Published in final edited form as:

Epidemiology. 2014 January ; 25(1): 134–138. doi:10.1097/EDE.0000000000000003.

Estimating the per-exposure effect of infectious disease interventions

Justin J. O'Hagan^{1,2}, Marc Lipsitch^{1,2,3}, and Miguel A. Hernán¹

¹Department of Epidemiology, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115, USA

²Center for Communicable Disease Dynamics, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115, USA

³Department of Immunology and Infectious Diseases, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115, USA

Abstract

The average effect of an infectious disease intervention (e.g., a vaccine) varies across populations with different degrees of exposure to the pathogen. As a result, many investigators favor a per-exposure effect measure that is considered independent of the population level of exposure and that can be used in simulations to estimate the total disease burden averted by an intervention across different populations. However, while per-exposure effects are frequently estimated, the quantity of interest is often poorly defined, and assumptions in its calculation are typically left implicit. In this paper, we build upon work by Halloran and Struchiner (*Epidemiology*. 1995; 6:172-151) to develop a formal definition of the per-exposure effect and discuss conditions necessary for its unbiased estimation. With greater care paid to the parameterization of transmission models, their results can be better understood and can thereby be of greater value to decision-makers.

Estimates of the impact of vaccines, behavior modification, or other risk factors for infectious diseases are often not transportable from the study population to other populations, even when the results are internally valid and the distribution of effect modifiers is the same across populations. This is because a person's risk of infection depends on the intervention status of others in their contact network, a phenomenon that is referred to as “interference” in the causal inference literature and is related to the concepts of “herd immunity”, “indirect effects”, and “dependent happenings” as discussed by infectious disease epidemiologists.^{1–3} Transmission-dynamic models, which are increasingly used by policy-makers,⁴ attempt to solve this problem by using simulations to predict how an intervention would affect an infection's spread in a new population. To do so, such models require the effect of interventions on risk of infection and infectiousness, and information on the contact network in the new population. Specifically, these simulations require as an input parameter an estimate of the effect of interventions on exposure to the pathogen of interest (a “per-exposure effect”) — a difficult quantity to estimate in empirical studies when exposure data are generally missing or limited.

The definition and estimation of per-exposure effects have received little attention in the causal inference literature. As a result, the assumptions made when estimating per-exposure

Corresponding author: Justin J. O'Hagan, Kresge 506, Harvard School of Public Health, 677 Huntington Avenue, Boston MA 02115, USA, Phone: +1-617-888-4599, jjo058@mail.harvard.edu.

Conflicts of Interest: The authors report no conflicts of interest.

effects in practice are often unclear, and transmission models are commonly parameterized with dubious estimates of per-exposure effects.⁵⁻⁷ For example, modelers assume, often implicitly, that the per-exposure effect is constant over time, that the average per-exposure effect does not depend on the number of previous exposures to the pathogen, that there is no unmeasured heterogeneity in infection risk across study participants, and that there is no confounding for exposure to infection.

Halloran and Struchiner⁸ showed that defining causal effects for infectious diseases can involve the intervention of interest (e.g., a vaccine), as well as an intervention on exposure to infection. The authors also suggested that per-exposure effects require further study. The present paper builds upon their work to (1) develop a formal definition of the per-exposure effect, (2) show that estimates of the per-exposure effect in most studies are obtained under unrealistic assumptions, (3) discuss previously presented methods that relax some of the assumptions, and (4) discuss how even valid per-exposure effect estimates may not be suitable for parameterizing transmission models. The application of causal inference principles to the parameterization of infectious disease models bridges two methodologic approaches that have little overlap in the training of epidemiologists.⁵

The average causal effect

To motivate later discussion of a per-exposure effect, we first define the average causal effect of an intervention in an “ideal” double-blind randomized controlled trial (RCT) with perfect adherence, no loss to follow-up, and no measurement error. For each trial participant, let A be an indicator for intervention (e.g., 1, vaccinated; 0, unvaccinated) at time $t=0$ (i.e., baseline), and Y_t an indicator for infection (1, yes; 0, no) at t time units (e.g., months) after baseline. The trial is restricted to persons who are not infected at baseline, i.e., $Y_0=0$. We define $\Pr[Y_t=1|A=a]$ as the proportion of participants who became infected by t among those receiving the intervention value $A=a$. The intent-to-treat analysis compares $\Pr[Y_t=1|A=1]$ versus $\Pr[Y_t=1|A=0]$, which is equivalent to a per-protocol analysis in our example since we assume perfect adherence.

