



Universiteit
Leiden
The Netherlands

Bias of time-varying exposure effects due to time-varying covariate measurement strategies

Vries, B.B.L.P. de; Groenwold, R.H.H.

Citation

Vries, B. B. L. P. de, & Groenwold, R. H. H. (2022). Bias of time-varying exposure effects due to time-varying covariate measurement strategies. *Pharmacoepidemiology And Drug Safety*, 31(1), 22-27. doi:10.1002/pds.5328

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/3307297>

Note: To cite this publication please use the final published version (if applicable).

ORIGINAL ARTICLE

WILEY

Bias of time-varying exposure effects due to time-varying covariate measurement strategies

Bas B. L. Penning de Vries¹  | Rolf H. H. Groenwold^{1,2}

¹Department of Clinical Epidemiology, Leiden University Medical Center, Leiden, The Netherlands

²Department of Biomedical Data Sciences, Leiden University Medical Center, Leiden, The Netherlands

Correspondence

Bas B. L. Penning de Vries, Department of Clinical Epidemiology, Leiden University Medical Center, PO Box 9600, 2300 RC, Leiden, The Netherlands.
Email: b.b.l.penning_de_vries@lumc.nl

Funding information

NWO-Vidi project 917.16.430

Abstract

Purpose: In studies of effects of time-varying drug exposures, adequate adjustment for time-varying covariates is often necessary to properly control for confounding. However, the granularity of the available covariate data may not be sufficiently fine, for example when covariates are measured for participants only when their exposure levels change.

Methods: To illustrate the impact of choices regarding the frequency of measuring time-varying covariates, we simulated data for a large target trial and for large observational studies, varying in covariate measurement design. Covariates were measured never, on a fixed-interval basis, or each time the exposure level switched. For the analysis, it was assumed that covariates remain constant in periods of no measurement. Cumulative survival probabilities for continuous exposure and non-exposure were estimated using inverse probability weighting to adjust for time-varying confounding, with special emphasis on the difference between 5-year event risks.

Results: With monthly covariate measurements, estimates based on observational data coincided with trial-based estimates, with 5-year risk differences being zero. Without measurement of baseline or post-baseline covariates, this risk difference was estimated to be 49% based on the available observational data. With measurements on a fixed-interval basis only, 5-year risk differences deviated from the null, to 29% for 6-monthly measurements, and with magnitude increasing up to 35% as the interval length increased. Risk difference estimates diverged from the null to as low as -18% when covariates were measured depending on exposure level switching.

Conclusion: Our simulations highlight the need for careful consideration of time-varying covariates in designing studies on time-varying exposures. We caution against implementing designs with long intervals between measurements. The maximum length required will depend on the rates at which treatments and covariates change, with higher rates requiring shorter measurement intervals.

KEYWORDS

inverse probability weighting, post-baseline covariate measurement, simulation, study design, time-varying confounding

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Pharmacoepidemiology and Drug Safety* published by John Wiley & Sons Ltd.

Key Points

- Adequate information on both baseline and time-varying covariates is important for confounding control in observational studies on the effects of time-varying drug exposures.
- Using simulated data, it was illustrated that considerable bias might arise when time-varying covariates are measured on a fixed-interval bases with long intervals between measurements or when subjects' covariates are assumed to remain constant in periods where their exposure levels remain constant.
- Whether or not data on time-varying covariates are collected with every exposure level switch, measurement strategies with long intervals between measurements are discouraged.

1 | INTRODUCTION

In many pharmacoepidemiologic studies, the use of the drugs that are investigated may change over time. In case of such time-varying exposures, the exposure effect can be defined in different ways. For example, one could contrast initiating drug treatment at a particular point in time (irrespective of whether the use is continued) with not initiating, or continuous drug use with continuous non-use. While analyses of point interventions (e.g., a single-dose vaccination) require adjustment for confounding at baseline only, for analyses of a time-varying exposure, information on time-varying covariates might be required to mitigate bias due to time-varying confounding. However, the granularity of the available information about the time-varying covariates may not be sufficiently fine to adequately control for confounding.

