
When the Oracle Misleads: Modeling the Consequences of Using Observable Rather than Potential Outcomes in Risk Assessment Instruments

Alan Mishler

Department of Statistics & Data Science
Carnegie Mellon University
Pittsburgh, PA
amishler@stat.cmu.edu

Niccolò Dalmaso

Department of Statistics & Data Science
Carnegie Mellon University
Pittsburgh, PA
ndalmass@stat.cmu.edu

1 Introduction

Machine learning is increasingly widely used to support decision making in domains as diverse as healthcare, criminal justice, and consumer finance. In particular, predictive models are often used to estimate the risk of a negative outcome such as death, recidivism, or default on a loan [1, 2, 3, 4, 5]. Scores from these Risk Assessment Instruments (RAIs) are made available to decision makers, such as doctors, judges, or loan officers, who may take them into account when deciding whether or not to admit a patient to a hospital, release a defendant on bail, or issue a loan to an applicant.

When the decision maker’s goal is to reduce the risk of the predicted outcome, they are naturally concerned with *potential outcomes*, the outcomes that would occur under each available decision. When these outcomes correspond to an intervention that actually takes place, they are observable; otherwise, they are counterfactual. (Many authors use “counterfactual outcomes” as a synonym for potential outcomes.) RAIs are typically trained on observational data, in which outcomes are affected by historical decisions, and they are typically designed to predict exclusively observable outcomes. Hence, these RAIs can only be sensibly understood as predicting the risk of an outcome *under the historical decision process that generated the data*; they are not generally appropriate for helping decision makers decide among different courses of action.

Although RAIs based on potential outcomes have been proposed in the context of medicine [6, 7] and recidivism prediction [8], RAIs designed to predict observable outcomes are in widespread use. While many of the limitations of such RAIs have been acknowledged [9, 10], and problems associated with particular RAIs have been investigated [11, 12], there does not appear to be a general mathematical model that provides insight into how and why such RAIs can lead users astray.

In this work, we aim to fill this gap, showing how RAIs based on observable outcomes can lead to *worse* outcomes, i.e., more severe departures from an optimal treatment regime, than before the RAI was introduced. This has nothing to do with the quality of prediction; it can occur even when (1) the oracle predictor is available and (2) there is no unmeasured confounding. We describe several dangerous properties of these RAIs and illustrate their suboptimality with a simple example.

2 Setup: RAIs and optimal treatment regimes

We anchor the problem in the context of medicine, but the results generalize to any domain where an estimated risk is used to drive decision making designed to mitigate that risk.

Suppose that at time $t = 0$ we have random variables drawn from a counterfactual distribution $(U, X, A, Y^0, Y^1, Y) \sim \mathbb{Q}_0$, where $U \in \mathbb{R}^{p'}$ is a set of unobserved confounders, $X \in \mathbb{R}^p$ is a set of observed covariates, $A \in \{0, 1\}$ is a binary treatment or intervention decision, and $Y \in \{0, 1\}$ is

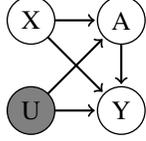


Figure 1: Causal graph at time $t = 0$, with unobserved confounders U .

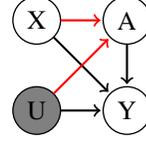


Figure 2: Causal graph at time $t = 1$, with possibly changed treatment decision process.

an outcome, with $Y = 1$ indicating an adverse event such as patient death. Y^0 and Y^1 denote the potential outcomes under treatment decisions $A = 0, 1$. Let \mathbb{P}_0 denote the marginal distribution of the observable vector (X, A, Y) at $t = 0$. We use \mathbb{E}_t and \mathbb{P}_t to denote expectations and probabilities at time t , but when these do not change over time we drop the subscript and use \mathbb{E} and \mathbb{P} .

