 3168452 (2018).

- [23] Matthew DeMichele, Megan Comfort, Shilpi Misra, Kelle Barrick, and Peter Baumgartner. 2018. The intuitive-override model: Nudging judges toward pretrial risk assessment instruments. *Available at SSRN 3168500* (2018).
- [24] Sarah L Desmarais, Samantha A Zottola, Sarah E Duhart Clarke, and Evan M Lowder. 2020. Predictive validity of pretrial risk assessments: A systematic review of the literature. *Criminal Justice and Behavior* (2020), 0093854820932959.
- [25] William Dieterich, Christina Mendoza, and Tim Brennan. 2016. COMPAS Risk Scales: Demonstrating Accuracy Equity and Predictive Parity. (2016).
- [26] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science advances* 4, 1 (2018), eaao5580.
- [27] Bradley Efron. 1971. Forcing a sequential experiment to be balanced. *Biometrika* 58, 3 (1971), 403–417.
- [28] Birte Englich, Thomas Mussweiler, and Fritz Strack. 2005. The last word in court—A hidden disadvantage for the defense. *Law and Human Behavior* 29, 6 (2005), 705–722.
- [29] Birte Englich, Thomas Mussweiler, and Fritz Strack. 2006. Playing dice with criminal sentences: The influence of irrelevant anchors on experts' judicial decision making. *Personality and Social Psychology Bulletin* 32, 2 (2006), 188–200.
- [30] Birte Englich and Thomas Mussweiler. 2001. Sentencing under uncertainty: Anchoring effects in the courtroom 1. *Journal of applied social psychology* 31, 7 (2001), 1535–1551.
- [31] Nicholas Epley and Thomas Gilovich. 2005. When effortful thinking influences judgmental anchoring: differential effects of forewarning and incentives on self-generated and externally provided anchors. *Journal of Behavioral Decision Making* 18, 3 (2005), 199–212.
- [32] Jonathan St BT Evans, David E Over, et al. 2004. *If: Supposition, pragmatics, and dual processes*. Oxford University Press, USA.
- [33] Riccardo Fogliato, Alexandra Chouldechova, and Max G'Sell. 2020. Fairness Evaluation in Presence of Biased Noisy Labels. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 2325–2336.
- [34] Karën Fort, Gilles Adda, and K Bretonnel Cohen. 2011. Amazon mechanical turk: Gold mine or coal mine? *Computational Linguistics* 37, 2 (2011), 413–420.
- [35] Francesca Gino and Don A Moore. 2007. Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making* 20, 1 (2007), 21–35.
- [36] Mark A Gluck and Gordon H Bower. 1988. From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General* 117, 3 (1988), 227.
- [37] Glenn Grant. 2019. Report to the Governor and the Legislature. *New Jersey Courts*. (2019).
- [38] 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*. ACM, 90–99.
- [39] Ben Green and Yiling Chen. 2019. The principles and limits of algorithm-in-the-loop decision making. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–24.
- [40] Nina Grgić-Hlača, Christoph Engel, and Krishna P Gummadi. 2019. Human decision making with machine assistance: An experiment on bailing and jailing. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–25.
- [41] Chris Guthrie, Jeffrey J Rachlinski, and Andrew J Wistrich. 2000. Inside the judicial mind. *Cornell L. Rev* 86 (2000), 777.
- [42] Jordan Hyatt and Steven L Chanenson. 2016. The use of risk assessment at sentencing: Implications for research and policy. *Villanova Law/Public Policy Research Paper* 2017-1040 (2016).
- [43] Jongbin Jung, Sharad Goel, Jennifer Skeem, et al. 2020. The limits of human predictions of recidivism. *Science advances* 6, 7 (2020), eaaz0652.
- [44] Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*. World Scientific, 99–127.
- [45] Alex Kale, Matthew Kay, and Jessica Hullman. 2020. Visual reasoning strategies for effect size judgments and decisions. *IEEE Transactions on Visualization and Computer Graphics* (2020).
