

Unit Selection with Nonbinary Treatment and Effect

Ang Li¹, Judea Pearl²

¹Department of Computer Science, Florida State University

²Cognitive Systems Laboratory, Department of Computer Science, University of California, Los Angeles
angli@cs.fsu.edu, judea@cs.ucla.edu

Abstract

The unit selection problem aims to identify a set of individuals who are most likely to exhibit a desired mode of behavior or to evaluate the percentage of such individuals in a given population, for example, selecting individuals who would respond one way if encouraged and a different way if not encouraged (Li and Pearl 2019). Using a combination of experimental and observational data, Li and Pearl solved the binary unit selection problem (binary treatment and effect) by deriving tight bounds on the “benefit function,” which is the payoff/cost associated with selecting an individual with given characteristics. This paper extends the benefit function to the general form such that the treatment and effect are not restricted to binary. We then propose an algorithm to test the identifiability of the nonbinary benefit function and an algorithm to compute the bounds of the nonbinary benefit function using experimental and observational data.

Introduction

Several areas of industry, marketing, and health science face the unit selection dilemma. For example, in customer relationship management (Berson, Smith, and Thearling 1999; Lejeune 2001; Hung, Yen, and Wang 2006; Tsai and Lu 2009), it is useful to determine the customers who are going to leave but might reconsider if encouraged to stay. Due to the high expense of such initiatives, management is forced to limit inducement to customers who are most likely to exhibit the behavior of interest. As another example, companies are interested in identifying users who would click on an advertisement if and only if it is highlighted in online advertising (Yan et al. 2009; Bottou et al. 2013; Li et al. 2014; Sun et al. 2015). The challenge in identifying these users stems from the fact that the desired response pattern is not observed directly but rather is defined counterfactually in terms of what the individual would do under hypothetical unrealized conditions. For example, when we observe that a user has clicked on a highlighted advertisement, we do not know whether they would click on that same advertisement if it were not highlighted.

The binary benefit function for the unit selection problem was defined by Li and Pearl (Li and Pearl 2019) (we will call this Li-Pearl’s model), and it properly captures the nature of

the desired behavior. Using a combination of experimental and observational data, Li and Pearl derived tight bounds of the benefit function. The only assumption is that the treatment has no effect on the population-specific characteristics. Inspired by the idea of Mueller, Li, and Pearl (Mueller, Li, and Pearl 2022) and Dawid et al. (Dawid, Musio, and Murta 2017) that the bounds of probabilities of causation could be narrowed using covariates information, Li and Pearl (Li and Pearl 2022c) narrowed the bounds of the benefit function using covariates information and their causal structure. However, the abovementioned studies are based on binary treatment and effect. Recently, researchers have shown interest in developing bounds for probabilities of causation with nonbinary treatment and effect. Zhang, Tian, and Bareinboim (Zhang, Tian, and Bareinboim 2022), as well as Li and Pearl (Li and Pearl 2022a), proposed nonlinear programming-based solutions to compute the bounds of nonbinary probabilities of causation numerically. Li and Pearl (Li and Pearl 2022b) provided the theoretical bounds of nonbinary probabilities of causation. The benefit function is a linear combination of probabilities of causation; therefore, in this paper, we focus on discovering the bounds of any benefit function without restricting them to binary treatment and effect.

Consider the following motivating scenario: a clinical study is conducted to test the effectiveness of a vaccine. The treatments include vaccinated and unvaccinated. The outcomes include uninfected, asymptomatic infected, and infected in a severe condition. The benefited individuals include the following: the individual who would be infected in a severe condition if unvaccinated and would be asymptomatic infected if vaccinated, the individual who would be infected in a severe condition if unvaccinated and would be uninfected if vaccinated, and the individual who would be asymptomatic infected if unvaccinated and would be uninfected if vaccinated. The harmed individuals include the following: the individual who would be asymptomatic infected if unvaccinated and would be infected in a severe condition if vaccinated, the individual who would be uninfected if unvaccinated and would be infected in a severe condition if vaccinated, and the individual who would be uninfected if unvaccinated and would be asymptomatic infected if vaccinated. All others are unaffected individuals. The researcher performing the clinical study has collected both experimental and observational data. The researcher then wants to know the expected difference

between benefited and harmed individuals to emphasize the effectiveness of the vaccine.

We cannot apply Li-Pearl’s model because we have two treatments and three outcomes. In this paper, we extend Li-Pearl’s benefit function to general form without restricting them to binary treatment and effect. We will provide an algorithm to test the identifiability of the nonbinary benefit function and an algorithm to compute the bounds of the non-binary benefit function using experimental and observational data.

Preliminaries

In this section, we review Li and Pearl’s binary benefit function of the unit selection problem (Li and Pearl 2019), and the theoretical bounds of the probabilities of causation recently proposed by Li and Pearl (Li and Pearl 2022b).