Now suppose that iid data drawn from \mathbb{P}_0 is used to construct a predictor $s(X)$ of Y given X . For example, suppose that $s(X) = \hat{\mathbb{E}}_0[Y|X]$. This predictor is made available to decision makers in the form of an RAI, as a ‘‘risk score,’’ giving rise at time $t = 1$ to new distributions $(U, X, A, Y^0, Y^1, Y) \sim \mathbb{Q}_1$ and $(X, A, Y) \sim \mathbb{P}_1$. We make the following assumptions at all time points t :

1. $Y = AY^1 + (1 - A)Y^0$ (*Consistency*)
2. $\mathbb{P}_t[\mathbb{P}_t(0 < \pi_t(X) < 1)] = 1$ (*Positivity*)
3. $A \perp\!\!\!\perp Y^a | X, U$, for $a \in \{0, 1\}$ (*No confounders beyond X and U*)
4. $(U, X, Y^0, Y^1)_{\mathbb{Q}_t} \stackrel{d}{=} (U, X, Y^0, Y^1)_{\mathbb{Q}_{t+1}}$ (*Only the treatment and outcome change after the RAI is introduced.*)
5. $0 < \mathbb{P}_t(Y^1 < Y^0) < 1$ (*Treatment sometimes helps and sometimes hurts overall. For example, hospitalization can expose patients to dangers such as MRSA or medical errors.*)

Note that U is not observable by the researchers who construct $s(X)$, but it may include variables that are available to doctors at the time they render a treatment decision. That is, the treatment decision process may change in light of the new RAI, but the RAI does not otherwise affect patient outcomes or the distribution of covariates. Causal graphs representing the change in the decision process from time 0 to time 1 are given in Figures 1 and 2.

Given all possible treatment decision functions $\mathcal{D} = \{d : (X, U) \mapsto \{0, 1\}\}$, it is easy to show that the optimal treatment regime with respect to the expectation of Y is

$$d^{\text{opt}}(X, U) := \arg \min_{d \in \mathcal{D}} \mathbb{E}[Y^{d(X, U)}] = \mathbb{1}\{\mathbb{E}[Y^1 | X, U] < \mathbb{E}[Y^0 | X, U]\}. \quad (1)$$

where the expectations in this expression do not change over time, as a consequence of Assumption 4. Given that $s(X)$ is not designed as an estimator of d^{opt} , the questions of interest are:

1. When is $\mathbb{E}_1[Y] \leq \mathbb{E}_0[Y]$, as desired? That is, when does the RAI make things better, or at least not worse?
2. How far is $\mathbb{E}_1[Y]$ from $\mathbb{E}_1[Y^{d^{\text{opt}}}]$, the optimal outcome?

We are also interested in versions of these questions where the quantities are conditional on (X, U) . For example, we wish to know when outcomes get better or worse differentially for patients from different demographic groups, which could cause the RAI to be considered unfair.

3 RAIs can make things worse

Let $\pi_t(X) = \mathbb{P}_t(A = 1 | X, U)$ denote the treatment propensity at time t , with $\Gamma(X, U) := \pi_1(X, U) - \pi_0(X, U)$, and let $\mu^a(X, U) = \mathbb{E}[Y | X, U, A = a]$ denote the outcome regression functions, for $a \in \{0, 1\}$. μ^a also does not change over time, per assumptions 3 and 4. We have:

$$\Delta := \mathbb{E}_1[Y] - \mathbb{E}_0[Y] = \mathbb{E}\{\Gamma(X, U)(\mu^1(X, U) - \mu^0(X, U))\} \quad (2)$$

(See the derivation in the Appendix). It is easy to see that Δ can be positive, meaning that more patients die after the introduction of the RAI, and that even if it is negative, outcomes could worsen for

particular strata of (X, U) . For example, consider a subpopulation for whom $\mu^1(X, U) < \mu^0(X, U)$ and $\Gamma(X, U) < 0$. These could be patients who historically benefited from hospitalization and were hospitalized at high rates, so that $\mathbb{E}_0[Y|X]$, their likelihood of death in the training data, is small. The apparent low risk could prompt doctors to reduce the rate at which they hospitalize these patients, causing death rates to rise. Of course, if Δ is positive, then $\mathbb{E}_1[Y] - \mathbb{E}[Y^{d^{\text{opt}}}]$ will be positive as well.

