

---

# ADDRESSING THE NULL PARADOX IN EPIDEMIC MODELS: CORRECTING FOR COLLIDER BIAS IN CAUSAL INFERENCE

---

A PREPRINT

**Heejong Bong**  
Department of Statistics  
University of Michigan  
Ann Arbor, MI 48109  
hbong@umich.edu

**Valérie Ventura**    **Larry Wasserman**  
Department of Statistics & Data Science  
Carnegie Mellon University  
Pittsburgh, PA 15213  
{vventura, larry}@stat.cmu.edu

October 15, 2024

## ABSTRACT

We address the null paradox in epidemic models, where standard methods estimate a non-zero treatment effect despite the true effect being zero. This occurs when epidemic models mis-specify how causal effects propagate over time, especially when covariates act as colliders between past interventions and latent variables, leading to spurious correlations. Standard approaches like maximum likelihood and Bayesian methods can misinterpret these biases, inferring false causal relationships. While semi-parametric models and inverse propensity weighting offer potential solutions, they often limit the ability of domain experts to incorporate epidemic-specific knowledge. To resolve this, we propose an alternative estimating equation that corrects for collider bias while allowing for statistical inference with frequentist guarantees, previously unavailable for complex models like SEIR.

**Keywords** Null paradox, Causal inference, Epidemic models, Inverse propensity weighting, Estimating equations

## 1 Introduction

In this paper we consider the problem of inferring the causal effects of time-varying interventions in epidemics. Our goal is to examine epidemic models through the lens of causal inference. In particular, we study the consequences of the common practice of adding an intervention  $A$  directly into an epidemic model, for example, by letting the model for the reproduction number depend on  $A$ . The term *intervention* can refer to control measures, treatments, public health policies or spontaneous changes in population behavior such as reduced mobility. Such interventions often change over time, often depending on the state of the epidemic. For example, we may want to estimate the effect of vaccinations, masks or social mobility on the number of infections or number of hospitalizations.

As an example, consider the usual SIR model due to [Kermack and McKendrick \(1927\)](#) which is given by three differential equations

$$\begin{aligned}\frac{dS_t}{dt} &= -\frac{\beta I_t S_t}{N_t} \\ \frac{dI_t}{dt} &= \frac{\beta I_t S_t}{N_t} - \gamma I_t \\ \frac{dR_t}{dt} &= \gamma I_t\end{aligned}\tag{1}$$

for  $t > 0$ , where  $S_t$  is the number of susceptibles at time  $t$ ,  $I_t$  is the number of new infections at  $t$ ,  $R_t$  is the number of removed and  $N_t = S_t + I_t + R_t$  is the total population size. In many cases, we do not observe  $(S_t, I_t, R_t)$  but rather we observe another variable  $Y_t$  which could be reported cases, hospitalizations, deaths, etc. This requires a further

model  $p(y_t|I_t)$  relating  $Y_t$  to infections  $I_t$ . The negative binomial distribution is a common choice. To include an intervention  $A_t$  we could, for example, replace  $\beta$  with  $\beta_t = \beta e^{\delta A_t}$  where  $\delta$  is the parameter that modulates the effect of  $A_t$  on the subsequent number of infections (and therefore on  $Y_t$ ).

Another example is the discretized SEIR model (Lekone and Finkenstädt, 2006; Gibson and Renshaw, 1998; Mode and Sleeman, 2000) which models susceptibles  $S_t$ , exposed  $E_t$ , infecteds  $I_t$  and removed  $R_t$  by

$$\begin{aligned} S_{t+h} &= S_t - B_t \\ E_{t+h} &= E_t + B_t - C_t \\ I_{t+h} &= I_t + C_t - Y_t \\ R_{t+h} &= R_t + Y_t \end{aligned} \tag{2}$$

where  $S_t + E_t + I_t + R_t = N_t$ ,  $B_t \sim \text{Binomial}(S_t, p_{B,t})$ ,  $C_t \sim \text{Binomial}(E_t, p_C)$ ,  $Y_t \sim \text{Binomial}(I_t, p_Y)$ , and

$$p_{B,t} = 1 - \exp\left\{-\frac{\eta_t}{N} h I_t\right\}, \quad p_C = 1 - e^{-\eta h}, \quad p_Y = 1 - e^{-\gamma h}.$$

In this model,  $h$  represents the time interval (for example  $h = 1$  day),  $B_t$  denotes the number of susceptibles who become infected,  $C_t$  is the number of cases, and  $Y_t$  is the number of removed cases. The parameters are the transmission rate  $\beta_t$ , the mean incubation period  $1/\eta$  and the mean infectious period  $1/\gamma$ . To include an intervention  $A_t$  we could, for example, replace  $\eta_t$  with  $\eta(\bar{A}_t; \beta) = e^{-\beta_1 - \beta_A A_t}$ .

We will refer to models that are modified to include an intervention as *augmented epidemic models*. In the language of causal inference, these are models for *counterfactuals*. That is, they are models that are used to specify the effect of the intervention on the outcome. We might use such a model to answer a question like: how many cases would there be if we locked down for 2 weeks?

In the epidemic modeling literature, the parameters of these models are usually estimated by interpreting the model as a model for the data generating process. Then one can use maximum likelihood or Bayesian methods. But data generating models and counterfactual models are, in general, not the same thing. They coincide if there are no confounding variables but otherwise they can be quite different. In that case maximum likelihood estimates can be biased.

In the causal literature, two methods are available for correctly estimating the parameters of the counterfactual model. The first (and simplest) is to use an estimating equation instead of likelihood based methods. The second is to construct a joint distribution that is consistent with the counterfactual model using, for example, the Evans-Didelez (2023) frugal parameterization. This permits likelihood based inference but requires many more modeling assumptions. Our main goal is to explain why using estimating equations to estimate the parameters of the model can have advantages over the common practice of using the maximum likelihood estimate or Bayes estimate.

**Paper Outline** In Sections 2 and 3 we provide brief backgrounds for causal inference and for common types of epidemic models, including the semi-mechanistic model and SEIR model of main interest here. Specifically we apply the  $g$ -formula to derive the mean causal effect of an intervention when using these models. In Section 4 we combine causal inference and epidemic models and describe how their parameters are estimated. In Section 5 we explain why the ubiquitous  $g$ -null paradox phenomenon arises and how our estimation approach provides remedies to it. Finally, we present empirical examples based on simulated and observational data in Section 6 and conclude in Section 7.

## 2 Background on Causal Inference

Putting aside epidemic models for a moment, we now review some background on causal inference.

First, consider a single outcome  $Y$  and a binary treatment  $A \in \{0, 1\}$ . The *counterfactual*  $Y(a)$  is the value the outcome  $Y$  would take if the intervention  $A$  were set to  $a$ . Thus, we now have four random variables  $(A, Y, Y(0), Y(1))$  where  $Y(0)$  is the value  $Y$  would have if  $A = 0$  and  $Y(1)$  is the value  $Y$  would have if  $A = 1$ . These counterfactuals are linked to the observed data  $(A, Y)$  by the equation  $Y = Y(A)$ . If  $A = 1$  then  $Y = Y(1)$  but  $Y(0)$  is unobserved. If  $A = 0$  then  $Y = Y(0)$  but  $Y(1)$  is unobserved. Many causal questions are quantified by these counterfactuals. For example,  $\mathbb{E}[Y(1)] - \mathbb{E}[Y(0)]$  is used to quantify the causal effect of the treatment. It can be shown that if there are no confounding variables — variables that affect  $Y$  and  $A$  — then  $P(Y \leq y|A = a) = P(Y(a) \leq y)$  so that the distribution of the counterfactual  $Y(a)$  is the same as the conditional distribution of the observable  $Y$ . But if there are confounding variables  $X$  then it can be shown (under three conditions described below) that

$$P(Y(a) \leq y) = \int P(Y \leq y|A = a, X = x)dP(x)$$

and in general,  $P(Y \leq y|A = a) \neq P(Y(a) \leq y)$ .