Let $Y_t^{a=1}$ and $Y_t^{a=0}$ be counterfactual outcomes indicating a person's infection status at time t had their intervention status been set to 1 (vaccinated) or 0 (unvaccinated), respectively. Note that this representation of a person's counterfactual outcome makes no reference to the intervention status of others in the population. We define $\Pr[Y_t^{a=1}=1]$ as the proportion of participants who would have become infected by t had all participants in the study received the intervention value a .

The average causal effect of A on Y_t involves a contrast between $\Pr[Y_t^{a=1}=1]$ and $\Pr[Y_t^{a=0}=1]$. For example, the average causal effect measured on the risk ratio scale is

$\frac{\Pr[Y_t^{a=1}=1]}{\Pr[Y_t^{a=0}=1]}$. For noncommunicable diseases, the causal risk ratio is expected to equal the

intent-to-treat risk ratio $\frac{\Pr[Y_t=1|A=1]}{\Pr[Y_t=1|A=0]}$ in ideal RCTs, but this is not true for infectious diseases. Here, the outcome of an index participant may depend on the intervention received by others because others' being vaccinated may prevent them from becoming infected and thus transmitting infection to the index person.⁷ Interference would occur if, for example, a vaccinated trial participant would escape infection if 100% of the study population were vaccinated, but would become infected if not all her contacts were vaccinated. Such

interference implies that the intent-to-treat risk ratio $\frac{\Pr[Y_t=1|A=1]}{\Pr[Y_t=1|A=0]}$ is not expected to equal

the causal risk ratio $\frac{\Pr[Y_t^{a=1}=1]}{\Pr[Y_t^{a=0}=1]}$. Moreover, the intent-to-treat risk ratio may not be transportable to other populations because interference depends on the specific network of contacts in a population.

These transportability problems can, in principle, be overcome through the use of a per-exposure effect measure.^{8,9} We first propose a definition of a per-exposure effect based on Halloran and Struchiner's work,⁸ and then describe the assumptions commonly made to estimate it.

The average per-exposure effect

To define the average per-exposure effect of an intervention we must introduce notation for exposure to the infectious agent. Let E_t be the number of exposures to infection experienced at time t (e.g., $E_t=1$ denotes that the person experienced a single exposure to infection at time t). For simplicity, we assume that all exposures are discrete and equivalent. We use overbars to denote history; for example, $\bar{e}_t = (0, 0, \dots, 0)$ indicates no exposure to the infectious agent between time 0 and t . Because exposure to the infectious agent is a necessary condition for infection, Y_{t+1} is always zero in the absence of exposure through t , i.e., $Y_{t+1}=0$ if $\bar{e}_t=0$ for any value of t . In practice, the time t should be chosen according to the question of interest and to ensure that a sufficient time has elapsed for the diagnostic test to discriminate between infected and uninfected persons.

We define $\Pr[Y_{t+1}^{a,\bar{e}_t}=1]$ as the proportion of participants who would have become infected by t had everyone received the intervention value a and exposure history $\bar{e}_t$ through time t .

For example, $\Pr[Y_{t+1}^{a=1,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1]$ is the proportion of people that would have been infected if everyone (1) had been vaccinated at baseline, (2) had been exposed to infection at time t , and (3) had not been exposed to infection between time 0 and $t-1$. This quantity can be equivalently written as the conditional probability

$\Pr[Y_{t+1}^{a=1,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1 | Y_t^{a=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=0]$, which is a discrete-time hazard. Similarly, $\Pr[Y_{t+1}^{a=0,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1 | Y_t^{a=0,\bar{e}_{t-1}=\bar{0}_{t-1}}=0]$ is the hazard if nobody had been vaccinated.

