On the causal interpretation of rate-change methods

van Aalst, Robertus; Thommes, Edward; Postma, Maarten; Chit, Ayman; Dahabreh, Issa J

Published in:
American Journal of Epidemiology

DOI:
10.1093/aje/kwaa122

Publication date:
2020

On the causal interpretation of rate-change methods: the prior event rate ratio and rate difference

Robertus van Aalst, Edward Thommes, Maarten Postma, Ayman Chit, and Issa J. Dahabreh

Correspondence to Mr. Robertus van Aalst, Vaccine Epidemiology and Modelling, Sanofi Pasteur, 1 Discovery Drive, Swiftwater, PA 18370 (e-mail rob.vanaalst@sanofi.com)

Author affiliations: Department of Health Sciences, University Medical Center Groningen, University of Groningen, Groningen, the Netherlands (Robertus van Aalst and Maarten Postma); Vaccine Epidemiology and Modelling, Sanofi Pasteur, Swiftwater, Pennsylvania (Robertus van Aalst, Edward Thommes, and Ayman Chit); Department of Mathematics & Statistics, University of Guelph, Guelph, Canada (Edward Thommes); Unit of PharmacoTherapy, -Epidemiology & -Economics (PTE2), University of Groningen, Department of Pharmacy, Groningen, the Netherlands (Maarten Postma); Department of Economics, Econometrics & Finance, University of Groningen, Faculty of Economics & Business, Groningen, the Netherlands (Maarten Postma); Leslie Dan Faculty of Pharmacy, University of Toronto, Toronto, Canada (Ayman Chit); Center for Evidence Synthesis in Health, Brown University School of Public Health, Providence, Rhode Island (Issa J. Dahabreh); and Departments of Health Services, Policy & Practice and Epidemiology, Brown University School of Public Health, Providence, Rhode Island (Issa J. Dahabreh).

This study was supported in part by Sanofi Pasteur.

Conflict of interest: Robertus van Aalst, Edward Thommes and Ayman Chit are employees of Sanofi Pasteur. Maarten Postma has received grants and honoraria from various pharmaceutical companies, including those developing, producing and marketing influenza vaccines (in particular GSK, Astra Zeneca, Seqirus, and Sanofi Pasteur). Issa J. Dahabreh has received consulting fees from Sanofi Pasteur for providing advice on a separate applied project that used PERR methods, but received no funding for the current study.

Running head: **Rate-change methods: PERR and PERD**

On the causal interpretation of rate-change methods: the prior event rate ratio and rate difference

Robertus van Aalst, MSc*[1-2], Edward Thommes, PhD[2-3], Maarten Postma, PhD[1,4-5], Ayman Chit, PhD[2,6], and Issa J. Dahabreh, MD, ScD[7-8]

[1]Department of Health Sciences, University Medical Center Groningen, University of Groningen, Groningen, the Netherlands

[2]Vaccine Epidemiology and Modelling, Sanofi Pasteur, Swiftwater, Pennsylvania, USA

[3]Department of Mathematics & Statistics, University of Guelph, Guelph, ON, Canada

[4]Unit of PharmacoTherapy, -Epidemiology & -Economics (PTE2), University of Groningen, Department of Pharmacy, Groningen, the Netherlands

[5]Department of Economics, Econometrics & Finance, University of Groningen, Faculty of Economics & Business, Groningen, the Netherlands

[6]Leslie Dan Faculty of Pharmacy, University of Toronto, Toronto, Ontario, Canada.

[7]Center for Evidence Synthesis in Health, Brown University School of Public Health, Providence, RI, USA

[8]Departments of Health Services, Policy & Practice and Epidemiology, Brown University School of Public Health, Providence, RI, USA

June 17, 2020

*Address for correspondence: Robertus van Aalst; Sanofi Pasteur, 1 Discovery Drive, Swiftwater, PA 18370; email: rob.vanaalst@sanofi.com

On the causal interpretation of rate-change methods: the prior event rate ratio and rate difference

Running head: Rate-change methods: PERR and PERD

Word count: abstract = 193; main text = 3,247.

Abbreviations that appear in the text: HR = hazard ratio; IRD = incidence rate difference; IRR = incidence rate ratio; PEHR = prior event ratio of hazard ratios; PERD = prior event rate difference; PERR = prior event rate ratio.

Abstract

A growing number of studies use data before and after treatment initiation in groups exposed to different treatment strategies to estimate "causal effects" using a ratio measure called the prior event rate ratio (PERR). Here, we offer a causal interpretation for PERR and its additive scale analog, the prior event rate difference (PERD). We show that causal interpretation of these measures requires untestable rate-change assumptions about the relationship between (1) the change of the counterfactual rate before and after treatment initiation in the treated group under hypothetical intervention to implement the control treatment; and (2) the change of the factual rate before and after treatment initiation in the control group. The rate-change assumption is on the multiplicative scale for PERR, but on the additive scale for PERD; the two assumptions hold simultaneously under testable, but unlikely, conditions. Even if investigators can pick the most appropriate scale, the relevant rate-change assumption may not hold exactly, so we describe sensitivity analysis methods to examine how assumption violations of different magnitudes would affect study results. We illustrate the methods using data from a published study of proton pump inhibitors and pneumonia.