- [46] Nathan Kallus and Angela Zhou. 2019. Assessing Disparate Impact of Personalized Interventions: Identifiability and Bounds. In *Advances in Neural Information Processing Systems*. 3426–3437.
- [47] Jon Kleinberg, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2017. Human decisions and machine predictions. *The quarterly journal of economics* 133, 1 (2017), 237–293.
- [48] Jonathan J Koehler. 1996. The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and brain sciences* 19, 1 (1996), 1–17.
- [49] Max Kuhn and Hadley Wickham. 2020. *tidymodels: Easily Install and Load the 'Tidymodels' Packages*. <https://CRAN.R-project.org/package=tidymodels> R package version 0.1.0.
- [50] Vivian Lai and Chenhao Tan. 2019. On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 29–38.
- [51] Himabindu Lakkaraju, Jon Kleinberg, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2017. The selective labels problem: Evaluating algorithmic predictions in the presence of unobservables. In *Proceedings of the 23rd ACM SIGKDD*

- International Conference on Knowledge Discovery and Data Mining*. 275–284.
- [52] Isaac M Lipkus, Greg Samsa, and Barbara K Rimer. 2001. General performance on a numeracy scale among highly educated samples. *Medical decision making* 21, 1 (2001), 37–44.
 - [53] Keri Mallari, Kori Inkpen, Paul Johns, Sarah Tan, Divya Ramesh, and Ece Kamar. 2020. Do I Look Like a Criminal? Examining how Race Presentation Impacts Human Judgement of Recidivism. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–13.
 - [54] Charles F Manski. 2004. Measuring expectations. *Econometrica* 72, 5 (2004), 1329–1376.
 - [55] Winter Mason and Siddharth Suri. 2012. Conducting behavioral research on Amazon’s Mechanical Turk. *Behavior research methods* 44, 1 (2012), 1–23.
 - [56] Cate Metz and Adam Satariano. 2020. An Algorithm that Grants Freedom, or Takes it Away. <https://www.nytimes.com/2020/02/06/technology/predictive-algorithms-crime.html>
 - [57] Thomas Mussweiler and Fritz Strack. 2000. Numeric judgments under uncertainty: The role of knowledge in anchoring. *Journal of experimental social psychology* 36, 5 (2000), 495–518.
 - [58] Daniel S Nagin et al. 2013. Deterrence: A review of the evidence by a criminologist for economists. *Annual Review of Economics* 5, 1 (2013), 83–105.
 - [59] Richard E Nisbett and Timothy D Wilson. 1977. Telling more than we can know: verbal reports on mental processes. *Psychological review* 84, 3 (1977), 231.
 - [60] Anthony O’Hagan, Caitlin E Buck, Alireza Daneshkhah, J Richard Eiser, Paul H Garthwaite, David J Jenkinson, Jeremy E Oakley, and Tim Rakow. 2006. *Uncertain judgements: eliciting experts’ probabilities*. John Wiley & Sons.
 - [61] Pennsylvania Commission on Sentencing. 2020. An Overview of the Sentence Risk Assessment Instrument. https://sentencing.umn.edu/sites/sentencing.umn.edu/files/pennsylvania_overview_of_the_sentencing_risk_assessment_instrument_2017.pdf
 - [62] The Pennsylvania Commission on Sentencing. 2020. Judicial System General Provisions. Part VIII Criminal Sentencing. Chapter 305. Sentence Risk Assessment Instrument. <http://pcs.la.psu.edu/guidelines/sentence-risk-assessment-instrument/sentence-risk-assessment-instrument-guideline-text>
 - [63] Daniel Paravisini and Antoinette Schoar. 2013. *The incentive effect of scores: Randomized evidence from credit committees*. Technical Report. National Bureau of Economic Research.
 - [64] Pennsylvania Commission on Sentencing. 2020. Sentencing Guidelines and Implementation Manuals. <http://pcs.la.psu.edu/guidelines/sentencing-guidelines-and-implementation-manuals>
 - [65] Forough Poursabzi-Sangdeh, Daniel G Goldstein, Jake M Hofman, Jennifer Wortman Vaughan, and Hanna Wallach. 2021. Manipulating and measuring model interpretability. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–52.