For ease of exposition, we now restrict our attention to a special case of the above scenario, wherein $U = \emptyset$, so that there is no unmeasured confounding, and $s(X) = \mathbb{E}[Y|X]$, so we have access to the MSE-minimal oracle predictor. We suppose that once the RAI is introduced, decisions are made deterministically according to a threshold rule $d(x) = \mathbb{1}\{s(X) \geq \theta\}$ for some $\theta \in [0, 1]$. That is, doctors hospitalize patients iff their estimated risk is at or above θ . We illustrate with a toy example.

3.1 Toy example

We assume a single covariate $X \sim \text{Unif}(0, 1)$, representing a marker of disease severity. We let both the treatment propensity and the risk of non-treatment increase in X , with $\pi_0(X) = \mathbb{E}[Y^0|X] = X$, and we let the risk of treatment be $\mathbb{E}[Y^1|X] = (0.7 - X)^2$. This represents a situation in which treatment is beneficial on average above a certain level of X but harmful otherwise.

Figure 3 (a) shows the two conditional expectations $\mathbb{E}[Y^0|X], \mathbb{E}[Y^1|X]$. The optimal treatment rule here is $d^{\text{opt}}(X) = \mathbb{1}\{X \geq 0.22\}$, indicated by the dashed line. The rule that is actually implemented at time $t = 1$ is $d(X) = \mathbb{1}\{\mathbb{E}_0[Y|X] \geq \theta\}$ for the chosen threshold θ . Figure 3 (b) shows the mean difference in outcomes Δ from time 0 to time 1 as a function of θ . (See the Appendix for a derivation.) This difference is around $1/3$, regardless of the θ chosen, indicating that more patients die as a result of the RAI. (In this scenario, $s(X)$ is bounded in $[0, 0.30]$, so we only show thresholds in this range.)

The reason that all values of θ lead to worse outcomes is that θ corresponds to a threshold for $s(X) = \mathbb{E}_0[Y|X]$ rather than a threshold for X . In Figure 3 (c) and (d), the vertical purple block represents the optimal treatment group $\{X \geq 0.22\}$, while the overlapping horizontal green block represents the group $\{\mathbb{E}_0[Y|X] \geq \theta\}$ that is actually treated under $d(X)$. Panel (c) shows the effect of choosing $\theta = 0.22$, the optimal threshold for X : we would fail to provide treatment to the group $\{X \geq 0.67\}$, indicated in red. This happens to be the group with the highest values of $\mathbb{E}[Y^0|X]$, i.e., the worst outcomes under no treatment. Conversely, Figure (d) shows the results of selecting the cutoff such that all those who would receive treatment under $d^{\text{opt}}(X)$ also receive treatment under $d(X)$: we wrongly treat the group $\mathbb{E}[Y^1|X] > \mathbb{E}[Y^0|X]$, again indicated in red.

These same problems can obviously arise in more complex scenarios, for example when $U \neq \emptyset$, when X is high dimensional, and when the relationship between X and the outcome is complex. In particular, we identify three properties of $s(X)$ that can give rise to these and other problems.

3.2 $s(X)$ doesn't map nicely to a quantity of interest like $\mathbb{E}[Y^0|X], \mathbb{E}[Y^1|X]$, or $d^{\text{opt}}(X)$

Even though it is designed to predict outcomes under a historical treatment decision process, the RAI could usefully inform a new decision process if it bore some readily apprehensible relationship with a potential outcome-based quantity of interest. For example, if $s(X)$ were monotonic in $d^{\text{opt}}(X)$, then doctors might be able to use $s(X)$ to get closer to $d^{\text{opt}}(X)$, even without an explicit awareness of this relationship. In general, however, the relationship between $s(X)$ and any potential outcome-based quantity can be arbitrarily complex.