Now consider observed time series data on a subject of the form

$$(X_1, A_1, Y_1), \dots, (X_T, A_T, Y_T)$$

where  $A_t$  is some intervention at time  $t$ ,  $Y_t$  is the outcome of interest at time  $t$  and  $X_t$  refers to potential confounding variables which might affect  $A_t$  and  $Y_t$ . We use overbars to represent histories such as  $\bar{A}_t = (A_1, \dots, A_t)$ . Again we introduce the counterfactual  $Y_t(\bar{a}_t)$  which is the value  $Y_t$  would have had the treatment sequence been  $\bar{a}_t = (a_1, \dots, a_t)$  rather than the actual observation  $\bar{A}_t = (A_1, \dots, A_t)$ . For example, suppose that  $a_t = 1$  means that the subject wore a mask and  $a_t = 0$  means that the subject didn't wear a mask. Then  $Y_t(0, 0, \dots, 0)$  is the outcome at time  $t$  if the subject never wore a mask. (In some literature,  $\mathbb{E}[Y_T(\bar{a}_T)]$  is denoted by  $\mathbb{E}[Y_T|\text{do}(\bar{A}_T = \bar{a}_T)]$ .) Causal inference requires three conditions:

(A1) No interference: if  $\bar{A}_t = \bar{a}_t$  then  $Y_t(\bar{a}_t) = Y_t$ .

(A2) Positivity:  $\pi(a_t|\bar{x}_{t-1}, \bar{a}_{t-1}) > \epsilon > 0$  for all values, where  $\pi(a_t|\bar{x}_{t-1}, \bar{a}_{t-1})$  is the density of  $A_t$  given the past.

(A3) No unmeasured confounding: the variable  $Y_t(\bar{a}_t)$  is independent of  $A_t$  given the past measured variables.

Assumption (A1) means that the observed  $Y_t$  is equal to the counterfactual  $Y_t(\bar{a}_t)$  if the observed treatment sequence  $\bar{A}_t$  happens to equal  $\bar{a}_t$ . This means a subject's outcome is affected by their treatment but not affected by another subject's treatment. Assumption (A2) means that, conditional on the past, every subject has nonzero probability of receiving treatment at any level. Assumption (A3) means that we have measured all important confounding variables, which are variables that affect the treatment and the outcome.

In this setting, and under assumptions (A1)-(A3), [Robins \(1986\)](#) proved that

$$\mathbb{E}[Y_t(\bar{a}_t)] = \psi(\bar{a}_t)$$

where

$$\psi(\bar{a}_t) = \int \dots \int \mathbb{E}[Y_t|\bar{x}_t, \bar{y}_{t-1}, \bar{a}_t] \prod_{s=1}^t p(x_s, y_s|h_s) dx_s dy_s \quad (3)$$

where  $h_s = (\bar{x}_{s-1}, \bar{a}_{s-1}, \bar{y}_{s-1})$  is the history before time  $s$ . Equation (3) is known as the *g-formula*. Note that, in general,

$$\mathbb{E}[Y_t(\bar{a}_t)] \neq \mathbb{E}[Y_t|\bar{A}_t = \bar{a}_t]$$

which is the difference between causation (the left hand side) and correlation (the right hand side). In what follows, we will often write  $\psi(\bar{a}_t; \theta)$  where  $\theta$  denotes any parameters that are involved. The *g-formula* above is for the mean but there are similar expressions for densities, cdf's, quantiles etc. In particular, let  $p^*(y_t; \bar{a}_t)$  denote the density of counterfactual  $Y_t(\bar{a}_t)$  evaluated at a point  $y_t$ . Then

$$p^*(y_t; \bar{a}_t) = \int \dots \int p(y_t|\bar{x}_t, \bar{y}_{t-1}, \bar{a}_t) \prod_{s=1}^t p(x_s, y_s|h_s) dx_s dy_s. \quad (4)$$

Again, the causal density  $p^*(y_t; \bar{a}_t)$  should not be confused with the usual conditional  $p(y_t|\bar{a}_t)$ . The causal effect in Eq. (3) is the mean of  $p^*(y_t; \bar{a}_t)$ ,

$$\mathbb{E}[Y_t(\bar{a}_t)] = \int \dots \int \mu(\bar{x}_t, \bar{y}_{t-1}, \bar{a}_t) \prod_{s=1}^t p(x_s, y_s|h_s) dx_s dy_s$$

where  $\mu(\bar{x}_t, \bar{y}_{t-1}, \bar{a}_t) = \mathbb{E}[Y_t|\bar{x}_t, \bar{y}_{t-1}, \bar{a}_t]$ .

The *g-formula* has a graphical interpretation. Starting with a directed graph  $G$  such as Fig. 1(a), form a new graph  $G^*$  in which all arrows pointing into any  $A_s$ ,  $s \leq t$ , are removed and in which any  $A_s$  is fixed at a value  $a_s$ ; see Fig. 1(b). Equation (4) is then the marginal density for  $Y_t$  corresponding to the density in the graph  $G^*$ .

It is common in the causal inference literature for the user to specify a simple, interpretable model called a *marginal structural model* (MSM). For example, one might take

$$\psi(\bar{a}_t; \theta) = \theta_0 + \theta_1 \sum_{s=1}^t a_s \quad (5)$$

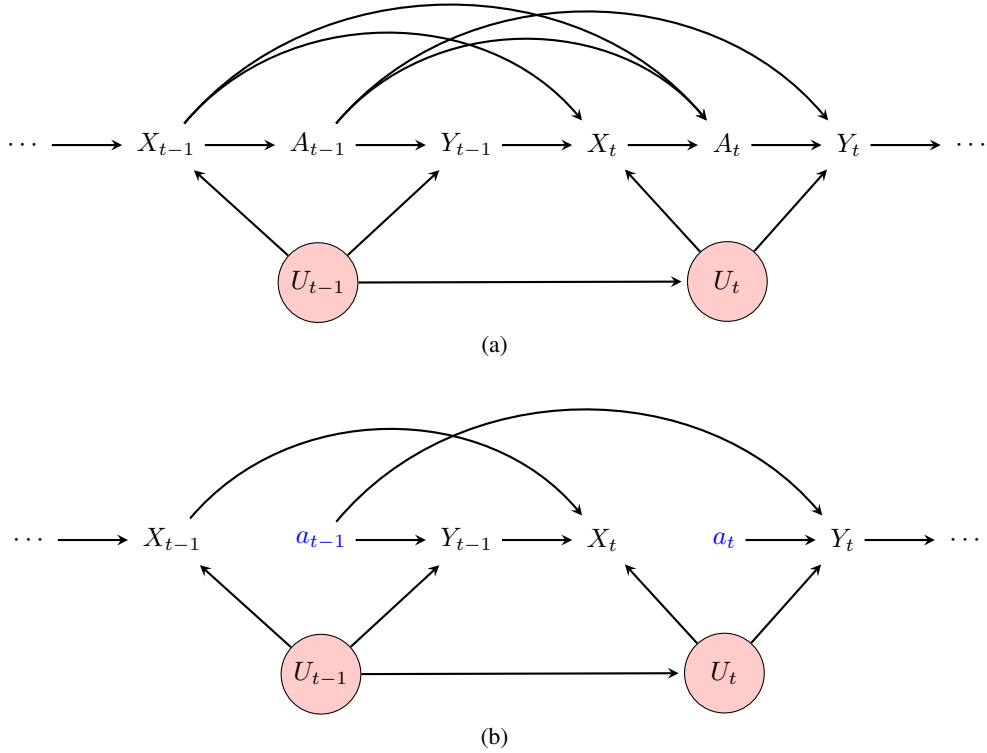


Figure 1: (a) The directed graph of the type of models we consider. Latent variables are indicated with pink nodes. Importantly, we allow unobserved variables  $U_t$  as long as they do not directly affect  $A_t$ . (b) The intervention graph after setting  $\bar{A}_t = \bar{a}_t$ .

which says that the mean of  $Y_t(\bar{a}_t)$  is a linear function of the cumulative dose  $\sum_{s=1}^t a_s$ . This is not a model for the data generating process, and it is possible to estimate the parameters of a MSM without completely specifying the data generating process using an estimating equation, as we explain later. Instead, this approach is akin to specifying a regression model for the effect of  $\bar{a}_t$  on  $Y_t$ , except that this must be done having observed  $(\bar{A}_t, Y_t)$ . We never observe the counterfactuals  $Y_t(\bar{a}_t)$  except for the one observed case  $\bar{a}_t = \bar{A}_t$ . This means that such models are not data driven but user specified, and they are not typically evaluated for goodness-of-fit.

