

Potential Outcome Rankings for Counterfactual Decision Making

Yuta Kawakami, Jin Tian

Mohamed bin Zayed University of Artificial Intelligence, UAE
{Yuta.Kawakami, Jin.Tian}@mbzuai.ac.ae

Abstract

Counterfactual decision-making in the face of uncertainty involves selecting the optimal action from several alternatives using causal reasoning. Decision-makers often rank expected potential outcomes (or their corresponding utility and desirability) to compare the preferences of candidate actions. In this paper, we study new counterfactual decision-making rules by introducing two new metrics: the probabilities of potential outcome ranking (PoR) and the probability of achieving the best potential outcome (PoB). PoR reveals the most probable ranking of potential outcomes for an individual, and PoB indicates the action most likely to yield the top-ranked outcome for an individual. We then establish identification theorems and derive bounds for these metrics, and present estimation methods. Finally, we perform numerical experiments to illustrate the finite-sample properties of the estimators and demonstrate their application to a real-world dataset.

Introduction

Decision-making in the face of uncertainty involves selecting the optimal action from several alternatives $\{1, \dots, K\}$ using logical reasoning or empirical evidence. Decision theory traces its origins to the work of Blaise Pascal, particularly Pascal’s Wager and his foundational contributions to probability theory (Pascal 1670). This line of reasoning aligns with the principles of expected utility theory, which later became a cornerstone of modern decision theory (Bernoulli 1954; Jeffrey 1990; Knight and Jones 2002). Today, AI is reshaping decision-making processes across a wide range of domains—from healthcare and finance to business strategy and everyday personal choices (Lai et al. 2023). Causal inference, on the other hand, examines the cause-and-effect relationship between an action (or decision) (X) and its outcome (or corresponding utility, desirability) (Y) with potential outcomes (POs) Y_x (Pearl 2009). When these two areas intersect, counterfactual decision-making emerges as a method for determining optimal action by integrating causal reasoning techniques (Lewis 1976; Gibbard and Harper 1981; Pearl 2009; Dawid 2021).

Using causal reasoning, decision-makers often focus on evaluating the *rank of the expected POs* (RoE), given by the

following counterfactual relationship:

$$\mathbb{E}[Y_{r_1}] > \mathbb{E}[Y_{r_2}] > \dots > \mathbb{E}[Y_{r_K}]. \quad (1)$$

Causal reasoning enables us to assess the expected POs, when they cannot be directly observed or evaluated. Based on the RoE in (1), decision-makers judge r_1 as the best action and r_K as the least favorable action. Many approaches rely on RoE-based decision-making, including methods in statistics (Student 1908; Bartholomew 1959; Page 1963; McCaffrey et al. 2013), in causal inference (Lopez and Gutman 2017), and reinforcement learning (Sutton, Barto et al. 1998; Auer, Cesa-Bianchi, and Fischer 2002).

Despite the widespread use of RoE-based decision-making, the action that maximizes the expected PO does not always maximize the PO for each individual. Similarly, the action with the lowest expected PO is not necessarily the least favorable choice for an individual. RoE is not a perfect decision-making criterion and may lead to undesirable choices, depending on the preferences of the decision-maker. To accommodate a broader range of decision-making preferences, alternative rules are desired.

In economics, discrete choice models (DCMs) offer an alternative framework for modeling *individual* choice behavior (McFadden 1974; Ben-Akiva and Lerman 1985; Anderson, De Palma, and Thisse 1992; McFadden 2001; Train 2009). DCMs are used to describe and predict the probability that an individual selects a particular alternative from available options. The utility $Y(k)$ that a person obtains by choosing the alternative k is modeled as a function of observable characteristics of the alternatives and the decision-maker, plus an unobservable random component. The core assumption is that individuals are rational and choose the alternative that maximizes their perceived utility. Researchers typically evaluate the probability that a specific alternative, say r_1 , is chosen, which can be expressed as:

$$\mathbb{P}(\arg \max_{k=1, \dots, K} \{Y(k)\} = r_1). \quad (2)$$

Sometimes the interest lies in the probability of a specific ranking of alternatives (r_1, \dots, r_K) , i.e.,

$$\mathbb{P}(Y(r_1) > Y(r_2) > \dots > Y(r_K)). \quad (3)$$

It is important to clarify that $Y(k)$ (utility) in this context does not represent POs in the causal inference sense, and DCMs do not inherently constitute counterfactual reasoning.

DCMs aim to model observed choices based on prevailing conditions, and to forecast what choices they will make if those conditions persist.

Inspired by the DCM framework, we formulate new types of counterfactual decision-making rules using POs Y_k and study their reasoning using historical experimental data. We ask the following counterfactual questions for individualized decision-making:

(Question 1). “What is the most probable rank of POs for an individual?”

(Question 2). “Which action is most likely to yield the best outcome for an individual?”

To address Question 1, we propose the following *probability of potential outcome ranking* (PoR) metric for ranks of actions $\mathbf{r} = (r_1, \dots, r_K)$:

$$S_{PoR}(\mathbf{r}) \triangleq \mathbb{P}(Y_{r_1} > Y_{r_2} > \dots > Y_{r_K}). \quad (4)$$

To address Question 2, we propose the following *probability of achieving the best potential outcome* (PoB) metric for actions $r_1 \in \{1, \dots, K\}$:

$$S_{PoB}(r_1) \triangleq \mathbb{P}(\arg \max_{k=1, \dots, K} \{Y_k\} = r_1). \quad (5)$$

Questions 1 and 2 are answered by optimizing the PoR and PoB metrics, respectively.

While the standard RoE metric $S_{RoE}(k) = \mathbb{E}[Y_k]$ evaluates an action based on the population average of its POs, our proposed metrics differ fundamentally. The PoB metric is based on the proportion of individuals for whom the action yields the best PO. The PoR metric does not directly evaluate individual actions, but assesses the rankings of all possible actions based on the proportion of individuals exhibiting the corresponding order of POs. PoR allows for the identification of the most probable rank of POs, and one could then select the best action as the one ordered first in this rank. PoR enables more nuanced decisions, particularly useful when the sequence or relative performance of actions matters, beyond just identifying a single optimal one.

Evaluating the PoR and PoB metrics poses significantly greater challenges than evaluating RoE, as it requires counterfactual reasoning that lies in the third layer of Pearl’s causal hierarchy (Bareinboim et al. 2022; Pearl and Mackenzie 2018), whereas RoE reasoning belongs to the second layer. This paper offers three key contributions: we establish identification theorems for PoR and PoB under the widely used rank invariance assumption (Heckman, Smith, and Clements 1997); we derive bounds for these metrics in the absence of this assumption; and we present corresponding estimation methods. Finally, we conduct numerical experiments to demonstrate the properties of the estimators and present an application using a real-world dataset.