One special case of where this issue may arise is where researchers choose to measure covariates for study subjects only when their exposure levels have changed since the last measurement. If exposure levels do not change, covariate levels are (implicitly) assumed to remain constant, which is an implementation of a method generally known as last-observation-carried-forward (LOCF). The accurateness of the observed covariate data may then depend on the observed exposure history. In studies of antidepressant use and the risk of hip fracture, for example, comorbidities and use of comedication were assessed only at baseline and whenever patients switched exposure level or after every 6 months in the absence of switching.^{1,2}

In this paper, we investigate the impact of various covariate measurement designs on the estimation of time-varying exposure effects in observational studies with time-varying confounding. We illustrate, by way of simulation, the potential for bias of inverse-probability-weighting (IPW) estimators under static designs of fixed-interval covariate measurement and under dynamic designs with covariates being measured depending on the observed exposure history. IPW estimators are considered as these are increasingly used for estimating causal effects of time-varying exposures, can accommodate exposure-covariate feedback,³ and readily allow for 'adjusted' survival curves to be created.⁴

2 | METHODS

We first simulated data for a hypothetical study, the 'target trial', which if implemented on the theoretical population of interest would

readily allow us to identify the exposure effect of interest.⁵ In practice, it is not always possible to implement a target trial, but we use it here as a means to clarify the exposure effect of interest and we simulate from it to give a reference against which to compare results from analyses that are based on simulated data for observational studies. We considered multiple observational studies, each with the same data-generating mechanism but with different covariate measurement designs to evaluate their impact. Having simulated data, we then estimated the survival curves for the period of 5 years, using a weighting approach (described below) that was designed to keep treatment arms comparable throughout follow-up in terms of measured covariates. For each of the trial and observational studies, we first generated data on a single sample of $n = 150\,000$ individuals, which is sufficiently large to allow us to ignore sampling variability and regard differences between the survival curves as measures of the impact of the measurement designs on the large sample bias of the IPW estimators. The results corresponding to this single simulation run are described in detail below. In Appendix S2, we summarise the results of 5000 independent simulation runs for samples sizes 150 000, 10 000, 1000, 100. R code for this simulation is provided as Appendix S1.

2.1 | Set-up

The target trial has the following key design elements: (1) study participants (subjects who satisfy the eligibility criteria) are randomised at a well-defined baseline time point t_0 to either being issued a drug prescription ($A_0 = 1$)—say, a prescription for a daily dose of some antidepressant drug for the next one-month period—or to not being issued the prescription ($A_0 = 0$) at t_0 ; (2) participants are then followed over time until the occurrence of an event (e.g., the first hip fracture or death if the subject dies without having sustained a hip fracture during follow-up) or the administrative study end, whichever comes first; (3) provided event-free survival is long enough, study participants in the ($A_0 = 1$)-group are issued a further prescription after every month since t_0 and those in the ($A_0 = 0$)-group do not receive a prescription during follow-up. For a given subject, we define A_k to be the indicator variable that takes the value of 1 if the subject is on a one-month prescription on month k ; $A_k = 0$ otherwise. We further define Y to be the amount of follow-up time between baseline and the subject's (first) event and let Y_k be that part of Y that relates to month k . We stipulate that study

participants are event-free at the start of the study and that subjects do not get lost to follow-up before the administrative study end, which we stipulated to be 5 years (or $K = 60$ months) post-baseline.

The observational studies differ from the target trial in the following ways only: (1) the decision to allocate a subject to $A_0 = 1$ versus $A_0 = 0$ is not made by randomisation; (2) the decisions to renew prescriptions for subjects in the ($A_0 = 1$)-group or to never issue a prescription throughout the follow-up period for those in the ($A_0 = 0$)-group are not determined by their baseline allocations A_0 . Rather, for month $k = 0, 1, \dots$, the decision to set exposure A_k to 0 or 1 is based only on past exposure history ($A_j: j < k$) and certain binary covariates L_k . In this observational setting, subjects can switch at the start of each month between exposure levels 'being on prescription' (or 'exposed') versus 'not being on prescription' (or 'not exposed'). In variations on this setting, covariate data were measured according to one of the following measurement designs: (1) covariates were not measured at all, thus precluding any adjustment for confounding and effectively forcing us to implement a 'crude' estimator; (2) covariates were measured on a monthly basis, which is sufficient for identification of our target quantity; (3) covariates were measured on a 6-monthly basis starting at baseline; (4) covariates were measured when the respective subject's exposure level switched; (5) covariates were measured with an exposure level switch and at a 6-monthly basis in the absence of exposure level switching. We also considered variations on designs (3) and (5) where, instead of 6 months, the fixed measurement interval have a length of 2, 3, 9, 12, ..., or 60 months. Where design (3) means that measurement times are known before the start of follow-up, designs (4) and (5) are dynamic in the sense that whether or not a subject's covariate level is measured depends on the subject's time-varying variables.

