Journal of the Royal Statistical Society Series A: Statistics in Society, 2023, **186**, 167–189
https://doi.org/10.1093/jrsssa/qnad010
Advance access publication 2 February 2023
Discussion Paper

ROYAL STATISTICAL SOCIETY DATA | EVIDENCE | DECISIONS
Journal of the Statistics Society
Series A
Statistics in Society

Experimental evaluation of algorithm-assisted human decision-making: application to pretrial public safety assessment*

Kosuke Imai[1,2], Zhichao Jiang[3], D. James Greiner[4], Ryan Halen[5] and Sooahn Shin[1]

[1]Department of Government, Harvard University, 1737 Cambridge Street, Cambridge, MA 02138, USA
[2]Department of Statistics, Harvard University, 1 Oxford Street, Cambridge, MA 02138, USA
[3]School of Mathematics, Sun Yat-sen University, Guangzhou, Guangdong 510275, China
[4]Harvard Law School, 1525 Massachusetts Avenue, Griswold 504, Cambridge, MA 02138, USA
[5]Harvard Law School, 1607 Massachusetts Avenue, Third Floor, Cambridge, MA 02138, USA

Address for correspondence: Kosuke Imai, Department of Government, Harvard University, 1737 Cambridge Street, Cambridge, MA 02138, USA; Department of Statistics, Harvard University, 1 Oxford Street, Cambridge, MA 02138, USA.
Email: imai@harvard.edu

*Read before The Royal Statistical Society at a meeting organized by the Discussion Meetings Committee, the Statistics and the Law Section, and the Data Ethics and Governance Section, on Tuesday, 8th February 2022, Dr. S. Coleman in the Chair.

Abstract

Despite an increasing reliance on fully-automated algorithmic decision-making in our day-to-day lives, humans still make consequential decisions. While the existing literature focuses on the bias and fairness of algorithmic recommendations, an overlooked question is whether they improve human decisions. We develop a general statistical methodology for experimentally evaluating the causal impacts of algorithmic recommendations on human decisions. We also examine whether algorithmic recommendations improve the fairness of human decisions and derive the optimal decision rules under various settings. We apply the proposed methodology to the first-ever randomized controlled trial that evaluates the pretrial Public Safety Assessment in the United States criminal justice system. Our analysis of the preliminary data shows that providing the PSA to the judge has little overall impact on the judge's decisions and subsequent arrestee behaviour.

Keywords: algorithmic fairness, causal inference, principal stratification, randomized experiments, recommendation systems, sensitivity analysis

1 Introduction

A growing body of literature has suggested the potential superiority of algorithmic decision-making over purely human choices across a variety of tasks (e.g., Hansen & Hasan, 2015; He et al., 2015). Although some of this evidence is decades old (e.g., Dawes et al., 1989), it has recently gained significant public attention for the spectacular defeats of humanity's best in cerebral games (e.g., Silver et al., 2018). Yet, even in contexts where research has warned of human frailties, we humans still make many consequential decisions for a variety of reasons including the preservation of human agency and accountability.

The desire for a human decision-maker as well as the precision and efficiency of algorithms have led to the adoption of hybrid systems involving both. By far the most popular system uses algorithmic recommendations to inform human decision-making. Such algorithm-assisted human decision-making has been deployed in many aspects of our daily lives, including medicine, hiring, credit lending, investment decisions, and online shopping. And of particular interest, algorithmic

Received: January 25, 2021. **Revised:** October 2, 2021. **Accepted:** October 7, 2021
© (RSS) Royal Statistical Society 2023. All rights reserved. For permissions, please e-mail: journals.permissions@oup.com

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

recommendations are increasingly of use in the realm of evidence-based public policy-making. A prominent example, studied in this paper, is the use of risk assessment instruments in the criminal justice system that are designed to improve incarceration rulings and other decisions made by judges.

While there exists a fast-growing literature in computer science that studies the bias and fairness of algorithms (see Chouldechova & Roth, 2020, for a review and many references therein), an overlooked question is whether such algorithms help humans make better decisions (see e.g., Green & Chen, 2019, for an exception). In this paper, we develop a general methodological framework for experimentally evaluating the impacts of algorithmic recommendations on human decision-making. We conducted the first-ever real-world field experiment by providing, for a randomly selected cases, information from a system consisting of Public Safety Assessment (PSA) risk scores and a recommendation from a Decision-Making Framework (DMF) to a judge who makes an initial release decision. We evaluate whether the PSA-DMF system (which for brevity we refer to as the PSA hereafter) helps judges achieve their goal of preventing arrestees from committing a new crime or failing to appear in court while avoiding an unnecessarily harsh decision.

Using the concept of principal stratification from the causal inference literature (e.g., Ding & Lu, 2017; Frangakis & Rubin, 2002), we propose the evaluation quantities of interest, identification assumptions, and estimation strategies. We also develop sensitivity analyses to assess the robustness of empirical findings to the potential violation of a key identification assumption (see also Hirano et al., 2000; Jiang et al., 2016; Mattei et al., 2013; Schwartz et al., 2011). In addition, we examine whether algorithmic recommendations improve the fairness of human decisions, using the concept of principal fairness that, unlike other fairness criteria, accounts for how the decision in question affects individuals (Imai & Jiang, 2023). Finally, we consider how the data from an experimental evaluation can be used to inform an optimal decision rule and assess the optimality of algorithmic recommendations and human decisions (see Ben-Michael et al., 2021, for a methodological framework for learning an optimal algorithmic recommendation). Although we describe and apply the proposed methodology in the context of evaluating the PSA, it is directly applicable or extendable to many other settings of algorithm-assisted human decision-making.

The use of risk assessment scores, which serves as the main application of the current paper, has played a prominent role in the literature on algorithmic fairness since the controversy over the potential racial bias of COMPAS risk assessment score used in the United States (US) criminal justice system (see e.g., Angwin et al., 2016; Dieterich et al., 2016; Dressel & Farid, 2018; Flores et al., 2016). With few exceptions, however, much of this debate focused on the accuracy and fairness properties of risk assessment scores itself rather than how they affect judges' decisions (see e.g., Berk et al., 2018; Kleinberg et al., 2018; Rudin et al., 2020, and references therein). Even studies that directly estimate the impacts of risk assessment scores on judges' decisions are based on either observational data or hypothetical vignettes in surveys (e.g., Albright, 2019; Berk, 2017; Garrett & Monahan, 2020; Green & Chen, 2019; Miller & Maloney, 2013; Skeem et al., 2020; Stevenson, 2018; Stevenson & Doleac, 2021).

We contribute to this literature by demonstrating how to evaluate the use of risk assessment scores experimentally when humans are ultimate decision-makers. To the best of our knowledge, this is the first real-world randomized controlled trial (RCT) that evaluates the impacts of algorithmic risk assessment scores on judges' decisions in the criminal justice system (see also the Manhattan Bail Project and Philadelphia Bail Experiment that evaluated the effects of bail guidelines on judges' decisions several decades ago Ares et al., 1963; Goldkamp & Gottfredson, 1984, 1985). Using the concept of principal stratification from causal inference literature, the proposed methodology allows us to evaluate the effects of the PSA on judges' decisions separately for the subgroups of arrestees with different levels of risks.

Based on the preliminary data from our experiment (complete data will not be available for some time), we find that the provision of the PSA has little overall impact on the judge's decisions across three outcomes we examine: failure to appear (FTA), new criminal activity (NCA), and new violent criminal activity (NVCA). Our analysis, however, provides some suggestive evidence that the PSA may make the judge's decisions more lenient for female arrestees regardless of their risk levels, while it encourages the judge to make stricter decisions for male arrestees who are deemed to be risky. In terms of fairness, the PSA appears to increase an existing gender difference while having no substantial impact on any racial differences in judges' decisions. Finally, we use the

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

experimental data to learn about the optimal decision rule that minimizes the prevalence of negative outcomes (FTA, NCA, and NVCA) while avoiding unnecessarily harsh decisions. Our analysis suggests that the PSA's recommendations may be unnecessarily severe unless a jurisdiction considers the costs of FTA, NCA, and NVCA to be sufficiently high. This might suggest that incarceration decisions themselves, whether PSA-informed or otherwise, are also unnecessarily severe.

2 Experimental evaluation of pretrial public safety assessment

In this section, we briefly describe our field experiment after providing some background about the use of the PSA in the US criminal justice system. Additional details about our experiment are given in Greiner et al. (2020).

2.1 Background

The US criminal justice apparatus consists of thousands of diverse systems. Some are similar in the decision points they feature, as an individual suspected of a crime travels from the investigation to sentencing. Common decision points include whether to stop and frisk an individual in a public place, whether to arrest or issue a citation to an individual suspect of committing a crime, whether to release the arrestee while they await the disposition of any charges against them (the subject of this paper), what charge(s) to be filed against the individual, whether to find the defendant guilty of those charges, and what sentence to impose on a defendant found guilty.

At present, human judges make all these decisions. In theory, algorithms could inform any of them and could even make some of these decisions without human involvement. To date, algorithmic outputs have appeared most frequently in two settings: (i) at the 'first appearance' hearing, during which a judge decides whether to release an arrestee pending disposition of any criminal charges and (ii) at sentencing, in which the judge imposes a punishment on a defendant found guilty. The first of these two motivates the present paper, but the proposed methodology is applicable or extendable to other settings.

We describe a typical first appearance hearing. The key decision the judge must make at a first appearance hearing is whether to release the arrestee pending disposition of any criminal charges and, if the arrestee is to be released, what conditions to impose. Almost all jurisdictions allow the judge to release the arrestee with only a promise to reappear at subsequent court dates. In addition, because arrestees have not yet been adjudicated guilty of any charge at the time of a pretrial hearing, there exists a consensus that pretrial incarceration is to be avoided unless the risks associated with release are sufficiently high.

Judges deciding whether to release arrestees ordinarily consider two risk factors among a variety of other concerns; the risk that the arrestee will fail to appear at subsequent court dates, and the risk that the arrestee will engage in new criminal activity (NCA) before the case is resolved (e.g., 18 U.S.C. § 3142(e)(1)). Jurisdiction laws vary regarding how these two risks are to be weighed. Some jurisdictions direct judges to consider both simultaneously along with other factors (e.g., Ariz. Const. art. II, § 22, Iowa Code § 811.2(1)(a)), while others focus on only FTA risk (e.g., N.Y. Crim. Proc. Law § 510.30(2)(a)). Despite these variations, NCA and FTA are constant and prominent in the debate over the first appearance decisions.

Concerns about the consequential nature of the first appearance decision have led to the development of the PSA, which is ordinarily offered as an input to first appearance judges. Predisposition risk assessment instruments take various forms, but most focus on classifying arrestees according to FTA and NCA risks. They are generally constructed by fitting a statistical model to a training dataset based on the past first appearance hearings and the subsequent incidences (or lack thereof) of FTA and NCA. The hope is that providing such instruments will improve the assessment of FTA and NCA risks and thereby lead to better decisions. The goal of this paper is to develop a general methodological framework for evaluating the impact of providing the PSA to judges at first appearance hearings using an RCT, to which we now turn.