In this paper, we use the language of counterfactuals in structural model semantics, as given in (Galles and Pearl 1998; Halpern 2000). we use $Y_x = y$ to denote the counterfactual sentence “Variable Y would have the value y , had X been x ”. For simplicity purposes, in the rest of the paper, we use y_x to denote the event $Y_x = y$, $y_{x'}$ to denote the event $Y_{x'} = y$, y'_x to denote the event $Y_x = y'$, and $y'_{x'}$ to denote the event $Y_{x'} = y'$. We assume that experimental data will be summarized in the form of the causal effects such as $P(y_x)$ and observational data will be summarized in the form of the joint probability function such as $P(x, y)$. If not specified, the variable X stands for treatment, and the variable Y stands for effect.

Individual behavior was classified into four response types: labeled complier, always-taker, never-taker, and defier. Suppose the benefit of selecting one individual in each category are $\beta, \gamma, \theta, \delta$ respectively (i.e., the benefit vector is $(\beta, \gamma, \theta, \delta)$). Li and Pearl defined the objective function of the unit selection problem as the average benefit gained per individual. Suppose x and x' are binary treatments, y and y' are binary outcomes, and c are population-specific characteristics, the objective function (i.e., benefit function) is following (If the goal is to evaluate the average benefit gained per individual for a specific population c , argmax_c can be dropped.):

$$\text{argmax}_c \beta P(y_x, y'_{x'} | c) + \gamma P(y_x, y_{x'} | c) + \theta P(y'_{x'}, y'_{x'} | c) + \delta P(y'_{x'}, y_{x'} | c).$$

Using a combination of experimental and observational data, Li and Pearl established the most general tight bounds on this benefit function (which we refer to as Li-Pearl’s Theorem in the rest of the paper). The only constraint is that the population-specific characteristics are not a descendant of the treatment.

Li and Pearl (Li and Pearl 2022b) provided eight theorems to compute bounds for any type of probability of causation with nonbinary treatment and effect. Suppose variable X has m values and Y has n values, the following (Equations 1 to 8) probabilities of causation are bounded. Besides, if the probabilities of causation are conditioned on a population-specific variable c that is not affected by X , then all the theorems still hold (we provided the extended theorems from Li and Pearl in the appendix).

$$P(y_{i_{x_j}}, y_{i_i}), \quad (1)$$

$$s.t., 1 \leq i \leq n, 1 \leq j \leq m,$$

$$P(y_{i_{x_j}}, y_{i_k}), \quad (2)$$

$$s.t., 1 \leq i, k \leq n, 1 \leq j \leq m, i \neq k$$

$$P(y_{i_{x_j}}, x_k), \quad (3)$$

$$s.t., 1 \leq i \leq n, 1 \leq j, k \leq m, j \neq k$$

$$P(y_{i_{x_j}}, y_{i_k}, x_p), \quad (4)$$

$$s.t., 1 \leq i, k \leq n, 1 \leq j, p \leq m, j \neq p$$

$$P(y_{i_1 x_{j_1}}, \dots, y_{i_k x_{j_k}}), \quad (5)$$

$$s.t., 1 \leq i_1, \dots, i_k \leq n, 1 \leq j_1, \dots, j_k \leq m, j_1 \neq \dots \neq j_k$$

$$P(y_{i_1 x_{j_1}}, \dots, y_{i_k x_{j_k}}, x_p), \quad (6)$$

$$s.t., 1 \leq i_1, \dots, i_k \leq n, 1 \leq j_1, \dots, j_k, p \leq m,$$

$$j_1 \neq \dots \neq j_k \neq p$$

$$P(y_{i_1 x_{j_1}}, \dots, y_{i_k x_{j_k}}, y_q), \quad (7)$$

$$s.t., 1 \leq i_1, \dots, i_k, q \leq n, 1 \leq j_1, \dots, j_k \leq m,$$

$$j_1 \neq \dots \neq j_k$$

$$P(y_{i_1 x_{j_1}}, \dots, y_{i_k x_{j_k}}, x_p, y_q), \quad (8)$$

$$s.t., 1 \leq i_1, \dots, i_k, q \leq n, 1 \leq j_1, \dots, j_k, p \leq m,$$

$$j_1 \neq \dots \neq j_k \neq p.$$

The benefit function is a linear combination of the probabilities of causation; therefore, we define the general benefit function for the unit selection problem based on Li and Pearl’s results.