2.2 | Data-generating mechanism

To simulate longitudinal data for a setting with time-varying confounding we used a variation on the approach described by Havercroft and Didelez⁶ and Young and Tchetgen Tchetgen.⁷ The data-generating mechanisms for the target trial and observational studies are described in the Appendix A and produce data that are consistent with the directed acyclic graphs (DAGs) of Figure 1. In the trial setting (left panel of Figure 1), the absence of arrows going into the exposure variables reflects the absence of (time-varying) confounding. In the target trial, post-baseline exposures are fully determined by the baseline level of exposure, which takes the value of 1 for half of subjects (i.e., exposure status does not change over time). In the observational study, however, approximately 40% of subjects will have switched exposure level by the end of follow-up in each of the arms that are defined by baseline exposure level.

2.3 | Defining and estimating the exposure effect

We define the exposure effect of interest as a contrast between continuous exposure ($A_j = 1$ for $j = 0, 1, \dots$) versus continuous non-

exposure ($A_j = 0$ for $j = 0, 1, \dots$). In particular, we suppose that the interest lies with a contrast between the 5-year event-free survival probabilities that we would observe had everyone received continuous exposure versus continuous non-exposure; that is, a contrast that is identified in the target trial as

$$\Pr(Y \geq 60 | A_0 = 1) \text{ versus } \Pr(Y \geq 60 | A_0 = 0).$$

As indicated by the absence of a directed path of arrows from the exposure variables to the outcome variables in the DAG for the target trial, the difference between these two survival probabilities is zero.

To account for time-varying confounding in the observational studies, we implemented IPW by applying a crude (Kaplan–Meier) estimator to an artificial data set where, given a time during follow-up, a subject received a weight of zero if the subject had experienced an exposure level switch by that time and otherwise a weight equal to the reciprocal of the product of the estimated probabilities of their observed exposure levels until that time given the respective measured exposure and covariate histories. That is, for $a = 0, 1$, a subject's weight for month k was

$$w_k = \prod_{j=0}^k \frac{1}{\Pr(A_j = a | Y \geq j, A_0 = \dots = A_{j-1} = a, L_0, \dots, L_j)}$$

if the subject received exposure level a in months 0 through k (i.e., $A_0 = \dots = A_k = a$). Subjects were censored (i.e., received a weight of zero) from the time at which they switched to another exposure level. Apart from the covariate measurement design, the validity of the approach also rests on the correct specification of the model for the conditional treatment probabilities. To ensure correct specification for the reference measurement design (1), we assumed that the exposure A_k given survival and past exposure and covariate levels was Bernoulli distributed with mean equal to

$$\Pr(A_k = 1 | Y \geq 1, A_0, \dots, A_{k-1}, L_0, \dots, L_k) = \frac{\exp[\alpha_0 + \alpha_1 I(k=0) + \alpha_2 A_{k-1} + \alpha_3 L_k]}{1 + \exp[\alpha_0 + \alpha_1 I(k=0) + \alpha_2 A_{k-1} + \alpha_3 L_k]}$$

for some $\alpha_0, \alpha_1, \alpha_2, \alpha_3$, which were estimated by a pooled logistic regression under this model. Throughout, variables that were unobserved by measurement design were handled with LOCF.

3 | RESULTS

Figure 2 shows the estimated survival curves for the 'always treat' and 'never treat' protocols. Consistent with the absence of a directed path from the exposure variables to the outcome variables in the DAGs of Figure 1, the trial-based estimates of the survival curves overlap (Figure 2, panel A). Where we observed a 5-year event risk of 31% in both arms of the target trial, in the observational setting, we observed a risk of 64% and 15% in those who do and those who do not receive a treatment prescription at baseline, respectively, giving a risk difference of 49% (panel B). With monthly covariate





FIGURE 1 Directed acyclic graphs representing the data-generating mechanism for the first 2 months of the target trial (left) and observational study (right). Here, U represents a unmeasured common cause of the measured covariates L_{k-1}, L_k and outcome variables Y_{k-1}, Y_k . The absence of directed paths from exposure variables A_{k-1}, A_k to outcome variables Y_{k-1}, Y_k reflects the absence of a causal exposure-outcome effect