The hazard ratio (HR) $\frac{\Pr[Y_{t+1}^{a=1,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1 | Y_t^{a=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=0]}{\Pr[Y_{t+1}^{a=0,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1 | Y_t^{a=0,\bar{e}_{t-1}=\bar{0}_{t-1}}=0]}$ measures the effect at the first exposure after baseline. We define the average per-exposure effect of an intervention at time t as

$$\left(1 - \frac{\Pr[Y_{t+1}^{a=1,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1 | Y_t^{a=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=0]}{\Pr[Y_{t+1}^{a=0,e_t=1,\bar{e}_{t-1}=\bar{0}_{t-1}}=1 | Y_t^{a=0,\bar{e}_{t-1}=\bar{0}_{t-1}}=0]} \right)$$

Unlike the average causal effect described in the previous section, the per-exposure effect is independent of both the distribution of exposure after baseline and the pattern of interference in the study population. This is because the per-exposure effect is defined under a hypothetical exposure intervention in which everybody gets exposed at time t only. On the other hand, like the average causal effect of the previous section, the per-exposure effect may vary over time t (e.g., a vaccine's per-exposure effect may wane) and may depend on

population-specific baseline characteristics that modify the infection risk upon exposure (e.g., age or proportion of people exposed to the infection before baseline).

Our notation makes explicit Halloran and Struchiner's paradigm⁸ that a per-exposure effect measure requires considering a hypothetical joint intervention on treatment A and exposure E_t during the follow-up, and not just an intervention on A conditional on E_t . We elected to define the average per-exposure effect as the effect at the first exposure after baseline, in order to simplify the example and because the literature on intervention effects conditional upon exposure has commonly used a single exposure as the relevant unit.⁹ However, researchers could estimate the effect of an intervention after any specified number of exposures where the timing of each exposure is specified.

Estimating the average per-exposure effect

A valid estimate of a per-exposure intervention effect at time t could be obtained if the investigator controlled not only the intervention A at baseline, but also the exposure E_t during the follow-up. Studies in which investigators expose study participants to the pathogen of interest, known as “challenge studies,” are rare, and their use in humans is restricted to treatable or non-life-threatening infections.^{10–15} In most studies, people are not deliberately exposed to infection and participants can experience heterogeneous exposure patterns (many participants may not be exposed at any time). Furthermore, exposure to the pathogen of interest is typically unknown to the investigators.

Yet, despite exposures to the pathogen being unknown, per-exposure effect measures are estimated in practice. A common estimator is one minus the intention-to-treat average hazard ratio,¹⁶ which is a weighted average of the time-specific hazard ratios

$\left(\frac{\Pr[Y_{t+1}=1|Y_t=0, A=1]}{\Pr[Y_{t+1}=1|Y_t=0, A=0]} \right)$ across all times t , and is often estimated in practice via a Cox model. The following conditions are required for this estimator to provide an unbiased estimate of the per-exposure effect measure:

Condition 1: The average per-exposure effect is constant during the follow-up for all t

Condition 2: The average per-exposure effect does not depend on the number of previous exposures to the pathogen during the follow-up

Condition 1 could be easily relaxed if one minus the time-specific hazard ratio at time t , i.e.,

$1 - \frac{\Pr[Y_{t+1}=1|Y_t=0, A=1]}{\Pr[Y_{t+1}=1|Y_t=0, A=0]}$ were a valid estimator of the per-exposure effect measure at time t . In that case, we would need to calculate only the time-specific hazard ratio rather than a weighted average over time points. Condition 2, on the other hand, cannot be easily relaxed because a person's exposure history (E_{t-1}) is usually unknown, and thus one cannot estimate the per-exposure effect by restricting the calculations at each time t to those not

exposed before t , i.e., $\left(1 - \frac{\Pr[Y_{t+1}=1|Y_t=0, A=1, E_t=1, \bar{E}_{t-1}=\bar{0}_{t-1}]}{\Pr[Y_{t+1}=1|Y_t=0, A=0, E_t=1, \bar{E}_{t-1}=\bar{0}_{t-1}]} \right)$. Therefore, a valid estimation of the per-exposure effect measure is often unattainable.

However, the situation is even more problematic. Even if all participants' exposure histories were known, two additional conditions are required to estimate the average per-exposure effect using the intent-to-treat hazard ratio without bias. In the following sections we examine these two conditions using causal diagrams.