2.2 The experiment

We conducted a field RCT in Dane county, Wisconsin, to evaluate the impacts of PSA provision on judges' decisions. The PSA consists of three scores—two six-point scores separately summarizing

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

FTA and NCA risks as well as a binary score for the risk of NVCA. These scores are based on the weighted indices of nine factors drawn from criminal history information, primarily prior convictions and FTA, and a single demographic factor, age. Notably, gender and race are not used to compute the PSA. The weights are calculated using past data. A Decision-Making Framework (DMF) combines information from the three PSA scores with other considerations to produce an overall recommendation to the judge, which the judge may accept or modify or ignore as they see fit. The details about the construction of the PSA and other relevant information are available at https://advancingpretrial.org/psa/factors/

The field operation was straightforward. In this county, a court employee assigned each matter a case number sequentially as it entered the system. No one but this clerk was aware of the pending matter numbers, so manipulation of the number by charging assistant district attorneys was not possible. Employees of the Clerk's office scanned online record systems to calculate the PSA for all cases. If the last digit of the case number was even, these employees made the PSA (specifically, a printout of the PSA scores, the DMF recommendation, and the supporting criminal history and age information) available to the judge. Otherwise, no PSA was made available. Thus, the provision of the PSA to judges was essentially randomized. Indeed, the comparison of the observed covariate distributions suggests that this scheme produced groups comparable on background variables (Greiner et al., 2020).

The judge presiding over the first appearance hearing by law was to consider the risk of FTA and NCA, along with other factors, including ties to the community as prescribed by statute. The judge could order the arrestee released with or without bail of varying amount. The judge could also condition release on compliance with certain conditions such as monitoring, but for the sake of simplicity, we focus on bail decisions and ignore other conditions in this paper.

When making decisions, the judge also had information other than the PSA and its inputs. In all cases, the judge had a copy of an affidavit sworn to by a police officer recounting the circumstances of the incident that led to the arrest. The defense attorney sometimes informed the judge of the following regarding the arrestee's connections to the community: length of time lived there, employment there, and family living there. When available, this information ordinarily stemmed from an arrestee interview conducted earlier by a paralegal. The assistant district attorney sometimes provided additional information regarding the circumstances of the arrest or criminal history. Given the lack of access to this additional information, we develop a sensitivity analysis to address a potential unobserved confounding bias.

2.3 The data

The field operation design called for approximately a 30-month treatment assignment period (from the middle of 2017 until the end of 2019) followed by the collection of data on FTA, NCA, NVCA, and other outcomes for a period of two years after randomization. At the time of this writing, we have outcome data from a 12-month follow-up of each first appearance event that occurred in the first 12 months of randomization. The 30-month randomization period has expired, and we will report the results of our comprehensive analysis of a full dataset in the future. Furthermore, although some arrestees had multiple cases during the study period, this paper focuses only on the first of the first appearance hearings for any individual arrestee. This leads to a total of 1891 cases for our analysis, of which 40.0% (38.8%) are white male arrestees and 13.0 are white female arrestees (non-white male and female arrestees account for 38.8% and 8.1%, respectively).

Based on the empirical distribution of bail amounts and expert's opinion, we categorize the judge's decisions into three ordinal categories: signature bond, small cash bond (less than $1,000), and large cash bond (greater than or equal to $1,000). A signature bond requires an arrestee to sign a promise to return to the court for trial but does not require any payment to be released. Cash bonds require an arrestee to deposit money with the court to obtain release. Table 1 summarizes the joint distribution of treatment assignment (PSA provision), the judge's decisions (three ordinal categories), and three binary outcomes. We observe that in about three quarters of cases, the judge imposed signature bonds, while in the remaining cases, the judge imposed bail. For the outcome variables, slightly less than 30% of arrestees commit FTA or NCA whereas the proportion of those who commit NVCA is only about 6%.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Table 1. The Joint Distribution of Treatment Assignment, Judge's Decisions, and Outcomes. The table shows the number of cases in each category with the corresponding percentage in parentheses. Only about 20% of all arrestees are female. Few cases result in NVCA (new violent criminal activity), while FTA (failure to appear in court) and NCA (new criminal activity) occur in slightly above 25% each. A majority of decisions are signature bonds rather than cash bonds.

	***no* PSA (Control group)**			**PSA (Treatment group)**			
	Signature	**Cash bond**		**Signature**	**Cash bond**		
	bond	**≤$1,000**	**>$1,000**	**bond**	**≤$1,000**	**>$1,000**	**Total (%)**
Non-white Female	64	11	6	67	6	0	154
	(3.4)	(0.6)	(0.3)	(3.5)	(0.3)	(0.0)	(8.1)
White Female	91	17	7	104	17	10	246
	(4.8)	(0.9)	(0.4)	(5.5)	(0.9)	(0.5)	(13.0)
Non-white Male	261	56	49	258	53	57	734
	(13.8)	(3.0)	(2.6)	(13.6)	(2.8)	(3.0)	(38.8)
White Male	289	48	44	276	54	46	757
	(15.3)	(2.5)	(2.3)	(14.6)	(2.9)	(2.4)	(40.0)
FTA committed	218	42	16	221	45	16	558
	(11.5)	(2.2)	(0.8)	(11.7)	(2.4)	(0.8)	(29.4)
not committed	487	90	90	484	85	97	1333
	(25.8)	(4.8)	(4.8)	(25.6)	(4.5)	(5.1)	(70.6)
NCA committed	211	39	14	202	40	17	523
	(11.2)	(2.1)	(0.7)	(10.7)	(2.1)	(0.9)	(27.7)
not committed	494	93	92	503	90	96	1368
	(26.1)	(4.9)	(4.9)	(26.6)	(4.8)	(5.1)	(72.4)
NVCA committed	36	10	3	44	10	6	109
	(1.9)	(0.5)	(0.2)	(2.3)	(0.5)	(0.3)	(5.7)
not committed	669	122	103	661	120	107	1782
	(35.4)	(6.5)	(5.4)	(35.0)	(6.3)	(5.7)	(94.3)
Total	705	132	106	705	130	113	1891
	(37.3)	(7.0)	(5.6)	(37.3)	(6.9)	(6.0)	(100)

2.4 The overall impact of PSA provision on judge's decisions

Figure 1 presents the distribution of the judge's decisions given each of the PSA scores among the cases in the treatment (top panel) and control (bottom panel) groups. The overall difference in the conditional distribution between the two groups is small though there are some differences in some subgroups (see Online Supplementary Material, Appendix S1). The PSA scores for FTA and NCA are ordinal, ranging from 1 (safest) to 6 (riskiest), whereas the PSA score for NVCA is binary, 0 (safe) and 1 (risky). We also plot the DMF recommendation, which aggregates these three PSA scores as well as other information such as types of charges. The DMF recommendation has four categories (signature bond, modest cash bond, moderate cash bond, and cash bond with maximum conditions), but we dichotomize it into signature or cash bond given its skewed empirical distribution.

In general, we observe a positive association between the PSA scores and judge's decisions, implying that a higher PSA score is associated with a harsher decision. We also find that for FTA and NCA, the most likely scores are in the medium range, while the vast majority of NVCA cases were classified as no elevated risk. For NCA and FTA, the judge's decisions varied little when the PSA score took a value in the lower range. For the DMF recommendation, the judge is far more likely to give a signature bond for the cases that are actually recommended for a signature bond.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

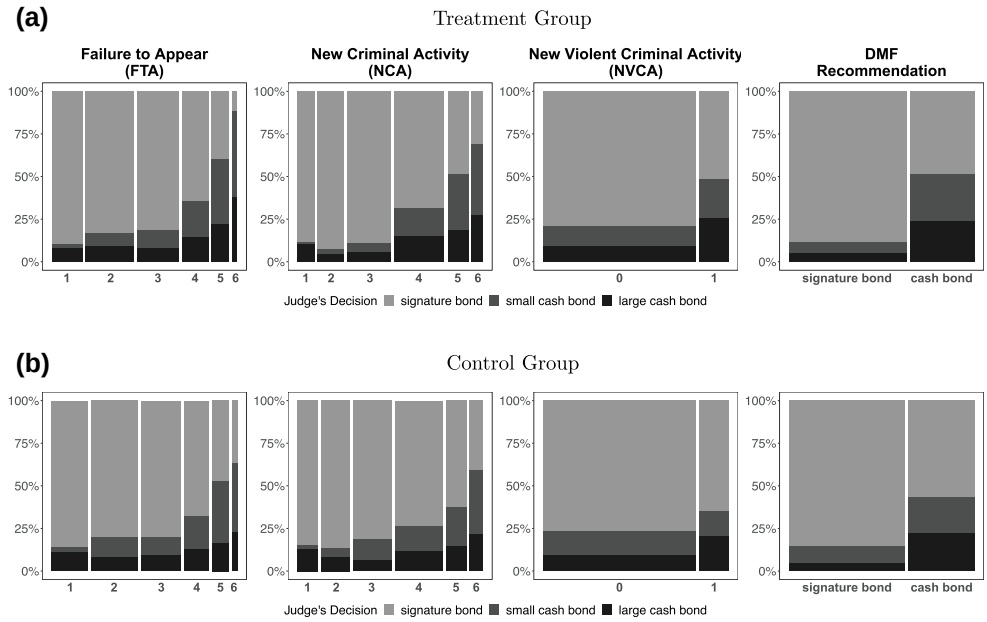

Figure 1. The distribution of the Judge's Decisions given the Pretrial Public Safety Assessment (PSA) among the Cases in the Treatment (Top Panel) and Control (Bottom Panel) Groups. There are three PSA scores, two of which are ordinal—FTA and NCA—while the other is dichotomous—NVCA. The judge's decision is coded as a three-category ordinal variable based on the type and amount of bail: a signature bond, a small cash bond (less than $1,000), and a large cash bond (greater than or equal to $1,000). The DMF recommendation is presented as a binary variable: signature or cash bond. The width of each bar is proportional to the number of cases for each value of the corresponding PSA score. There exists a positive correlation between PSA scores and the severity of the judge's decisions in both treatment and control groups. (a) Treatment group. (b) Control group.

Figure 2 presents the estimated average causal effect of PSA provision on the judge's decisions (left plot) and three outcomes of interest (right plot). We use the difference-in-means estimator and display the 95% confidence intervals as well as the point estimates. We do not compute separate estimates for white females and non-white females because we have too few female arrestees (see Table 1). The results imply that PSA provision, on average, has little effect on the judge's decisions. In addition, the average effects of PSA provision on the three outcomes are also largely ambiguous although there is suggestive evidence that it may slightly increase NVCA among female arrestees. In Online Supplementary Material, Appendix S2.1, we also explore the average causal effects of PSA provision across different age groups. We find some suggestive causal effects for the group of 29–35 year old arrestees.

Although these results show whether PSA provision leads to a harsher or more lenient decision (and whether it increases or decreases the proportions of negative outcomes), they are not informative about whether it helps judges make better decisions. In the current context, a primary goal of the judge is to make lenient decisions in low-risk cases and less lenient decisions in high-risk cases. If the PSA is helpful, therefore, its provision should encourage the judge to impose small or no bail on safe cases and impose a greater amount of bail on risky cases (we formally define 'safe' and 'risky' cases below). This demands the study of an important causal heterogeneity by distinguishing among cases with different risk levels. In addition, we may also be interested in knowing how PSA provision affects the gender and racial fairness of judges' decisions. Thus, the goal of the remainder of the paper is to develop statistical methods that directly address these and other questions.