Counterfactual Formulation of the Unit Selection Problem

Based on Li and Pearl (Li and Pearl 2019), the objective is to find a set of characteristics c that maximizes the benefit associated with the resulting mixture of different response types of individuals. Let X denotes the treatment with m values and Y denotes the effect with n values. Therefore, we have n^m different response types (i.e., one response type means assigning one effect to each of the treatments). Suppose the benefit of selecting an individual are $(\alpha_1, \dots, \alpha_{n^m})$ (we call $(\alpha_1, \dots, \alpha_{n^m})$ as benefit vector). Our objective, then, should be to find c that maximizes the following expression (If the goal is to evaluate the average benefit gained per individual for a specific population c , argmax_c can be dropped):

$$\begin{aligned} \text{argmax}_c \quad & \alpha_1 P(y_{1x_1}, y_{1x_2}, \dots, y_{1x_m} | c) + \\ & \alpha_2 P(y_{1x_1}, y_{1x_2}, \dots, y_{2x_m} | c) + \dots \\ & \alpha_n P(y_{1x_1}, y_{1x_2}, \dots, y_{nx_m} | c) + \dots \\ & \alpha_{n^{m-1}+1} P(y_{2x_1}, y_{1x_2}, \dots, y_{1x_m} | c) + \dots \\ & \alpha_{n^m} P(y_{nx_1}, y_{nx_2}, \dots, y_{nx_m} | c). \end{aligned}$$

Note that c can be interpreted as the population-specific variable, the only assumption is that the treatment X has no effect on the population-specific variable. Recall from Li and Pearl’s paper (Li and Pearl 2019), the benefit vector is provided by the decision-maker who uses the model.

In the next section, we will provide an algorithm that could check whether a given benefit function with the benefit vector is identifiable with purely experimental data (i.e., we can find the exact value of the benefit function rather than bounds). If it is not identifiable, we will then provide an algorithm that computes the bounds of the benefit function given the benefit vector using experimental and observational data.

Main Results

Identifiability of Benefit Function

Recall that in the binary case, the conditions of identifiable are gain equality (i.e., $\beta + \delta = \gamma + \theta$) or monotonicity (i.e., $P(y_{x'}, y_x) = 0$) (Li and Pearl 2019). Here, it is complicated in nonbinary cases; therefore, we provide an algorithm to test whether a given benefit function with the benefit vector is identifiable with purely experimental data.

Theorem 1. *Suppose variables X has m values x_1, \dots, x_m and Y has n values y_1, \dots, y_n . Then the benefit function $f(c)$ is identifiable if Algorithm 1 returns (*True, res*), and *res* is the value of the benefit function.*

$$\begin{aligned} f(c) = & \alpha_1 P(y_{1x_1}, y_{1x_2}, \dots, y_{1x_m} | c) + \\ & \alpha_2 P(y_{1x_1}, y_{1x_2}, \dots, y_{2x_m} | c) + \dots \\ & \alpha_n P(y_{1x_1}, y_{1x_2}, \dots, y_{nx_m} | c) + \dots \\ & \alpha_{n^{m-1}+1} P(y_{2x_1}, y_{1x_2}, \dots, y_{1x_m} | c) + \dots \\ & \alpha_{n^m} P(y_{nx_1}, y_{nx_2}, \dots, y_{nx_m} | c). \end{aligned}$$

The correctness of the algorithm simply follow the fact that $\sum_{n^{m-1} \text{ terms}} P(\dots, y_{ix_j}, \dots | c) = P(y_{ix_j} | c)$. Therefore, if there exist such n^{m-1} terms in the benefit function, then we can obtain an equivalent benefit function by replacing one of the n^{m-1} terms with experimental data $P(y_{ix_j} | c)$. We exhausted all equivalent benefit functions to check if we could replace all the counterfactual terms with experimental data (i.e., identifiable).

For example, consider $m = n = 2$ and the benefit function:

$$\begin{aligned} & 7P(y_{1x_1}, y_{1x_2} | c) + 2P(y_{1x_1}, y_{2x_2} | c) + \\ & 4P(y_{2x_1}, y_{1x_2} | c) - P(y_{2x_1}, y_{2x_2} | c) \\ = & 7P(y_{1x_1} | c) - 5P(y_{1x_1}, y_{2x_2} | c) + \\ & 4P(y_{2x_1}, y_{1x_2} | c) - P(y_{2x_1}, y_{2x_2} | c) \\ = & 7P(y_{1x_1} | c) - 5P(y_{1x_1}, y_{2x_2} | c) + \\ & 4P(y_{2x_1} | c) - 5P(y_{2x_1}, y_{2x_2} | c) \\ = & 7P(y_{1x_1} | c) - 5P(y_{2x_2}) + 4P(y_{2x_1} | c). \end{aligned}$$

Bounds of Benefit Function

If Algorithm 1 returns false, we then need to compute the bounds of the benefit function using experimental and observational data. We first obtain the bounds of the probabilities of causation, $P(y_{1x_1}, y_{1x_2}, \dots, y_{1x_m} | c), \dots, P(y_{1x_1}, y_{1x_2}, \dots, y_{nx_m} | c), \dots, P(y_{2x_1}, y_{1x_2}, \dots, y_{1x_m} | c), \dots, P(y_{nx_1}, y_{nx_2}, \dots, y_{nx_m} | c)$, by Li and Pearl's theorems (Li and Pearl 2022b). We then have the following theorem.