Condition3: No unmeasured heterogeneity in infection risk

As an example, Figure 1 depicts a double-blind RCT of a Chlamydia vaccine in women. Let A be an indicator for vaccination at baseline and Y_{t+1} an indicator for Chlamydia infection at $t+1$. E_t is the number of exposures to infection experienced at t , U represents the unmeasured factors that are responsible for the differential susceptibility to infection (e.g., being a sex worker increases the risk of bacterial vaginosis, a known risk factor for Chlamydia infection¹⁷). Because participants were blinded to their vaccine status, there is no arrow from A to exposure. Because A was randomly assigned, there is no arrow from U to A , i.e., no confounding for the effect of A .

Consider the per-exposure effect at time $t+2$. The hazard at $t+2$ is the probability of infection at $t+2$ among those who were not infected at $t+1$. We denote this conditioning on $Y_{t+1}=0$ by placing a box around Y_{t+1} in Figure 1. Because Y_{t+1} is a common effect of A and U , conditioning on $Y_{t+1}=0$ will generally create an association between the intervention A and Y_{t+2} by opening the path $A \rightarrow \boxed{Y_{t+1}} \leftarrow U \rightarrow Y_{t+2}$. That is, the A - Y_{t+2} association is a mixture of the causal effect of A on Y_{t+2} , and the selection bias induced by conditioning on Y_{t+1} .

As a result, the time-specific hazard ratio (HR) $\frac{\Pr[Y_{t+2}=1|Y_{t+1}=0, A=1]}{\Pr[Y_{t+2}=1|Y_{t+1}=0, A=0]}$ may differ from 1 even if A has no direct effect on Y_{t+2} ,^{18,19} and the quantity $\left(1 - \frac{\Pr[Y_{t+2}=1|Y_{t+1}=0, A=1]}{\Pr[Y_{t+2}=1|Y_{t+1}=0, A=0]}\right)$ is generally a biased estimator of the per-exposure effect at time $t+2$.

The bias described here is a form of the built-in selection bias in hazards analyses.^{18–20} The impact of this bias needs to be more widely appreciated in the infectious disease literature, and its impact on per-exposure effect estimates needs to receive greater attention.^{21–23} While there has been some treatment of this issue from a statistical perspective,^{21,24,25} the causal structure of this bias for per-exposure effects has not been described.

Condition 4: No confounding for the effect of exposure

The per-exposure effect is a joint effect of A (e.g., vaccination) and E_t (exposure to infection), and therefore its unbiased estimation requires no unmeasured confounding for the effect of both A and E_t at all times t . In an RCT no confounding is expected for A , which is the randomized intervention, but there may be unmeasured confounding for E_t , which is not controlled by the investigator. Figure 2 is the same as Figure 1 except that there are arrows from the unmeasured U to the exposure variables, i.e., there is confounding due to U for the effect of E_t on Y_{t+1} .

This confounding may result in a biased estimate of per-exposure effect when using the time-specific HRs, but also when using the average HR: because the average HR is a weighted average of the time-specific HRs, the average HR over certain periods may still differ from 1 even if the average per-exposure effect is null.

Interestingly, confounding by U may occur even if there is no direct arrow from U to Y_{t+1} in the absence of vaccination. For example, suppose that, in the absence of vaccination, sex workers and other people have the same per-exposure risk of infection but, in the presence of vaccination, the per-exposure risk is higher among sex workers (e.g., due to genital inflammation, which is more common among sex workers and which could inhibit the immune response to the vaccine). In this setting of “co-action” between U and E_t , the causal diagram still includes a direct arrow from U to Y_{t+1} and U needs to be adjusted for when estimating the per-exposure effect in the study population.²⁶ The eAppendix presents a

numerical example of this bias. Notably, no confounding would arise in the presence of $U-E_t$ co-action if U were independent of exposure.⁹

Note that, unlike Struchiner and Halloran,⁹ we use here a definition of confounding that is not based on collapsibility.^{27,28}

Discussion

We have provided a definition of the average per-exposure effect of infectious disease interventions, and have described four conditions necessary for its unbiased estimation from ideal RCT data. Except for condition 1, which can be easily met by estimating time-specific effects, these conditions are unlikely to be met in practice. Condition 2 generally requires measurement of each person's exposure history during the follow-up, which is rarely, if ever, available. Conditions 3 and 4 require that, even in an ideal RCT, unbiased estimation of the per-exposure effect via an intent-to-treat HR requires the unrealistic assumption that risk factors U for infection either do not exist or are adjusted for. As a result, standard analyses will not generally provide valid estimates of the average per-exposure effect of an intervention.