Key words: rate-change methods; prior event rate ratio; prior event rate difference; observational studies; confounding; causal inference.

Introduction

Many recent pharmacoepidemiologic studies (e.g., [1–11]) collect outcome data before and after treatment initiation in two groups of individuals, each exposed to a different strategy (henceforth, we refer to these groups as the treated and control groups, even though both can be receiving active treatment; and we refer to the strategies they are exposed to as the treatment and control strategy, respectively). These studies estimate "treatment effects" on outcomes that can occur multiple times by using a measure called the *prior event rate ratio* (PERR). Though modeling details differ across applications, a common thread is that the ratio of event rates before and after treatment initiation in the control group is used as a proxy for what the ratio of event rates before and after treatment initiation would have been in the treated group, under hypothetical intervention to implement the control strategy. Because PERR relies on assumptions about the change of the rate before and after treatment we refer to it as a "rate-change" method. It is often informally claimed that a causal interpretation of PERR does not require the absence of confounding by unmeasured time-fixed variables [12, 13].

Here, we define the target causal quantity of PERR analyses, and formalize the requirements for endowing the analyses with a causal interpretation. We show how PERR analysis can be viewed as a form of "difference-in-differences" analysis [14] on the multiplicative scale. We also describe an analog of PERR on the additive scale, the *prior event rate difference* (PERD), which connects with the econometric literature on difference-in-differences methods. For both PERR and PERD analyses, we show that identification of the target causal quantities requires strong and untestable *rate-change assumptions* about the relationship between (1) the counterfactual change of the rate before and after treatment in the treated group under intervention to implement the control strategy; and (2) the factual change of the rate before and after treatment in the control group. The rate-change assumption is on the multiplicative scale for PERR, but on the additive scale for PERD. We show that these assumptions can hold simultaneously only under testable, but unlikely, conditions. Even if investigators can decide which rate-change assumption is most appropriate, that assumption is unlikely to hold exactly; to address possible violations, we describe sensitivity analysis methods that can be used to examine the degree to which violations of assumptions might affect study results. We illustrate the methods using data from a recently published study of the effect of proton pump inhibitors on pneumonia.

Study design and data

Study design

Suppose that two groups of individuals are exposed to two different treatment strategies. For example, we might want to compare outcomes among individuals in two healthcare plans that are subject to different reimbursement policies after a given date. Or, we might want to compare outcomes among individuals who meet some eligibility criteria and who receive recommendations to initiate two different treatments for the same condition. We refer to the time of policy implementation or treatment initiation as time zero. We focus on outcomes that can be assessed both before time zero (during the pretreatment period), and after time zero (during the posttreatment period).

Illustrative example: For concreteness, in the remainder of the paper we will consider a recent pharmacoepidemiologic study that examined the effect of proton pump inhibitor prescription on the risk of community acquired pneumonia [5]. The authors used UK-based primary care electronic health records to identify a treated group that received a proton pump inhibitor prescription and a matched control group that did not. We will use numerical data from that study to illustrate different methods. Our main objective is to discuss the methods in general terms; we do not take any position on the validity of this particular study.

Observed data

The study design described in the previous section provides adequate data to estimate the rate of the outcome among the treated and control groups, during the pretreatment and posttreatment periods. Specifically, for each of these periods and each of the treatment groups we observe the number of events that occurred and the person-time under follow-up for the treated and control groups, during the pretreatment and posttreatment periods (Table 1).

To introduce some notation for the population parameters underlying the data of Table 1, let $r_{\text{pre}}(A = a)$ denote the pretreatment event rate and $r_{\text{post}}(A = a)$ denote the posttreatment event rate among individuals who received treatment $A = a$ at time zero. In our example, the random variable A denotes "proton pump

inhibitor prescription" (1 if received at time zero; 0 if not received). The study design can be used to identify the population rate parameters in Table 2 and that the data in Table 1 can be used to estimate those parameters.

Causal quantities of interest

To define causal quantities of interest and state identifiability conditions, we need additional notation for counterfactual incidence rates, that is, rates that would be observed under hypothetical interventions to implement a particular treatment strategy, possibly contrary to fact [15, 16]. Let $r_{\text{post}}^{a=0}(A=1)$ be the counterfactual (potential) posttreatment event rate had we intervened to implement the control strategy $a=0$ to the treated group; $r_{\text{post}}^{a=1}(A=1)$ be the counterfactual posttreatment event rate had we intervened to implement the treatment strategy $a=1$ in the treated group; and $r_{\text{pre}}^{a=0}(A=1)$ be the counterfactual pretreatment event rate had we intervened to implement the control strategy $a=0$ in the treated group.

We are now ready to define two causal quantities of interest, both of which pertain to the individuals in the treated group (i.e., they are similar to the average treatment effect on the treated [17]). The first quantity is the causal incidence rate ratio (IRR) among the treated,

$$\text{IRR}_{\text{causal}}(A=1) \equiv \frac{r_{\text{post}}^{a=1}(A=1)}{r_{\text{post}}^{a=0}(A=1)}.$$

The second quantity is the causal incidence rate difference (IRD) among the treated,

$$\text{IRD}_{\text{causal}}(A=1) \equiv r_{\text{post}}^{a=1}(A=1) - r_{\text{post}}^{a=0}(A=1).$$

Our next task is to consider the conditions under which these causal quantities can be identified from the observed data.