3 The proposed evaluation methodology

In this section, we describe the proposed methodology for experimentally evaluating the impacts of algorithmic recommendations on human decision-making. Although we refer to our specific application throughout, the proposed methodology can be applied or extended to other settings, in

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

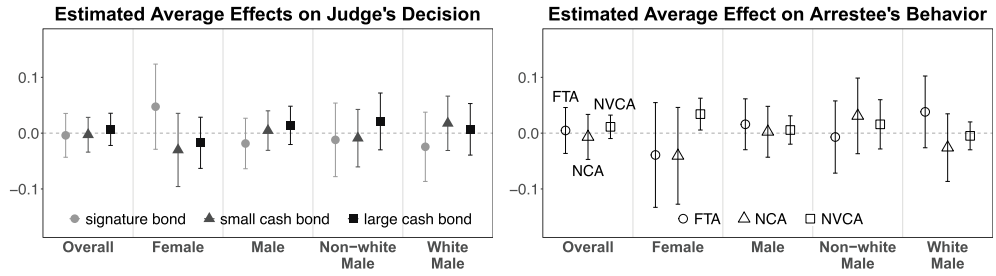

Figure 2. Estimated average causal effects of PSA provision on the judge's decisions and outcome variables. The results are based on the difference-in-means estimator. The vertical bars represent the 95% confidence intervals. In the left plot, we report the estimated effects of PSA provision on the judge's decision to charge a signature bond (solid circles), a small cash bail ($1,000 dollars or less; solid triangles), and a large cash bail (greater than $1,000; solid squares). In the right plot, we report the estimated effects of PSA provision on the three different outcome variables: FTA (open circles), NCA (open triangles), and NVCA (open squares). PSA provision appears to have little overall effect on the judge's decision and arrestee's behaviour, on average, though it may slightly increase NVCA among female arrestees.

which humans make decisions using algorithmic recommendations as an input. We will begin by considering a binary decision and then extend our methodology to an ordinal decision in Section 3.4.

3.1 The setup

Let Z_i be a binary treatment variable indicating whether the PSA is presented to the judge of case $i = 1, 2, \ldots, n$. We use D_i to denote the binary detention decision made by the judge to either detain ($D_i = 1$) or release ($D_i = 0$) the arrestee prior to the trial. In addition, let Y_i represent the binary outcome: we code all our outcomes—NCA, NVCA, and FTA—as binary variables. For example, $Y_i = 1$ ($Y_i = 0$) implies that the arrestee of case i commits (does not commit) an NCA. Finally, we use $\mathbf{X}_i$ to denote a vector of observed pre-treatment covariates for case i. They include age, gender, race, and prior criminal history.

We adopt the potential outcomes framework of causal inference and assume the stable unit treatment value assumption (SUTVA) (Rubin, 1990). In particular, we assume no interference among cases, implying that the treatment assignment for one case does not influence the judge's decision and outcome variable in another case. This assumption is reasonable in our analysis because we focus only on first arrests and do not analyse cases with subsequent arrests. Online Supplementary Material, Appendix S3 provides the empirical evidence in support of this assumption.

Let $D_i(z)$ be the potential value of the pretrial detention decision if case i is assigned to the treatment condition $z \in \{0, 1\}$. Furthermore, $Y_i(z, d)$ represents the potential outcome under the scenario, in which case i is assigned to the treatment condition z and the judge makes the decision $d \in \{0, 1\}$. Then, the observed decision is given by $D_i = D(Z_i)$, whereas the observed outcome is denoted by $Y_i = Y_i(Z_i, D_i(Z_i))$.

Throughout this paper, we maintain the following three assumptions, all of which we believe are reasonable in our application. First, because the treatment assignment is essentially randomized, the following independence assumption is automatically satisfied.

Assumption 1 (Randomization of the Treatment Assignment).

$$\{D_i(z), Y_i(z, d), \mathbf{X}_i\} \perp\!\!\!\perp Z_i \quad \text{for } z \in \{0, 1\} \text{ and all } d.$$

Second, we assume that the provision of the PSA influences the outcome only through the judge's decision. Because an arrestee would not care and, perhaps, would not even know whether the judge is presented with the PSA at their first appearance, it is reasonable to assume that their behaviour, be it NCA, NVCA, or FTA, is not affected directly by the treatment assignment.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Assumption 2 (Exclusion Restriction).

$$Y_i(z, d) = Y_i(z', d) \quad \text{for } z, z' \in \{0, 1\} \text{ and all } i, d.$$

Under Assumption 2, we can simplify our notation by writing $Y_i(z, d)$ as $Y_i(d)$. A potential violation of this assumption is that the PSA may directly influence the judge's decision about release conditions, which can in turn affect the outcome. The extension of the proposed methodology to multi-dimensional decisions, including the bail amount and monitoring conditions, is left for future research.

Finally, we assume that the judge's decision monotonically affects the outcome. Thus, for NCA (NVCA), the assumption implies that each arrestee is no less likely to commit a new (violent) crime if released. If FTA is the outcome of interest, this assumption implies that an arrestee is no more likely to appear in court if released. The assumption is reasonable because being held in custody of a court makes it difficult to engage in NCA, NVCA, and FTA.

Assumption 3 (Monotonicity).

$$Y_i(1) \leq Y_i(0) \quad \text{for all } i.$$

3.2 Causal quantities of interest

We define causal quantities of interest using principal strata that are determined by the joint values of potential outcomes, i.e., $(Y_i(1), Y_i(0)) = (y_1, y_0)$, where $y_1, y_0 \in \{0, 1\}$ (Frangakis & Rubin, 2002). Since Assumption 3 eliminates one principal stratum, $(Y_i(1), Y_i(0)) = (1, 0)$, there are three remaining principal strata. The stratum $(Y_i(1), Y_i(0)) = (0, 1)$ consists of those who would engage in NCA (NVCA or FTA) only if they are released. We call members of this stratum as 'preventable cases' because keeping those arrestees in custody would prevent the negative outcome (NCA, NVCA, or FTA). The stratum $(Y_i(1), Y_i(0)) = (1, 1)$ is called 'risky cases' and corresponds to those who always engage in NCA (NVCA or FTA) regardless of the judge's decision. In contrast, the stratum $(Y_i(1), Y_i(0)) = (0, 0)$ represents 'safe cases' in which the arrestees would never engage in NCA (NVCA or FTA) regardless of the detention decision.

We are interested in examining how PSA provision influences the judge's detention decisions across different types of cases. We define the following three average principal causal effects (APCE),

$$\text{APCEp} = \mathbb{E}\{D_i(1) - D_i(0) \mid Y_i(1) = 0, Y_i(0) = 1\}, \tag{1}$$

$$\text{APCEr} = \mathbb{E}\{D_i(1) - D_i(0) \mid Y_i(1) = 1, Y_i(0) = 1\}, \tag{2}$$

$$\text{APCEs} = \mathbb{E}\{D_i(1) - D_i(0) \mid Y_i(1) = 0, Y_i(0) = 0\}. \tag{3}$$

If the PSA is helpful, its provisions should make the judge more likely to detain the arrestees of preventable cases. That is, the principal causal effect on the detention decision for the preventable cases (APCEp) should be positive. In addition, the PSA should encourage the judge to release the arrestees of the safe cases, implying that the principal causal effect for the safe cases (APCEs) should be negative. The desirable direction of the principal causal effect for risky cases (APCEr) depends on various factors including the societal costs of holding the arrestees of this category in custody.

3.3 Nonparametric identification

We consider the nonparametric identification of the principal causal effects defined above. The following theorem shows that under the aforementioned assumptions, these effects can be identified up to the marginal distributions of $Y_i(d)$ for $d = 0, 1$.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Theorem 1 (Identification). Under Assumptions 1, 2, and 3,

$$APCEp = \frac{\Pr(Y_i = 1 \mid Z_i = 0) - \Pr(Y_i = 1 \mid Z_i = 1)}{\Pr\{Y_i(0) = 1\} - \Pr\{Y_i(1) = 1\}},$$
$$APCEr = \frac{\Pr(D_i = 1,\ Y_i = 1 \mid Z_i = 1) - \Pr(D_i = 1,\ Y_i = 1 \mid Z_i = 0)}{\Pr\{Y_i(1) = 1\}},$$
$$APCEs = \frac{\Pr(D_i = 0,\ Y_i = 0 \mid Z_i = 0) - \Pr(D_i = 0,\ Y_i = 0 \mid Z_i = 1)}{1 - \Pr\{Y_i(0) = 1\}}.$$

Proof is given in Online Supplementary Material, Appendix S4.2. Because $\Pr\{Y_i(d)\}$ is not identifiable without additional assumptions, we cannot estimate the causal effects based on Theorem 1. The denominators of the expressions on the right-hand side of Theorem 1, however, are positive under Assumption 3. As a result, the signs of the causal effects are identified from Theorem 1, which allows us to draw qualitative conclusions. In addition, the theorem implies that the sign of APCEp is the opposite of the sign of the average causal effect on the outcome. This is intuitive because if the provision of the PSA increases the probability of NCA (NVCA or FTA), then the judge must have released more arrestees for preventable cases.

Furthermore, we can obtain the nonparametric bounds on these causal quantities by bounding $\Pr\{Y_i(d) = y\}$ that appears in the denominators. From Assumption 1 and the law of total probability

$$\begin{aligned}\Pr\{Y_i(d) = 1\} &= \Pr\{Y_i(d) = 1 \mid Z_i = z\} \\ &= \Pr(Y_i = 1 \mid D_i = d, Z_i = z)\Pr(D_i = d \mid Z_i = z) \\ &\quad + \Pr\{Y_i(d) = 1 \mid D_i = 1 - d, Z_i = z\}\Pr(D_i = 1 - d \mid Z_i = z)\end{aligned}$$

for z, $d = 0$, 1. Under Assumption 3, the bounds on the unidentifiable terms are $\Pr\{Y_i = 1 \mid D_i = 1, Z_i = z\} \leq \Pr\{Y_i(0) = 1 \mid D_i = 1, Z_i = z\} \leq 1$ and $0 \leq \Pr\{Y_i(1) = 1 \mid D_i = 0, Z_i = z\} \leq \Pr\{Y_i = 1 \mid D_i = 0, Z_i = z\}$. This yields the following bounds on $\Pr\{Y_i(d) = 1\}$,

$$\max_z \Pr(Y_i = 1, D_i = 1 \mid Z_i = z) \leq \Pr\{Y_i(1) = 1\} \leq \min_z \Pr(Y_i = 1 \mid Z_i = z),$$
$$\max_z \Pr(Y_i = 1 \mid Z_i = z) \leq \Pr\{Y_i(0) = 1\} \leq 1 - \max_z \Pr(Y_i = 0, D_i = 0 \mid Z_i = z).$$

For point identification, we consider the following unconfoundedness assumption, which states that conditional on a set of observed pre-treatment covariates $\mathbf{X}_i$ and PSA provision, the judge's decision is independent of the potential outcomes.

Assumption 4 (Unconfoundedness).

$$Y_i(d) \perp\!\!\!\perp D_i \mid \mathbf{X}_i = \mathbf{x}, Z_i = z,$$

where we also assume $0 < \Pr(D_i = d \mid \mathbf{X}_i = \mathbf{x}, Z_i = z) < 1$ for $z \in \{0, 1\}$, and all $\mathbf{x} \in \mathcal{X}$ and d.

Assumption 4 holds if $\mathbf{X}_i$ contains all the information the judge has access to when making the detention decision under each treatment condition. As noted in Section 2.2, however, the judge may receive and use additional information regarding whether the arrestee has a job or a family in the jurisdiction, or perhaps regarding the length of time the arrestee has lived in the jurisdiction. If these factors have an impact on both the judge's decisions and arrestee's behaviours, then the

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

assumption is unlikely to be satisfied. Later, we address this issue by developing a sensitivity analysis for the potential violation of Assumption 4 (see Section 3.5).