Challenge studies provide investigators with precise exposure data and, in principle, are ideal for estimating per-exposure effects, but assumptions are still required. For example, Hudgens and Gilbert^{29,30} developed an analysis method for challenge study data that estimates a per-exposure effect, but it requires assumptions regarding the intervention's mechanism of action and the distribution of susceptibility to infection across persons. Since challenge studies can be used only for easily treated infections, alternative methods are needed to obtain less biased estimates from standard studies. We suggest the following methods — none of which is guaranteed to give unbiased estimates. Which approach is most appropriate for a given situation will depend on the availability of data and whether the methods' respective assumptions are met.

First, some study designs allow the opportunity to collect at least partial data on the number and timing of exposures, which can be used for a better estimation of the per-exposure effect. One example is a study of serodiscordant couples, in which one member of a sexual partnership is known to be HIV-infected.^{31,32}

Second, selection bias due to heterogeneity in infection risk may be reduced by collecting data on known risk factors for infection and adjusting for them in the analysis (e.g., as covariates in the Cox models)³³ or by restricting the analysis to more homogeneous participants (e.g., those with similar levels of known risk factors). Interestingly, enrolling a more heterogeneous population in an effort to estimate a more representative effect measure may introduce greater bias because of greater violations of conditions 2-4. When possible, crossover designs for reversible interventions (e.g., prophylactic treatment) may be less susceptible to selection bias.²²

Third, the follow-up period could be shortened to mitigate violations of Conditions 1-3. However, restricting the analysis to a follow-up shorter than time t prevents the estimation of per-exposure effects at times greater than t , which limits generalizability if the intervention's effect varies with time or number of exposures (e.g., for some pathogens exposure may confer partial immunity, even in exposed persons with no apparent infection^{34,35}).

Fourth, mathematical models could be used to assess the plausibility of possible per-exposure effect estimates. For example, the estimates could be used to assess how the model fits to data of an infection's prevalence over time.^{36,37}

Last, sensitivity analyses are suitable for describing the possible range of results when data are missing on potentially important sources of bias. These sensitivity analyses range from simple bias formulas to mathematical models that simulate aspects of transmission, data collection, and analysis.³⁸⁻⁴⁰

Modelers also need to consider possible interactions for the effects of exposure or the intervention, even in the absence of confounding, because the predicted population impact will differ if transmission models are parameterized with an average effect versus subgroup-specific effects.^{41,42} Estimating interactions is difficult, however, because studies are rarely powered to estimate subgroup effects and because many interaction variables may be unknown.

Observational studies and randomized trials are rarely designed with mathematical models of disease transmission in mind, despite the growing importance of such models in understanding infectious disease data and informing policy-making. It would be useful if the study planning process included consideration of data needs for transmission modeling, such as collecting additional data on risk factors and interaction variables, and powering studies to allow testing for important differences in effects across subgroups. Such additional data collection could greatly enhance the utility of empirical studies to inform transmission models.

In summary, the common practice of parameterizing transmission models with standard estimates from epidemiologic studies or randomized trials should be re-evaluated. Simulation studies are warranted to characterize the direction and magnitude of the biases described here. When such estimates are used to parameterize cost-effectiveness, transmission-dynamic, or other parameters used for decision-making, the results need to be interpreted cautiously, and the range of sensitivity analyses expanded to encompass not just the statistical uncertainty in the estimates, but also their possible bias. Acknowledging and exploring such assumptions should be seen as a strength of a study, not a limitation.

Acknowledgments

Sources of Funding: This work was supported by Grant Number U54GM088558 from the National Institute of General Medical Sciences (JOH and ML) and National Institutes of Health grant R01 AI102634 (MAH).