Identification

Why exchangeability-based methods might not work

The most commonly used approaches [17] for identifying $\text{IRR}_{\text{causal}}(A = 1)$ and $\text{IRD}_{\text{causal}}(A = 1)$ rest on exchangeability (ignorability) assumptions between the treatment groups [18]. Specifically, these methods require that the counterfactual event rate in the treated group under intervention to implement the control strategy is equal to the factual rate in the control group, typically within strata defined by baseline covariates. That is to say, the usual approaches require that, conditional on covariates, the observed post-treatment event rate in the control group is a good proxy for the counterfactual event rate for the treated group, under intervention to implement the control strategy. This assumption is often questionable in pharmacoepidemiologic studies because it requires that baseline covariates are sufficiently information-rich to remove all confounding.

In our illustrative example, we might be suspicious of the assumption that all confounding factors are sufficiently captured in the observational data. For example, comorbid conditions were categorized on the basis of the Charlson score, obtained using diagnostic codes extracted from electronic health records. For many chronic diseases such information does not differentiate between different severity levels or reflect how well disease is controlled by treatment. Limitations like these might explain why the authors themselves considered that treatment effect estimates obtained from methods that require exchangeability of the treatment groups conditional on baseline time-fixed covariates were likely affected by residual confounding [5].

A number of recent publications suggest that rate-change methods can overcome these limitations by allowing identification of the causal quantities of interest even in the presence of confounding by unmeasured baseline variables. To our knowledge, these arguments have not been couched in explicitly causal terms, and we undertake the task in the next section.

Identification of the causal rate ratio by PERR

Identifiability conditions: The following identifiability conditions are sufficient for identifying $\text{IRR}_{\text{causal}}(A = 1)$.

1. Consistency among the actually treated: $r^{a=1}_{\text{post}}(A=1) = r_{\text{post}}(A=1)$; among the treated group, the counterfactual event rate under intervention to assign treatment is equal to the factual rate.

2. Hypothetical intervention to implement the control strategy does not affect the pretreatment event rate among the treated: $r^{a=0}_{\text{pre}}(A=1) = r_{\text{pre}}(A=1)$; the factual pretreatment event rate among the treated equals the counterfactual event rate of the same group under intervention to implement the control strategy.

3. Common rate-change assumption on the multiplicative scale:

$$\frac{r^{a=0}_{\text{post}}(A=1)}{r^{a=0}_{\text{pre}}(A=1)} = \frac{r_{\text{post}}(A=0)}{r_{\text{pre}}(A=0)};$$

the ratio of the counterfactual posttreatment and pretreatment event rates among the treated under intervention to implement the control strategy equals the ratio of the factual posttreatment and pretreatment event rates among the control group.

4. Positivity of the treatment probability: $1 > \Pr[A=1] > 0$, so that, in large samples, we observe individuals in both the treated and untreated groups.

5. Positivity of event rates: for all treatments $a \in \{0,1\}$, $r_{\text{pre}}(A=a) > 0$ and $r_{\text{post}}(A=a) > 0$.

In addition to these conditions, we assume that all subjects can be observed from the start of the pretreatment period until the end of the posttreatment period. Extensions to address identification in the presence of drop-out or competing events are possible but beyond the scope of this paper.

Reasoning about the identifiability conditions: Conditions *1* through *3*, listed above, make up the core of the PERR method and cannot be verified using observed data (i.e., they are untestable). Reasoning about the conditions requires background knowledge and can be informed by results of other studies (e.g., research about treatment preferences or the impact of time-varying factors on the outcome). Web Appendix 1.3 offers a brief discussion of potential violations of assumptions *1* through *3*.

Identification of the causal rate ratio by PERR: As we show in Web Appendix 1, under identifiability conditions *1* through *5*, the causal incidence rate ratio among the treated $\text{IRR}_{\text{causal}}(A=1)$ is identifiable by

the population PERR, defined as

$$\text{PERR} \equiv \frac{r_{\text{post}}(A=1)/r_{\text{post}}(A=0)}{r_{\text{pre}}(A=1)/r_{\text{pre}}(A=0)}. \quad (1)$$

Identification of the causal rate difference

Astute readers will have perceived the similarity between PERR and the assumptions needed for it to have a causal interpretation, and so-called "difference-in-differences" methods [14]. The connection stems from the fact that both PERR and difference-in-differences methods use the change between the posttreatment and pretreatment period in the "average" factual outcome among the control group as a proxy for the change between the posttreatment and pretreatment period in the counterfactual outcome among the treated group, under intervention to implement the control strategy. The difference is that in conventional difference-in-differences analyses the assumption is usually made on the additive scale (i.e., change over time is expressed as a difference), whereas in PERR analyses the assumption is made on the multiplicative scale (i.e., change over time is expressed as a ratio). The connection between PERR and difference-in-differences analyses suggests that with slight modification of the identifiability conditions for the causal incidence rate ratio, we should be able to identify the causal incidence rate difference.