To derive the identification result, consider the following principal scores (Ding & Lu, 2017), which represent in our application the population proportion (conditional on $\mathbf{X}_i$) of preventable, risky, and safe cases, respectively,

$$
\begin{aligned}
e_P(\mathbf{x}) &= \Pr\{Y_i(1) = 0,\ Y_i(0) = 1 \mid \mathbf{X}_i = \mathbf{x}\},\\
e_R(\mathbf{x}) &= \Pr\{Y_i(1) = 1,\ Y_i(0) = 1 \mid \mathbf{X}_i = \mathbf{x}\},\\
e_S(\mathbf{x}) &= \Pr\{Y_i(1) = 0,\ Y_i(0) = 0 \mid \mathbf{X}_i = \mathbf{x}\}.
\end{aligned}
$$

Under Assumptions 2, 3, and 4, we can identify the principal scores as,

$$
\begin{aligned}
e_P(\mathbf{x}) &= \Pr\{Y_i = 1 \mid D_i = 0, \mathbf{X}_i = \mathbf{x}\} - \Pr\{Y_i = 1 \mid D_i = 1, \mathbf{X}_i = \mathbf{x}\},\\
e_R(\mathbf{x}) &= \Pr\{Y_i = 1 \mid D_i = 1, \mathbf{X}_i = \mathbf{x}\},\\
e_S(\mathbf{x}) &= \Pr\{Y_i = 0 \mid D_i = 0, \mathbf{X}_i = \mathbf{x}\}.
\end{aligned}
$$

The next theorem shows that we can identify the APCE as the difference in the weighted average of judge's decisions between the treatment and control groups.

Theorem 2 (Identification under Unconfoundedness). Under Assumptions 1, 2, 3, and 4, APCEp, APCEr and APCEs are identified as,

$$
\begin{aligned}
APCEp &= \mathbb{E}\{w_P(\mathbf{X}_i)D_i \mid Z_i = 1\} - \mathbb{E}\{w_P(\mathbf{X}_i)D_i \mid Z_i = 0\},\\
APCEr &= \mathbb{E}\{w_R(\mathbf{X}_i)D_i \mid Z_i = 1\} - \mathbb{E}\{w_R(\mathbf{X}_i)D_i \mid Z_i = 0\},\\
APCEs &= \mathbb{E}\{w_S(\mathbf{X}_i)D_i \mid Z_i = 1\} - \mathbb{E}\{w_S(\mathbf{X}_i)D_i \mid Z_i = 0\},
\end{aligned}
$$

where

$$
w_P(\mathbf{x}) = \frac{e_P(\mathbf{x})}{\mathbb{E}\{e_P(\mathbf{X}_i)\}}, \quad w_R(\mathbf{x}) = \frac{e_R(\mathbf{x})}{\mathbb{E}\{e_R(\mathbf{X}_i)\}}, \quad w_S(\mathbf{x}) = \frac{e_S(\mathbf{x})}{\mathbb{E}\{e_S(\mathbf{X}_i)\}}.
$$

Proof is given in Online Supplementary Material, Appendix S4.2. Although Ding and Lu (2017) also identify principal causal effects using principal scores, they consider principal strata based on an intermediate variable. In contrast, we are interested in the causal effects on the decision within each principal stratum defined by the potential values of the outcome.

In some situations, we might consider the following strong monotonicity assumption instead of Assumption 3.

Assumption 5 (Strong Monotonicity).

$$
Y_i(1) = 0 \quad \text{for all } i.
$$

The assumption implies that the detention decision prevents NCA, NVCA, or FTA. The assumption is plausible for FTA, but may not hold for NCA/NVCA in some cases. In our data, for example, we find some NCA and NVCA among the incarcerated arrestees.

Under Assumption 5, the risky cases do not exist and hence the APCEr is not defined. This leads to the following identification result.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Theorem 3 (Identification under Strong Monotonicity). Under Assumptions 1, 2, and 5,

$$APCEp = \frac{\Pr(D_i = 0,\, Y_i = 1 \mid Z_i = 0) - \Pr(D_i = 0,\, Y_i = 1 \mid Z_i = 1)}{\Pr\{Y_i(0) = 1\}},$$
$$APCEs = \frac{\Pr(D_i = 0,\, Y_i = 0 \mid Z_i = 0) - \Pr(D_i = 0,\, Y_i = 0 \mid Z_i = 1)}{\Pr\{Y_i(0) = 0\}}.$$

Proof is given in Online Supplementary Material, Appendix S4.4. As in Theorem 1, the APCEp and APCEs depend on the distribution of $Y_i(0)$, which is not identifiable. However, as before, the sign of each effect is identifiable.

For point identification, we invoke the unconfoundedness assumption. Note that under the strong monotonicity assumption, Assumption 4 is equivalent to a weaker conditional independence relation concerning only one of the two potential outcomes,

$$Y_i(0) \perp\!\!\!\perp D_i \mid \mathbf{X}_i, Z_i = z$$

for $z = 0, 1$. We now present the identification result.

Theorem 4 (Identification under Unconfoundedness and Strong Monotonicity). Under Assumptions 1, 2, 4 and 5,

$$APCEp = \mathbb{E}\{w_P(\mathbf{X}_i) D_i \mid Z_i = 1\} - \mathbb{E}\{w_P(\mathbf{X}_i) D_i \mid Z_i = 0\},$$
$$APCEs = \mathbb{E}\{w_S(\mathbf{X}_i) D_i \mid Z_i = 1\} - \mathbb{E}\{w_S(\mathbf{X}_i) D_i \mid Z_i = 0\},$$

where

$$w_P(\mathbf{x}) = \frac{e_P(\mathbf{x})}{\mathbb{E}\{e_P(\mathbf{X}_i)\}}, \quad w_S(\mathbf{x}) = \frac{e_S(\mathbf{x})}{\mathbb{E}\{e_S(\mathbf{X}_i)\}}.$$

Proof is straightforward and hence omitted. While the identification formulas are identical to those in Theorem 2, under Assumption 5, we can simply compute the principal score as $e_S(\mathbf{x}) = \Pr(Y_i = 0 \mid D_i = 0, \mathbf{X}_i = \mathbf{x})$ and set $e_P(\mathbf{x}) = 1 - e_S(\mathbf{x})$.

3.4 Ordinal decision

We generalize the above identification results to an ordinal decision. In our application, this extension is important as the judge's release decision often is based on different amounts of cash bail or varying levels of supervision of an arrestee. We first generalize the monotonicity assumption (Assumption 3) by requiring that a decision with a greater amount of bail is no less likely to make an arrestee engage in NCA (NVCA or FTA). The assumption may be reasonable, for example, because a greater amount of bail is expected to imply a greater probability of being held in custody. The assumption could be violated if arrestees experience financial strain in an effort to post bail, causing them to commit NCA (NVCA or FTA).

Formally, let D_i be an ordinal decision variable where $D_i = 0$ is the least amount of bail, and $D_i = 1, \ldots, k$ represents a bail of increasing amount, i.e., $D_i = k$ is the largest bail amount. Then, the monotonicity assumption for an ordinal decision is given by,

Assumption 6 (Monotonicity with Ordinal Decision)

$$Y_i(d_1) \leq Y_i(d_2) \quad \text{for } d_1 \geq d_2.$$

To generalize the principal strata introduced in the binary decision case, we define the decision with the least amount of bail that prevents an arrestee from committing NCA (NVCA or FTA) as

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

follows,

$$R_i = \begin{cases} \min\{d : Y_i(d) = 0\} & \text{if } Y_i(k) = 0, \\ k+1 & \text{if } Y_i(k) = 1. \end{cases}$$

We may view R_i as an ordinal measure of risk with a greater value indicating a higher degree of risk. When D_i is binary, R_i takes one of the three values, $\{0, 1, 2\}$, representing safe, preventable, and risky cases, respectively. Thus, R_i generalizes the principal strata to the ordinal case under the monotonicity assumption.

Now, we define the principal causal effects in the ordinal decision case. Specifically, for $r = 1, \ldots, k$ (excluding the cases with $r = 0$ and $r = k + 1$), we define the average principal causal effect of the PSA on the judge's decisions as a function of this ordinal risk measure,

$$\text{APCEp}(r) = \Pr\{D_i(1) \geq r \mid R_i = r\} - \Pr\{D_i(0) \geq r \mid R_i = r\}. \tag{4}$$

Since the arrestees with $R_i = r$ would not commit NCA (NVCA or FTA) under the decision with $D_i \geq r$, APCEp(r) represents a reduction in the proportion of NCA (NVCA or FTA) that is attributable to PSA provision among the cases with $R_i = r$. Thus, the expected proportion of NCA (NVCA or FTA) that would be reduced by the PSA is given by,

$$\sum_{r=1}^{k} \text{APCEp}(r) \cdot \Pr(R_i = r).$$

This quantity equals the overall Intention-to-Treat (ITT) effect of PSA provision on NCA (NVCA or FTA).

Furthermore, the arrestees with $R_i = 0$ would never commit a new crime regardless of the judge's decisions. We may, therefore, be interested in estimating the increase in the proportion of the most lenient decision for these safest cases. This generalizes the APCEs to the ordinal decision case,

$$\text{APCEs} = \Pr\{D_i(1) = 0 \mid R_i = 0\} - \Pr\{D_i(0) = 0 \mid R_i = 0\}.$$

For the cases with $R_i = k + 1$ that would always result in a new criminal activity, a desirable decision may depend on a number of factors. Note that if we assume the strong monotonicity, i.e., $Y_i(k) = 0$ for all i, then such cases do not exist.

Like the APCEs, the APCEp(r) can be expressed as a function of the average principal causal effect (APCE) for each decision $d = 0, 1, 2, \ldots, k$. This generalized APCE is given by,

$$\text{APCE}(d, r) = \Pr\{D_i(1) = d \mid R_i = r\} - \Pr\{D_i(0) = d \mid R_i = r\}. \tag{5}$$

In our empirical analysis, we estimate this causal quantity, which has the same identification conditions.

The identification of these principal causal effects requires the knowledge of the distribution of R_i. Fortunately, under the unconfoundedness and monotonicity assumptions (Assumptions 4 and 6), this distribution is identifiable conditional on $\mathbf{X}_i$,

$$\begin{aligned} e_r(\mathbf{x}) &= \Pr(R_i = r \mid \mathbf{X}_i = \mathbf{x}) \\ &= \Pr(R_i \geq r \mid \mathbf{X}_i = \mathbf{x}) - \Pr(R_i \geq r+1 \mid \mathbf{X}_i = \mathbf{x}) \\ &= \Pr\{Y_i(r-1) = 1 \mid \mathbf{X}_i = \mathbf{x}\} - \Pr\{Y_i(r) = 1 \mid \mathbf{X}_i = \mathbf{x}\} \\ &= \Pr\{Y_i = 1 \mid D_i = r-1, \mathbf{X}_i = \mathbf{x}\} - \Pr\{Y_i = 1 \mid D_i = r, \mathbf{X}_i = \mathbf{x}\}, \quad \text{for } r = 1, \ldots, k, \\ e_{k+1}(\mathbf{x}) &= \Pr\{Y_i(k) = 1 \mid \mathbf{X}_i = \mathbf{x}\} = \Pr\{Y_i = 1 \mid D_i = k, \mathbf{X}_i = \mathbf{x}\}, \\ e_0(\mathbf{x}) &= \Pr\{Y_i(0) = 0 \mid \mathbf{X}_i = \mathbf{x}\} = \Pr\{Y_i = 0 \mid D_i = 0, \mathbf{X}_i = \mathbf{x}\}. \end{aligned} \tag{6}$$

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Since $e_r(\mathbf{x})$ cannot be negative for each r, this yields a set of testable conditions for Assumptions 4 and 6. This statement is also true in the binary decision case.