3.3 Expertise can make things worse

The more skilled doctors are at time $t = 0$, the worse the system can get at time $t = 1$. As an extreme example, if doctors are already behaving according to the optimal policy at time $t = 0$, then necessarily, $\mathbb{E}_1[Y] \geq \mathbb{E}_0[Y]$. Alternatively, suppose that there are two medical systems \mathbb{P}_0 and \mathbb{P}_0^* that are identical in their distribution of (X, U, Y^0, Y^1) . Imagine that they're also identical in terms of A , except that in system \mathbb{P}_0^* doctors are more skilled at identifying who needs to be hospitalized:

$$\mathbb{P}_0^*(A = 1 | d^{\text{opt}}(X, U) = 1) > \mathbb{P}_0(A = 1 | d^{\text{opt}}(X, U) = 1)$$

Then, under a threshold decision rule, we have that $\mathbb{E}_0^*[Y] < \mathbb{E}_0[Y]$ but $\mathbb{E}_1^*[Y] > \mathbb{E}_1[Y]$, so, perversely, people in system \mathbb{P}_0^* are **better off** than people in system \mathbb{P} at time 0 and **worse off** at time 1.

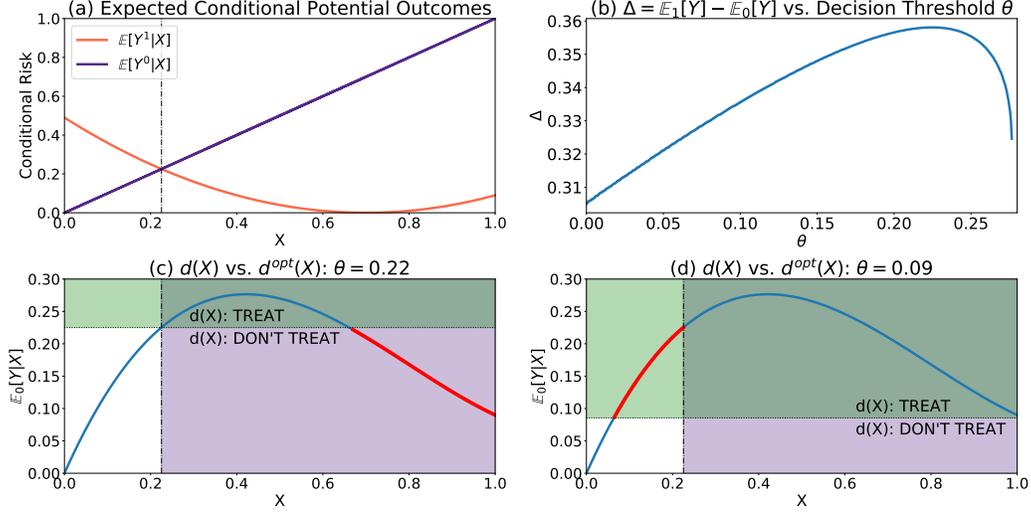


Figure 3: (a) Conditional expectations in the toy example. (b) Behavior of $\Delta = \mathbb{E}_1[Y] - \mathbb{E}_0[Y]$ as a function of the cutoff θ . (c) and (d) show groups treated at time 1 under $d(X)$ and $d^{opt}(X)$, for two possible values of θ . The optimal treatment group is $\{X > 0.22\}$, in purple. The group treated under $d(X)$ is $\{\mathbb{E}_0[Y|X] \geq \theta\}$, in green. Red lines indicate groups that are harmed under $d(X)$ as a result of not receiving or receiving treatment, respectively.

Time t	Treatment decision	$\mathbb{E}[Y_t X]$	$\mathbb{E}_t[Y X]$ relative to θ
0	Treat with probability $\pi_0(X)$	$\mathbb{E}_0[Y X]$	$> \theta$
1	Treat all	$\mathbb{E}[Y^1 X]$	$< \theta$
2	Treat none	$\mathbb{E}[Y^0 X]$	$> \theta$
3	Treat all	$\mathbb{E}[Y^1 X]$	$< \theta$
4	Treat none	$\mathbb{E}[Y^0 X]$	$> \theta$
...			