### 3 Epidemic Causal Models

While semi-parametric marginal structural models such as Eq. (5) have been applied to epidemic data (Bonvini et al., 2022), they often limit the ability of domain experts to incorporate epidemic-specific knowledge. Instead, we interpret the augmented epidemic model as defining a model for the counterfactual  $Y_t(\bar{a}_t)$ . It seems, in fact, that this is how augmented models are intended to be interpreted. For example, it is common to do scenario predictions from an epidemic model by fixing  $\bar{a}_t$  and simulating from the model. In other words, we use the augmented model to answer questions like: what would happen if we set  $\bar{A}_t$  to be  $\bar{a}_t$ ? This is precisely what a counterfactual is. Throughout the paper, we maintain the assumption that the augmented epidemic model defines a model for the counterfactuals  $Y_t(\bar{a}_t)$ .

As always, it should be understood that the model  $\psi(\bar{a}_t; \theta)$  is at best an approximation. More precisely, we are estimating the projection of  $\mathbb{E}[Y_t(\bar{a}_t)]$  onto the model  $\mathcal{M} = \{\psi(\bar{a}_t; \theta) : \theta \in \Theta\}$ .

Now we discuss the details of obtaining  $\psi(\bar{a}_t; \theta)$  from an augmented epidemic model. We start with an example where it can be obtained in closed form.

*Example 1 (Semi-mechanistic Model).* We consider a version of the semi-mechanistic epidemic model from [Bhatt et al. \(2022\)](#): for parameter  $\beta \equiv (\beta_1, \beta_A)$ ,

$$\begin{aligned}\mathbb{E}[I_t | \bar{A}_t, \bar{I}_{t-1}, \bar{Y}_{t-1}] &= R(\bar{A}_t, \beta) \sum_{s < t} g_{t-s} I_s \\ \mathbb{E}[Y_t | \bar{A}_t, \bar{I}_t, \bar{Y}_{t-1}] &= \alpha_t \sum_{s < t} \pi_{t-s} I_s\end{aligned}\quad (6)$$

where

$$R(\bar{A}_t, \beta) = \frac{K}{1 + \exp(\beta_1 + \beta_A \bar{A}_t)}, \quad (7)$$

$K$  is the maximum transmission rate,  $\pi$  is the infection to death distribution and  $g$  is the generating distribution. A different version takes

$$\mathbb{E}[I_t | \bar{A}_t, \bar{I}_{t-1}, \bar{Y}_{t-1}] = \sum_{s < t} e^{\beta_1 + \beta_A \bar{A}_s} g_{t-s} I_s. \quad (8)$$

We call these the multiplicative and exponential versions. For fixed values of  $\bar{A}_t$ ,  $\alpha_t$ ,  $g$  and  $\pi$ , this model is a pair of linear equations and is an example of a *structural equation model* (SEM).

Meaningful dynamics in this model requires some positive infections  $I_t$  prior to time  $t = 1$ . [Bhatt et al. \(2022\)](#) assumed  $I_t = 0$  for  $t \leq -T_0$  and  $I_t = e^\mu$  for  $t = -T_0 + 1, \dots, 0$ , where  $\mu$  is a parameter to be estimated and  $T_0 = 6$ . Let  $\bar{I}_0 \equiv (I_{-T_0+1}, \dots, I_0)$  indicate those *seeding values* in the infection process. (We still exclude those seeding values in the definition of  $\bar{I}_t = (I_1, \dots, I_t)$ .)

For the exponential model define

$$\Lambda^e = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ g_1 e^{\beta_1 + \beta_A a_1} & 0 & \cdots & 0 \\ g_2 e^{\beta_1 + \beta_A a_1} & g_1 e^{\beta_1 + \beta_A a_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ g_{t-1} e^{\beta_1 + \beta_A a_1} & g_{t-2} e^{\beta_1 + \beta_A a_2} & \cdots & 0 \end{pmatrix}, \quad \Lambda_0^e = \begin{pmatrix} g_{T_0} e^{\beta_1 + \beta_A a_{-T_0+1}} & \cdots & g_1 e^{\beta_1 + \beta_A a_0} \\ g_{T_0+1} e^{\beta_1 + \beta_A a_{-T_0+1}} & \cdots & g_2 e^{\beta_1 + \beta_A a_0} \\ g_{T_0+2} e^{\beta_1 + \beta_A a_{-T_0+1}} & \cdots & g_3 e^{\beta_1 + \beta_A a_0} \\ \vdots & \vdots & \vdots \\ g_{T_0+t-1} e^{\beta_1 + \beta_A a_{-T_0+1}} & \cdots & g_t e^{\beta_1 + \beta_A a_0} \end{pmatrix},$$

and for the multiplicative model define

$$\Lambda^m = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ g_1 R(\bar{a}_2, \beta) & 0 & \cdots & 0 \\ g_2 R(\bar{a}_3, \beta) & g_1 R(\bar{a}_3, \beta) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ g_{t-1} R(\bar{a}_t, \beta) & g_{t-2} R(\bar{a}_t, \beta) & \cdots & 0 \end{pmatrix}, \quad \Lambda_0^m = \begin{pmatrix} g_{T_0} R(\bar{a}_1, \beta) & \cdots & g_1 R(\bar{a}_1, \beta) \\ g_{T_0+1} R(\bar{a}_2, \beta) & \cdots & g_2 R(\bar{a}_2, \beta) \\ g_{T_0+2} R(\bar{a}_3, \beta) & \cdots & g_3 R(\bar{a}_3, \beta) \\ \vdots & \vdots & \vdots \\ g_{T_0+t-1} R(\bar{a}_t, \beta) & \cdots & g_t R(\bar{a}_t, \beta) \end{pmatrix}.$$

Finally, define

$$\Pi = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ \pi_1 \alpha_2 & 0 & \cdots & 0 \\ \pi_2 \alpha_3 & \pi_1 \alpha_3 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \pi_{t-1} \alpha_t & \pi_{t-2} \alpha_t & \cdots & 0 \end{pmatrix}, \quad \Pi_0 = \begin{pmatrix} \pi_{T_0} \alpha_1 & \cdots & \pi_1 \alpha_1 \\ \pi_{T_0+1} \alpha_2 & \cdots & \pi_2 \alpha_2 \\ \pi_{T_0+2} \alpha_3 & \cdots & \pi_3 \alpha_3 \\ \vdots & \vdots & \vdots \\ \pi_{T_0+t-1} \alpha_t & \cdots & \pi_t \alpha_t \end{pmatrix}.$$

Then the marginal structural model  $\psi(\bar{a}_t; \theta)$  is given in a closed form as follows.

*Theorem 3.1.* For the exponential model,

$$\mathbb{E}[I_t(\bar{a}_t)] = [(id - \Lambda^e)^{-1} \Lambda_0^e \bar{I}_0]_t \quad (9)$$

and

$$\mathbb{E}[Y_t(\bar{a}_t)] = [\{\Pi(id - \Lambda^e)^{-1} \Lambda_0^e + \Pi_0\} \bar{I}_0]_t, \quad (10)$$

where the subscript  $t$  represents the  $t$ -th element of the outcome vector. For the multiplicative model, the expressions are the same except that  $\Lambda^m$  and  $\Lambda_0^m$  replace  $\Lambda^e$  and  $\Lambda_0^e$ .