Finally, we formally present the identification result for the ordinal decision case.

Theorem 5 (Identification with Ordinal Decision)

Under Assumptions 1, 2, 4 and 6, the APCE is identified by

$$APCEp(r) = \mathbb{E}\{w_r(\mathbf{X}_i)\mathbf{1}(D_i \geq r) \mid Z_i = 1\} - \mathbb{E}\{w_r(\mathbf{X}_i)\mathbf{1}(D_i \geq r) \mid Z_i = 0\},$$
$$APCEs = \mathbb{E}\{w_0(\mathbf{X}_i)\mathbf{1}(D_i = 0) \mid Z_i = 1\} - \mathbb{E}\{w_0(\mathbf{X}_i)\mathbf{1}(D_i = 0) \mid Z_i = 0\},$$

where $w_r(\mathbf{x}) = e_r(\mathbf{x})/\mathbb{E}\{e_r(\mathbf{X}_i)\}$ and $\mathbf{1}()$ is the indicator function.

Proof is given in Online Supplementary Material, Appendix S4.5.

3.5 Sensitivity analysis

The unconfoundedness assumption, which enables the nonparametric identification of causal effects, may be violated when researchers do not observe some information that is used by the judge and is predictive of arrestees' behaviour. As noted in Section 2.2, the length of time the arrestee has lived in the community may represent an example of such unobserved confounders. It is important, therefore, to develop a sensitivity analysis for the potential violation of the unconfoundedness assumption (Assumption 4).

We propose a parametric sensitivity analysis (see Online Supplementary Material, Appendix S8 for nonparametric sensitivity analysis). We consider the following bivariate ordinal probit model for the observed judge's decision D and the latent risk measure R_i,

$$D_i^*(z) = \beta_Z z + \mathbf{X}_i^\top \beta_X + z\mathbf{X}_i^\top \beta_{ZX} + \epsilon_{i1}, \tag{7}$$

$$R_i^* = \mathbf{X}_i^\top \alpha_X + \epsilon_{i2}, \tag{8}$$

where

$$\begin{pmatrix} \epsilon_{i1} \\ \epsilon_{i2} \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right),$$

and

$$D_i(z) = \begin{cases} 0 & D^*(z) \leq \theta_{z1} \\ 1 & \theta_{z1} < D_i^*(z) \leq \theta_{z2} \\ \vdots & \vdots \\ k-1 & \theta_{z,k-1} < D_i^*(z) \leq \theta_{zk} \\ k & \theta_{zk} < D_i^*(z) \end{cases}, \quad R_i = \begin{cases} 0 & R_i^* \leq \delta_0 \\ 1 & \delta_0 < R_i^* \leq \delta_1 \\ \vdots & \vdots \\ k & \delta_{k-1} < R_i^* \leq \delta_k \\ k+1 & \delta_k < R_i^* \end{cases}.$$

The error terms $(\epsilon_{i1}, \epsilon_{i2})$ are assumed to follow a bivariate normal distribution. Under this model, ρ represents a sensitivity parameter since $\rho = 0$ implies Assumption 4. If the value of ρ is known, then the other coefficients, i.e., β_X, α_X and β_Z, can be estimated, which in turn enables the estimation of the APCE. In the literature, Frangakis et al. (2002), Barnard et al. (2003), and Forastiere et al. (2016) also model the distribution of principal strata using the ordinal probit model.

Because R_i is a latent variable, the estimation of this model is not straightforward. In our empirical application, we conduct a Bayesian analysis to estimate the causal effects (see e.g., Hirano et al., 2000; Jiang et al., 2016; Mattei et al., 2013; Schwartz et al., 2011, for other applications of Bayesian sensitivity analysis). Online Supplementary Material, Appendix S5 presents the details of the Bayesian estimation. We also perform a frequentist analysis, based on Theorem 2, that does

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

not require an outcome model, assessing the robustness of the results to the outcome model (though we assume $\rho = 0$).

3.6 Fairness

Next, we discuss how the above causal effects relate to the fairness of the judge's decision. In particular, Imai and Jiang (2023) introduce the concept of 'principal fairness'. The basic idea is that within each principal stratum a fair decision should not depend on protected attributes (race, gender, etc.). Imai and Jiang (2023) provide a detailed discussion about how principal fairness is related to the existing definitions of fairness (see also Chouldechova & Roth, 2020; Corbett-Davies et al., 2017, and references therein). Although Coston et al. (2020) consider the potential outcomes framework, they only focus on one potential outcome $Y_i(0)$ rather than the joint potential outcomes $(Y_i(1), Y_i(0))$.

Formally, let $A_i \in \mathcal{A}$ be a protected attribute such as race and gender. We first consider a binary decision. We say that decisions are fair on average with respect to A_i if it does not depend on the attribute within each principal stratum, i.e.,

$$\Pr\{D_i = 1 \mid A_i,\, Y_i(1) = y_1,\, Y_i(0) = y_0\} = \Pr\{D_i = 1 \mid Y_i(1) = y_1,\, Y_i(0) = y_0\} \tag{9}$$

for all $y_1, y_0 \in \{0, 1\}$. We can generalize this definition to the ordinal case as,

$$\Pr(D_i \geq d \mid A_i,\, R_i = r) = \Pr(D_i \geq d \mid R_i = r)$$

for $1 \leq d \leq k$ and $0 \leq r \leq k + 1$.

The degree of fairness for principal stratum $R_i = r$ can be measured using the maximal deviation among the distributions for different groups,

$$\Delta_r(z) = \max_{a,a',d} \left| \Pr\{D_i(z) \geq d \mid A_i = a,\, R_i = r\} - \Pr\{D_i(z) \geq d \mid A_i = a',\, R_i = r\} \right| \tag{10}$$

for $z = 0, 1$. By estimating $\Delta_r(z)$, we can use the experimental data to examine whether or not the provision of the PSA improves the fairness of the judge's decisions. Specifically, if PSA provision improves the fairness of judge's decisions for the principal stratum r, we should have $\Delta_r(1) \leq \Delta_r(0)$.

3.7 Optimal decision rule

The discussion so far has focused on estimating the impacts of algorithmic recommendations on human decisions. We now show that the experimental data can also be used to derive an optimal decision rule given a certain objective. In addition, by comparing human decisions and algorithmic recommendations with optimal decision rules, we can evaluate their efficacy. In our application, one goal is to prevent as many NCAs (NVCAs or FTAs) as possible while avoiding unnecessarily harsh initial release decisions. To achieve this, we must carefully weigh the cost of negative outcomes and that of unnecessarily harsh decisions. Once these costs are specified as part of the utility function, one can empirically assess this tradeoff using the experimental data.

Formally, let δ be the judge's decision based on $\mathbf{X}_i$, which may include the PSA. We consider a deterministic decision rule, i.e., $\delta(\mathbf{x}) = d$ if $\mathbf{x} \in \mathcal{X}_d$ where $\mathcal{X}_d$ is a non-overlapping partition of the covariate space $\mathcal{X}$ with $\mathcal{X} = \bigcup_{r=0}^{k} \mathcal{X}_r$ and $\mathcal{X}_r \cap \mathcal{X}_{r'} = \emptyset$. We consider the utility function of the following form,

$$U_i(\delta) = \begin{cases} -c_0 & \delta(\mathbf{X}_i) < R_i \\ 1 & \delta(\mathbf{X}_i) = R_i \\ 1 - c_1 & \delta(\mathbf{X}_i) > R_i \end{cases},$$

where c_0 and c_1 represent the cost of an NCA (NVCA or FTA) and that of an unnecessarily harsh decision, respectively. Under this setting, preventing an NCA (NVCA or FTA) with the most lenient decision ($\delta(\mathbf{X}_i) = R_i$) yields the utility of one, while we incur the cost c_1 for an unnecessarily harsh decision ($\delta(\mathbf{X}_i) > R_i$), leading to the net utility of $1 - c_1$.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

The relative magnitude of these two cost parameters, c_0 and c_1, may depend on the consideration of various factors, including the potential harm to the public and arrestees caused by the negative outcomes and unnecessarily harsh decisions, respectively. When $c_0 = c_1 = 0$, for example, $U_i(\delta)$ reduces to $\mathbf{1}\{\delta(\mathbf{X}_i) \geq R_i\}$, which is non-zero only if the decision is sufficiently harsh so that it prevents the negative outcome. The optimal decision under this utility is the most stringent decision, i.e., $\delta(\mathbf{X}_i) = k$, for all cases. If $c_0 = 2$ and $c_1 = 1$, the resulting utility function implies that the cost of NCA (NVCA or FTA) is twice as large as that of an unnecessarily harsh decision.

We derive the optimal decision rule δ^* that maximizes the expected utility,

$$\delta^* = argmax_\delta \mathbb{E}\{U_i(\delta)\}.$$

For $r = 0, \ldots, k+1$ and $d = 0, \ldots, k$, we can write,

$$\mathbb{E}[\mathbf{1}\{\delta(\mathbf{X}_i) = d,\, R_i = r\}] = \mathbb{E}\{\mathbf{1}(\mathbf{X}_i \in \mathcal{X}_d,\, R_i = r)\} = \mathbb{E}\big\{\mathbf{1}(\mathbf{X}_i \in \mathcal{X}_d) \cdot e_r(\mathbf{X}_i)\big\}.$$

Thus, we can express the expected utility as,

$$\begin{aligned}
&\mathbb{E}\{U_i(\delta)\} \\
&= \sum_{r=0}^{k+1} \Bigg(\sum_{d>r} (1-c_1)\mathbb{E}[\mathbf{1}\{\delta(\mathbf{X}_i) = d,\, R_i = r\}] + \sum_{d=r} \mathbb{E}[\mathbf{1}\{\delta(\mathbf{X}_i) = d,\, R_i = r\}] \\
&\quad - \sum_{d<r} c_0 \mathbb{E}[\mathbf{1}\{\delta(\mathbf{X}_i) = d,\, R_i = r\}] \Bigg) \qquad (11) \\
&= \sum_{r=0}^{k+1} \Bigg[\sum_{d\geq r} \mathbb{E}\big\{\mathbf{1}(\mathbf{X}_i \in \mathcal{X}_d) \cdot e_r(\mathbf{X}_i)\big\} - c_0 \sum_{d<r} \mathbb{E}\big\{\mathbf{1}(\mathbf{X}_i \in \mathcal{X}_d) \cdot e_r(\mathbf{X}_i)\big\} \\
&\quad - c_1 \sum_{d>r} \mathbb{E}\big\{\mathbf{1}(\mathbf{X}_i \in \mathcal{X}_d) \cdot e_r(\mathbf{X}_i)\big\} \Bigg] \\
&= \sum_{d=0}^{k} \mathbb{E}\left[\mathbf{1}(\mathbf{X}_i \in \mathcal{X}_d) \left\{ \sum_{r\leq d} e_r(\mathbf{X}_i) - c_0 \cdot \sum_{r>d} e_r(\mathbf{X}_i) - c_1 \cdot \sum_{r<d} e_r(\mathbf{X}_i) \right\} \right]. \qquad (12)
\end{aligned}$$

This yields the following optimal decision,

$$\delta^*(\mathbf{x}) = \underset{d\in\{0,\ldots,k\}}{\operatorname{argmax}}\, g_d(\mathbf{x}) \quad \text{where } g_d(\mathbf{x}) = \sum_{r\leq d} e_r(\mathbf{x}) - c_0 \cdot \sum_{r>d} e_r(\mathbf{x}) - c_1 \cdot \sum_{r<d} e_r(\mathbf{x}). \qquad (13)$$

We can, therefore, use the experimental estimate of $e_r(\mathbf{x})$ to learn about the optimal decision.