References

1. Halloran ME, Haber M, Longini IM, Struchiner CJ. Direct and indirect effects in vaccine efficacy and effectiveness. *Am J Epidemiol.* 1991; 133:323–331. [PubMed: 1899778]
2. Hudgens MG, Halloran ME. Toward Causal Inference With Interference. *J Am Stat Assoc.* 2008; 103:832–842. [PubMed: 19081744]
3. Tchetgen EJT, VanderWeele TJ. On causal inference in the presence of interference. *Stat Methods Med Res.* 2012; 21:55–75. [PubMed: 21068053]
4. Garnett GP, Cousens S, Hallett TB, Steketee R, Walker N. Mathematical models in the evaluation of health programmes. *Lancet.* 2011; 378:515–525. [PubMed: 21481448]
5. Bellan SE, Pulliam JRC, Scott JC, Dushoff J. How to make epidemiological training infectious. *PLoS Biol.* 2012; 10:e1001295. [PubMed: 22509129]
6. Grad YH, Miller JC, Lipsitch M. Cholera modeling: challenges to quantitative analysis and predicting the impact of interventions. *Epidemiology.* 2012; 23:523–530. [PubMed: 22659546]
7. Pitzer VE, Basta NE. Linking data and models: the importance of statistical analyses to inform models for the transmission dynamics of infections. *Epidemiology.* 2012; 23:520–522. [PubMed: 22659545]
8. Halloran ME, Struchiner CJ. Causal inference in infectious diseases. *Epidemiology.* 1995; 6:142–151. [PubMed: 7742400]

9. Struchiner CJ, Halloran ME. Randomization and baseline transmission in vaccine field trials. *Epidemiol Infect.* 2007; 135:181–194. [PubMed: 17291359]
10. Cohen MB, Giannella RA, Bean J, et al. Randomized, controlled human challenge study of the safety, immunogenicity, and protective efficacy of a single dose of Peru-15, a live attenuated oral cholera vaccine. *Infect Immun.* 2002; 70:1965–1970. [PubMed: 11895960]
11. Roestenberg M, McCall M, Hopman J, et al. Protection against a malaria challenge by sporozoite inoculation. *N Engl J Med.* 2009; 361:468–477. [PubMed: 19641203]
12. Killingley B, Enstone J, Booy R, et al. Potential role of human challenge studies for investigation of influenza transmission. *The Lancet Infectious Diseases.* 2011; 11:879–886. [PubMed: 21798808]
13. Atmar RL, Bernstein DI, Harro CD, et al. Norovirus vaccine against experimental human Norwalk Virus illness. *N Engl J Med.* 2011; 365:2178–2187. [PubMed: 22150036]
14. Harro C, Chakraborty S, Feller A, et al. Refinement of a human challenge model for evaluation of enterotoxigenic *Escherichia coli* vaccines. *Clin Vaccine Immunol.* 2011; 18:1719–1727. [PubMed: 21852546]
15. Sun W, Eckels KH, Putnak JR, et al. Experimental Dengue Virus Challenge of Human Subjects Previously Vaccinated with Live-Attenuated Tetravalent Dengue Virus Vaccines. *J Infect Dis.* 2012
16. Halloran, ME.; Longini, IM.; Struchiner, CJ. Design and analysis of vaccine studies. Springer Verlag; 2009.
17. Gallo MF, Macaluso M, Warner L, et al. Bacterial vaginosis, gonorrhea, and chlamydial infection among women attending a sexually transmitted disease clinic: a longitudinal analysis of possible causal links. *Ann Epidemiol.* 2012; 22:213–220. [PubMed: 22192490]
18. Greenland S. Absence of confounding does not correspond to collapsibility of the rate ratio or rate difference. *Epidemiology.* 1996; 7:498–501. [PubMed: 8862980]
19. Hernán MA. The hazards of hazard ratios. *Epidemiology.* 2010; 21:13–15. [PubMed: 20010207]
20. Hernán MA, Hernández-Díaz S, Robins JM. A structural approach to selection bias. *Epidemiology.* 2004; 15:615–625. [PubMed: 15308962]
21. Halloran ME, Longini IM, Struchiner CJ. Estimability and interpretation of vaccine efficacy using frailty mixing models. *Am J Epidemiol.* 1996; 144:83–97. [PubMed: 8659489]
22. Auvert B, Sitta R, Zarca K, Mahiane SG, Pretorius C, Lissouba P. The effect of heterogeneity on HIV prevention trials. *Clinical Trials.* 2011; 8:144–154. [PubMed: 21478327]
23. O'Hagan JJ, Hernán MA, Walensky RP, Lipsitch M. Apparent declining efficacy in randomized trials: examples of the Thai RV144 HIV vaccine and South African CAPRISA 004 microbicide trials. *AIDS.* 2012; 26:123–126. [PubMed: 22045345]
24. Longini IM Jr, Halloran ME. A frailty mixture model for estimating vaccine efficacy. *Applied Statistics.* 1996:165–173.
25. Haber M, Ndikuyeze A. Estimation of individual and population vaccination effectiveness from time-to-event data. *Stat Med.* 1998; 17:2617–2623. [PubMed: 9839352]
26. Flanders WD, Johnson CY, Howards PP, Greenland S. Dependence of confounding on the target population: a modification of causal graphs to account for co-action. *Ann Epidemiol.* 2011 Sep; 21(9):698–705. [PubMed: 21737305]
27. Greenland S, Robins JM. Identifiability, exchangeability, and epidemiological confounding. *Int J Epidemiol.* 1986; 15:413–419. [PubMed: 3771081]
28. Greenland S, Robins JM, Pearl J. Confounding and collapsibility in causal inference. *Statistical Science.* 1999:29–46.
29. Hudgens MG, Gilbert PB. Assessing vaccine effects in repeated low-dose challenge experiments. *Biometrics.* 2009; 65:1223–1232. [PubMed: 19397589]
30. Hudgens MG, Gilbert PB, Mascola JR, Wu CD, Barouch DH, Self SG. Power to detect the effects of HIV vaccination in repeated low-dose challenge experiments. *J Infect Dis.* 2009; 200:609–613. [PubMed: 19591571]
31. Wawer MJ, Gray RH, Sewankambo NK, et al. Rates of HIV-1 transmission per coital act, by stage of HIV-1 infection, in Rakai, Uganda. *J Infect Dis.* 2005; 191(9):1403–9. [PubMed: 15809897]