Illustrating Example

We provide an illustrative example that highlights the differences between the three metrics RoE, PoR, and PoB.

Educational Scenario. In educational research, the causal effect of education plays a critical role practically. Specifically, imagine three classes (A, B, and C) designed to enhance students’ overall knowledge proficiency. Decision-makers, such as students, teachers, or parents, need to

Student ID	A	B	C	Rankings
1	30	40	50	A < B < C
2	40	50	60	A < B < C
3	50	60	70	A < B < C
4	30	50	40	A < C < B
5	40	60	50	A < C < B
6	60	70	40	C < A < B
7	50	60	40	C < A < B
8	100	5	0	C < B < A
Average of potential score	50	49.375	43.75	-
Proportion of being the best score	0.125	0.5	0.375	-

Table 1: An example of students’ potential scores. We highlight the maximum value in each row.

choose one course from among them. Table 1 presents an example of the students’ potential scores. Traditionally, decision-makers use the RoE metric and compare the expected scores across the available classes. Class A exhibits the highest expected score, followed by Class B. Based on this comparison, one concludes that Class A is the most beneficial, Class B is a bit less effective, and Class C is the least effective. However, one may reasonably conclude that the optimal ranking is Class C first, followed by Class B, and Class A last, based on the observation that this is the most probable rank (37.5% chance). This conclusion corresponds to the PoR metric. Finally, some other decision-makers might prefer to enroll in Class B based on the PoB metric, since B is most likely to yield the highest score (in 50% of the students).

As illustrated in this example, our PoB and PoR criteria yield different decision-making rules, potentially leading decision-makers to choose different actions. Similar situations arise in medical contexts (such as selecting a type of surgery to improve patient life expectancy), in economic contexts (such as choosing a policy to improve economic performance in each municipality), and in marketing contexts (such as selecting an advertising strategy to influence the attitudes of customers).

Notations and Backgrounds

We represent each variable with a capital letter (X) and its realized value with a small letter (x). Let $\mathbb{I}(x)$ be an indicator function that takes 1 if x is true and 0 if x is false. Denote Ω_Y be the domain of Y , $\mathbb{E}[Y]$ be the expectation of Y , and $F_Y(y) = \mathbb{P}(Y \leq y)$ be the cumulative distribution function (CDF) of a continuous variable Y . The inverse of the CDF is defined as $F_Y^{-1}(p) = \inf\{y; F_Y(y) \geq p\}$ for $p \in [0, 1]$.

Structural causal models (SCM). We use the language of SCMs as our basic semantic and inferential framework (Pearl 2009). An SCM \mathcal{M} is a tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, \mathbb{P}_{\mathbf{U}} \rangle$, where \mathbf{U} is a set of exogenous (unobserved) variables following a joint distribution $\mathbb{P}_{\mathbf{U}}$, and \mathbf{V} is a set of endogenous (observable) variables whose values are determined by structural functions $\mathcal{F} = \{f_{V_i}\}_{V_i \in \mathbf{V}}$ such that $v_i := f_{V_i}(\mathbf{pa}_{V_i}, \mathbf{u}_{V_i})$ where $\mathbf{PA}_{V_i} \subseteq \mathbf{V}$ and $U_{V_i} \subseteq \mathbf{U}$. An atomic intervention of setting a set of endogenous variables \mathbf{X} to constants \mathbf{x} , denoted by $do(\mathbf{x})$, replaces the original equations of \mathbf{X} by $\mathbf{X} := \mathbf{x}$ and induces a *sub-model* $\mathcal{M}_{\mathbf{x}}$. We denote the potential outcome (PO) Y under intervention $do(\mathbf{x})$

by $Y_x(\mathbf{u})$, which is a solution of Y given $\mathbf{U} = \mathbf{u}$ in the sub-model \mathcal{M}_x .

Identification of Counterfactuals. Researchers often consider the following SCM, \mathcal{M} :

$$Y := f_Y(X, U^Y), X := f_X(U^X), \quad (6)$$

where U^Y and U^X are latent exogenous variables. Y is a continuous variable and X is a discrete variable taking values in $\Omega_X = \{1, 2, \dots, K\}$. We assume the following:

Assumption 1 (Continuity and Differentiability of CDF of Potential Outcome). $F_{Y_x}(y)$ is continuous and differentiable for any $x \in \Omega_X$ and $y \in \Omega_Y$.

Assumption 2 (Exogeneity). $Y_x \perp\!\!\!\perp X$ for any $x \in \Omega_X$.

Remark. Assumption 2 is used to identify $\mathbb{P}(Y_x)$, under which we have $\mathbb{P}(Y_x) = P(Y|x)$. Alternatively, we may assume *conditional exogeneity* given covariates W , i.e., $Y_x \perp\!\!\!\perp X|W$ for any $x \in \Omega_X$, in which case $\mathbb{P}(Y_x) = \int P(Y|x, w)\mathbf{p}(w)dw$, where \mathbf{p} denotes probability density function. Another possibility is that the interventional distribution $\mathbb{P}(Y_x)$ is directly available from experimental data. All identification and bounding results in the paper are expressed in terms of $F_{Y_x}(y) = F_{Y|X=x}(y)$ under Assumption 2. However, these results remain valid under conditional exogeneity (by replacing $F_{Y|X=x}(y)$ with $\int F_{Y|X=x, w}(y)\mathbf{p}(w)dw$) or when experimental data are available (by replacing $F_{Y|X=x}(y)$ with $F_{Y_x}(y)$).

We have the following proposition:

Proposition 1 (Identifiability of RoE). *Under SCM \mathcal{M} and Assumption 2, the expected PO $\mathbb{E}[Y_x]$ is identifiable by $\mathbb{E}[Y|X = x]$.*

Heckman, Smith, and Clements (1997) showed that for binary treatment $\Omega_X = \{1, 2\}$, the value of the unobserved PO is identified in terms of the observed one under a rank-invariant assumption as follows.

Assumption 3 (Rank invariance assumption). (Heckman, Smith, and Clements 1997) $F_{Y_1}(y_1) = F_{Y_2}(y_2)$ holds for almost every subject whose POs are $(Y_1, Y_2) = (y_1, y_2)$.

Proposition 2 (Identification of the unobserved PO for two POs). (Heckman, Smith, and Clements 1997) *Under SCM \mathcal{M} and Assumptions 1, 2, and 3, given CDF $F_{Y|X=1}$ and $F_{Y|X=2}$, for almost every subject whose POs are $(Y_1, Y_2) = (y_1, y_2)$, the value of y_2 is identifiable by*

$$y_2 = F_{Y|X=2}^{-1}(F_{Y|X=1}(y_1)). \quad (7)$$

The function $F_{Y|X=2}^{-1}(F_{Y|X=1}(\cdot))$ is called the shift function in (Doksum 1974; Doksum and Sievers 1976).