Policy makers could derive the optimal decision rule by using the above result and then adopt this rule as the recommendation for judges. However, this may not be useful if the judge decides to follow the algorithmic recommendation selectively for some cases or ignore it altogether. Instead, we may wish to construct PSA scores that maximize the optimality of the judge's decision. Unfortunately, the derivation of such an optimal PSA score is challenging since the PSA scores were not directly randomized in our experiment. We tackle this problem in a separate paper (Ben-Michael et al., 2021). In Online Supplementary Material, Appendix S6, we also consider the optimal provision of the PSA given the same goal considered above (i.e., prevent as many NCAs (NVCAs or FTAs) as possible with a minimal amount of bail).

4 Empirical analysis

In this section, we apply the proposed methodology to the data from the field RCT described in Section 2.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

4.1 Preliminaries

As explained in Section 2.3, we use the ordinal decision variable with three categories—the signature bond ($D_i = 0$), the bail amount of \$1,000 or less ($D_i = 1$), and the bail amount of greater than \$1,000 ($D_i = 2$). Given this ordinal decision, we label the principal strata as safe ($R_i = 0$), easily preventable ($R_i = 1$), preventable ($R_i = 2$), and risky cases ($R_i = 3$).

We fit the Bayesian model defined in Equations (7) and (8) with a diffuse prior distribution as specified in Online Supplementary Material, Appendix S5, separately for each of three binary outcome variables—FTA, NCA, and NVCA. The model incorporates the following pre-treatment covariates: gender (male or female), race (white or non-white), the interaction between gender and race, age, and several indicator variables regarding the current and past charges. It also includes a binary variable for the presence of pending charge (felony, misdemeanor, or both) at the time of offense, four binary variables for current charges (non-violent misdemeanor, violent misdemeanor, non-violent felony, and violent felony), a four-level ordinal variable for the DMF recommendation, three variables for prior conviction (binary variables for misdemeanor and felony as well as a four-level ordinal variable for violent conviction), a binary variable for prior sentence to incarceration, and two variables for prior FTAs (a three-level ordinal variable for FTAs from past two years, and a binary variable for FTAs from over two years ago).

We use the Gibbs sampling and run five Markov chains of 100,000 iterations each with random starting values independently drawn from the prior distribution. Based on the Gelman-Rubin statistic for convergence diagnostics, we retain the second half of each chain and combine them to be used for our analysis. Online Supplementary Material, Appendix S5 presents the computational details including the Gibbs sampling algorithm we use.

We begin by computing the estimated population proportion of each principal stratum based on Equation (6). Figure 3 presents the results. We find that for FTA, the overall proportion of safe cases (blue) is estimated to be 67%, whereas those of easily preventable (black), preventable (red), and risky (brown) cases are 6%, 7%, and 20%, respectively. A similar pattern is observed for FTA and NCA across different racial and gender groups, while the estimated overall proportion of safe cases is even higher for NVCA, exceeding 90%.

4.2 Average principal causal effects

Figure 4 presents the estimated APCE of PSA provision on the three ordinal decision categories, separately for each of the three outcomes and each principal stratum (see Equation (5)). The overall and subgroup-specific results are given for each of the four principal strata—safe (blue), easily preventable (black), preventable (red), and risky (brown) cases. For a given principal stratum, we present the estimated APCE on each decision category—signature bond (circle), small cash bond (triangle), and large cash bond (square). The left column of each panel shows that PSA provision has little overall impact on the judge's decision across four principal strata for FTA and NCA. There is a suggestive, but inconclusive, evidence that PSA provision leads to an overall harsher decision for NVCA among easily preventable, preventable, and risky cases.

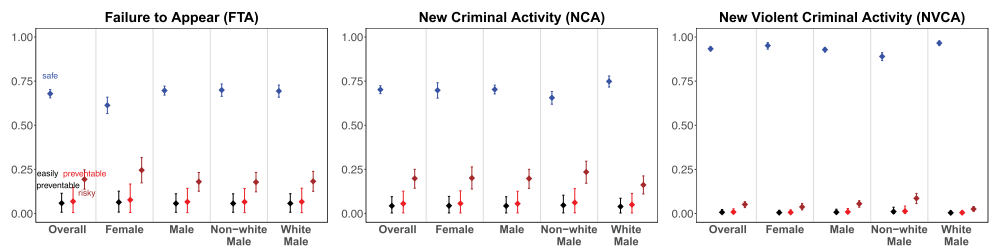

Figure 3. Estimated proportion of each principal stratum. Each plot represents the result using one of the three outcome variables (FTA, NCA, and NVCA), where the blue, black, red, and brown diamonds represent the estimates for safe, easily preventable, preventable, risky cases, respectively. The solid vertical lines represent the 95% Bayesian credible intervals. The results show that a vast majority of cases are safe across subgroups and across different outcomes. The proportion of safe cases is estimated to be especially high for NVCA.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

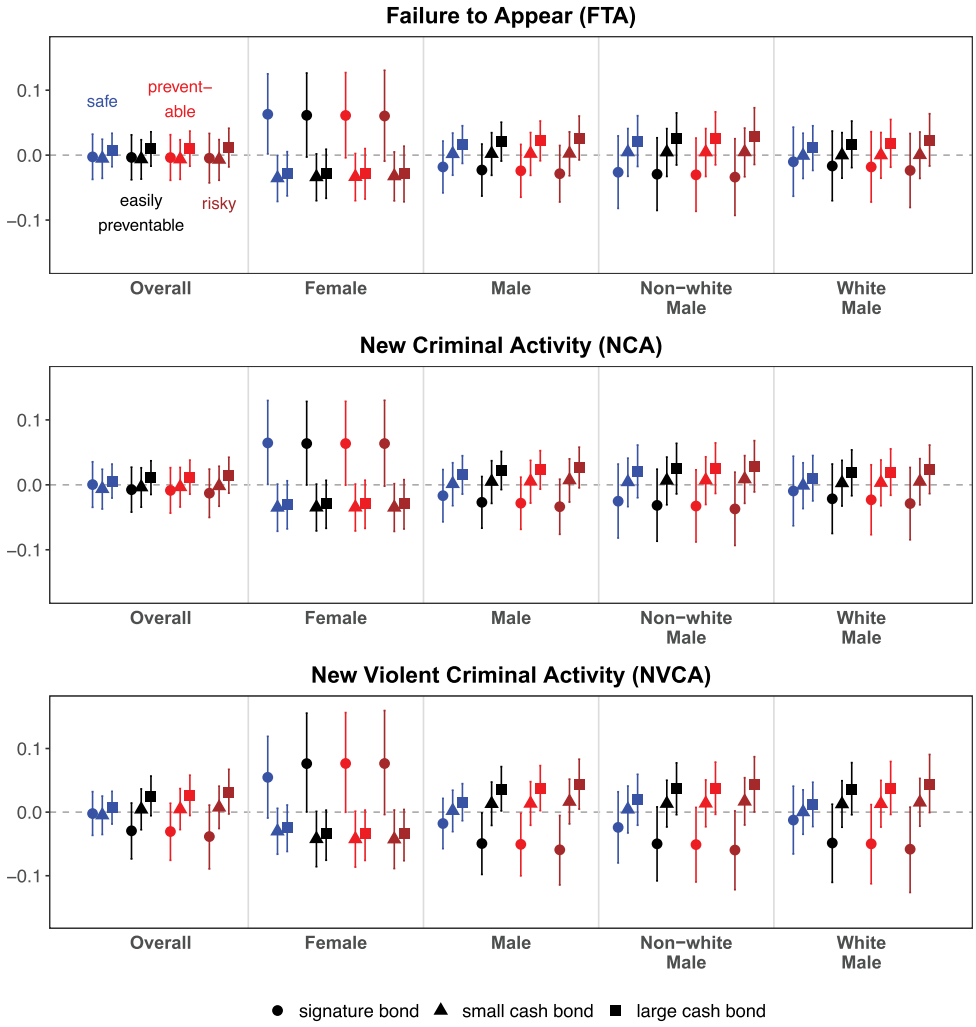

Figure 4. Estimated average principal causal effects (APCE) of PSA provision on the judge's decision. Each panel presents the overall and subgroup-specific results for a different outcome variable. Each column within a panel shows the estimated APCE of PSA provision for safe (blue), easily preventable (black), preventable (red), and risky (brown) cases. For each of these principal strata, we report the estimated APCE on the judge's decision to impose a signature bond (circles), a small cash bail amount of 1,000 dollars or less (triangles), and a large cash bail amount of greater than 1,000 (squares). The vertical line for each estimate represents the Bayesian 95% credible interval. The results show that PSA provision may make the judge's decision more lenient for female arrestees regardless of their risk levels. PSA provision may also encourage the judge to make harsher decisions for male arrestees with a greater risk level though the effect sizes are relatively small.

We also present the estimated APCE for different gender and racial groups in the remaining columns of each panel. We find potentially suggestive evidence that PSA provision may make it more likely for the judge to impose signature bonds (circles) on female arrestees instead of cash bonds (triangles and squares) across three outcomes. Interestingly, for all outcomes, this pattern appears to hold for any of the four principal strata, implying that PSA provision might not help the judge distinguish different risk levels of female arrestees. Our analysis also finds that for NVCA, PSA provision may lead to a harsher decision for easily preventable, preventable, and risky cases among male arrestees while it has little effect on the safe cases. This suggests that PSA provision may help distinguish different risk levels among male arrestees, resulting in improved decisions at least in terms of the original goal of the PSA. There is no discernible racial difference in these effects.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

In Online Supplementary Material, Appendix S2.2, we explore the estimated APCE for different age groups. We find that PSA provision may lead to a harsher decision for arrestees of the 29–35 years old group across three outcomes. This pattern appears to generally hold across all principal strata, though for NVCA, the effects are more pronounced for easily preventable, preventable, and risky cases. In addition, our analysis yields suggestive evidence that across all outcomes, PSA provision may make the judge's decision more lenient for the oldest (46 years old or above) group. This appears to be true across all three outcomes except that for NVCA the effect may exist only for safe cases. We reiterate, however, that these results are based on preliminary data and the effect sizes are relatively small.

We conduct two robustness analyses. Firstly, we perform a frequentist analysis that is based on Theorem 2 and does not assume an outcome model. The results are shown in Online Supplementary Material, Appendix S7 and are largely consistent with those shown here. As expected, the estimation uncertainty of the frequentist analysis, which makes less stringent assumptions than Bayesian analysis, is greater. Secondly, we conduct parametric sensitivity analyses using the methods described in Section 3.5 (see Online Supplementary Material, Appendix S9). We set the value of correlation parameter ρ to 0.05, 0.1, and 0.3 and examine how the estimated APCE changes. The results (see Online Supplementary Material, Figures S13–S15) are largely consistent across different values of ρ although the effects for females tend to exhibit a large degree of estimation uncertainty especially when the correlation is high and particularly for NVCA. This is not surprising. There are only a small number of female arrestees and only a handful of NVCA events corresponding to them.

4.3 Gender and racial fairness

We now examine the impacts of PSA provision on gender and racial fairness. Specifically, we evaluate the principal fairness of PSA provision as discussed in Section 3.6. We use gender (female vs. male) and race (white male vs. non-white male) separately as a protected attribute and analyse whether or not the provision of the PSA improves the fairness of the judge's decision in terms of the protected attribute. While the gender analysis is based on the entire sample, the racial analysis is based on the male sample only due to the limited sample size for females.