Table 1: Treatment decisions and mean outcomes in stratum X over time, under the deterministic decision rule that treats patients at time t iff $\mathbb{E}_{t-1}[Y|X] > \theta$ for some threshold θ , and assuming that $\mathbb{E}_0[Y^1|X] < \theta$, $\mathbb{E}_0[Y^0|X] > \theta$ and $\mathbb{E}[Y_{(0)}|X] > \theta$.

3.4 The procedure is unstable under iteration

Imagine that we iterate the process of gathering data from the system, developing a predictor, and implementing the threshold-based decision rule above. This seems like a plausible occurrence, in that as RAIs get integrated into more and more systems, necessarily any future data gathered from those systems will reflect the influence of those tools.

For time points $t = 1, 2, \dots$, we have

$$\mathbb{E}_t[Y|X] = \mathbb{1}\{\mathbb{E}_{t-1}[Y|X] > \theta\} \mathbb{E}[Y^1|X] + (1 - \mathbb{1}\{\mathbb{E}_{t-1}[Y|X] > \theta\}) \mathbb{E}[Y^0|X]$$

Suppose we have some X for which $\mathbb{E}_0[Y^1|X] < \theta$, $\mathbb{E}_0[Y^0|X] > \theta$ and $\mathbb{E}[Y_{(0)}|X] > \theta$. Then we'll have the situation depicted in Table 1, in which the treatment decision for this stratum just alternates at different time points. Ideally, as more and more data is gathered from a system, a decision procedure gets closer and closer to optimal. In this scenario, however, the treatment decision is the optimal one only at odd time points, while at even time points it's precisely the opposite.

4 Conclusion

Decision makers choosing among different courses of action are naturally interested in the risk associated with each option. RAIs are in widespread use in many domains, but they are typically designed to predict outcomes under the historical decision process that generated the training data, rather than predicting potential outcomes under the available courses of action. This makes them

generally unsuitable for informing future treatment or intervention decisions that are designed to reduce risk. Although previous work has proposed using potential outcome-based predictors in certain contexts, there has been little formal modeling of the consequences of current practice. Here, we show how RAIs based on observable rather than potential outcomes can plausibly lead to worse outcomes overall or for specific demographic groups than before their introduction, making them potentially both dangerous and unfair.

References

- [1] Konstantina Kourou, Themis P. Exarchos, Konstantinos P. Exarchos, Michalis V. Karamouzis, and Dimitrios I. Fotiadis. Machine learning applications in cancer prognosis and prediction. *Computational and Structural Biotechnology Journal*, 13:8–17, 2015.
- [2] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 1721–1730, New York, NY, USA, 2015. ACM.
- [3] Andres Colubri, Tom Silver, Terrence Fradet, Kalliroi Retzepi, Ben Fry, and Pardis Sabeti. Transforming Clinical Data into Actionable Prognosis Models: Machine-Learning Framework and Field-Deployable App to Predict Outcome of Ebola Patients. *PLOS Neglected Tropical Diseases*, 10(3):e0004549, March 2016.
- [4] Tim Brennan, William Dieterich, and Beate Ehret. Evaluating the Predictive Validity of the Compas Risk and Needs Assessment System. *Criminal Justice and Behavior*, 36(1):21–40, January 2009.
- [5] Amir E. Khandani, Adlar J. Kim, and Andrew W. Lo. Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11):2767–2787, November 2010.
- [6] Peter Schulam and Suchi Saria. Reliable Decision Support using Counterfactual Models. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 1697–1708. Curran Associates, Inc., 2017.
- [7] Uri Shalit, Fredrik D Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, page 10, 2017.
- [8] Alan Mishler. Modeling Risk and Achieving Algorithmic Fairness Using Potential Outcomes. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society - AIES '19*, pages 555–556, Honolulu, HI, USA, 2019. ACM Press.
- [9] Jonathan H. Chen and Steven M. Asch. Machine Learning and Prediction in Medicine — Beyond the Peak of Inflated Expectations. *The New England journal of medicine*, 376(26):2507–2509, June 2017.
- [10] Michael Veale, Max Van Kleek, and Reuben Binns. Fairness and Accountability Design Needs for Algorithmic Support in High-Stakes Public Sector Decision-Making. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–14, Montreal QC, Canada, 2018. ACM Press.
- [11] Andrey A. Povyakalo, Eugenio Alberdi, Lorenzo Strigini, and Peter Ayton. How to Discriminate between Computer-Aided and Computer-Hindered Decisions: A Case Study in Mammography. *Medical Decision Making*, 33(1):98–107, January 2013.
- [12] Kristian Lum and William Isaac. To predict and serve? *Significance*, 13(5):14–19, October 2016.