Parameter ψ in (Fay et al. 2018). Fay et al. (2018) have studied the causal parameter

$$\psi \triangleq \mathbb{P}(Y_2 > Y_1), \quad (8)$$

which can be considered as a special case of PoR when $K = 2$. They provide bounds for ψ under Assumption 2, i.e., $\psi_L \leq \psi \leq \psi_U$, where

$$\psi_L = \sup_y \{F_{Y|X=1}(y) - F_{Y|X=2}(y)\}, \quad (9)$$

$$\psi_U = 1 - \sup_y \{F_{Y|X=2}(y) - F_{Y|X=1}(y)\}. \quad (10)$$

The parameter ψ has been actively studied in causal inference (Fan and Park 2010; Huang et al. 2016; Lu, Zhang, and Ding 2019; Gabriel, Sachs, and Jensen 2024).

Decision-Making Using PoR and PoB

In this section, we formally introduce new counterfactual decision-making rules inspired by DCM: PoR and PoB.

Probabilities of potential outcome ranking (PoR). Let $\mathbf{R} = (R_1, R_2, \dots, R_K)$ denote a permutation of elements in $\{1, 2, \dots, K\}$, e.g., $\mathbf{R} = (1, 2, \dots, K)$. We denote the set of $K!$ elements that \mathbf{R} takes as \mathfrak{R} . We then define the PoR metric as follows:

Definition 1 (PoR). *The probabilities of potential outcome ranking (PoR) of $\mathbf{r} = (r_1, \dots, r_K)$, for each $\mathbf{r} \in \mathfrak{R}$, are defined as*

$$S_{PoR}(\mathbf{r}) \triangleq \mathbb{P}(\mathbf{R} = \mathbf{r}) \triangleq \mathbb{P}(Y_{r_1} > Y_{r_2} > \dots > Y_{r_K}). \quad (11)$$

The vector \mathbf{R} can be regarded as a new discrete random variable with domain \mathfrak{R} , whose distribution is given by Eq. (11). Since Y_x is a continuous random variable (Assumption 1), the probability of POs being tied is zero. Throughout this paper, we do not consider cases in which POs are tied.

Decision-makers may seek the ranking with the largest PoR, i.e.,

$$\mathbf{r}^{PoR} \triangleq \arg \max_{\mathbf{r} \in \mathfrak{R}} S_{PoR}(\mathbf{r}). \quad (12)$$

\mathbf{r}^{PoR} answers (**Question 1**). In concrete terms, decision-makers that employ the PoR metric conclude r_k^{PoR} as the k -th preferred action, for $k = 1, \dots, K$. As an example, we present the six patterns ($3! = 6$) of ranking for POs when $K = 3$ in Table 2. PoR assesses the likelihood of each ranking pattern.

Probability of achieving the best potential outcome (PoB). Some decision-makers may focus on only the action that achieves the largest POs. For example, in Table 2, action 1 obtains the largest PO in patterns (A) and (B), that is, $R_1 = 1$; $R_1 = 2$ in patterns (C) and (D); and $R_1 = 3$ in patterns (E) and (F). We define the probability of an action achieving the best potential outcome (PoB) as the marginalized form of PoR.

Definition 2 (PoB). *We define the probability of action r_1 achieving the best potential outcome (PoB), for each $r_1 \in \{1, 2, \dots, K\}$, as*

$$S_{PoB}(r_1) \triangleq \mathbb{P}(R_1 = r_1) = \sum_{\mathbf{r} \in \mathfrak{R}; R_1=r_1} \mathbb{P}(\mathbf{R} = \mathbf{r}). \quad (13)$$

Alternatively,

$$S_{PoB}(r_1) = \mathbb{P}(\arg \max_{k=1, \dots, K} \{Y_k\} = r_1). \quad (14)$$

Decision-makers seek the action that has the largest PoB:

$$r_1^{PoB} \triangleq \arg \max_{r_1=1, 2, \dots, K} S_{PoB}(r_1). \quad (15)$$

r_1^{PoB} answers (**Question 2**). Decision-makers employing PoB conclude r_1^{PoB} as the most preferred action. Note that r_1^{PoB} may not be equal to r_1^{PoR} , which is the most preferred action by \mathbf{r}^{PoR} .

Ranking Patterns	R_1	R_2	R_3
(A)	1	2	3
(B)	1	3	2
(C)	2	1	3
(D)	2	3	1
(E)	3	1	2
(F)	3	2	1

Table 2: Six patterns of rankings for POs (Y_1, Y_2, Y_3).

Example 1 (Homogeneous causal effects). We consider SCM $Y = X + U^Y$ with $\mathbb{E}[U^Y] = 0$ and $K = 3$. We have $Y_k = k + U^Y$, and the individual causal effect (ICE) $Y_3 - Y_2 = Y_2 - Y_1 = 1$ for every subject, meaning that the causal effect is homogeneous. The expectations of Y_1, Y_2 , and Y_3 are 1, 2, and 3, respectively. $Y_3 > Y_2 > Y_1$ holds for every subject. Then $\mathbb{P}(\mathbf{R} = (3, 2, 1)) = 1$ and PoR of all other ranking patterns is equal to 0. Therefore, $\mathbf{r}^{PoR} = (3, 2, 1)$. Finally, $S_{PoB}(3) = \mathbb{P}(R_1 = 3) = 1$ with PoB of all other actions being 0; therefore $r_1^{PoB} = 3$. We conclude that action 3 is the optimal choice from the perspectives of all three metrics RoE, PoR, and PoB. In general, the three metrics lead to the same optimal choice when the causal effect is homogeneous, meaning $Y_x - Y_{x-1}$ is constant.