Identifiability conditions: We replace identifiability condition *3* with the following alternative:

3. Common rate-change assumption on the additive scale:*

$$r_{\text{post}}^{a=0}(A=1) - r_{\text{pre}}^{a=0}(A=1) = r_{\text{post}}(A=0) - r_{\text{pre}}(A=0);$$

the difference between the counterfactual posttreatment and pretreatment event rate among the treated group under intervention to implement the control strategy, is equal to the difference of the factual posttreatment and pretreatment event rate among the control group.

Furthermore, we retain assumptions *1*, *2*, and *4*, but not assumption *5* (i.e., we no longer need to assume that the event rates are strictly positive).

Identification of the causal rate difference by PERD: As we show in Web Appendix 2, under conditions *1*, *2*, *3**, and *4* the causal incidence rate difference, $\text{IRD}_{\text{causal}}(A=1)$, is identified by PERD,

$$\text{PERD} \equiv \Big\{r_{\text{post}}(A=1) - r_{\text{post}}(A=0)\Big\} - \Big\{r_{\text{pre}}(A=1) - r_{\text{pre}}(A=0)\Big\}. \tag{2}$$

Identification results for the causal hazard ratio, using an identification strategy similar to that for the causal incidence rate ratio and difference, are presented in Web Appendix 3.

Estimation

Using data from Table 1 we obtain the estimated rates in Table 3. Using these estimated rates we can estimate PERR and PERD by their sample analogs:

$$\widehat{\text{PERR}} \equiv \frac{\widehat{r}_{\text{post}}(A=1)/\widehat{r}_{\text{post}}(A=0)}{\widehat{r}_{\text{pre}}(A=1)/\widehat{r}_{\text{pre}}(A=0)} = 0.89$$

and

$$\widehat{\text{PERD}} \equiv \Big\{\widehat{r}_{\text{post}}(A=1) - \widehat{r}_{\text{post}}(A=0)\Big\} - \Big\{\widehat{r}_{\text{pre}}(A=1) - \widehat{r}_{\text{pre}}(A=0)\Big\} = -3.83 \text{ events/1,000 person-years.}$$

Figure 1 is graphical representation of the identification strategies underlying PERR and PERD and a visual summary of the estimation results in the illustrative example.

Because our analyses are based on published data, we cannot obtain estimates of sampling variability that account for possible dependence between the pretreatment and posttreatment rates in the treatment and control groups. Such estimates, however, are straightforward to obtain when primary data are available, using standard estimating equation methods [19] and the delta method [20].

Should we use PERR or PERD?

At this point, it is natural to wonder how to decide between using PERR or PERD and whether it is reasonable to apply both methods on the same data as a form of stability analysis. As we show in Web Appendix 4, if the conditions that are needed to endow both PERR and PERD with a causal interpretation hold simultaneously, then it has to be that $r_{\text{post}}(A=0) = r_{\text{pre}}(A=0)$ or $r_{\text{pre}}(A=1) = r_{\text{pre}}(A=0)$. These conditions are testable using the observed data and thus if one of them is rejected by the data with high confidence, then we could infer that at least one of the conditions needed for both PERR and PERD to have a causal interpretation does not hold. Such rejection does not imply that either PERR or PERD have a causal interpretation; it only implies that at most one of PERR or PERD can have a causal interpretation in the particular application. In such cases, it is not sensible to conduct both PERR and PERD analyses on the same data. In fact, because neither of the observed data conditions, $r_{\text{post}}(A=0) = r_{\text{pre}}(A=0)$ or $r_{\text{pre}}(A=1) = r_{\text{pre}}(A=0)$, is likely to hold in practice, it will usually be necessary to rely on background knowledge to choose on which scale a rate-change condition is most likely to hold and, thus, to decide whether PERR or PERD analysis is most appropriate.

Sensitivity analysis

Even if investigators can pick the most appropriate scale to work on, the relevant rate-change assumption is unlikely to hold exactly. In other words, the validity of PERR or PERD analyses depends critically on the relevant rate-change assumption, which is not testable from the data and will often be controversial in applications. Thus, investigators should perform sensitivity analyses to examine how assumption violations of different magnitudes might affect study results.

Violation of the common rate-change assumption on the multiplicative scale

Suppose that the common rate-change assumption on the multiplicative scale does not hold, so that

$$\frac{r^{a=0}_{\text{post}}(A=1)}{r^{a=0}_{\text{pre}}(A=1)} \neq \frac{r_{\text{post}}(A=0)}{r_{\text{pre}}(A=0)}.$$

When that is the case, the identification results presented above do not hold and the PERR analysis does not have a causal interpretation. A convenient way to parameterize the violations of the assumption is to assume that

$$\frac{r^{a=0}_{\text{post}}(A=1)}{r^{a=0}_{\text{pre}}(A=1)} = \frac{r_{\text{post}}(A=0)}{r_{\text{pre}}(A=0)} \times u,$$

where u is a positive bias parameter. Under this sensitivity analysis model, we can express the causal incidence rate ratio among the treated as a function of u,

$$\text{IRR}_{\text{causal}}(u, A=1) \equiv \frac{r_{\text{post}}(A=1)/r_{\text{post}}(A=0)}{r_{\text{pre}}(A=1)/r_{\text{pre}}(A=0)} \times \frac{1}{u} = \text{PERR} \times \frac{1}{u}. \tag{3}$$

When $u = 1$, we recover the result as when the rate-change condition holds exactly; when $0 < u < 1$, PERR underestimates the causal incidence rate ratio by a factor of $1/u$ (i.e., the causal incidence rate ratio is $1/u$ times larger than PERR); when $u > 1$, PERR overestimates the causal rate ratio by the same factor.