 - [66] Drazen Prelec. 1998. The probability weighting function. *Econometrica* (1998), 497–527.
 - [67] R Core Team. 2019. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>
 - [68] Manish Raghavan, Solon Barocas, Jon Kleinberg, and Karen Levy. 2020. Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 469–481.
 - [69] Elissa M Redmiles, Sean Kross, and Michelle L Mazurek. 2019. How well do my results generalize? comparing security and privacy survey results from mturk, web, and telephone samples. In *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1326–1343.
 - [70] Catherine A Roster, Lorenzo Lucianetti, and Gerald Albaum. 2015. Exploring slider vs. categorical response formats in web-based surveys. *Journal of Research Practice* 11, 1 (2015), D1–D1.
 - [71] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the US Child Welfare System. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
 - [72] Lisa M Schwartz, Steven Woloshin, William C Black, and H Gilbert Welch. 1997. The role of numeracy in understanding the benefit of screening mammography. *Annals of internal medicine* 127, 11 (1997), 966–972.
 - [73] Nicholas Scurich and Richard S John. 2011. The effect of framing actuarial risk probabilities on involuntary civil commitment decisions. *Law and human behavior* 35, 2 (2011), 83–91.
 - [74] Mark Sendak, Madeleine Clare Elish, Michael Gao, Joseph Futoma, William Ratliff, Marshall Nichols, Armando Bedoya, Suresh Balu, and Cara O’Brien. 2020. "The human body is a black box" supporting clinical decision-making with deep learning. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 99–109.
 - [75] Jennifer Skeem and Christopher Lowenkamp. 2020. Using algorithms to address trade-offs inherent in predicting recidivism. *Behavioral Sciences & the Law* (2020).
 - [76] Jennifer L Skeem and Christopher T Lowenkamp. 2016. Risk, race, and recidivism: Predictive bias and disparate impact. *Criminology* 54, 4 (2016), 680–712.
 - [77] Megan Stevenson. 2018. Assessing risk assessment in action. *Minn. L. Rev.* 103 (2018), 303.

- [78] Sarah Tan, Julius Adebayo, Kori Inkpen, and Ece Kamar. 2018. Investigating Human+ Machine Complementarity for Recidivism Predictions. *arXiv preprint arXiv:1808.09123* (2018).
- [79] Robert Tibshirani. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* 58, 1 (1996), 267–288.
- [80] Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science* 185, 4157 (1974), 1124–1131.
- [81] Michelle Vaccaro and Jim Waldo. 2019. The Effects of Mixing Machine Learning and Human Judgment. *Queue* 17, 4 (2019), 30.
- [82] Vivianne HM Visschers, Ree M Meertens, Wim WF Passchier, and Nanne NK De Vries. 2009. Probability information in risk communication: a review of the research literature. *Risk Analysis: An International Journal* 29, 2 (2009), 267–287.
- [83] Hadley Wickham, Mara Averick, Jennifer Bryan, Winston Chang, Lucy D’Agostino McGowan, Romain François, Garrett Grolemund, Alex Hayes, Lionel Henry, Jim Hester, Max Kuhn, Thomas Lin Pedersen, Evan Miller, Stephan Milton Bache, Kirill Müller, Jeroen Ooms, David Robinson, Dana Paige Seidel, Vitalie Spinu, Kohske Takahashi, Davis Vaughan, Claus Wilke, Kara Woo, and Hiroaki Yutani. 2019. Welcome to the tidyverse. *Journal of Open Source Software* 4, 43 (2019), 1686. <https://doi.org/10.21105/joss.01686>
- [84] Timothy D Wilson. 2004. *Strangers to ourselves*. Harvard University Press.
- [85] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the Effect of Accuracy on Trust in Machine Learning Models. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 279.
- [86] Yunfeng Zhang, Q Vera Liao, and Rachel KE Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in AI-assisted decision making. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 295–305.

Received January 2021; revised April 2021; revised July 2021; accepted July 2021