Algorithm 1: Check identifiability of the benefit function

Input: a , the benefit function, where $a[i]$ is a $m + 1$ tuple that stands for i th term in the benefit function. If the i th term is $\alpha_i P(y_{i_1 x_1}, y_{i_2 x_2}, \dots, y_{i_m x_m} | c)$, then $a[i] = (\alpha_i, i_1, i_2, \dots, i_m)$.

$d[1, \dots, m][1, \dots, n]$, the experimental data, where $d[i][j] = P(y_{j x_i} | c)$.

e , the adjusted value of the benefit function, 0 for initial call.

The initial call of the algorithm is $IBF(a[1, \dots, n^m], d[1, \dots, m][1, \dots, n], 0)$, where $a[1, \dots, n^m]$ corresponding to the original benefit function.

All lists in this algorithm start with index 1.

Output: (identifiable, value), a tuple, where identifiable = True if the given benefit function is identifiable and value is the value of the benefit function.

Function $IBF(a, d, e)$:

```

1: mark = True;
2: l = length(a);
3: // Base case, if all benefit vector equals to (0, ..., 0), then
   the input benefit function is identifiable, and its value
   equals to the adjusted value.
4: for i = 1 to l do
5:   if a[i][1] ≠ 0 then
6:     mark = False;
7:     break;
8:   end if
9: end for
10: if mark == True then
11:   Return(True, e);
12: end if
13: // Build an equivalent benefit function by the fact that
   if ∃2 ≤ r ≤ (m + 1) s.t., a[j1][r] = ... = a[jnm-1][r],
   then the sum of these nm-1 terms without coefficients
   is equal to P(ya[j1][r]xr), we then recursively solve the
   equivalent benefit function.
14: for every nm-1 pair in a, (a[j1], ..., a[jnm-1]), s.t., there
   ∃2 ≤ r ≤ (m + 1) s.t., a[j1][r] = ... = a[jnm-1][r] do
15:   for k = 1 to nm-1 do
16:     na = a;
17:     nc = e + a[jk][1] * d[r - 1][a[j1][r]];
18:     for t = 1 to nm-1 do
19:       if t ≠ k then
20:         na[jt][1] = na[jt][1] - na[jk][1];
21:       end if
22:     end for
23:     Remove(na[jk]);
24:     res = IBF(na, d, nc);
25:     if res[0] == True then
26:       Return res
27:     end if
28:   end for
29: end for
30: Return(False, e);

```

	Vaccinated	Unvaccinated
Uninfected	52	329
Asymptomatic	512	58
Severe Condition	36	213
Overall	600	600

Table 1: Experimental data of the clinical study. Here, 600 people were forced to take the vaccine and 600 people were forced to take no vaccine.

	Vaccinated	Unvaccinated
Uninfected	14	121
Asymptomatic	933	65
Severe Condition	6	61
Overall	953	247

Table 2: Observational data of the clinical study. Here, 1200 people were free to access the vaccine. 953 people chose to take the vaccine and 247 people chose to take no vaccine.

Theorem 2. Suppose variables X has m values x_1, \dots, x_m and Y has n values y_1, \dots, y_n . Then the bounds of the benefit function $f(c)$ is obtained by Algorithm 2.

$$\begin{aligned}
f(c) = & \alpha_1 P(y_{1x_1}, y_{1x_2}, \dots, y_{1x_m} | c) + \\
& \alpha_2 P(y_{2x_1}, y_{2x_2}, \dots, y_{2x_m} | c) + \dots \\
& \alpha_n P(y_{nx_1}, y_{nx_2}, \dots, y_{nx_m} | c) + \dots \\
& \alpha_{n^{m-1}+1} P(y_{2x_1}, y_{1x_2}, \dots, y_{1x_m} | c) + \dots \\
& \alpha_{n^m} P(y_{nx_1}, y_{nx_2}, \dots, y_{nx_m} | c).
\end{aligned}$$

Again, the correctness of the algorithm simply follows the fact that $\sum_{n^{m-1} \text{ terms}} P(\dots, y_{ix_j}, \dots | c) = P(y_{ix_j} | c)$. We exhausted all equivalent benefit functions and took the maximum of all the lower bounds and the minimum of all the upper bounds of equivalent benefit functions.

Example: Effectiveness of a Vaccine

Recall the motivating example at the beginning; a clinical study is conducted to test the effectiveness of a vaccine. The treatments include vaccinated and unvaccinated. The outcomes include uninfected, asymptomatic infected, and infected in a severe condition. The researcher of the clinical study has collected both experimental and observational data.

Task 1

The researcher wants to know the expected difference between benefited and harmed individuals to emphasize the effectiveness of the vaccine.