Because u cannot be identified using the observed data, sensitivity analysis can be conducted by repeating the analyses while using a sufficiently broad set of values. Figure 2 panel A shows the results of sensitivity analysis for our illustrative example, varying u from 0.8 to 1.2 and estimating all quantities in equation (3) by their sample analogs.

Violation of the common rate-change assumption on the additive scale

Similarly, when the common rate-change assumption on the additive scale does not hold, we can write

$$r^{a=0}_{\text{post}}(A=1) - r^{a=0}_{\text{pre}}(A=1) \neq r_{\text{post}}(A=0) - r_{\text{pre}}(A=0);$$

and parameterize the violations of the assumption as

$$r^{a=0}_{\text{post}}(A=1) - r^{a=0}_{\text{pre}}(A=1) = r_{\text{post}}(A=0) - r_{\text{pre}}(A=0) - d,$$

where d is a bias parameter represents the magnitude of the assumption violation. Under this sensitivity analysis model, we can express the causal incidence rate difference as a function of d:

$$\text{IRD}_{\text{causal}}(d, A=1) \equiv \left\{r_{\text{post}}(A=1) - r_{\text{post}}(A=0)\right\} - \left\{r_{\text{pre}}(A=1) - r_{\text{pre}}(A=0)\right\} - d = \text{PERD} - d. \quad (4)$$

When $d > 0$ the PERD overestimates the causal rate difference by d; conversely, if $d < 0$ the PERD underestimates the causal rate difference by the same amount. Figure 2 panel B shows the results of sensitivity analysis for our illustrative example, varying d from -5 to 5 events per 1,000 person-years and estimating all quantities in equation (4) by their sample analogs.

Discussion

This paper describes the assumptions needed for the causal interpretation of PERR, an increasingly popular rate-change method that addresses confounding by unmeasured time-fixed covariates when the event of interest is recurring. We show that PERR can be viewed as a form of difference-in-differences analysis on the multiplicative scale and show how an analog of PERR on the additive scale corresponds to the usual difference-in-differences approach that is popular in econometrics. Interestingly, we show that adopting the assumptions needed to endow both PERR and PERD estimates with a causal interpretation, has testable implications for the observed event rates that are unlikely to hold exactly in applications. Because the rate-change assumptions for both PERR and PERD analyses are unlikely to hold simultaneously, use of such analyses in applications requires substantive knowledge about the underlying data generating mechanism to identify the scale on which a rate-change identifiability condition is likely to hold. This result relates to issues that arise in difference-in-differences analyses when using different transformations of the outcome [21]; our result involves assumptions about the relationship between the expectation of the factual and counterfactual outcomes (i.e., the event rates) on different scales, but not transformations of the outcome itself.

Even if investigators are able to select the most appropriate scale, the relevant rate-change assumption is unlikely to hold exactly so we sketch simple sensitivity analysis methods for PERR and PERD. A benefit of our approach to sensitivity analysis is that investigators need not have detailed background knowledge

about the unmeasured time-varying confounding variables or their relationship with the observed data [22]. A potential limitation is that the elicitation of the range of the bias parameters from domain experts may be challenging.

For simplicity, we assumed no censoring or competing risks, and that incidence rates were reasonable measures of incidence. These assumptions will often be tenable over short periods of time, but in practical applications more refined methods will be needed to address censoring, competing risks, or allow more flexible modeling of time. Furthermore, we did not discuss conditioning on baseline (time-fixed covariates). As shown in Web Appendices 1.4 and 2.3, however, such conditioning can be handled easily with minor modifications of our approach (in fact, the main text of our paper can be read as pertaining to a single stratum of covariates that is sufficiently narrow as to justify the required causal assumptions, but not too narrow so as to violate the positivity conditions). Most methodological literature on rate-change methods on the multiplicative scale has focused on the hazard ratio measure, rather than the incidence rate ratio measure. Conceptually, our identification analysis applies with little modification to the hazard ratio measure (or the less commonly used odds ratio measure) and, when conditioning on covariates, we avoid difficulties due to the non-collapsibility of some effect measures, such as hazard or odds ratios. Nevertheless, for completeness, we considered identification of the causal hazard ratio among the treated by the prior event hazard ratio (PEHR) using proportional hazards models (Web Appendix 3).