Figure 5 presents the results for gender (top panel) and racial (bottom panel) fairness across the principal strata and separately for each of the three outcomes. Each column within a given plot presents $\Delta_r(z)$ defined in equation (10), which represents the maximal subgroup difference in the judge's decision probability distribution within the same principal stratum $R_i = r$ under the provision of the PSA ($z = 1$) compared to no use of the PSA ($z = 0$). In this application, the maximal difference always occurs at $d = 1$, allowing us to interpret $\Delta_r(z)$ as the difference in the probability of imposing a cash bond ($D \geq 1$) rather than a signature bond. We also present the estimated difference caused by PSA provision in the two maximal subgroup differences, i.e., $\Delta_r(1) - \Delta_r(0)$. If this difference is estimated to be positive, then PSA provision reduces the fairness of judge's decisions by increasing the maximal subgroup difference.

We find that PSA provision might worsen the gender fairness of the judge's decisions. When the PSA is provided, the maximal gender difference in the judge's decision probability is on average greater than that when it is not provided. The effect is particularly large and statistically significant for NVCA and for preventable, easily preventable, and risky cases. This is consistent with our finding that especially for NVCA, PSA provision might make the judge's decision more lenient for female arrestees while it leads to a harsher decision for male arrestees among preventable, easily preventable, and risky cases. Thus, PSA provision appears to increase disparate decision-making across gender.

PSA provision, however, does not have a statistically significant impact on the racial fairness of the judges' decisions among male arrestees. For instance, in the principal stratum of safe cases, we find that PSA provision does not affect the maximal difference in the judge's decision probability (between non-white males and white males). This suggests that in terms of principal fairness, the PSA may not alter any existing racial difference in the judge's decisions.

4.4 Using optimal decision to evaluate the DMF recommendation

Finally, we evaluate the DMF recommendation by comparing it with the optimal decision under different values of the costs. For simplicity, we consider a binary decision: signature or cash bond.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

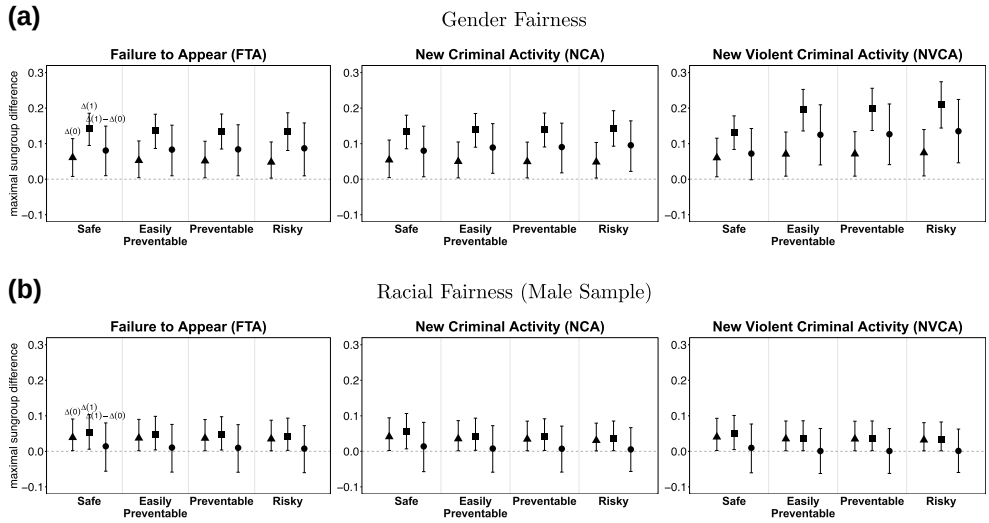

Figure 5. Gender and racial fairness of the judge's decisions. Within each plot, we show three estimates separately for each principal stratum—the maximal subgroup difference in the judge's decision probability of imposing a cash bond with PSA provision (squares; $\Delta(1)$) and without it (triangles; $\Delta(0)$) as well as the difference between them (circles; $\Delta(1) - \Delta(0)$). The vertical solid lines represent the 95% Bayesian credible intervals. A positive value of the difference would imply that the PSA reduces the fairness of the judge's decisions. For the gender analysis (top panel), even without the PSA, the judge seems to be more likely to impose a cash bond on male arrestees when compared to female arrestees with the same risk levels. PSA provision appears to increase this tendency. For the race analysis (bottom panel), PSA provision has little impact across all outcomes and risk levels. We reiterate, however, that these results are based on the preliminary data. (a) Gender fairness. (b) Racial fairness (male sample).

As discussed in Section 3.7, given a specific pair of cost parameters (c_0, c_1) and the experimental estimate of $e_r(\mathbf{x})$ for $r = 0, 1, 2$, we can compute the optimal decision for each case according to Equation (13). We then obtain the estimated proportion of cases, for which a cash bond is optimal. We repeat this process for a grid of different values for the cost of a negative outcome (c_0; FTA, NCA, and NVCA) and that of an unnecessarily harsh decision (c_1).

The top panel of Figure 6 presents the results for the cases whose DMF recommendation is a signature bond. In contrast, the bottom panel of the figure shows the results for the other cases (i.e., the DMF recommendation is a cash bond). In each plot, a darker grey region represents a greater proportion of cases, for which a cash bond is optimal. The results suggest that unless the cost of a negative outcome is much higher than the cost of an unnecessarily harsh decision, imposing a signature bond is the optimal decision for a vast majority of cases.

We also find that for all three outcomes, a cash bond is optimal for a greater proportion of cases when the DMF recommendation is indeed a cash bond. However, this difference is small, suggesting that the DMF recommendation is only mildly informative. Similar results are found even if we separately examine three PSA scores (see Figure S16 in Online Supplementary Material, Appendix S10).

4.5 Comparison between the judge's decisions and DMF recommendations

Lastly, we compare the judge's actual decision with the DMF recommendation in terms of the expected utility given in Equation (12). The top panel of Figure 7 represents the results for the treatment group (i.e., judge's decisions with the PSA), whereas the bottom panel represents those for the control group (i.e., judge's decisions without the PSA). A darker grey area indicates that the expected utility for the judge's decision is estimated to be greater than the DMF recommendation. Most of these estimates are statistically significant (see Online Supplementary Material, Figure S11 for more details). Therefore, unless the cost of a negative outcome is much greater than the cost of an unnecessarily harsh decision, the judge's decision (with or without the PSA scores) yields a greater expected utility than the DMF recommendation. This is especially true for NVCA.

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

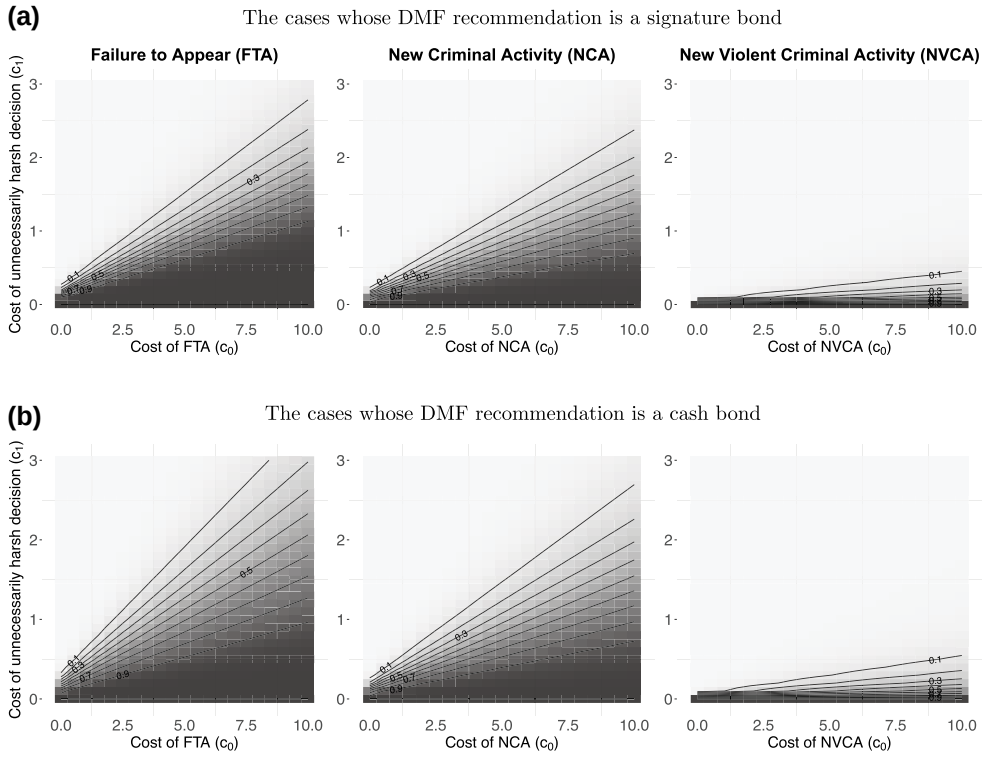

Figure 6. Estimated proportion of cases for which cash bond is optimal. Each column represents the results based on one of the three outcomes (FTA, NCA, and NVCA). The top (bottom) panel shows the results for the cases whose DMF recommendation is a signature (cash) bond. In each plot, the contour lines represent the estimated proportions of cases for which a cash bond is optimal, given the cost of an unnecessarily harsh decision (c_1; *y*-axis) and that of a negative outcome (c_0; *x*-axis). A dark grey area represents a greater proportion of such cases. The results show that regardless of DMF recommendation, a signature bond is optimal unless the cost of a negative outcome is much greater than the cost of an unnecessarily harsh decision. (a) The cases whose DMF recommendation is a signature bond. (b) The cases whose DMF recommendation is a cash bond.

Altogether, our analysis implies that the DMF recommendations may be unnecessarily harsher than the judge's decisions.

5 Concluding remarks

In today's data-rich society, many human decisions are guided by algorithmic recommendations. While some of these algorithmic-assisted human decisions may be trivial and routine (e.g., online shopping and movie suggestions), others that are much more consequential include judicial and medical decision-making. As algorithmic recommendation systems play increasingly important roles in our lives, we believe that a policy-relevant question is how such systems influence human decisions and how the biases of algorithmic recommendations interact with those of human decisions. These questions necessitate the empirical evaluation of the impacts of algorithmic recommendations on human decisions.