Example 2 (Heterogeneous causal effects). We consider SCM $Y = XU^Y$ with $K = 3$, $\mathbb{E}[U^Y] = 0$, $\mathbb{P}(U^Y = 0) = 0$, and $\mathbb{P}(U^Y > 0) = 0.6$. We have $Y_k = kU^Y$, and ICE $Y_3 - Y_2 = Y_2 - Y_1 = U^Y$, which varies between subjects. The expectations of Y_1, Y_2 , and Y_3 are all 0. We have $\mathbb{P}(\mathbf{R} = \mathbf{r}) = \mathbb{P}(Y_{r_1} > Y_{r_2} > Y_{r_3}) = \mathbb{P}(r_1 U^Y > r_2 U^Y > r_3 U^Y)$, $\mathbb{P}(r_1 U^Y > r_2 U^Y > r_3 U^Y | U^Y > 0) = \mathbb{I}(r_1 > r_2 > r_3)$, and $\mathbb{P}(r_1 U^Y > r_2 U^Y > r_3 U^Y | U^Y < 0) = \mathbb{I}(r_3 > r_2 > r_1)$. Then, $\mathbb{P}(\mathbf{R} = \mathbf{r}) = \mathbb{I}(r_1 > r_2 > r_3) \mathbb{P}(U^Y > 0) + \mathbb{I}(r_3 > r_2 > r_1) \mathbb{P}(U^Y < 0)$. We obtain $\mathbb{P}(\mathbf{R} = (3, 2, 1)) = 0.6$, $\mathbb{P}(\mathbf{R} = (1, 2, 3)) = 0.4$, and the PoR of all other ranking patterns are equal to 0. Therefore, $\mathbf{r}^{PoR} = (3, 2, 1)$. Finally, $S_{PoB}(3) = 0.6$, $S_{PoB}(1) = 0.4$, and $S_{PoB}(2) = 0$. We conclude that the three actions are tied from the perspective of RoE; action 3 is the optimal choice, followed by action 2, with action 1 ranked last from the perspective of PoR; and action 3 is the optimal choice, followed by action 1, with action 2 ranked last from the perspective of PoB.

Remark (RoE, PoR, PoB). While PoR is constructed to assess the preference rankings of actions, one can interpret r_1^{PoR} , the first element in \mathbf{r}^{PoR} , as the optimal action. On the other hand, although RoE and PoB are designed to select a single best action, they can be interpreted to measure the quality of a ranking $\mathbf{r} = (r_1, \dots, r_K)$ of actions by the following metrics:

$$S_{RoE}(\mathbf{r}) = \mathbb{I}(\mathbb{E}[Y_{r_1}] \geq \mathbb{E}[Y_{r_2}] \geq \dots \geq \mathbb{E}[Y_{r_K}]), \quad (16)$$

and

$$S_{PoB}(\mathbf{r}) = \mathbb{I}(S_{PoB}(r_1) \geq S_{PoB}(r_2) \geq \dots \geq S_{PoB}(r_K)), \quad (17)$$

respectively. In other words, RoE and PoB select the best ranking of actions as $\mathbf{r}^{RoE} \triangleq \arg \max_{\mathbf{r} \in \mathfrak{R}} S_{RoE}(\mathbf{r})$ and $\mathbf{r}^{PoB} \triangleq \arg \max_{\mathbf{r} \in \mathfrak{R}} S_{PoB}(\mathbf{r})$, respectively.

Remark (PoC). The probabilities of causation (PoC) are a family of counterfactual probabilities quantifying whether one event was the real cause of another in a given scenario (Pearl 1999; Tian and Pearl 2000). Personalized decision-making based on PoC with discrete treatments and outcomes has been explored by Mueller and Pearl (2023) and Li and Pearl (2024). PoC with continuous outcomes were examined by Kawakami, Kuroki, and Tian (2024). PoC expressions like $\mathbb{P}(Y_{x_0} < y_1 \leq Y_{x_1} \dots < y_p \leq Y_{x_p})$ in (Kawakami, Kuroki, and Tian 2024) may appear similar to PoR. The key difference is that PoR does not have thresholds like y_i for POs. For binary Y and X ($K = 2$) taking values in $\{1, 2\}$, our PoR and PoB both reduce to the probability of necessity and sufficiency (PNS) (Pearl 1999; Tian and Pearl 2000). Specifically, $S_{PoR}(2, 1) = S_{PoB}(2) = \mathbb{P}(Y_2 > Y_1) = \mathbb{P}(Y_2 = 2, Y_1 = 1)$.

Identification and Estimation of PoR and PoB

Next, we investigate the identification challenges associated with PoR and PoB. By Proposition 1, the RoE metric is identifiable under the exogeneity assumption. However, the identification of PoR and PoB is more challenging and requires additional assumptions.

Identification of the values of the unobserved POs. We extend the identification result by Heckman, Smith, and Clements (1997) given in Proposition 2 to the setting of multiple actions.

Assumption 4 (Rank invariance assumption for multiple actions). $F_{Y_1}(y_1) = F_{Y_2}(y_2) = \dots = F_{Y_K}(y_K)$ holds for almost every subject whose POs are $(Y_1, Y_2, \dots, Y_K) = (y_1, y_2, \dots, y_K)$.

This assumption means that a subject whose PO under action 1 lies in the top q percent also has POs in the top q percent under all other actions $k = 2, \dots, K$. Then, given Y_1 , the values of other unobserved POs are identifiable as follows:

Lemma 1 (Identification of the unobserved POs). *Under SCM \mathcal{M} and Assumptions 1, 2, and 4, given CDF $F_{Y|X=k}$ for $k = 1, \dots, K$, for almost every subject whose POs are $(Y_1, \dots, Y_K) = (y_1, \dots, y_K)$, the values of y_2, \dots, y_K are identifiable by*

$$y_k = F_{Y|X=k}^{-1}(F_{Y|X=1}(y_1)). \quad (18)$$

Identification of PoR and PoB. We apply the identification result in Lemma 1 to the identification of PoR and PoB.

Theorem 1 (Identification of PoR). *Under SCM \mathcal{M} and Assumptions 1, 2, and 4, $S_{PoR}(\mathbf{r})$ is identified by $\sigma(\mathbf{r})$ for any $\mathbf{r} \in \mathfrak{R}$, given by*

$$\sigma(\mathbf{r}) = \mathbb{P}(Y > F_{Y|X=r_2}^{-1}(F_{Y|X=r_1}(Y)) > \dots > F_{Y|X=r_K}^{-1}(F_{Y|X=r_1}(Y)) | X = r_1). \quad (19)$$

Theorem 2 (Identification of PoB). *Under SCM \mathcal{M} and Assumptions 1, 2, and 4, $S_{PoB}(r_1)$ is identified by $\eta(r_1)$ for any $r_1 \in \{1, \dots, K\}$, given by*

$$\eta(r_1) = \mathbb{P}(Y > F_{Y|X=r_2}^{-1}(F_{Y|X=r_1}(Y)),$$

$$Y > F_{Y|X=r_3}^{-1}(F_{Y|X=r_1}(Y)), \dots, \\ Y > F_{Y|X=r_K}^{-1}(F_{Y|X=r_1}(Y)) \Big| X = r_1, \quad (20)$$

where (r_2, \dots, r_K) is an arbitrary permutation of elements in $\{1, 2, \dots, K\} \setminus \{r_1\}$.

When $K = 2$, Theorem 1 establishes the identification of the causal parameter ψ in (8), which was not addressed in (Fay et al. 2018).