Our main message is that a causal interpretation of rate-change methods, even in a simplified scenario, requires strong and untestable assumptions. As is well-known [17], conventional methods that rely on conditioning on baseline covariates to address confounding, also require strong and untestable assumptions, most notably that measured covariates are adequate to render the treatment and control groups exchangeable. Arguably, interest in rate-change methods, and PERR in particular, is motivated by the belief that the variables in routinely collected data sources do not meet this stringent requirement. Whether this belief is true or not needs to be examined on a case-by-case basis. Similarly, whether PERR or PERD methods are a viable alternative depends critically on the plausibility of their respective rate-change assumptions, and thus also needs to be examined on a case-by-case basis. *A priori* preference for rate-change methods over exchangeability-based methods or vise-versa is not defensible without reference to a particular scientific

question; thus, researchers might want to consider strategies to detect potential violations of exchangeability assumptions (e.g., using negative control outcomes [23]) or rate-change assumptions (e.g., by using external data on time-varying factors that affect the outcome and may differentially affect the treated and control groups).

We hope that our results will help researchers to make informed choices when considering alternative identification strategies.

Acknowledgments

Author affiliations: Department of Health Sciences, University Medical Center Groningen, University of Groningen, Groningen, the Netherlands (Robertus van Aalst and Maarten Postma); Vaccine Epidemiology and Modelling, Sanofi Pasteur, Swiftwater, Pennsylvania (Robertus van Aalst, Edward Thommes, and Ayman Chit); Department of Mathematics & Statistics, University of Guelph, Guelph, Canada (Edward Thommes); Unit of PharmacoTherapy, -Epidemiology & -Economics (PTE2), University of Groningen, Department of Pharmacy, Groningen, the Netherlands (Maarten Postma); Department of Economics, Econometrics & Finance, University of Groningen, Faculty of Economics & Business, Groningen, the Netherlands (Maarten Postma); Leslie Dan Faculty of Pharmacy, University of Toronto, Toronto, Canada (Ayman Chit); Center for Evidence Synthesis in Health, Brown University School of Public Health, Providence, Rhode Island (Issa J. Dahabreh); and Departments of Health Services, Policy & Practice and Epidemiology, Brown University School of Public Health, Providence, Rhode Island (Issa J. Dahabreh). This study was supported in part by Sanofi Pasteur. Robertus van Aalst, Edward Thommes and Ayman Chit are employees of Sanofi Pasteur. Maarten Postma has received grants and honoraria from various pharmaceutical companies, including those developing, producing and marketing influenza vaccines (in particular GlaxoSmithKline, Astra Zeneca, Seqirus, and Sanofi Pasteur). Issa J. Dahabreh has received consulting fees from Sanofi Pasteur for providing advice on a separate applied project that used prior event rate ratio methods, but received no funding for the current study.

References

[1] Mark G Weiner, Dawei Xie, and Richard L Tannen. Replication of the scandinavian simvastatin survival study using a primary care medical record database prompted exploration of a new method to address unmeasured confounding. *Pharmacoepidemiology and Drug Safety*, 17(7):661–670, 2008.

[2] Richard L Tannen, Mark G Weiner, and Dawei Xie. Replicated studies of two randomized trials of angiotensin-converting enzyme inhibitors: further empiric validation of the 'prior event rate ratio'to adjust for unmeasured confounding by indication. *Pharmacoepidemiology and Drug Safety*, 17(7):671–685, 2008.

[3] Richard L Tannen, Mark G Weiner, and Dawei Xie. Use of primary care electronic medical record database in drug efficacy research on cardiovascular outcomes: comparison of database and randomised controlled trial findings. *BMJ*, 338:b81, 2009.

[4] Richard Tannen, Dawei Xie, Xingmei Wang, Menggang Yu, and Mark G Weiner. A new "comparative effectiveness" assessment strategy using the thin database: comparison of the cardiac complications of pioglitazone and rosiglitazone. *Pharmacoepidemiology and Drug Safety*, 22(1):86–97, 2013.

[5] Fatmah Othman, Colin J Crooks, and Timothy R Card. Community acquired pneumonia incidence before and after proton pump inhibitor prescription: population based study. *BMJ*, 355:i5813, 2016.

[6] Adam J Streeter, Nan Xuan Lin, Louise Crathorne, Marcela Haasova, Christopher Hyde, David Melzer, and William E Henley. Adjusting for unmeasured confounding in nonrandomized longitudinal studies: a methodological review. *Journal of Clinical Epidemiology*, 87:23–34, 2017.

[7] Yinong Young-Xu, Robertus Van Aalst, Salaheddin M Mahmud, Kenneth J Rothman, Julia Thornton Snider, Daniel Westreich, Vincent Mor, Stefan Gravenstein, Jason KH Lee, Edward W Thommes, et al. Relative vaccine effectiveness of high-dose versus standard-dose influenza vaccines among veterans health administration patients. *The Journal of Infectious Diseases*, 217(11):1718–1727, 2018.

[8] Michael Dennis, Laura Shine, Ann John, Amanda Marchant, Joanna McGregor, Ronan A Lyons, and Sinead Brophy. Risk of adverse outcomes for older people with dementia prescribed antipsychotic medication: a population based e-cohort study. *Neurology and Therapy*, 6(1):57–77, 2017.