FIGURE 2 Estimated event-free survival curves for ‘always treat’ and ‘never treat’ protocols based on target trial (panel A) and observational study (B through F) with varying covariate measurement designs: no covariate measurement B, continuous to monthly covariate measurement C, 6-monthly covariate measurement D, covariate measurement only with covariate level switching E, and with exposure switching and 6-monthly in periods without switching F






FIGURE 3 Estimated 2- and 5-year event risk differences comparing ‘always treat’ versus ‘never treat’ protocols. Estimates derive from observational studies with varying covariate measurement designs. Panel A gives the estimates for fixed-interval measurement; panel B gives the estimates for covariate measurement with exposure switching and with fixed-length intervals in periods without switching

measurement, IPW resulted in survival curves that virtually coincide with those of the trial (panel C), for which we found a risk difference of zero. Six-monthly measurements (panel D), however, brought the curves closer to those of the no measurement setting (panel B), that is, in the ‘direction of confounding’. The 5-year risks with 6-monthly measurements were estimated to be 50% and 21%, respectively, giving a risk-difference of 29%. In Figure 3, panel A, it is shown that the estimated risk differences at 2 and 5 years increase with the interval measurement length, until they reach a plateau of approximately 20% and 35%, respectively. When the interval length was set equal to the maximum follow-up duration (60 months), only baseline covariates were measured, which resulted in an estimated 5-year risk difference that was approximately 15% points closer to the target than that of no covariate measurement at all (Figure 2, panel B). When we implemented measurement design (4), the estimated 5-year risk difference flipped to the other side of the null, -14% (panel E), with 5-year risks estimated to be 27% and 41% for the ‘always treat’ and ‘never treat’ protocols, respectively. For design (5), we observed a 5-year risk difference of -5% , somewhere between the results of design (3) and (4) (panel F). With increasingly large measurement intervals within periods of no switching, the estimated 2-year risk difference steadily decreased to approximately -15% (Figure 3, panel B). The estimated 5-year risk was also -15% with 60 months between measurements in periods of no switching, equal to the observed risk of design (4), as expected. However, it was lowest, approximately -18% , with an interval length of around 30 months.

The bias estimates of the survival curves and 5-year risk difference that were derived by averaging across 5000 independent samples of sizes 150 000, 10 000 and 1000 are nearly identical to the corresponding estimates described above (see Supplementary Table and Figures). For sample size 100, however, we observed substantial (small sample) bias for all measurement designs, even in the reference observational setting with full/monthly covariate measurement.

4 | DISCUSSION

We used simulation to study and illustrate the potential for bias due to measurement design choices in the estimation of the effects of

time-varying exposures. The potential for bias in settings with static or fixed-interval covariate measurement designs has recently been illustrated already.⁸ We additionally showed that bias might arise in settings where decisions to measure are driven by observed values of the time-varying exposure.

As expected, in our simulations, fixed-interval measurement resulted in bias in the direction of confounding, bias that is attributable to residual confounding. Interestingly, we found bias in the opposite direction when we implemented measurement designs where covariates were measured preferentially with exposure level switches. Together with LOCF, these measurement designs introduced a form of differential misclassification, which may result in bias even in the absence of confounding.⁹ Researchers familiar with DAGs might be alerted by the presence of colliders in the DAG that encodes part of the misclassification mechanism. For example, on the DAG of the right panel of Figure 1, the differential misclassification of L_1 can be represented by adding a measured version of L_1 with incoming arrows from L_0 , L_1 , A_0 and A_1 . The measured variable can then be seen to be a collider on the path from A_1 to Y_1 via L_1 and U . By conditioning on the collider (and not the unmeasured variable L_1 or U), the path is opened, potentially leading to collider-stratification bias.³ In addition to adequate measurement of the time-varying covariates, the validity of IPW rests on the correct specification of the model for the distribution of the treatment variables given survival and past covariate and exposure levels. It is possible that the biases that we observed are partly due to model misspecification.

We considered a specific and relatively simple setting with a single, binary covariate, no censoring before the administrative study end and an interest in static rather than dynamic treatment strategies. These features are not required for valid inference with IPW.³ However, the magnitude and direction of bias in other settings may differ from those observed in the current study. We stress that the bias that was observed in our simulation does not depend critically on the choice of IPW as a means to control for time-varying confounding. The choices regarding the frequency of covariate measurements will likely also affect other methods, including the commonly applied Cox’ regression analysis with time-varying covariates. The extent to which such choices impact a particular study are obviously context-specific. For example, it will likely depend on the rate at which subjects cross

over between treatment arms as well as on the extent to which covariates are subject to change over time.