Estimation of PoR and PoB. We present an estimation method for PoR and PoB following the approach proposed in (Doksum 1974; Doksum and Sievers 1976). We denote the samples of Y corresponding to $X = k$ (or samples from the intervention $do(X = k)$ in the setting of given experimental data) as

$$\mathcal{D}^{(k)} = \{Y_1^{(k)}, \dots, Y_{N^{(k)}}^{(k)}\} \quad (21)$$

for $k = 1, 2, \dots, K$, where $N^{(k)}$ is the sample size of $\mathcal{D}^{(k)}$. We sort dataset $\mathcal{D}^{(k)}$ and denote

$$\mathcal{D}^{(k)} = \{Y_{(1)}^{(k)}, \dots, Y_{(N^{(k)})}^{(k)}\} \quad (22)$$

where $Y_{(1)}^{(k)} \leq \dots \leq Y_{(N^{(k)})}^{(k)}$. Then, our estimates of $\sigma(\mathbf{r})$ are given by

$$\hat{\sigma}(\mathbf{r}) \\ = \frac{1}{N^{(r_1)}} \sum_{i=1}^{N^{(r_1)}} \mathbb{I}\left(Y_i^{(r_1)} > Y_{\left(\lfloor \frac{N^{(r_2)} \hat{F}_{Y|X=r_2}(Y_i^{(r_1)}) \rfloor + 1\right)}^{(r_2)}\right) \\ > \dots > Y_{\left(\lfloor \frac{N^{(r_K)} \hat{F}_{Y|X=r_K}(Y_i^{(r_1)}) \rfloor + 1\right)}^{(r_K)}\right), \quad (23)$$

where

$$\hat{F}_{Y|X=r_k}(y) = \frac{1}{N^{(r_k)}} \sum_{j=1}^{N^{(r_k)}} \mathbb{I}(Y_j^{(r_k)} \leq y). \quad (24)$$

Our estimates of $\eta(r_1)$ are given by

$$\hat{\eta}(r_1) \\ = \frac{1}{N^{(r_1)}} \sum_{i=1}^{N^{(r_1)}} \mathbb{I}\left(Y_i^{(r_1)} > Y_{\left(\lfloor \frac{N^{(r_2)} \hat{F}_{Y|X=r_2}(Y_i^{(r_1)}) \rfloor + 1\right)}^{(r_2)}\right) \\ > Y_{\left(\lfloor \frac{N^{(r_3)} \hat{F}_{Y|X=r_3}(Y_i^{(r_1)}) \rfloor + 1\right)}^{(r_3)}, \dots, \\ > Y_{\left(\lfloor \frac{N^{(r_K)} \hat{F}_{Y|X=r_K}(Y_i^{(r_1)}) \rfloor + 1\right)}^{(r_K)}\right). \quad (25)$$

The computational complexities for $\hat{\sigma}(\mathbf{r})$ and $\hat{\eta}(r_1)$ are $\mathcal{O}\left(\sum_{k=1}^K N^{(k)} \log N^{(k)}\right)$, as each sorting operation on $\mathcal{D}^{(k)}$ requires $\mathcal{O}(N^{(k)} \log N^{(k)})$ computations. They are consistent estimators as follows:

Theorem 3. Under SCM \mathcal{M} and Assumptions 1, 2, and 4, $\hat{\sigma}(\mathbf{r})$ and $\hat{\eta}(r_1)$ converge in probability to $S_{PoR}(\mathbf{r})$ and $S_{PoB}(r_1)$, respectively, at the rate $\mathcal{O}_p(\sum_{k=1}^K N^{(k)-1/2})$.

When $N^{(k)} = N/K$ for any k , the rate becomes $\mathcal{O}_p(K^{3/2}N^{-1/2})$, which worsens as K increases.

Bounding PoR and PoB

The exogeneity (Assumption 2) is considered reasonable for experimental data by many researchers. In contrast, some researchers consider the rank invariance assumption (Assumption 4) overly restrictive; therefore, we derive bounds of PoR and PoB without relying on Assumption 4.

Since $S_{PoR}(\mathbf{r})$ is given as

$$\mathbb{P}((Y_{r_1} > Y_{r_2}) \wedge (Y_{r_2} > Y_{r_3}) \wedge \dots \wedge (Y_{r_{K-1}} > Y_{r_K})), \quad (26)$$

$S_{PoR}(\mathbf{r})$ is bounded as follows from the Fréchet inequalities (Fréchet 1935, 1960):

$$\max\left\{\sum_{k=1}^{K-1} \mathbb{P}(Y_{r_k} > Y_{r_{k+1}}) - K + 2, 0\right\} \leq S_{PoR}(\mathbf{r}) \\ \leq \min_{k=1, \dots, K-1} \{\mathbb{P}(Y_{r_k} > Y_{r_{k+1}})\}. \quad (27)$$

Then, we have the following results using the bounds of ψ given in Eqs. (9) and (10):

Theorem 4 (Bounds of PoR). Under SCM \mathcal{M} and Assumptions 1 and 2, for any $\mathbf{r} \in \mathfrak{R}$, $S_{PoR}(\mathbf{r})$ is bounded by $\sigma_L(\mathbf{r}) \leq S_{PoR}(\mathbf{r}) \leq \sigma_U(\mathbf{r})$, where

$$\sigma_L(\mathbf{r}) = \max\left\{\sum_{k=1}^{K-1} \sup_y \{F_{Y|X=r_{k+1}}(y) - F_{Y|X=r_k}(y)\} - K + 2, 0\right\}, \quad (28)$$

$$\sigma_U(\mathbf{r}) = \min_{k=1, \dots, K-1} \left\{1 - \sup_y \{F_{Y|X=r_k}(y) - F_{Y|X=r_{k+1}}(y)\}\right\}. \quad (29)$$

The bounds in Theorem 4 for $K = 2$ reduce to the bounds of ψ . Similarly, since $S_{PoB}(r_1)$ is given as

$$\mathbb{P}((Y_{r_1} > Y_{r_2}) \wedge (Y_{r_1} > Y_{r_3}) \wedge \dots \wedge (Y_{r_1} > Y_{r_K})), \quad (30)$$

$S_{PoB}(r_1)$ is bounded as follows from the Fréchet inequalities:

$$\max\left\{\sum_{k=1}^{K-1} \mathbb{P}(Y_{r_1} > Y_{r_{k+1}}) - K + 2, 0\right\} \leq S_{PoB}(r_1) \\ \leq \min_{k=1, \dots, K-1} \{\mathbb{P}(Y_{r_1} > Y_{r_{k+1}})\}. \quad (31)$$

Then, we have the following results:

Theorem 5 (Bounds of PoB). Under SCM \mathcal{M} and Assumptions 1 and 2, for any $r_1 = 1, 2, \dots, K$, $S_{PoB}(r_1)$ is bounded by $\eta_L(r_1) \leq S_{PoB}(r_1) \leq \eta_U(r_1)$, where

$$\eta_L(r_1) = \max\left\{\sum_{k=1}^{K-1} \sup_y \{F_{Y|X=r_{k+1}}(y) - F_{Y|X=r_1}(y)\} - K + 2, 0\right\}, \quad (32)$$

$$\eta_U(r_1) = \min_{k=1, \dots, K-1} \left\{1 - \sup_y \{F_{Y|X=r_1}(y) - F_{Y|X=r_{k+1}}(y)\}\right\}. \quad (33)$$

These bounds of PoR and PoB are not sharp except for the case of $K = 2$.