[9] Jan Zirk-Sadowski, Jane A Masoli, Joao Delgado, Willie Hamilton, W David Strain, William Henley, David Melzer, and Alessandro Ble. Proton-pump inhibitors and long-term risk of community-acquired pneumonia in older adults. *Journal of the American Geriatrics Society*, 66(7):1332–1338, 2018.

[10] Richard L Tannen, Kurt T Barnhart, and Joshua C Rubin. Advantages of large medical record database for outcomes research: Insights into post-menopausal hormone therapy. *Learning Health Systems*, page e10193, 2019.

[11] Sinead Brophy, Jonathan Kennedy, Fabiola Fernandez-Gutierrez, Ann John, Robert Potter, Christine Linehan, and Michael Kerr. Characteristics of children prescribed antipsychotics: analysis of routinely collected data. *Journal of Child and Adolescent Psychopharmacology*, 28(3):180–191, 2018.

[12] M. J. Uddin, R. H. Groenwold, T. P. van Staa, A. de Boer, S. V. Belitser, A. W. Hoes, K. C. Roes, and O. H. Klungel. Performance of prior event rate ratio adjustment method in pharmacoepidemiology: a simulation study. *Pharmacoepidemiology and Drug Safety*, 24(5):468–77, 2015.

[13] Nan Xuan Lin and William Edward Henley. Prior event rate ratio adjustment for hidden confounding in observational studies of treatment effectiveness: a pairwise cox likelihood approach. *Statistics in Medicine*, 35(28):5149–5169, 2016.

[14] Joshua David Angrist, Jörn-Steffen Pischke, and Jörn-Steffen Pischke. *Mostly harmless econometrics: an empiricists companion*. Princeton University Press, Princeton, NJ, 2013.

[15] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.

[16] James M Robins and Sander Greenland. Causal inference without counterfactuals: comment. *Journal of the American Statistical Association*, 95(450):431–435, 2000.

[17] Miguel A Hernán and James M Robins. *Causal inference (forthcoming)*. Chapman & Hall/CRC, Boca Raton, FL, 2020.

[18] Sander Greenland and James M Robins. Identifiability, exchangeability, and epidemiological confounding. *International Journal of Epidemiology*, 15(3):413–419, 1986.

[19] Scott L Zeger, Kung-Yee Liang, and Paul S Albert. Models for longitudinal data: a generalized estimating equation approach. *Biometrics*, 44(4):1049–1060, 1988.

[20] George Casella and Roger L Berger. *Statistical inference*, volume 2. Duxbury Pacific Grove, CA, 2002.

[21] Paul R. Rosenbaum. *Observation and experiment: an introduction to causal inference*. Harvard University Press, Cambridge, MA, 2019.

[22] James M Robins, Andrea Rotnitzky, and Daniel O Scharfstein. Sensitivity analysis for selection bias and unmeasured confounding in missing data and causal inference models. In *Statistical models in epidemiology, the environment, and clinical trials*, pages 1–94. Springer, 2000.

[23] Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology (Cambridge, Mass.)*, 21(3):383, 2010.

Tables

Table 1: Data from the illustrative example.

Treatment group	Pretreatment		Posttreatment	
	Events	Person-years	Events	Person-years
$A = 1$	9,642	155,341	8,727	142,110
$A = 0$	4,298	157,783	4,516	148,504

Table 2: Treatment group and period-specific population rate parameters.

Treatment group	Pretreatment	Posttreatment
$A = 1$	$r_{\text{pre}}(A = 1)$	$r_{\text{post}}(A = 1)$
$A = 0$	$r_{\text{pre}}(A = 0)$	$r_{\text{post}}(A = 0)$

Table 3: Treatment group and period-specific population rate estimates (expressed as events per 1,000 person-years).

Treatment group	Pretreatment	Posttreatment
$A = 1$	$\widehat{r}_{\text{pre}}(A = 1) = 62.07$	$\widehat{r}_{\text{post}}(A = 1) = 61.41$
$A = 0$	$\widehat{r}_{\text{pre}}(A = 0) = 27.24$	$\widehat{r}_{\text{post}}(A = 0) = 30.41$

Figures

Figure 1: Graphical depiction of identification using the prior event rate ratio (PERR, panel A) and the prior event rate difference (PERD, panel B).