Let X denotes vaccination with x_1 being vaccinated and x_2 being unvaccinated and Y denotes outcome, where y_1 denotes uninfected, y_2 denotes asymptomatic infected, and y_3 denotes infected in a severe condition. The experimental and observational data of the clinical study are summarized in Tables 1 and 2.

Based on the clinical study, the researcher of the vaccine claimed that the vaccine is effective in controlling the se-

Algorithm 2: Compute the bounds of the benefit function

Input: a , the benefit function, where $a[i]$ is a $m + 1$ tuple that stands for i th term in the benefit function. If the i th term is $\alpha_i P(y_{i_1 x_1}, y_{i_2 x_2}, \dots, y_{i_m x_m} | c)$, then $a[i] = (\alpha_i, i_1, i_2, \dots, i_m)$.

lb , the lower bound of all possible terms obtained from Li-Pearl's theorems, where $lb[(i_1, i_2, \dots, i_m)]$ is the lower bound of $P(y_{i_1 x_1}, y_{i_2 x_2}, \dots, y_{i_m x_m} | c)$.

ub , the upper bound of all possible terms obtained from Li-Pearl's theorems, where $ub[(i_1, i_2, \dots, i_m)]$ is the upper bound of $P(y_{i_1 x_1}, y_{i_2 x_2}, \dots, y_{i_m x_m} | c)$.

e , the adjusted value of the benefit function, 0 for initial call.

The initial call of the algorithm is $BBF(a[1, \dots, n^m], lb, ub, 0)$, where $a[1, \dots, n^m]$ corresponding to the original benefit function.

All lists in this algorithm start with index 1.

Output: (lo, up), lower and upper bound of the benefit function.

Function $BBF(a, lb, ub, e)$:

```

1:  $l = \text{length}(a)$ ;
2: // Base case, compute the bounds.
3:  $up = e, lo = e$ ;
4: for  $i = 1$  to  $l$  do
5:   if  $a[i][1] < 0$  then
6:      $lo = lo + a[i][1] * ub[(a[i][2], \dots, a[i][m + 1])]$ ;
7:      $up = up + a[i][1] * lb[(a[i][2], \dots, a[i][m + 1])]$ ;
8:   else
9:      $lo = lo + a[i][1] * lb[(a[i][2], \dots, a[i][m + 1])]$ ;
10:     $up = up + a[i][1] * ub[(a[i][2], \dots, a[i][m + 1])]$ ;
11:   end if
12: end for
13: // Build an equivalent benefit function by the fact that
    if  $\exists 2 \leq r \leq (m + 1) \text{ s.t., } a[j_1][r] = \dots = a[j_{n^{m-1}}][r]$ ,
    then the sum of these  $n^{m-1}$  terms without coefficients
    is equal to  $P(y_{a[j_1][r] x_r})$ , we then recursively solve the
    equivalent benefit function.
14: for every  $n^{m-1}$  pair in  $a$ ,  $(a[j_1], \dots, a[j_{n^{m-1}}])$ , s.t., there
     $\exists 2 \leq r \leq (m + 1) \text{ s.t., } a[j_1][r] = \dots = a[j_{n^{m-1}}][r]$  do
15:   for  $k = 1$  to  $n^{m-1}$  do
16:      $na = a$ ;
17:      $nc = e + a[j_k][1] * d[r - 1][a[j_1][r]]$ ;
18:     for  $t = 1$  to  $n^{m-1}$  do
19:       if  $t \neq k$  then
20:          $na[j_t][1] = na[j_t][1] - na[j_k][1]$ ;
21:       end if
22:     end for
23:      $Remove(na[j_k])$ ;
24:      $res = BBF(na, lb, ub, nc)$ ;
25:      $lo = \max\{lo, res[0]\}$ 
26:      $up = \min\{up, res[1]\}$ 
27:   end for
28: end for
29: Return  $(lo, up)$ ;

```

vere condition, and the number of severe condition patients dropped from 213 to only 36.

Now consider the expected difference between benefited and harmed individuals. Recall the benefited individuals include the individual who would be infected in a severe condition if unvaccinated and would be asymptomatic infected if vaccinated, the individual who would be infected in a severe condition if unvaccinated and would be uninfected if vaccinated, and the individual who would be asymptomatic infected if unvaccinated and would be uninfected if vaccinated. The harmed individuals include the individual who would be asymptomatic infected if unvaccinated and would be infected in a severe condition if vaccinated, the individual who would be uninfected if unvaccinated and would be infected in a severe condition if vaccinated, and the individual who would be uninfected if unvaccinated and would be asymptomatic infected if vaccinated. All others are unaffected individuals. In order to maximize the difference between benefited and harmed individuals; therefore, we assign 1 to benefited individuals, assign -1 to harmed individuals, and 0 to all others in the benefit vector. The objective function (i.e., benefit function) is then

$$\begin{aligned} f(c) = & 0P(y_{1x_1}, y_{1x_2}|c) + P(y_{1x_1}, y_{2x_2}|c) + \\ & P(y_{1x_1}, y_{3x_2}|c) - P(y_{2x_1}, y_{1x_2}|c) + \\ & 0P(y_{2x_1}, y_{2x_2}|c) + P(y_{2x_1}, y_{3x_2}|c) - \\ & -P(y_{3x_1}, y_{1x_2}|c) - P(y_{3x_1}, y_{2x_2}|c) + \\ & 0P(y_{3x_1}, y_{3x_2}|c). \end{aligned}$$