Estimation of Bounds. We can estimate the bounds in Eqs. (28), (29), (32), and (33) by plugging the estimates of conditional CDF $\hat{F}_{Y|X=r_k}(y)$ given in (24). The bound estimators are

$$\hat{\sigma}_L(\mathbf{r}) = \max \left\{ \sum_{k=1}^{K-1} \text{s\hat{u}p}_y \{ \hat{F}_{Y|X=r_{k+1}}(y) - \hat{F}_{Y|X=r_k}(y) \} - K + 2, 0 \right\}, \quad (34)$$

$$\hat{\sigma}_U(\mathbf{r}) = \min_{k=1, \dots, K-1} \left\{ 1 - \text{s\hat{u}p}_y \{ \hat{F}_{Y|X=r_k}(y) - \hat{F}_{Y|X=r_{k+1}}(y) \} \right\}, \quad (35)$$

$$\hat{\eta}_L(r_1) = \max \left\{ \sum_{k=1}^{K-1} \text{s\hat{u}p}_y \{ \hat{F}_{Y|X=r_{k+1}}(y) - \hat{F}_{Y|X=r_1}(y) \} - K + 2, 0 \right\}, \quad (36)$$

$$\hat{\eta}_U(r_1) = \min_{k=1, \dots, K-1} \left\{ 1 - \text{s\hat{u}p}_y \{ \hat{F}_{Y|X=r_1}(y) - \hat{F}_{Y|X=r_{k+1}}(y) \} \right\}, \quad (37)$$

where $\text{s\hat{u}p}_y$ denotes the supremum operator evaluated over a finite grid. Specifically, for $\Omega_Y = [a, b]$, $\text{s\hat{u}p}_{y \in \Omega_Y} f(y) := \max_{i=1, \dots, M} f(y^i)$, where $y^i = a + \frac{i-1}{M-1}(b-a)$ for $i = 1, \dots, M$. The computational complexities for $\hat{\sigma}_L(\mathbf{r})$, $\hat{\sigma}_U(\mathbf{r})$, $\hat{\eta}_L(r_1)$, and $\hat{\eta}_U(r_1)$ are $\mathcal{O}(KM \sum_{k=1}^K N^{(k)})$. These estimators are consistent as follows:

Theorem 6. *Under SCM \mathcal{M} and Assumptions 1 and 2, if Ω_Y is bounded, $\hat{\sigma}_L(\mathbf{r})$, $\hat{\sigma}_U(\mathbf{r})$, $\hat{\eta}_L(r_1)$, and $\hat{\eta}_U(r_1)$ converge in probability to $\sigma_L(\mathbf{r})$, $\sigma_U(\mathbf{r})$, $\eta_L(r_1)$, and $\eta_U(r_1)$, respectively, at the rate $\mathcal{O}_p(1/M + \sum_{k=1}^K N^{(k)-1/2})$.*

Numerical Experiments

We illustrate the finite-sample properties of the estimators. To our knowledge, there are no existing baselines for estimating PoR or PoB when the number of actions $K \geq 3$.

Three Candidate Actions

First, we evaluate the performance of the PoR and PoB estimators in the case of three candidate actions $K = 3$.

Settings. We assume the following SCMs:

$$(A): Y := -(X-4)U^Y, U^Y \sim \text{Unif}(-0.5, 1), \quad (38)$$

$$(B): Y := -XU^Y + U'^Y,$$

$$U^Y \sim \mathcal{N}(-1, 1), U'^Y \sim \text{Unif}(-1, 1), \quad (39)$$

where X takes $\{1, 2, 3\}$ and $\mathbb{P}(X = k) = 1/3$ for any $k = 1, 2, 3$. Both SCMs satisfy Assumptions 1 and 2. SCM (A) satisfies the rank invariance assumption (Assumption 4), whereas SCM (B) violates it. We use SCM (A) to illustrate the performance of our PoR and PoB estimators, and use SCM (B) to illustrate the performance of the bounds estimators. We perform 100 simulation runs and report the mean along with the 95% confidence intervals (CI) for each statistic. We use sample size $N = KN^{(k)}$ with varying $N^{(k)} = N' = 30, 300, 3000$ for $k = 1, 2, 3$.

Results. We present in Table 3 the estimates for $S_{PoR}(1, 2, 3)$ and $S_{PoB}(1)$ under SCM (A), and the upper

bound (UB) and lower bound (LB) for $S_{PoR}(1, 2, 3)$ and $S_{PoB}(1)$ under SCM (B). The results under SCM (A) show that the estimates become more reliable as the sample size increases. Our model correctly predicts that the most probable ranking is (1, 2, 3), and the action that is the most likely to yield the best outcome is 1 (both probabilities are greater than 0.5). The results under SCM (B) show that the ground truth value lies within the computed bounds, and the CIs of estimates become narrower as the sample size increases. All the computed lower bounds are non-informative for this model.

Many Candidate Actions

Next, we examine the reliability of the estimated PoR and PoB under the rank invariance assumption as the number of actions K increases.

Setting. We assume the following SCM:

$$(C): Y := -XU^Y, U^Y \sim \text{Unif}(-1, 1), \quad (40)$$

where X takes $\{1, 2, \dots, K\}$ and $\mathbb{P}(X = k) = 1/K$ for any $k = 1, 2, \dots, K$. This SCM satisfies Assumptions 1, 2, and 4. We perform 100 simulation runs and report the mean along with the 95% CI for each estimate, using a fixed per action sample size of $N' = 3000$ and varying the number of actions $K \in \{3, 5, 10, 20\}$.

Results. Table 4 shows the absolute errors of estimated $P_{PoR}(1, 2, \dots, K)$ and $S_{PoB}(1)$ from their ground truths. As the number of candidate actions K increases, the estimates deviate further from the ground truth. We present curves of absolute errors of estimates vs. K in Figures 1 and 2 in Appendix B in (Kawakami and Tian 2025). The errors roughly increase linearly in terms of K , matching with the convergence rate results in Theorem 3.