*Proof.* Consider the intervened graph in Fig. 1(b) with  $\bar{A}_t$  set to  $\bar{a}_t$ . For this graph, we have

$$\bar{I}_t(\bar{a}_t) = \Lambda^e \bar{I}_t(\bar{a}_t) + \Lambda_0^e \bar{I}_0 + \epsilon,$$

for the exponential model (Eq. (8)), which, as mentioned above, is a linear structural equation model. Now

$$\bar{I}_t(\bar{a}_t) = (id - \Lambda^e)^{-1} \Lambda_0^e \bar{I}_0 + (id - \Lambda^e)^{-1} \epsilon$$

and hence, the last element of this vector is

$$\mathbb{E}[I_t(\bar{a}_t)] = [(id - \Lambda^e)^{-1} \Lambda_0^e \bar{I}_0]_t.$$

Subsequently,

$$\mathbb{E}[\bar{Y}_t(\bar{a}_t)] = \Pi \mathbb{E}[\bar{I}_t(\bar{a}_t)] + \Pi_0 \bar{I}_0 = \{[\Pi(id - \Lambda^e)^{-1} \Lambda_0^e + \Pi_0] \bar{I}_0\}_t.$$

The proof proceeds similarly for the multiplicative model (Eq. (6)), but with  $\Lambda^m$  and  $\Lambda_0^m$  in place of  $\Lambda^e$  and  $\Lambda_0^e$ .  $\square$

**The General Case.** In general, writing down  $\psi(\bar{a}_t; \theta)$  is not possible. Instead, we approximate it by Monte Carlo for each value of  $\theta$ . We fix  $\bar{A}_t$  fixed at  $\bar{a}_t$ . Now we simulate  $Y_1, \dots, Y_T$  from the augmented epidemic model parameterized with  $\theta$ . We repeat this simulation  $N$  times giving values  $Y_1^{(k)}, \dots, Y_T^{(k)}$  for  $k = 1, \dots, N$ . Then we set

$$\psi(\bar{a}_t; \theta) \approx \frac{1}{N} \sum_{k=1}^N Y_t^{(k)}. \quad (11)$$

This approach is very simple albeit computationally expensive.

## 4 Estimation

Now we discuss the estimation of the parameters  $\theta$ .

**No Confounding Case.** First suppose there are no confounding variables. In this case,  $P_\theta(Y_t(\bar{a}_t) \leq y) = P_\theta(Y_t \leq y | \bar{A}_t = \bar{a}_t)$ . Thus, the augmented model not only provides a model for the counterfactuals, it also provides a model for the observed data  $Y_1, \dots, Y_T$ . Specifically, it defines

$$G_{t,\theta}(y | \bar{A}_t = \bar{a}_t, \bar{Y}_t = \bar{y}_t) = P_\theta(Y_t \leq y | \bar{A}_t = \bar{a}_t, \bar{Y}_t = \bar{y}_t).$$

One way to estimate  $\theta$  is to maximize the likelihood

$$\mathcal{L}(\theta) = \prod_t p_\theta(y_t | \bar{a}_t, \bar{y}_{t-1})$$

where  $p_\theta(y_t | \bar{a}_t, \bar{y}_{t-1}) = G'_{t,\theta}(y | \bar{A}_t = \bar{a}_t, \bar{Y}_t = \bar{y}_t)$  is the density of  $Y_t$ . When  $p_\theta$  is not available in closed form we can estimate the likelihood using simulation based inference (SBI) (Gourieroux and Monfort, 1993). (One can also combine the likelihood with a prior and get the Bayes estimate.)

Alternatively, we can define  $\hat{\theta}$  to be the solution of the estimating equation

$$\frac{1}{T} \sum_t h_t(\bar{A}_t) (Y_t - \psi_\theta(\bar{a}_t)) W_t = 0$$

where

$$W_t = \prod_{s=1}^t \frac{\pi(A_s | \bar{A}_{s-1})}{\pi(A_s | \bar{A}_{s-1}, \bar{Y}_{s-1})},$$

$\pi(a_t | \bar{a}_{t-1}, \bar{y}_{t-1})$  is the density of  $A_t$  given the past and  $\psi_\theta(\bar{a}_t) = \mathbb{E}_\theta[Y_t(\bar{a}_t)]$ . The function  $h_t(\bar{A}_t)$  is an arbitrary function of  $\bar{A}_t$  whose choice affects the standard deviation of the estimator. For simplicity, we take  $h_t(\bar{A}_t) = 1$ .

In principle, the mle is optimal in parametric models and the estimator from the estimating equation might have larger variance than the mle. The advantage of using the estimating equation is that the estimator is robust to the presence of certain unobserved variables called phantoms; we explain this point in Section 5.

**Confounding Case.** Now suppose there are confounding variables  $X_1, \dots, X_T$ . These are variables that affect both the  $Y_t$ 's and the  $A_t$ 's. In this case, it is no longer true that the distribution of the counterfactuals  $Y_t(\bar{a}_t)$  is equal to the conditional distribution of the observables  $Y_t$ . In this case, the estimating equation still provides a valid way to estimate  $\theta$ .

We estimate  $\theta$  by solving the estimating equation

$$\sum_t (Y_t - \psi_\theta(\bar{a}_t)) W_t = 0 \quad (12)$$

where

$$W_t = \prod_{s=1}^t \frac{\pi(A_s | \bar{A}_{s-1})}{\pi(A_s | \bar{A}_{s-1}, \bar{Y}_{s-1}, \bar{X}_{s-1})}$$

and  $\pi(a_t | \bar{a}_{t-1}, \bar{y}_{t-1}, \bar{x}_{t-1})$  is the density of  $A_t$  given the past.

**Solving Estimating Equations.** To solve the estimating equations in Eq. (12), we apply Newton's method, which requires the computation of the first derivative of the left-hand side of the equation with respect to the parameter of interest. Specifically, since  $\beta$  is the key parameter in both the semi-parametric model (Eq. (6)) and the SEIR model (Eq. (2)), we provide detailed calculations for the derivative with respect to  $\beta$ .

*Example 2 (Derivative for the Semi-mechanistic Model).* In Example 1, we derived the closed-form expression for the marginal structural model  $\psi(\bar{a}_t; \theta)$  in the multiplicative semi-mechanistic model as:

$$\psi(\bar{a}_t; \theta) = [\{\Pi(id - \Lambda^m)^{-1} \Lambda_0^m + \Pi_0\} \bar{I}_0]_t.$$

For each element  $\beta_i$  of the parameter vector  $\beta$  (either  $\beta_1$  or  $\beta_A$ ), the first derivative of  $\Lambda^m$  with respect to  $\beta_i$  is given by:

$$\frac{\partial}{\partial \beta_i} \Lambda^m(t, s) = g_{t-s} \frac{\partial}{\partial \beta_i} R(\bar{a}_t, \beta) \mathbf{1}\{t > s\},$$

where  $\Lambda^m(t, s)$  is parameterized by the rate function  $R(\bar{a}_t, \beta)$ , and  $\mathbf{1}\{t > s\}$  is an indicator function. Similarly, the derivative of  $\Lambda_0^m$  with respect to  $\beta_i$  follows the same structure.

Using the identity from matrix calculus,  $\frac{\partial U^{-1}}{\partial x} = -U^{-1} \frac{\partial U}{\partial x} U^{-1}$ , we can express the derivative of  $\psi(\bar{a}_t; \theta)$  with respect to  $\beta_i$  as:

$$\begin{aligned} \frac{\partial}{\partial \beta_i} \psi(\bar{a}_t; \theta) &= \left[ \left\{ \Pi(id - \Lambda^m)^{-1} \frac{\partial \Lambda^m}{\partial \beta_i} (id - \Lambda^m)^{-1} \Lambda_0^m + \Pi(id - \Lambda^m)^{-1} \frac{\partial \Lambda_0^m}{\partial \beta_i} \right\} \bar{I}_0 \right]_t \\ &= \left[ \Pi(id - \Lambda^m)^{-1} \left\{ \frac{\partial \Lambda^m}{\partial \beta_i} \mathbb{E}[\bar{I}_t(\bar{a}_t)] + \frac{\partial \Lambda_0^m}{\partial \beta_i} \bar{I}_0 \right\} \right]_t. \end{aligned}$$

The derivative for exponential model is given similarly.