The experimental data in Table 1 provide the following estimates:

$$\begin{aligned} P(y_{1x_1}|c) &= 52/600 = 0.087 \\ P(y_{2x_1}|c) &= 512/600 = 0.853 \\ P(y_{3x_1}|c) &= 36/600 = 0.060 \\ P(y_{1x_2}|c) &= 329/600 = 0.548 \\ P(y_{2x_2}|c) &= 58/600 = 0.097 \\ P(y_{3x_2}|c) &= 213/600 = 0.355 \end{aligned}$$

The observational data Table 2 provide the following estimates:

$$\begin{aligned} P(x_1, y_1|c) &= 14/1200 = 0.012 \\ P(x_1, y_2|c) &= 933/1200 = 0.778 \\ P(x_1, y_3|c) &= 6/1200 = 0.005 \\ P(x_2, y_1|c) &= 121/1200 = 0.101 \\ P(x_2, y_2|c) &= 65/1200 = 0.054 \\ P(x_2, y_3|c) &= 61/1200 = 0.051 \end{aligned}$$

We plug the estimates and the benefit function into Theorem 1, the Algorithm 1 returns false (i.e., not identifiable by experimental data). We then plug the estimates and the benefit function into Theorem 2 to obtain the bounds

$$-0.228 \leq f(c) \leq -0.107$$

Thus, the expected difference between benefited and harmed individuals is at most -0.107 per individual. We can conclude that the vaccine is ineffective for the virus.

Task 2

The researcher of the clinic study claimed that the individual who would be infected in a severe condition if unvaccinated and would be uninfected if vaccinated and the individual who would be uninfected if unvaccinated and would be infected in a severe condition if vaccinated should be twice important than other individuals. Based on the clinical study, the number of severe condition patients dropped from 213 to only 36; therefore, the vaccine should be effective for the virus.

Now consider the expected difference between benefited and harmed individuals. The benefit vector should be the same except assigning 2 to the individual who would be infected in a severe condition if unvaccinated and would be uninfected if vaccinated and assigning -2 to the individual who would be uninfected if unvaccinated and would be infected in a severe condition if vaccinated.

The objective function (i.e., benefit function) is then

$$\begin{aligned} f(c) = & 0P(y_{1x_1}, y_{1x_2}|c) + P(y_{1x_1}, y_{2x_2}|c) + \\ & 2P(y_{1x_1}, y_{3x_2}|c) - P(y_{2x_1}, y_{1x_2}|c) + \\ & 0P(y_{2x_1}, y_{2x_2}|c) + P(y_{2x_1}, y_{3x_2}|c) - \\ & -2P(y_{3x_1}, y_{1x_2}|c) - P(y_{3x_1}, y_{2x_2}|c) + \\ & 0P(y_{3x_1}, y_{3x_2}|c). \end{aligned}$$

We plug the estimates and the benefit function into Theorem 1, the Algorithm 1 returns true (i.e., identifiable by experimental data) with value -0.167 . The benefit function can be simplified as follow:

$$\begin{aligned} f(c) = & 2P(y_{1x_1}|c) + P(y_{2x_1}|c) - \\ & -2P(y_{1x_2}|c) - P(y_{2x_2}|c) \\ = & -0.167. \end{aligned}$$

Thus, the expected difference between benefited and harmed individuals is exactly -0.167 per individual. We can conclude that the vaccine is still ineffective for the virus.

Simulated Results

In this section, we show the quality of the bounds of the benefit function obtained by Theorem 2 using four common benefit vectors.

First, we set $m = 2$ (i.e., X has two values) and $n = 3$ (i.e., Y has three values). We set the benefit vector to one of the most common ones, $(0, 1, 1, -1, 0, 1, -1, -1, 0)$, which is to evaluate the expected difference between benefited and harmed individuals. We randomly generated 1000 populations where each population consists of different fractions of nine response types of individuals. For each population, we then generated sample distributions (observational data and experimental data) compatible with the fractions of response types (see the appendix for the generating algorithm). The advantage of this generating process is that we have the real benefit value (because we know the fractions of the response types) for comparison. Each sample population represents a different instantiate of the population-specific characteristics C in the model. The generating algorithm ensures that the experimental data and observational data satisfy the general relation (i.e., $P(x, y|c) \leq P(y_x|c) \leq P(x, y|c) + 1 - P(x|c)$).