32. Cohen MS, Chen YQ, McCauley M, et al. Prevention of HIV-1 infection with early antiretroviral therapy. *N Engl J Med*. 2011; 365(6):493–505. [PubMed: 21767103]
33. Hauck WW, Anderson S, Marcus SM. Should we adjust for covariates in nonlinear regression analyses of randomized trials? *Control Clin Trials*. 1998; 19:249–256. [PubMed: 9620808]
34. Ogunjimi B, Smits E, Hens N, et al. Exploring the impact of exposure to primary varicella in children on varicella-zoster virus immunity of parents. *Viral Immunol*. 2011; 24:151–157. [PubMed: 21449725]
35. Hasselrot K, Säberg P, Hirbod T, et al. Oral HIV-exposure elicits mucosal HIV-neutralizing antibodies in uninfected men who have sex with men. *AIDS*. 2009; 23:329–333. [PubMed: 19114861]
36. Poole D, Raftery AE. Inference for deterministic simulation models: the Bayesian melding approach. *Journal of the American Statistical Association*. 2000:1244–1255.
37. Powers KA, Ghani AC, Miller WC, et al. The role of acute and early HIV infection in the spread of HIV and implications for transmission prevention strategies in Lilongwe, Malawi: a modelling study. *Lancet*. 2011; 378:256–268. [PubMed: 21684591]
38. Vanderweele TJ, Arah OA. Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. *Epidemiology*. 2011; 22(1):42–52. [PubMed: 21052008]
39. Maire N, Aponte JJ, Ross A, et al. Modeling a field trial of the RTS,S/AS02A malaria vaccine. *Am J Trop Med Hyg*. 2006; 75(2 Suppl):104–10. [PubMed: 16931821]
40. Boily MC, Anderson RM. Human immunodeficiency virus transmission and the role of other sexually transmitted diseases. Measures of association and study design. *Sex Transm Dis*. 1996; 23(4):312–32. [PubMed: 8836026]
41. Ball F. Deterministic and stochastic epidemics with several kinds of susceptibles. *Advances in applied probability*. 1985:1–22.
42. Miller JC. A note on the derivation of epidemic final sizes. *Bull Math Biol*. 2012; 74:2125–2141. [PubMed: 22829179]

Figure 1. Causal diagram for a double-blind randomized trial of a Chlamydia vaccine A and Chlamydia infection Y . E represents exposure to infection and U unmeasured risk factors for infection. The subscripts denote time period. For simplicity, only two time periods are shown.

Figure 2.
Causal diagram for a randomized trial like that in Figure 1 except that the risk factors U affect exposure E .