*Example 3 (Derivative for the SEIR Model).* The SEIR model does not admit a closed-form expression for the marginal structural model  $\psi(\bar{a}_t; \theta)$ . In Eq. (11), we proposed estimating this quantity through Monte Carlo approximation for each given  $\theta$ . Here, we describe how the derivative of this approximation can also be computed using Monte Carlo samples.

First, by applying the  $g$ -formula, we have:

$$\begin{aligned} \psi(\bar{a}_t; \theta) &= \mathbb{E}_\theta[Y_t(\bar{a}_t)] \\ &= \int \cdots \int \mathbb{E}[Y_t | \bar{B}_{t-1}, \bar{C}_{t-1}, \bar{Y}_{t-1}] \prod_{s=1}^{t-1} dP_{Y_s}(Y_s | \bar{B}_{s-1}, \bar{C}_{s-1}, \bar{Y}_{s-1}) \\ &\quad \times dP_{C_s}(C_s | \bar{B}_{s-1}, \bar{C}_{s-1}, \bar{Y}_{s-1}) \times dP_{B_s}(B_s | \bar{B}_{s-1}, \bar{C}_{s-1}, \bar{Y}_{s-1}), \end{aligned}$$

where  $P_{Y_s}$ ,  $P_{C_s}$ , and  $P_{B_s}$  are binomial random variables. The "number of trials" parameters for these variables depend on the conditioning terms  $\bar{B}_{s-1}$ ,  $\bar{C}_{s-1}$ , and  $\bar{Y}_{s-1}$ , with success probabilities denoted by  $p_Y$ ,  $p_C$ , and  $p_B$ , respectively. Importantly, only  $p_{B,s}$ , and consequently  $P_{B_s}$ , are parametrized by  $\beta$  through  $\eta(\bar{a}_s; \beta)$ .

Now, suppose we have Monte Carlo samples  $\{(\bar{B}_T^{(k)}, \bar{C}_T^{(k)}, \bar{Y}_T^{(k)}) : k = 1, \dots, N\}$  drawn under the parameter  $\beta$ . For any alternative parameter  $\beta'$ , we approximate the  $g$ -formula using importance sampling as follows:

$$\psi(\bar{a}_t; \beta') \approx \frac{1}{N} \sum_{k=1}^N \mathbb{E}[Y_t^{(k)} | \bar{B}_{t-1}^{(k)}, \bar{C}_{t-1}^{(k)}, \bar{Y}_{t-1}^{(k)}] \prod_{s=1}^{t-1} \frac{dP_{B_s} | \beta'}{dP_{B_s} | \beta} (B_s | \bar{B}_{s-1}^{(k)}, \bar{C}_{s-1}^{(k)}, \bar{Y}_{s-1}^{(k)}),$$

where  $\frac{dP_{B_s|\beta'}}{dP_{B_s|\beta}}$  is the Radon–Nikodym derivative. Note that when  $\beta' = \beta$ , the importance sampling estimator reduces to the Monte Carlo approximation for  $\psi(\bar{a}_t; \beta)$  as given in Eq. (11).

Next, to compute the derivative of  $\psi(\bar{a}_t; \beta)$  with respect to  $\beta$ , we recognize that the derivative of the Radon–Nikodym derivative  $\frac{dP_{B_s|\beta'}}{dP_{B_s|\beta}}$  at  $\beta' = \beta$  is simply the gradient of the log-likelihood, i.e.,  $\nabla_\beta \log\{f_{B_s|\beta}(B_s|\bar{B}_{s-1}^{(k)}, \bar{C}_{s-1}^{(k)}, \bar{Y}_{s-1}^{(k)})\}$ . By applying the chain rule, we obtain:

$$\nabla_\beta \psi(\bar{a}_t; \beta) \approx \frac{1}{N} \sum_{k=1}^N \mathbb{E}[Y_t^{(k)} | \bar{B}_{t-1}^{(k)}, \bar{C}_{t-1}^{(k)}, \bar{Y}_{t-1}^{(k)}] \sum_{s=1}^{t-1} \nabla_\beta \log\{f_{B_s|\beta}(B_s|\bar{B}_{s-1}^{(k)}, \bar{C}_{s-1}^{(k)}, \bar{Y}_{s-1}^{(k)})\}.$$

**Inference.** Under some conditions, we have that

$$\sqrt{T}(\hat{\theta} - \theta) \rightsquigarrow N(0, \Sigma)$$

for a positive definite matrix  $\Sigma$ . Moreover, we can estimate  $\Sigma$  consistently from the data. See, for example, [De Jong \(1997\)](#); [De Jong and Davidson \(2000\)](#); [Andrews \(1991, 1988\)](#); [Newey et al. \(1987\)](#). Then  $\hat{\theta}_j \pm z_{\alpha/2} \sqrt{\hat{\Sigma}_{j,j}}$  is an asymptotic  $1 - \alpha$  confidence interval for  $\theta_j$ .

## 5 Phantoms and the Null Paradox

An additional reason for recommending to use estimation equations is due to variables we refer to as *phantoms*. Phantoms are unobserved variables that affect the  $Y_t$ 's but not the  $A_t$ 's; they are not confounding variables. Phantoms are very likely to exist, and they pose a serious problem because they make likelihood based methods inconsistent ([Robins, 1986](#); [Robins and Wasserman, 1997a](#)).

It may be shown that phantoms do not alter the  $g$ -formula at all. But they cause problems for maximum likelihood and Bayesian methods. In particular, even if the  $A_t$ 's have no causal effect, the maximum likelihood estimator will show a non-zero effect. Robins called this the  $g$ -null paradox.

To see a specific example of this, suppose we model the observables in Fig. 2(a) as:

$$\begin{aligned} \mathbb{E}[U_1] &= 0, \mathbb{E}[A_0] = 0, \\ X_1 &= \alpha_U U_1 + \alpha_A A_0 + \epsilon_X, \\ A_1 &= \gamma_X X_1 + \gamma_A A_0 + \epsilon_A, \\ Y_1 &= \beta_A A_1 + \beta_X X_1 + \beta_U U_1 + \epsilon_Y, \end{aligned}$$

where  $\beta_A$  is the true causal effect of  $A$  on  $I$ . Now if we were to observe data  $Y_1$ , we would fit the model

$$Y_1 = \theta_A A_1 + \theta_X X_1 + \delta_Y,$$

to estimate  $\theta_A$ , the effect of  $A_1$  on  $Y_1$ . If  $Y_1$  depends on  $A_1$  then  $\theta_A$  must be nonzero. If  $A_1$  has no causal effect on  $Y_1$  then  $\theta_A$  must be zero. But what if both are true? What if (i)  $Y_1$  depends on  $A_1$  but (ii) there is no causal effect? The reason  $Y_1$  can be dependent on  $A_1$  even when there is no causal effect, is the presence of phantoms. Consider the directed graph in Fig. 2(b) from [Robins and Wasserman \(1997b\)](#), [Bonvini et al. \(2022\)](#) and [Bates et al. \(2022\)](#). Because there are no arrows from the  $A_0$  and  $A_1$  to  $Y_1$ , there is no causal effect and it is easy to show that the distribution of  $Y_1(a_0, a_1)$  does not depend on  $(a_0, a_1)$ . However, the variable  $X_1$  is a collider — it has two arrows converging. From standard directed graph calculations, this implies that  $Y_1$  is dependent on  $A_0$  and  $A_1$  given  $X_1$ . That is  $p(y_1|x_1, a_0, a_1)$  is a function of  $(a_0, a_1)$ . Thus we have precisely the situation where  $Y_1$  has conditional dependence on the treatment but there is no causal effect.