Figure 1: Bounds of the benefit function for 100 sample populations out of 1000 with the benefit vector $(0, 1, 1, -1, 0, 1, -1, -1, 0)$.

For a sample population i , let $[a_i, b_i]$ be the bounds of the benefit function from the proposed theorem. We summarized the following criteria for each population as illustrated in Figure 1:

- lower bound : a_i ;
- upper bound : b_i ;
- midpoint : $(a_i + b_i)/2$;
- real benefit : dot product of the benefit vector and the fractions of response types;

From Figure 1, it is clear that the proposed bounds obtained from Theorem 2 are a good estimation of the real benefit. The lower and upper bounds are closely around the real benefit, and the midpoints are almost identified with the real benefit. Besides, the average gap of the bounds, $\frac{\sum(b_i - a_i)}{1000}$, is 0.330, which is also small compared to the largest possible gap of 6.

Second, we set the benefit vector to another common one, $(-1, 1, 1, -1, -1, 1, -1, -1, -1)$, which is to evaluate the expected difference between benefited and unbenefited (i.e., unaffected and harmed) individuals. We again randomly generated 1000 populations where each population consists of different fractions of nine response types. The data-generating process and all other factors remain the same. We summarized the same criteria for each population as illustrated in Figure 2.

From Figure 2, it is clear that the proposed bounds obtained from Theorem 2 are a good estimation of the real benefit. The lower and upper bounds are closely around the real benefit, and the midpoints are almost identified with the real benefit. Besides, the average gap of the bounds, $\frac{\sum(b_i - a_i)}{1000}$, is 0.6520, which is also small compared to the largest possible gap of 9.

Third, we set the benefit vector to another common one, $(0, 1, 1, 0, 0, 1, 0, 0, 0)$, which is to evaluate the expected benefited individuals. We again randomly generated 1000 populations where each population consists of different fractions




Figure 2: Bounds of the benefit function for 100 sample populations out of 1000 with the benefit vector $(-1, 1, 1, -1, -1, 1, -1, -1, -1)$.

of nine response types. The data-generating process and all other factors still remain the same. We summarized the same criteria for each population as illustrated in Figure 3.

From Figure 3, it is clear that the proposed bounds obtained from Theorem 2 are a good estimation of the real benefit. The lower and upper bounds are closely around the real benefit, and the midpoints are almost identified with the real benefit. Besides, the average gap of the bounds, $\frac{\sum(b_i - a_i)}{1000}$, is 0.3284, which is also small compared to the largest possible gap of 3.

Lastly, we set the benefit vector to the last common one, $(0, 0, 0, -1, 0, 0, -1, -1, 0)$, which is to evaluate the expected harmed individuals (we set the benefit vector to -1 because we want to minimize the harmed individuals). We again randomly generated 1000 populations where each population consists of different fractions of nine response types. The data-generating process and all other factors still remain the same. We summarized the same criteria for each population as illustrated in Figure 4.

From Figure 4, it is clear that the proposed bounds obtained from Theorem 2 are a good estimation of the real benefit. The lower and upper bounds are closely around the real benefit, and the midpoints are almost identified with the real benefit. Besides, the average gap of the bounds, $\frac{\sum(b_i - a_i)}{1000}$, is 0.3266, which is also small compared to the largest possible gap of 3.

Discussion

We have shown that the proposed theorems are a good estimation of the non-binary benefit function using examples and simulated studies. One may concern about the computation complexity of Algorithms 1 and 2. They are for sure in exponential time. However, the m and n (i.e., values of X and Y) are usually small constants; therefore, we do not need to worry about too much.

Moreover, one may wonder why we need to investigate the bounds of the linear combination of the probabilities of




Figure 3: Bounds of the benefit function for 100 sample populations out of 1000 with the benefit vector $(0, 1, 1, 0, 0, 1, 0, 0, 0)$.




Figure 4: Bounds of the benefit function for 100 sample populations out of 1000 with the benefit vector $(0, 0, 0, -1, 0, 0, -1, -1, 0)$.

causation because Li and Pearl have already proposed the bounds for all the probabilities of causation (Li and Pearl 2022b). Note that the linear combination of the probabilities of causation is not a simple extension of the theoretical bounds for any single probability of causation. First, all single probabilities of causation are not identifiable with no further assumption. Here in this paper, the benefit function becomes identifiable to a point estimation with a certain relationship of the benefit vector detected by Algorithm 1. In other words, the benefit vector provides identifiability conditions. Second, say a probability of causation A has bounds $[a, b]$ and another probability of causation B has bounds $[c, d]$, then the bounds of $A + B$ is not simply $[a + c, b + d]$. We have to use the properties of probabilities of causation (i.e., in Algorithm 2) to obtain non-trivial bounds (the new bounds might be even narrower than both $[a, b]$ and $[c, d]$). Besides, compared to (Li and Pearl 2019), as the response types increased from 4 to exponential many, the identifiability conditions became non-unique, and the difficulty of obtaining the narrow bounds raised rapidly.