In this paper, we present a set of general statistical methods that can be used for the experimental evaluation of algorithm-assisted human decision-making. We applied these methods to the preliminary data from the first-ever randomized controlled trial for assessing the impacts of PSA provision on judges' pretrial decisions. There are several findings that emerge from our initial analysis. Firstly, we find that PSA provision has little overall impact on the judge's decisions. Secondly, we find potentially suggestive evidence PSA provision may encourage the judge to make more lenient decisions for female arrestees regardless of their risk levels while leading to more stringent decisions for males who are classified as risky. Thirdly, PSA provision appears to widen the existing

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

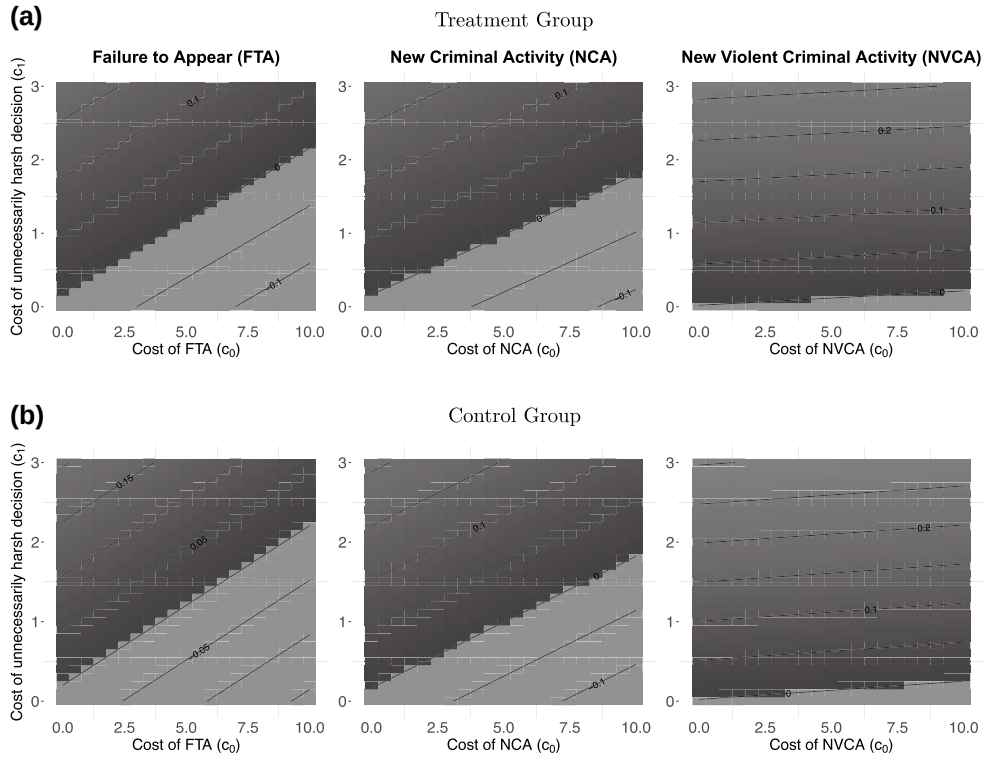

Figure 7. Estimated difference in the expected utility between judge's decisions and DMF recommendations for the treatment (top panel) and control (bottom panel) group. Each column represents the results based on one of the three outcomes, with a darker region indicating the values of the costs (the cost of a negative outcome and the cost of an unnecessarily harsh decision) for which the judge's decision yields a higher expected utility than the corresponding DMF recommendation. The results show that the judge's decision yields a higher expected utility than the DMF recommendation unless the cost of a negative outcome is much higher than that of an unnecessarily harsh decision. This pattern holds for all outcomes and is unchanged by the provision of the PSA. (a) Treatment group. (b) Control group.

gender difference of the judge's decisions against male arrestees, whereas it does not seem to alter decision-making across race among male arrestees. We caution, however, that these findings could be explained by other factors that are correlated with gender and race. Finally, we find that for a vast majority of cases, the optimal decision is to impose a signature bond rather than a cash bond unless the cost of a negative outcome is much higher than that of a decision that may result in unnecessary incarceration. This suggests that the PSA's recommendations may be harsher than necessary.

Acknowledgments

We thank Dean Knox and many participants at various seminars for helpful comments. The experiment analysed in this paper has been approved by Harvard University's Institutional Review Board (#16–1258).

Funding

We acknowledge partial financial support from Arnold Ventures, Cysco Systems, Inc. (CG# 2370386), the National Science Foundation (SES–2051196), and the Sloan Foundation (Economics Program; 2020-13946).

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Supplementary material

Supplementary material is available at *Journal of the Royal Statistical Society: Series A* online.

Conflict of interests: None declared.

References

Albright A. (2019). *If you give a judge a risk score: Evidence from Kentucky Bail Decisions* (Technical Report). Department of Economics, Harvard University.

Angwin J., Larson J., Mattu S., & Kirchner L. (2016). Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

Ares C. E., Rankin A., & Sturz H. (1963). The Manhattan Bail Project: An interim report on the use of pretrial parole. *New York University Law Review*, *38*(1), 67–95.

Barnard J., Frangakis C. E., Hill J. L., & Rubin D. B. (2003). Principal stratification approach to broken randomized experiments: A case study of school choice vouchers in New York City. *Journal of the American Statistical Association*, *98*(462), 299–323. https://doi.org/10.1198/016214503000071

Ben-Michael E., Greiner J., Imai K., & Jiang Z. (2021). *Safe policy learning through extrapolation: Application to pre-trial risk assessment* (Technical Report). arXiv:2109.11679.

Berk R. (2017). An impact assessment of machine learning risk forecasts on parole board decisions and recidivism. *Journal of Experimental Criminology*, *13*(2), 193–216. https://doi.org/10.1007/s11292-017-9286-2

Berk R., Heidari H., Jabbari S., Kearns M., & Roth A. (2018). Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, *50*(1), 3–44. https://doi.org/10.1177/0049124118782533

Chouldechova A., & Roth A. (2020). A snapshot of the frontiers of fairness in machine learning. *Communications of the ACM*, *63*(5), 82–89. https://doi.org/10.1145/3376898

Corbett-Davies S., Pierson E., Feller A., Goel S., & Huq A. (2017). Algorithmic decision making and the cost of fairness. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 797–806. https://doi.org/10.1145/3097983.3098095

Coston A., Mishler A., Kennedy E. H., & Chouldechova A. (2020). Counterfactual risk assessments, evaluation, and fairness. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 582–593. https://doi.org/10.1145/3351095.3372851

Dawes R. M., Faust D., & Meehl P. E. (1989). Clinical versus actuarial judgment. *Science*, *243*(4899), 1668–1674. https://doi.org/10.1126/science.2648573

Dieterich W., Mendoza C., & Brennan T. (2016). *COMPAS risk scales: Demonstrating accuracy equity and predictive parity*. Northpointe Inc. Research Department. http://go.volarisgroup.com/rs/430-MBX-989/images/ProPublica_Commentary_Final_070616.pdf

Ding P., & Lu J. (2017). Principal stratification analysis using principal scores. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, *79*(3), 757–777. https://doi.org/10.1111/rssb.12191

Dressel J., & Farid H. (2018). The accuracy, fairness, and limits of predicting recidivism. *Science Advances*, *4*(1), eaao5580. https://doi.org/10.1126/sciadv.aao5580

Flores A. W., Bechtel K., & Lowenkamp C. (2016). False positives, false negatives, and false analyses: A rejoinder to "Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks". *Federal Probation Journal*, *80*(2), 28–46.

Forastiere L., Mealli F., & VanderWeele T. J. (2016). Identification and estimation of causal mechanisms in clustered encouragement designs: Disentangling bed nets using Bayesian principal stratification. *Journal of the American Statistical Association*, *111*(514), 510–525. https://doi.org/10.1080/01621459.2015.1125788

Frangakis C. E., & Rubin D. B. (2002). Principal stratification in causal inference. *Biometrics*, *58*(1), 21–29. https://doi.org/10.1111/j.0006-341X.2002.00021.x

Frangakis C. E., Rubin D. B., & Zhou X.-H. (2002). Clustered encouragement designs with individual noncompliance: Bayesian inference with randomization, and application to advance directive forms. *Biostatistics*, *3*(2), 147–164. https://doi.org/10.1093/biostatistics/3.2.147

Garrett B. L., & Monahan J. (2020). Judging risk. *California Law Review*, *108*(2), 439–493. https://doi.org/10.15779/Z38B56D515

Goldkamp J. S., & Gottfredson M. R. (1984). *Judicial guidelines for bail: The Philadelphia experiment*. US Department of Justice, National Institute of Justice.

Goldkamp J. S., & Gottfredson M. R. (1985). *Policy guidelines for bail: An experiment in court reform*. Temple University Press.

Green B., & Chen Y. (2019). Disparate interactions: An algorithm-in-the-loop analysis of fairness in risk assessments. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 90–99. https://doi.org/10.1145/3287560.3287563

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023

Greiner D. J., Halen R., Stubenberg M., & Griffin Jr. C. L. (2020). *Randomized control trial evaluation of the implementation of the PSA-DMF system in Dane County, WI* (Technical Report). Access to Justice Lab, Harvard Law School.

Hansen J. H. L., & Hasan T. (2015). Speaker recognition by machines and humans: A tutorial review. *IEEE Signal Processing Magazine*, *32*(6), 74–99. https://doi.org/10.1109/MSP.2015.2462851

He K., Zhang X., Ren S., & Sun J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *2015 IEEE International Conference on Computer Vision (ICCV)*, 1026–1034. https://doi.org/10.1109/ICCV.2015.123

Hirano K., Imbens G. W., Rubin D. B., & Zhou X.-H. (2000). Assessing the effect of an influenza vaccine in an encouragement design. *Biostatistics*, *1*(1), 69–88. https://doi.org/10.1093/biostatistics/1.1.69

Imai K., & Jiang Z. (2023). Principal fairness for human and algorithmic decision-making. *Statistical Science*. https://doi.org/10.1214/22-STS872

Jiang Z., Ding P., & Geng Z. (2016). Principal causal effect identification and surrogate end point evaluation by multiple trials. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *78*(4), 829–848. https://doi.org/10.1111/rssb.12135

Kleinberg J., Lakkaraju H., Leskovec J., Ludwig J., & Mullainathan S. (2018). Human decisions and machine predictions. *Quarterly Journal of Economics*, *133*(1), 237–293. https://doi.org/10.1093/qje/qjx032

Mattei A., Li F., & Mealli F. (2013). Exploiting multiple outcomes in Bayesian principal stratification analysis with application to the evaluation of a job training program. *The Annals of Applied Statistics*, *7*(4), 2336–2360. https://doi.org/10.1214/13-AOAS674

Miller J., & Maloney C. (2013). Practitioner compliance with risk/needs assessment tools: A theoretical and empirical assessment. *Criminal Justice and Behavior*, *40*(7), 716–736. https://doi.org/10.1177/0093854812468883

Rubin D. B. (1990). Comments on "On the application of probability theory to agricultural experiments. Essay on principles. Section 9" by J. Splawa-Neyman translated from the Polish and edited by D. M. Dabrowska and T. P. Speed. *Statistical Science*, *5*(4), 472–480. https://doi.org/10.1214/ss/1177012032

Rudin C., Wang C., & Coker B. (2020). The age of secrecy and unfairness in recidivism prediction. *Harvard Data Science Review*, *2*(1), https://doi.org/10.1162/99608f92.6ed64b30

Schwartz S. L., Li F., & Mealli F. (2011). A Bayesian semiparametric approach to intermediate variables in causal inference. *Journal of the American Statistical Association*, *106*(496), 1331–1344. https://doi.org/10.1198/jasa.2011.ap10425

Silver D., Hubert T., Schrittwieser J., Antonoglou I., Lai M., Guez A., Lanctot M., Sifre L., Kumaran D., Graepel T., Lillicrap T., Simonyan K., & Hassabi D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, *362*(6419), 1140–1144. https://doi.org/10.1126/science.aar6404

Skeem J., Scurich N., & Monahan J. (2020). Impact of risk assessment on judges' fairness in sentencing relatively poor defendant. *Law and Human Behavior*, *44*(1), 51–59. https://doi.org/10.1037/lhb0000360

Stevenson M. (2018). Assessing risk assessment in action. *Minnesota Law Review*, *103*(1), 303.

Stevenson M., & Doleac J. L. (2021). *Algorithmic risk assessment in the hands of humans* (Technical Report). Social Science Research Network (SSRN). http://dx.doi.org/10.2139/ssrn.3489440

Downloaded from https://academic.oup.com/jrsssa/article/186/2/167/7024674 by guest on 20 May 2023