Application to Real-World Data

Dataset. We pick up an open medical dataset ‘‘coagulation’’ (Kropf et al. 2000) in (<https://cran.r-project.org/web/packages/SimComp/SimComp.pdf>). Three extracorporeal circulation systems in heart-lung machines were evaluated: treatments ‘‘H’’ and ‘‘B’’, alongside the standard ‘‘S’’. The study included 12 male adult patients each for treatments S and H, and 11 for treatment B. The analysis focused on laboratory parameters related to blood coagulation, using three primary endpoints expressed as the ratio of post- to pre-surgery values. Higher values indicate a more beneficial or favorable treatment effect. Kropf et al. (2000) compared the mean vectors of outcomes using their developed statistical hypothesis testing method. We focus on an endpoint referred to as ‘‘Thromb.count’’ among the three endpoints. Decision-makers, such as doctors, try to choose the most effective treatment to achieve better endpoint results. We denote POs as Y_S , Y_H , and Y_B .

Analysis. We adopt Assumption 2 (exogeneity), as Kropf et al. (2000) did not report the presence of any confounding variables. We impose Assumption 1, as the outcome variable is continuously distributed. We assume the rank invariant assumption (Assumption 4). The rank invariance assumption implies that a patient who belongs to the top q percent in

Setting	Measures	$N' = 30$	$N' = 300$	$N' = 3000$	True Value
(A)	$S_{PoR}(1, 2, 3)$	0.557 [0.149,0.918]	0.580 [0.421,0.720]	0.635 [0.560,0.638]	0.666
(A)	$S_{PoB}(1)$	0.561 [0.133,0.884]	0.645 [0.439,0.763]	0.650 [0.580,0.702]	0.666
(B)	UB of $S_{PoR}(1, 2, 3)$	0.813 [0.683,0.900]	0.857 [0.806,0.890]	0.873 [0.855,0.887]	0.500
(B)	LB of $S_{PoR}(1, 2, 3)$	0.000 [0.000,0.000]	0.000 [0.000,0.000]	0.000 [0.000,0.000]	0.500
(B)	UB of $S_{PoB}(1)$	0.767 [0.649,0.867]	0.763 [0.708,0.813]	0.767 [0.745,0.798]	0.500
(B)	LB of $S_{PoB}(1)$	0.000 [0.000,0.000]	0.000 [0.000,0.000]	0.000 [0.000,0.000]	0.500

Table 3: Results of experiments (three candidate actions): Mean estimates with 95% CI. UB: upper bound; LB: lower bound.

Setting	Measures	$K = 3$	$K = 5$	$K = 10$	$K = 20$
(C)	$S_{PoR}(1, 2, \dots, K)$	0.022 [0.001,0.067]	0.047 [0.007,0.104]	0.117 [0.049,0.210]	0.256 [0.147,0.391]
(C)	$S_{PoB}(1)$	0.015 [0.001,0.043]	0.015 [0.001,0.044]	0.016 [0.004,0.045]	0.023 [0.012,0.048]

Table 4: Results of experiments (many candidate actions): Absolute errors of the estimates from the truth, with 95% CIs.

terms of the endpoint under treatment S also belongs to the top q percent under treatments B and H . We conduct the bootstrapping (Efron 1979) to provide the means and 95% confidence intervals (CIs) for each estimator.

Results. We show the estimates of RoE, six patterns of PoR, and three PoB in Table 5. We also compute the bounds for each estimator and present the results in Appendix C in (Kawakami and Tian 2025). The bounds are relatively wide. Overall, due to the small sample size, the CIs of the estimates for different treatments have considerable overlap, and it’s difficult to draw reliable conclusions about the effectiveness ranking of the three treatments.

In the following, we analyze the results purely based on the mean values. From Table 5, all three metrics conclude that treatment B is the most effective. After B , RoE and PoB will (slightly) prefer treatment H over S . However, the most probable ranking by PoR is (B, S, H) , which prefers S over H . Our proposed PoR and PoB metrics constitute alternative frameworks for decision-making, which may yield action choices that in some cases differ significantly from those recommended by traditional RoE. Furthermore, they provide more nuanced insights into treatment effects. For example, PoR reveals that, for 8.1% of patients, treatment S is the most effective, followed by H , and then B . This indicates the existence of patients whose outcome ranking is completely opposite to that implied by RoE. As another example, although treatment B is the most effective on average, PoB reveals that B is the optimal choice for only 49% of patients, not for the other 51%. The information provided by PoR and PoB enables clinicians to better understand the benefits and risks associated with each treatment choice.

Conclusion

We propose two novel counterfactual decision-making rules. Decision-making based on PoR or PoB differs from that based on the traditional RoE. We do not claim that one approach is universally better than the other, as the choice often depends on the decision-maker’s perspective and preferences. An individual who prioritizes the expected outcome—possibly reflecting a more aggressive decision-making style—may prefer actions based on RoE, even when the probability of receiving the best individual outcome is

RoE, PoR, or PoB	Estimates
$\mathbb{E}[Y_B]$	0.994 [0.859,1.128]
$\mathbb{E}[Y_H]$	0.917 [0.757,1.087]
$\mathbb{E}[Y_S]$	0.873 [0.782,0.984]
$S_{PoR}(B, H, S)$	0.215 [0.000,0.500]
$S_{PoR}(B, S, H)$	0.277 [0.083,0.583]
$S_{PoR}(H, B, S)$	0.172 [0.000,0.467]
$S_{PoR}(H, S, B)$	0.123 [0.000,0.333]
$S_{PoR}(S, B, H)$	0.132 [0.000,0.333]
$S_{PoR}(S, H, B)$	0.081 [0.000,0.250]
$S_{PoB}(B)$	0.490 [0.182,0.818]
$S_{PoB}(H)$	0.295 [0.083,0.583]
$S_{PoB}(S)$	0.216 [0.000,0.500]

Table 5: Mean estimates with 95% CIs of the RoE, PoR, and PoB metrics. We highlight the highest values.

low. In contrast, a more conservative decision-maker may favor PoB, prioritizing the likelihood of achieving the best individual outcome, even if it means foregoing the action with the highest expected outcome. There may also be decision-makers who aim to balance different perspectives, incorporating considerations of all three RoE, PoR, and PoB into their choices.

Evaluating the PoR and PoB metrics poses greater challenges than evaluating RoE. Our identification results rely on the restrictive rank invariance assumption, which may limit their practical applications. We provide bounds that do not rely on this assumption, but they are not sharp in general. Deriving tight bounds for PoR and PoB remains an important direction for future research. We see personalized decision-making using PoB or PoR as a promising direction for future research, offering an alternative to the traditional RoE, such as in multi-armed bandit (MAB) or reinforcement learning problems.

Acknowledgements

The authors thank the anonymous reviewers for their time and thoughtful comments.