However, the model cannot represent this situation. Indeed, no finite dimensional parametric model can model (i) and (ii). The reason we cannot model (i) and (ii) simultaneously is that the model is not variation independent: dependence and causation are tied together in the parameterization of the model. The ML and Bayes estimates (which are estimating the Kullback-Leibler (KL) projection of the distribution onto the model) are driven strongly by the dependence between  $Y_1$  and  $A_1$ , rather than by the causal effect. So when both (i) and (ii) hold, both these the estimates will be nonzero even though there is no causal effect. In fact, because  $U_1 = -\frac{\alpha_A}{\alpha_U} A_0 + \frac{1}{\alpha_U} X_1 - \frac{1}{\alpha_U} \epsilon_X$ ,

$$\begin{aligned} Y_1 &= \beta_A A_1 + \left(\beta_X + \frac{1}{\alpha_U} \beta_U\right) X_1 - \frac{\alpha_A}{\alpha_U} \beta_U A_0 + \left(\epsilon_I - \frac{1}{\alpha_U} \epsilon_X\right) \\ &= \left(\beta_A - \frac{1}{\gamma_A} \frac{\alpha_A}{\alpha_U} \beta_U\right) A_1 + \left(\beta_X + \frac{1}{\alpha_U} \beta_U + \frac{\gamma_X}{\gamma_A} \frac{\alpha_A}{\alpha_U} \beta_U\right) X_1 + \left(\epsilon_I - \frac{1}{\alpha_U} \epsilon_X + \frac{1}{\gamma_A} \frac{\alpha_A}{\alpha_U} \beta_U \epsilon_A\right), \end{aligned} \tag{13}$$






Figure 2: The effect of phantoms. The latent variable  $U$  is not a confounder as it has no arrows to  $A_0$  or  $A_1$ . The variable  $X_1$  is a collider, meaning that two arrowheads meet at  $X_1$ . This implies that  $Y_1$  and  $(A_0, A_1)$  are dependent conditional on  $X_1$ . In (b), neither  $A_0$  nor  $A_1$  have a causal effect on  $Y_1$ , but the estimate of the parameters that relate  $Y_1$  to  $(A_0, A_1)$  in the epidemic model will be non-zero.

which no longer involves the phantom variable. Hence, the MLE of  $\theta_A$  estimates  $\beta_A - \frac{1}{\gamma_A} \frac{\alpha_A}{\alpha_U} \beta_U$ , which is biased for the true causal effect  $\beta_A$ . There is phantom bias.

To summarize, this problem is due to three things: phantoms, which enable (i) and (ii) to both be true (see Section 5), variation dependence, which is a property of the model, and the fact that the KL projection is driven by dependence.

We can illustrate the problem as follows. Let  $\theta$  represent some measure of conditional dependence and let  $\beta$  denote the causal effect. Then

$$\underbrace{\theta \neq 0 \text{ but } \beta = 0}_{\text{phantoms}} \implies \underbrace{\hat{\theta} \neq 0}_{\text{KL projection}} \implies \underbrace{\hat{\beta} \neq 0}_{\text{variation dependence}}.$$

The fact that in the real world we can have dependence but no causal effect and the model cannot represent this, means that the model is misspecified. If  $\Theta_0$  denotes the parameter values that correspond to no causal effect and  $\Theta_+$  denotes the parameter values that correspond to conditional dependence, we have that  $\Theta_0 \cap \Theta_+ = \emptyset$ . This problem is well known in causal inference. It is called the *g-null paradox* and affects any sequentially parameterized causal model (models that parameterize the distribution of each variable given the past). This was first pointed out by Robins (1986) and has received much attention since then; see, for example, Robins and Wasserman (1997b), Bates et al. (2022), Robins (2000), Babino et al. (2019) and Evans and Didelez (2024). It appears that the problem has gone unnoticed in the literature on modeling epidemics.

Our estimating approach using estimation equation in Section 4 suggests a remedy. For example, under the model of Fig. 2(a), if we use the estimating equation approach instead, the counterfactual from the model is:

$$\begin{aligned} \psi(a_0, a_1; \theta) &= \int (\theta_A a_1 + \theta_X x_1) p_{X_1}(x_1 | a_0) dx_1 \\ &= \theta_X \alpha_A a_0 + \theta_A a_1, \end{aligned}$$

which yields the estimating equation

$$\begin{aligned} &\mathbb{E} \left[ (I_1 - \psi(A_0, A_1; \theta)) h(A_0, A_1) \frac{\pi(A_0) \pi(A_1 | A_0)}{\pi(A_0) \pi(A_1 | X_1, A_0)} \right] \\ &= \int (\beta_A a_1 + \beta_X x_1 - \theta_X \alpha_A a_0 - \theta_A a_1) h(a_0, a_1) \pi(a_0, a_1) p_{X_1}(x_1 | a_0) dx_1 da_1 da_0 \\ &= \int (\beta_X \alpha_A a_0 + \beta_A a_1 - \theta_X \alpha_A a_0 - \theta_A a_1) h(a_0, a_1) \pi(a_0, a_1) da_1 da_0. \end{aligned}$$

Setting the estimating equation to zero for any  $h$  yields  $(\theta_0, \theta_1) = (\beta_0, \beta_1)$  (and  $\theta_X = \beta_X$ ). Hence the estimating equation yields an unbiased estimate of the causal effect  $\beta_A$ . There is no phantom bias. We have focused on the relationship between  $I_t$  and the  $A_t$ 's but similar comments apply to  $Y_t$ .

## 6 Examples

The first two examples illustrate that phantom variables induce bias in parameter estimated via ML, but that using estimating equations instead yields unbiased estimates. The last example is an analysis of the effect of a certain measure of mobility on Covid-19 deaths in 30 US states at the start of the pandemic.




Figure 3: Generating distribution  $g$  from [Bhatt et al. \(2022\)](#).

### 6.1 Semi-mechanistic model simulated data

We now turn to a more realistic epidemic model. We simulate epidemic time series consistent with the DAG in Fig. 1(a), as follows. For  $t = 1, \dots, 120$ , we

1. simulate phantoms  $U_t$  from a Gaussian random process with mean zero and covariance kernel  $\Sigma(t, s) = \phi^{|t-s|}$ , with  $\phi = 0.95$ ;
2. sample confounders  $X_t$  from a Gaussian distribution with variance  $\sigma^2$  and mean  $\epsilon_I + \zeta_U U_t + \zeta_X X_{t-1} + \zeta_A A_{t-1} + \zeta_I I_{t-1}$ , where  $(\epsilon_I, \zeta_U, \zeta_X, \zeta_A, \zeta_I) = (0, 0.2, 0, 1, 0)$  and  $\sigma^2 = 0.09$ ;
3. generate binary interventions  $A_t$  from a Bernoulli distribution with

$$\mathbb{P}(A_t = 1 \mid \bar{X}_t, \bar{A}_{t-1}, \bar{I}_{t-1}) = \frac{e^{\gamma_1 + \gamma_X X_t + \gamma_A A_{t-1} + \gamma_I I_{t-1}}}{1 + e^{\gamma_1 + \gamma_X X_t + \gamma_A A_{t-1} + \gamma_I I_{t-1}}},$$

where  $(\gamma_1, \gamma_X, \gamma_A, \gamma_I) = (-2.5, 0, 4, 0.001)$ ;

4. simulate an infection process  $I_t$  from a negative binomial distribution with “number of successes” parameter  $\nu = 10$  and mean parameter specified in Eq. (6) – the mean of the semi-mechanistic model with generating distribution  $g$  in Fig. 3 and

$$R(\bar{A}_t, \beta) = \frac{K}{1 + \exp(\beta_1 + \beta_U U_t + \beta_X X_t + \beta_A A_t)},$$

where  $K = 6.5$  is the maximum transmission rate. We consider 21 linearly spaced values of  $\beta_A$  in  $[-1, 0]$ , and for each value, we set  $(\beta_1, \beta_U, \beta_X) = (-\log(5.5) + 0.5 - \beta_A/2, 0.3, 0)$ , which has the effect of keeping  $R(\bar{A}_t, \beta)$  in the same ballpark for all values of  $\beta_A$ .

Because  $g$  has substantial support on up to 40 time points, we seed the time series as follows:  $I_t = 0$  for  $t \leq -40$  and  $I_t = e^\mu$  for  $t = -39, \dots, 0$ , where  $\mu = \log(100)$  ([Bhatt et al., 2022](#); [Bong et al., 2023](#)). (An alternative would be to simulate infections prior to  $t = 1$ , but we chose to proceed as we would with observed data.)