A Derivations

A.1 Equation (2): difference in mean outcome from time 0 to time 1

Recall that we define

$$\begin{aligned}\pi_t(X, U) &= \mathbb{E}_t[A|X, U] \\ \mu^a(X, U) &= \mathbb{E}[Y^a|X, U]\end{aligned}$$

for $t = 1, 2, \dots$ and $a \in \{0, 1\}$. Recall also that, per assumption 4, the distribution of the covariates (X, U) doesn't change over time, so functions of (X, U) don't change either.

For any time point t , we have

$$\begin{aligned}\mathbb{E}_t[Y] &= \mathbb{E}_t\{\mathbb{E}_t[AY^1 + (1 - A)Y^0|X, U]\} \\ &= \mathbb{E}_t\{\mathbb{E}_t[A|X, U]\mathbb{E}_t[Y^1|X, U] + (1 - \mathbb{E}_t[A|X, U])\mathbb{E}_t[Y^0|X, U]\} \\ &= \mathbb{E}\{\mathbb{E}_t[A|X, U]\mathbb{E}[Y^1|X, U] + (1 - \mathbb{E}_t[A|X, U])\mathbb{E}[Y^0|X, U]\} \\ &= \mathbb{E}\{\pi_t(X, U)\mu^1(X, U) + (1 - \pi_t(X, U))\mu^0(X, U)\}\end{aligned}$$

where the second equality follows because $A \perp\!\!\!\perp Y^a|X, U$ and the third equality follows from assumption 4. With $\Gamma(X, U) := \pi_1(X, U) - \pi_0(X, U)$, we have:

$$\begin{aligned}\mathbb{E}_1[Y] - \mathbb{E}_0[Y] &= \mathbb{E}\{\pi_1(X, U) - \pi_0(X, U)(\mu^1(X, U) - \mu^0(X, U))\} \\ &= \mathbb{E}\{\Gamma(X, U)(\mu^1(X, U) - \mu^0(X, U))\}\end{aligned}$$

A.2 Calculating Δ in the toy example

We have $U = \emptyset$, $X \sim \text{Unif}(0, 1)$, $\pi_0(X) = \mathbb{E}[Y^0|X] = X$, $\mathbb{E}[Y^1|X] = (0.7 - X)^2$, and $\pi_1(X) = \mathbb{1}\{\mathbb{E}_0[Y|X] \geq \theta\}$ for some chosen θ . Plugging these into the above yields

$$\begin{aligned}\Delta(\theta) = \mathbb{E}_1[Y] - \mathbb{E}_0[Y] &= \int (\mathbb{1}\{\mathbb{E}_0[Y|X] \geq \theta\} - X)((0.7 - X)^2 - X)d\mathbb{P}(X) \\ &= \int (\mathbb{1}\{\mathbb{E}_0[AY^1 + (1 - A)Y^0|X] \geq \theta\} - X)((0.7 - X)^2 - X)d\mathbb{P}(X) \\ &= \int (\mathbb{1}\{\pi_0(X)\mu^1(X) + (1 - \pi_0(X))\mu^0(X) \geq \theta\} - X)((0.7 - X)^2 - X)d\mathbb{P}(X) \\ &= \int (\mathbb{1}\{X(0.7 - X)^2 + (1 - X)X \geq \theta\} - X)((0.7 - X)^2 - X)d\mathbb{P}(X)\end{aligned}$$

where the second equality follows from the consistency assumption, and the third equality follows from the no unmeasured confounding assumption and assumption 4. This yields the curve in Figure 3(b).