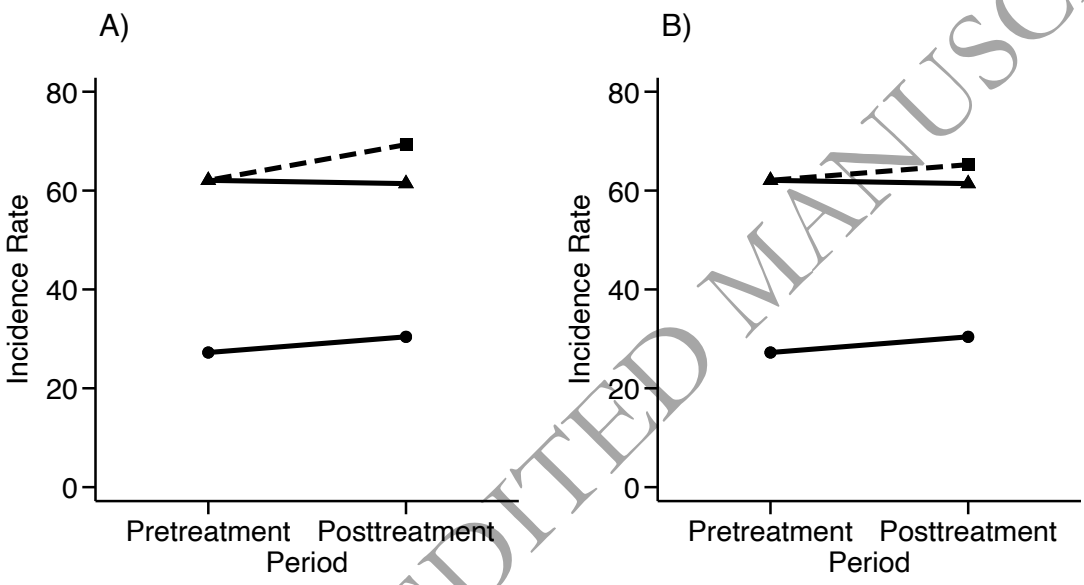

Observed incidence rates, denoted by circles for the control group and triangles for the treated group (connected by solid lines to highlight the change between the pretreatment to the posttreatment periods), are identical in panels A) and B). The observed incidence rate in the control group ($A = 0$) increased from 27.24 in the pretreatment to 30.41 in the posttreatment period. On the multiplicative scale, that is an increase by a factor of 1.12. Using PERR, in panel A), the counterfactual posttreatment rate in the treated group is estimated to be the observed pretreatment rate, 62.07, increased by the same factor: $62.07 \times 1.12 = 69.29$ (dashed lines highlight the change between the pretreatment rate and the counterfactual posttreatment rate). On the additive scale, the increase in the control group is 3.17. Using PERD, in panel B), the counterfactual posttreatment rate in the treated group is estimated to be the observed pretreatment rate, 62.07, increased by 3.17, that is $62.07 + 3.17 = 65.24$. All rates are expressed as events per 1,000 person-years.

Figure 2: Sensitivity analysis for PERR (panel A, over different values of u) and PERD (panel B, over different values of d).

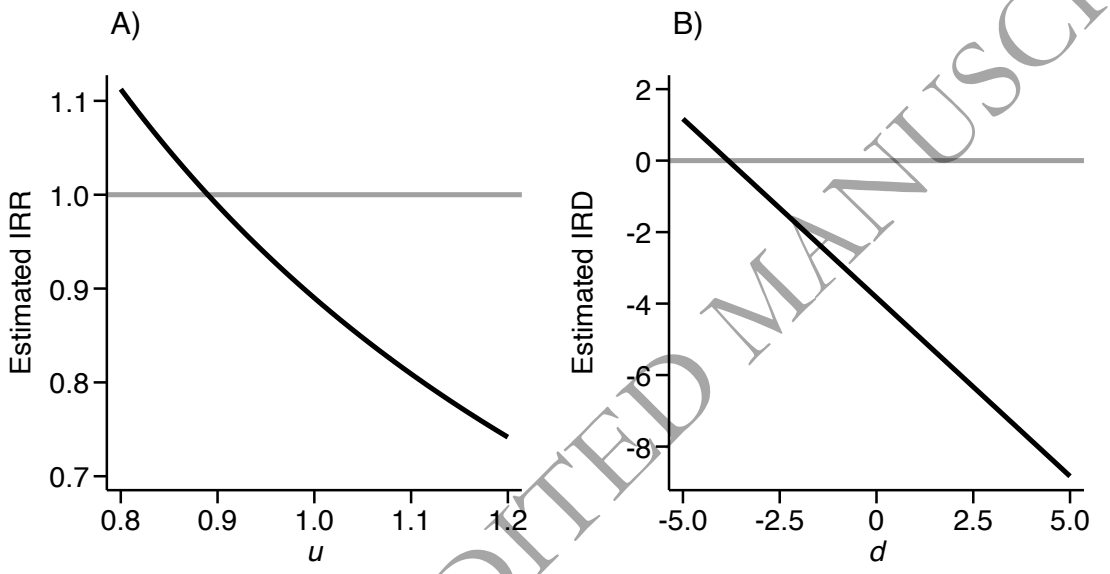

In the PERR analysis, values for u range between 0.8 and 1.2; $u = 1$ represents the primary analysis without adjustment for violations of the common rate-change assumption on the multiplicative scale, and estimates below the gray horizontal line denote benefit from proton pump inhibitors. In the PERD analysis, values for d range between -5 and 5 events per 1,000 person-years; $d = 0$ represents the primary analysis without adjustment for violations of the common rate-change assumption on the additive scale, and estimates below the gray horizontal line denote benefit from proton pump inhibitors. For the PERD analysis, incidence rate differences are expressed as differences in events per 1,000 person-years. In panel A, estimated IRR denotes the estimated $\text{IRR}_{\text{causal}}(A = 1)$; in panel B, estimated IRD denotes the estimated $\text{IRD}_{\text{causal}}(A = 1)$.