In conclusion, our simulations highlight the need for adequate measurement of time-varying covariates in observational studies on the effects of time-varying exposures. Researches should consider differential covariate misclassification as a possible source of bias when designing covariate measurement strategies. Whether or not covariates are measured with every exposure level switch, we caution against implementing measurement designs with long intervals between measurements, particularly when the impact of the design choices are poorly understood. The maximum interval length that is sufficient to yield negligible bias will depend on the rates at which treatments and covariates can change,⁸ with higher rates requiring shorter measurement intervals.

ACKNOWLEDGEMENTS

RHHG was funded by the Netherlands Organization for Scientific Research (NWO-Vidi project 917.16.430). The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding bodies.

CONFLICT OF INTEREST

The authors declare no conflict of interest.

ETHICS STATEMENT

No ethics committee approval was required for this study.

ORCID

Bas B. L. Penning de Vries  <https://orcid.org/0000-0001-9989-7732>

REFERENCES

1. Ali MS, Groenwold RHH, Belitser SV, et al. Methodological comparison of marginal structural model, time-varying cox regression, and propensity score methods: the example of antidepressant use and the risk of hip fracture. *Pharmacoepidemiol Drug Saf.* 2016;25:114-121.
2. Souverein PC, Abbing-Karahagopian V, Martin E, et al. Understanding inconsistency in the results from observational pharmacoepidemiological studies: the case of antidepressant use and risk of hip/femur fractures. *Pharmacoepidemiol Drug Saf.* 2016;25:88-102.
3. Hernán MA, Robins JM. *Causal Inference: What If*. Boca Raton, FL: Chapman & Hall/CRC; 2020.
4. Cole SR, Hernán MA. Adjusted survival curves with inverse probability weights. *Comput Methods Programs Biomed.* 2004;75(1):45-49.
5. Hernán MA, Robins JM. Using big data to emulate a target trial when a randomized trial is not available. *Am J Epidemiol.* 2016;183(8):758-764.
6. Havercroft WG, Didelez V. Simulating from marginal structural models with time-dependent confounding. *Stat Med.* 2012;31(30):4190-4206.
7. Young JG, Tchetgen Tchetgen EJ. Simulating from a known Cox MSM using standard parametric models for the g-formula. *Stat Med.* 2014;33(6):1001-1014.
8. Young JG, Vatsa R, Murray EJ, Hernán MA. Interval-cohort designs and bias in the estimation of per-protocol effects: a simulation study. *Trials.* 2019;20(1):552.
9. Webster-Clark M, Jonsson-Funk M, Stürmer T. Single-arm trials with external comparators and confounder misclassification: how adjustment can fail. *Med Care.* 2020;58(12):1116-1121.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Penning de Vries BBL, Groenwold RHH. Bias of time-varying exposure effects due to time-varying covariate measurement strategies. *Pharmacoepidemiol Drug Saf.* 2022;31(1):22-27. <https://doi.org/10.1002/pds.5328>

APPENDIX A.

Data-generating mechanism

For convenience of notation, define $L_k = A_k = 0$ if $k < 0$. Also, let I and expit denote the indicator and inverse logit function, respectively, so that $\text{expit}[x] = \exp[x]/(1 + \exp[x])$. For all 150 000 study participants, we generated data using independent runs of the following algorithm: (1) generate U from a uniform distribution over the interval $[0,1]$, initialise $Y = 0$ and set $k = 0$; (2) if $Y \geq k$, draw L_k from the Bernoulli distribution with parameter $\text{expit}[-4.25 + 0.25/(k = 0) + 6U + 0.5A_{k-1} + L_{k-1}]$ and draw A_k from the Bernoulli distribution with parameter p_k ; (3) draw T from the exponential distribution with rate $\exp[-9.5 + 7U]$ and increment Y with $Y_k = \min\{T, 1\}$; (4) if $k < K$, increment k with 1 and return to (2); stop otherwise. For the observational study, we defined $p_k = \text{expit}[-7 + 4I(k = 0) + 10A_{k-1} + 4L_k]$, whereas for the trial we defined $p_k = 0.5I(k = 0) + A_{k-1}I(k > 0)$.