Another concern is why we do not apply linear programming-based approaches to the proposed problem, such as in (Balke 1995). The reason is that such an approach would necessitate a linear programming formulation with an exponential number of variables (mn^m variables), rendering it impractical in terms of both computational time and accuracy. However, in contrast, even though our approach still involves exponential run-time, we have successfully reduced the objective function using Algorithm 2 to the pre-computed probabilities of causation as proposed by Li and Pearl (Li and Pearl 2022b).

Conclusion and Future Work

We demonstrated the formalization of the general benefit function with nonbinary treatment and effect. We provided the algorithm to compute the bounds of the general benefit function and the algorithm to check whether the benefit function is identifiable with purely experimental data. Examples and simulation results are provided to support the proposed theorems.

Future studies could assess the statistical properties of the proposed bounds. How tight would the bounds be? Does it sufficient to make decisions? Which data, experimental or observational, would affect the bounds more? How would the number of values in treatment and effect affect the quality of the bounds?

Another future direction could be to improve the bounds using covariate information as Li and Pearl (Li and Pearl 2019) did for the binary benefit function.

Acknowledgements

This research was supported in parts by grants from the National Science Foundation [#IIS-2106908 and #IIS-2231798], Office of Naval Research [#N00014-21-1-2351], and Toyota Research Institute of North America [#PO-000897].

References

- Balke, A. A. 1995. *Probabilistic counterfactuals: semantics, computation, and applications*. University of California, Los Angeles.
- Berson, A.; Smith, S.; and Thearling, K. 1999. *Building data mining applications for CRM*. McGraw-Hill Professional.
- Bottou, L.; Peters, J.; Quiñero-Candela, J.; Charles, D. X.; Chikering, D. M.; Portugaly, E.; Ray, D.; Simard, P.; and Snelson, E. 2013. Counterfactual reasoning and learning systems: The example of computational advertising. *The Journal of Machine Learning Research*, 14(1): 3207–3260.
- Dawid, P.; Musio, M.; and Murtas, R. 2017. The Probability of Causation. *Law, Probability and Risk*, (16): 163–179.
- Galles, D.; and Pearl, J. 1998. An axiomatic characterization of causal counterfactuals. *Foundations of Science*, 3(1): 151–182.
- Halpern, J. Y. 2000. Axiomatizing causal reasoning. *Journal of Artificial Intelligence Research*, 12: 317–337.
- Hung, S.-Y.; Yen, D. C.; and Wang, H.-Y. 2006. Applying data mining to telecom churn management. *Expert Systems with Applications*, 31(3): 515–524.
- Lejeune, M. A. 2001. Measuring the impact of data mining on churn management. *Internet Research*, 11(5): 375–387.
- Li, A.; and Pearl, J. 2019. Unit Selection Based on Counterfactual Logic. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, 1793–1799. International Joint Conferences on Artificial Intelligence Organization.
- Li, A.; and Pearl, J. 2022a. Bounds on causal effects and application to high dimensional data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5773–5780.
- Li, A.; and Pearl, J. 2022b. Probabilities of Causation with Non-binary Treatment and Effect. Technical Report R-516, Department of Computer Science, University of California, Los Angeles, CA.
- Li, A.; and Pearl, J. 2022c. Unit selection with causal diagram. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 5765–5772.
- Li, L.; Chen, S.; Kleban, J.; and Gupta, A. 2014. Counterfactual estimation and optimization of click metrics for search engines. *arXiv preprint arXiv:1403.1891*.
- Mueller, S.; Li, A.; and Pearl, J. 2022. Causes of effects: Learning individual responses from population data. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI-22)*, 2712–2718.
- Sun, W.; Wang, P.; Yin, D.; Yang, J.; and Chang, Y. 2015. Causal inference via sparse additive models with application to online advertising. In *AAAI*, 297–303.
- Tsai, C.-F.; and Lu, Y.-H. 2009. Customer churn prediction by hybrid neural networks. *Expert Systems with Applications*, 36(10): 12547–12553.
- Yan, J.; Liu, N.; Wang, G.; Zhang, W.; Jiang, Y.; and Chen, Z. 2009. How much can behavioral targeting help online advertising? In *Proceedings of the 18th international conference on World Wide Web*, 261–270. ACM.
- Zhang, J.; Tian, J.; and Bareinboim, E. 2022. Partial counterfactual identification from observational and experimental data. In *International Conference on Machine Learning*, 26548–26558. PMLR.