5. Finally, we simulate an observed time series  $Y_t$  – e.g. cases or deaths – according to Eq. (6) with  $\alpha_t = 1$  and  $\pi_t = \mathbf{1}\{t = 1\}$ , for simplicity, so that  $Y_t = \mathbb{E}[Y_t \mid \bar{I}_{t-1}, \bar{Y}_{t-1}, \bar{A}_t] \equiv I_{t-1}$  for all  $t$ .

For each  $\beta_A$ , we simulated 200 times series  $Y_t$  from the model, and for each, we obtained the MLEs of the  $\beta$ 's using the package `freqepid`, assuming  $Y_t$  negative binomial with semi-mechanistic model mean (Eq. (6)) and reproduction number  $R(\bar{A}_t, \beta) = \frac{K}{1 + \exp(\beta_1 + \beta_X X_t + \beta_A A_t)}$ . Fig. 4(a) shows the averages of the 200 MLEs of  $\beta_A$  for each true  $\beta_A$ , along with 95% confidence intervals. There is phantom bias.

This bias is absent in the results under the same simulation conditions, except with  $\beta_U$  set to zero, as shown in Fig. 7. This observation confirms that the bias in Fig. 4(a) is due to the presence of phantom variables.

Next, we estimated  $\beta_A$  using the estimating equation (Eq. (12)) from the same simulation. For this, we employed the multiplicative semi-mechanistic model with the rate function  $R(\bar{A}_t, \beta) = \frac{K}{1 + \exp(\beta_1 + \beta_A A_t)}$ . This model has closed-form expressions for both  $\psi_\theta(\bar{a}_t)$  and its derivative  $\nabla_\beta \psi_\theta(\bar{a}_t)$ , as derived in Examples Example 1 and Example 2. Using

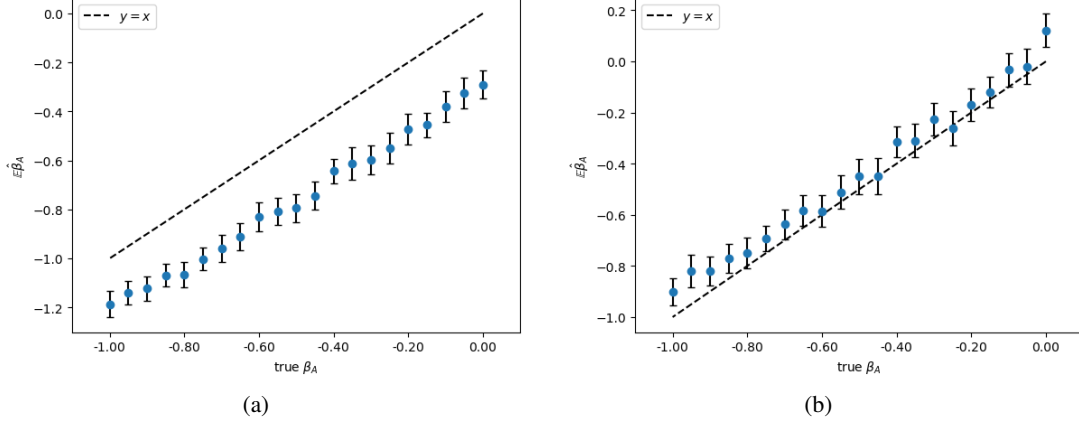


Figure 4: **(Left)** ML estimate of  $\beta_A$  averaged across 200 repeat simulations (blue dots) with 95%-confidence intervals (error bars), for a range of true  $\beta_A$  values. There is phantom bias. **(Right)** Corresponding estimating equation estimates. We used the estimating equation implied by the semi-mechanistic model (Eq. (9)).

these results, we solved the estimating equation via Newton’s method. To address confounding, we modeled the propensity score for  $A_t$  using a logistic regression based on  $A_{t-1}$ ,  $I_{t-1}$ , and  $X_t$ . From 200 i.i.d. estimates for each  $\beta_A$ , we computed 95% confidence intervals, and Fig. 4(b) presents the results for 21 different true values of  $\beta_A$ . These results demonstrate that the estimates obtained from the estimating equations are unbiased, even in the presence of phantom variables.

## 6.2 SEIR model simulated data

We demonstrate the same using the SEIR model in Eq. (2). First, we set  $I_0 = 100$ ,  $E_0 = 0$  and  $S_0 = N - I_0 - E_0$ , where the total population  $N$  was set 100,000. Then for  $t = 1, \dots, 120$ , we

1. simulate phantoms  $U_t$  from the same Gaussian random process as the previous example;
2. sample confounders  $X_t$  from a Gaussian distribution with variance  $\sigma^2$  and mean  $\epsilon_I + \zeta_U U_t + \zeta_X X_{t-1} + \zeta_A A_{t-1} + \zeta_Y Y_{t-1}$ , where  $(\epsilon_I, \zeta_U, \zeta_X, \zeta_A, \zeta_Y) = (0, 0.5, 0, 1, 0)$  and  $\sigma^2 = 0.09$ ;
3. generate binary interventions  $A_t$  from a Bernoulli distribution with

$$\mathbb{P}(A_t = 1 \mid \bar{X}_t, \bar{A}_{t-1}, \bar{I}_{t-1}) = \frac{e^{\gamma_1 + \gamma_X X_t + \gamma_A A_{t-1} + \gamma_Y Y_{t-1}}}{1 + e^{\gamma_1 + \gamma_X X_t + \gamma_A A_{t-1} + \gamma_Y Y_{t-1}}},$$

where  $(\gamma_1, \gamma_X, \gamma_A, \gamma_D) = (-2.5, 0, 4, 100/N)$ ;

4. simulate an exposure process  $B_t$  from a binomial distribution with number of trials  $S_{t-1}$  and success probability

$$p_{B,t} = 1 - \exp(-\eta_t I_{t-1}/N),$$

where  $\eta_t = \exp(-\beta_1 - \beta_U U_t - \beta_X X_t - \beta_A A_t - \beta_Y Y_{t-1})$ . We consider 21 linearly spaced values of  $\beta_A$  in  $[-1, 0]$ , and for each value, we set  $(\beta_1, \beta_U, \beta_X) = (1 - \beta_A/2, 0.3, 0)$ , which has the effect of keeping  $\eta_t$  in the same ballpark for all values of  $\beta_A$ .

Then,  $S_t = S_{t-1} - B_t$ .

5. simulate an infection process  $C_t$  from a binomial distribution with number of trials  $E_{t-1}$  and success probability  $p_C = 0.2$ .

Then,  $E_t = E_{t-1} + B_t - C_t$ .

6. Finally, we simulate an observed time series  $Y_t$  from a binomial distribution with number of trials  $I_{t-1}$  and success probability  $p_Y = 0.2$ .

Then,  $I_t = I_{t-1} + C_t - Y_t$ .

For each value of  $\beta_A$ , we simulated 200 time series  $Y_t$  from the SEIR model. Since we do not have a developed method for estimating the MLE in this setting, we instead used a regressive approach to estimate the  $\beta$  parameters. Specifically, for each iteration, we fit a binomial regression model:

$$B_t \sim \text{Binomial}(S_{t-1}, \exp(-\theta_1 - \theta_X X_t - \theta_A A_t + \log(I_{t-1}/N))),$$




Figure 5: **(Left)** ML estimate of  $\beta_A$  averaged across 200 repeat simulations (blue dots) with 95%-confidence intervals (error bars), for a range of true  $\beta_A$  values. There is phantom bias. **(Right)** Corresponding estimating equation estimates. We used the estimating equation implied by the SEIR model (Eq. (2)).

where  $\log(I_{t-1}/N)$  was included as a log-offset in the regression. This approximates the generation of the exposure process by using that  $1 - \exp(-\eta_t I_{t-1}/N) \approx \eta_t I_{t-1}/N$  when the right hand side is sufficiently small. Fig. 5(a) shows the average regressive estimates of  $\beta_A$  over the 200 simulations for each true value of  $\beta_A$ , along with the corresponding 95% confidence intervals. A phantom bias is evident in the estimates.