References

- Anderson, S. P.; De Palma, A.; and Thisse, J.-F. 1992. *Discrete choice theory of product differentiation*. MIT press.
- Auer, P.; Cesa-Bianchi, N.; and Fischer, P. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Mach. Learn.*, 47(2–3): 235–256.
- Bareinboim, E.; Correa, J. D.; Ibeling, D.; and Icard, T. 2022. *On Pearl’s Hierarchy and the Foundations of Causal Inference*, 507–556. New York, NY, USA: Association for Computing Machinery, 1 edition.
- Bartholomew, D. J. 1959. A Test Of Homogeneity for Ordered Alternatives. *Biometrika*, 46(1/2): 36–48.
- Ben-Akiva, M. E.; and Lerman, S. R. 1985. *Discrete choice analysis: theory and application to travel demand*, volume 9. MIT press.
- Bernoulli, D. 1954. Exposition of a New Theory on the Measurement of Risk. *Econometrica*, 22(1): 23–36.
- Dawid, P. 2021. Decision-theoretic foundations for statistical causality. *Journal of Causal Inference*, 9(1): 39–77.
- Doksum, K. 1974. Empirical Probability Plots and Statistical Inference for Nonlinear Models in the Two-Sample Case. *The Annals of Statistics*, 2(2): 267 – 277.
- Doksum, K. A.; and Sievers, G. L. 1976. Plotting with Confidence: Graphical Comparisons of Two Populations. *Biometrika*, 63(3): 421–434.
- Efron, B. 1979. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1): 1 – 26.
- Fan, Y.; and Park, S. S. 2010. Sharp Bounds on the Distribution of the Treatment Effect and Their Statistical Inference. *Econometric Theory*, 26(3): 931–951.
- Fay, M. P.; Brittain, E. H.; Shih, J. H.; Follmann, D. A.; and Gabriel, E. E. 2018. Causal estimands and confidence intervals associated with Wilcoxon-Mann-Whitney tests in randomized experiments. *Stat Med*, 37(20): 2923–2937.
- Fréchet, M. 1935. Généralisation du théorème des probabilités totales. *Fundamenta Mathematicae*, 25(1): 379–387.
- Fréchet, M. 1960. Sur les tableaux dont les marges et des bornes sont données. *Revue de l’Institut International de Statistique / Review of the International Statistical Institute*, 28(1/2): 10–32.
- Gabriel, E. E.; Sachs, M. C.; and Jensen, A. K. 2024. Sharp symbolic nonparametric bounds for measures of benefit in observational and imperfect randomized studies with ordinal outcomes. *Biometrika*, 111(4): 1429–1436.
- Gibbard, A.; and Harper, W. 1981. Counterfactuals and two kinds of expected utility. Ifs: Conditionals, Beliefs, Decision. *Chance, and Time*, 153–190.
- Heckman, J. J.; Smith, J.; and Clements, N. 1997. Making The Most Out Of Programme Evaluations and Social Experiments: Accounting For Heterogeneity in Programme Impacts. *The Review of Economic Studies*, 64(4): 487–535.
- Huang, E. J.; Fang, E. X.; Hanley, D. F.; and Rosenblum, M. 2016. Inequality in treatment benefits: Can we determine if a new treatment benefits the many or the few? *Biostatistics*, 18(2): 308–324.
- Jeffrey, R. 1990. *The Logic of Decision*. McGraw-Hill series in probability and statistics. University of Chicago Press.
- Kawakami, Y.; Kuroki, M.; and Tian, J. 2024. Probabilities of Causation for Continuous and Vector Variables. *Proceedings of the 40th Conference on Uncertainty in Artificial Intelligence (UAI-2024)*.
- Kawakami, Y.; and Tian, J. 2025. Potential Outcome Rankings for Counterfactual Decision Making. arXiv:2511.10776.
- Knight, F.; and Jones, D. 2002. *Risk, Uncertainty and Profit*. Warriors (Washington, D.C.). Beard Books.
- Kropf, S.; Hommel, G.; Schmidt, U.; Brickwedel, J.; and Jepsen, M. S. 2000. Multiple Comparisons of Treatments with Stable Multivariate Tests in a Two-Stage Adaptive Design, Including a Test for Non-Inferiority. *Biometrical Journal*, 42(8): 951–965.
- Lai, V.; Chen, C.; Smith-Renner, A.; Liao, Q. V.; and Tan, C. 2023. Towards a Science of Human-AI Decision Making: An Overview of Design Space in Empirical Human-Subject Studies. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’23, 1369–1385. New York, NY, USA: Association for Computing Machinery. ISBN 9798400701924.
- Lewis, D. 1976. Probabilities of Conditionals and Conditional Probabilities. *The Philosophical Review*, 85(3): 297–315.
- Li, A.; and Pearl, J. 2024. Probabilities of Causation with Nonbinary Treatment and Effect. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(18): 20465–20472.
- Lopez, M. J.; and Gutman, R. 2017. Estimation of causal effects with multiple treatments: a review and new ideas. *Statistical Science*, 432–454.
- Lu, J.; Zhang, Y.; and Ding, P. 2019. Sharp Bounds on the Relative Treatment Effect for Ordinal Outcomes. *Biometrics*, 76(2): 664–669.
- McCaffrey, D. F.; Griffin, B. A.; Almirall, D.; Slaughter, M. E.; Ramchand, R.; and Burgette, L. F. 2013. A tutorial on propensity score estimation for multiple treatments using generalized boosted models. *Stat Med*, 32(19): 3388–3414.
- McFadden, D. 1974. Conditional Logit Analysis of Qualitative Choice Behavior. *Frontiers in Econometrics*.
- McFadden, D. 2001. Economic choices. *American economic review*, 91(3): 351–378.
- Mueller, S.; and Pearl, J. 2023. Personalized decision making—a conceptual introduction. *Journal of Causal Inference*, 11(1): 20220050.
- Page, E. B. 1963. Ordered hypotheses for multiple treatments: a significance test for linear ranks. *Journal of the American Statistical Association*, 58(301): 216–230.
- Pascal, B. 1670. *Pensées*. London: Dent.
- Pearl, J. 1999. Probabilities Of Causation: Three Counterfactual Interpretations And Their Identification. *Synthese*, 121(1): 93–149.
- Pearl, J. 2009. *Causality: Models, Reasoning and Inference*. Cambridge University Press, 2nd edition.

- Pearl, J.; and Mackenzie, D. 2018. *The book of why: the new science of cause and effect*. Basic books.
- Student. 1908. THE PROBABLE ERROR OF A MEAN. *Biometrika*, 6(1): 1–25.
- Sutton, R. S.; Barto, A. G.; et al. 1998. *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.
- Tian, J.; and Pearl, J. 2000. Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*, 28(1): 287–313.
- Train, K. 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.