Next, we estimated  $\beta_A$  using the estimating equation approach (Eq. (12)) on the same simulated data. In this approach, the SEIR model was parameterized as  $\eta_t = \exp(-\beta_1 - \beta_A A_t)$ . Since this model lacks a closed-form expression for  $\psi_\theta(\bar{a}_t)$  and its derivative  $\frac{\partial}{\partial \beta} \psi_\theta(\bar{a}_t)$ , we employed Monte Carlo approximations to compute them as described in Eq. (11) and Example 3, and solved the estimating equation using Newton’s method. To account for confounding, we modeled the propensity score of  $A_t$  using a logistic regression model based on  $A_{t-1}$ ,  $I_{t-1}$ , and  $X_t$ . The 95% confidence intervals were derived from 200 independent estimates for each  $\beta_A$ . Fig. 5(b) presents the results for 21 true  $\beta_A$  values, showing that the estimates from the estimating equations are unbiased, even in the presence of phantom variables.

### 6.3 We Estimate the Causal Effect of Mobility on Covid-19 Transmission

We analyzed the effect of a mobility measure on Covid-19 death data for U.S. states, using the dataset described in Bong et al. (2023). The data are sourced from the Delphi repository at Carnegie Mellon University ([delphi.cmu.edu](https://delphi.cmu.edu)), and consist of daily observations from February 15 to August 1, 2020 (168 days). The dataset includes state-level records of Covid-19 deaths, denoted as  $Y_t$ , and a mobility measure, “proportion of full-time work” ( $A_t$ ), which represents the fraction of mobile devices that spent more than six hours at a location other than their home during daytime (using SafeGraph’s `full_time_work_prop`). We focused on the 30 states that reported more than 20 deaths on at least one day and truncated the time series 30 days prior to reaching a total of 10 accumulated deaths, following the procedure outlined in Bhatt et al. (2022). A preprocessing step was used to correct for the weekend effect, which shows fewer deaths reported on Saturdays and Sundays (see Bong et al. (2023) for further details).

Figure 6(a) presents the estimates of  $\beta_A$  from the estimating equation in Eq. (12), assuming the semi-parametric model (Eq. (6)) for the 30 states. The faint thick lines show the point estimates and 95% confidence intervals, calculated separately for each state’s data. These estimates can be improved by borrowing strength across states using a frequentist approach based on the robust empirical Bayes shrinkage method introduced by Armstrong et al. (2022), and extended to multivariate parameters by Bong et al. (2023). The dark thin lines represent the shrinkage-adjusted estimates. We also compare these results with the maximum likelihood (ML) estimates from Bong et al. (2023), based on the same semi-parametric model, assuming negative binomial (NB) distributed death counts. All 30 states show differences in the estimates within the 95% confidence intervals between the estimating equation and ML approaches. Specifically, in 24 out of 30 states, the estimates from the estimating equation were lower than those from ML. This result is statistically significant, assuming a binomial probability of 0.5 for the two methods yielding smaller estimates equally across all states ( $p < 0.001$ ). This suggests the potential presence of a phantom effect, where the ML estimates may overestimate the global causal effect of mobility on Covid-19 deaths.




Figure 6: Estimates and confidence intervals for  $\beta_A$  in Eq. (7) for 30 states using (a) estimating equation and (b) ML. The faint thick lines are the estimates and intervals before shrinkage and the dark thin lines after shrinkage.

## 7 Conclusion

Augmenting an epidemic model to include an intervention variable is quite common. We have seen that such a model is best viewed as a model for counterfactual random variables. This makes it possible to use the model to answer causal questions about the effect of an intervention on an epidemic. We have argued that the best way to estimate the parameters in such a model is to use an appropriate estimating equation. Maximum likelihood and Bayesian methods are less useful as they are susceptible to certain biases due to latent variables. In particular, these methods can lead to rejecting the null hypothesis that there is no causal effect, even when the null is true.

There are many other subtleties involved in causal inference in the context of modeling epidemics. Yet the literature on combining methods of modern causal inference with epidemic modeling is sparse. This paper is meant as a first step in this direction.

## References

- Donald WK Andrews. Laws of large numbers for dependent non-identically distributed random variables. *Econometric theory*, 4(3):458–467, 1988.
- Donald WK Andrews. An empirical process central limit theorem for dependent non-identically distributed random variables. *Journal of Multivariate Analysis*, 38(2):187–203, 1991.
- Timothy B Armstrong, Michal Kolesár, and Mikkel Plagborg-Møller. Robust empirical bayes confidence intervals. *Econometrica*, 90(6):2567–2602, 2022.
- Lucia Babino, Andrea Rotnitzky, and James Robins. Multiple robust estimation of marginal structural mean models for unconstrained outcomes. *Biometrics*, 75(1):90–99, 2019.
- Stephen Bates, Edward Kennedy, Robert Tibshirani, Valerie Ventura, and Larry Wasserman. Causal inference with orthogonalized regression: Taming the phantom. *arXiv preprint arXiv:2201.13451*, 2022.
- Samir Bhatt, Neil Ferguson, Seth Flaxman, Axel Gandy, Swapnil Mishra, and James A Scott. Semi-mechanistic bayesian modeling of covid-19 with renewal processes. *Journal of the Royal Statistical Society, Series A*, 2022.
- Heejong Bong, Valérie Ventura, and Larry Wasserman. Frequentist inference for semi-mechanistic epidemic models with interventions. *arXiv preprint arXiv:2309.10792*, 2023.
- Matteo Bonvini, Edward H. Kennedy, Valerie Ventura, and Larry Wasserman. Causal inference for the effect of mobility on COVID-19 deaths. *The Annals of Applied Statistics*, 16(4):2458 – 2480, 2022. doi: 10.1214/22-AOAS1599. URL <https://doi.org/10.1214/22-AOAS1599>.

- Robert M De Jong. Central limit theorems for dependent heterogeneous random variables. *Econometric Theory*, 13(3): 353–367, 1997.
- Robert M De Jong and James Davidson. Consistency of kernel estimators of heteroscedastic and autocorrelated covariance matrices. *Econometrica*, 68(2):407–423, 2000.
- Robin J Evans and Vanessa Didelez. Parameterizing and simulating from causal models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(3):535–568, 2024.
- Gavin J Gibson and Eric Renshaw. Estimating parameters in stochastic compartmental models using markov chain methods. *Mathematical Medicine and Biology: A Journal of the IMA*, 15(1):19–40, 1998.
- Christian Gourieroux and Alain Monfort. Simulation-based inference: A survey with special reference to panel data models. *Journal of Econometrics*, 59(1-2):5–33, 1993.
- William Ogilvy Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.
- Phenyo E Lekone and Bärbel F Finkenstädt. Statistical inference in a stochastic epidemic seir model with control intervention: Ebola as a case study. *Biometrics*, 62(4):1170–1177, 2006.
- Charles J Mode and Candace K Sleeman. *Stochastic processes in epidemiology: HIV/AIDS, other infectious diseases and computers*. World Scientific, 2000.
- Whitney K Newey, Kenneth D West, et al. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.
- James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986.
- James M Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology, the environment, and clinical trials*, pages 95–133. Springer, 2000.
- James M Robins and L Wasserman. Estimation of effects of sequential treatments by reparameterizing directed acyclic graphs. In *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, pages 409–420. Morgan Kaufmann, 1997a.
- James M Robins and Larry Wasserman. Estimation of effects of sequential treatments by reparameterizing directed acyclic graphs. In *Proceedings of the Thirteenth conference on Uncertainty in artificial intelligence*, pages 409–420, 1997b.

## Appendix




Figure 7: **Point estimates (blue dots) and 95%-confidence intervals (error bars) of  $\beta_A$  from the ML estimates without phantom variables .** For each true  $\beta_A$  value, the point estimate and confidence interval were obtained from 200 i.i.d. estimates  $\hat{\beta}_A$ . The ML estimates are unbiased when phantom variables are absent.