

Rejoinder: ‘Estimating time-varying causal excursion effects in mobile health with binary outcomes’

BY TIANCHEN QIAN

*Department of Statistics, University of California, Irvine,
Bren Hall 2019, Irvine, California 92697, U.S.A.*

tiancq1@uci.edu

HYESUN YOO

*Department of Statistics, University of Michigan,
323 West Hall, 1085 South University, Ann Arbor, Michigan 48109, U.S.A.*

yoohs@umich.edu

PREDRAG KLASNJA

*School of Information, University of Michigan,
4364 North Quad, 105 South State Street, Ann Arbor, Michigan 48109, U.S.A.*

klasnja@umich.edu

DANIEL ALMIRALL

*Department of Statistics, University of Michigan,
323 West Hall, 1085 South University, Ann Arbor, Michigan 48109, U.S.A.*

dalmiral@umich.edu

AND SUSAN A. MURPHY

*Department of Statistics, Harvard University,
1 Oxford Street, Cambridge, Massachusetts 02138, U.S.A.*

samurphy@fas.harvard.edu

1. INTRODUCTION

We would like to thank the editor for making this a discussion paper. We are honoured that these three groups of researchers agreed to be discussants. Reading their discussions has been an enjoyable and thought-provoking journey for us. We organize our rejoinder by topic, as some topics have been discussed by multiple discussants.

2. DEPENDENCE OF THE CAUSAL EXCURSION EFFECT ON TREATMENT PROTOCOL

In general, the causal excursion effect in (2) of the main manuscript depends on the randomization probability of past treatment assignments (\bar{A}_{t-1}), which we refer to as the treatment protocol. This feature distinguishes the excursion effect from most other causal effects in time-varying treatment settings, such as causal effects in marginal structural models or structural nested mean models.

Guo et al. (2021) and Zhang & Laber (2021) provided numerical examples to illustrate this feature of the excursion effect. Guo et al. (2021, Equation 6) considered as one example the following data-generating process:

$$Y_2(a_1, a_2) \sim \text{Ber}\{\exp(-a_2 + 2a_1 \cdot a_2)/4\},$$

where the prior treatment, A_1 , interacts with the present treatment, A_2 , to affect Y_2 . They showed that the magnitude and sign of the causal excursion effect of A_2 on Y_2 , which marginalizes over A_1 , depends on the randomization probability of A_1 . In the second example they considered, there is a three-way interaction between A_1 , X_2 and A_2 on Y_2 , and they showed that the effect modification by X_2 of the causal excursion effect of A_2 can vary depending on the randomization probability of A_1 . Zhang & Laber (2021) provided two examples where the current treatment interacts with a prior treatment as well as a time-varying covariate, and showed the dependence of the excursion effect on the randomization probability of prior treatments. Zhang & Laber (2021) wrote ‘a large negative causal excursion effect does not mean a treatment is harmful; rather it may mean that the protocol drove patients into states for which treatment appeared harmful.’ We very much agree. Indeed, the comparison between two excursions starting at time t is likely to greatly differ by the treatment protocol up to time t ; an example is to consider (i) a treatment protocol that is feasible to implement in practice, and (ii) a treatment protocol that is infeasible or excessively onerous. Below we argue that this consequence makes sense in a continual learning setting such as in mobile health.

First, note that this consequence of marginalization occurs in all marginal models; models that don’t condition on all relevant covariates. In Equation (6) of Guo et al. (2021) A_1 plays the role of a covariate in the analysis of the effect of A_2 on Y_2 . Another related setting is the primary analysis for a factor in a factorial design, where the other factors are averaged over. If there is interaction between factors, then the primary analysis result for one factor will depend on the distribution of the other factors.

As discussed in § 8 of the main manuscript, we view the dependence of the causal excursion effect on the randomization probability of past treatments, the treatment protocol, as a desirable feature. In our mind, the causal excursion effect makes most sense in a continual learning setting. Because the technology is evolving rapidly, mobile health intervention development/optimization requires repeated updating and improvement. We advocate that the treatment protocol at any given time, including when implemented in a micro-randomized trial, be sufficiently realistic that it can be implemented in practice. We mean realistic in the sense that considerations of burden and domain knowledge on effectiveness have been taken into account. Thus, at each stage researchers are interested in answering the question: how would one further improve the current treatment protocol? Causal excursion effects provide an approximation to a gradient in addressing this question; see § 3 below.

Nonetheless, the examples provided by Guo et al. (2021) and Zhang & Laber (2021) clearly illustrate the dependence of the causal excursion effect on the treatment protocol, and highlight the importance of interpreting the causal excursion effect in the context of the treatment protocol. We agree with the discussants that additional analyses, such as effect moderation of the causal excursion effect by prior treatments and time-varying covariates, are often useful to understand the dependence of the causal excursion effect on the treatment protocol. We consider them as secondary and exploratory analyses, which have been routinely carried out in micro-randomized trial analyses. For example, in HeartSteps V1 (Klasnja et al., 2018), a micro-randomized trial for developing push notifications to encourage regular physical activity among sedentary adults, after the primary analysis to assess the marginal excursion effect of the push notification, the researchers conducted a variety of secondary analyses, which included the effect modification by past treatments and past covariates, and how the marginal excursion effect changes over time. They found that the magnitude of the effect changes from positive to not different from 0 around mid-study, suggesting that burden and habituation due to past treatments affects the causal excursion effect. Additional qualitative studies, e.g., via exit interviews, suggested that the decreasing, with time, magnitude of the excursion effect might be due to habituation.

Zhang & Laber (2021) discussed two exploratory analysis approaches to assess the degree of dependence of causal excursion effect on treatment protocol. One approach is to use a partially observable Markov decision process to capture the non-Markov dynamics of patient histories, including past treatments, which

can be understood as constructing features that summarize the past effect modifiers. The other approach is to vary the randomization probabilities, or other features of the design, and assess how the causal excursion effect varies accordingly. These are interesting future research topics.

3. USING THE CAUSAL EXCURSION EFFECT TO MODIFY THE PROTOCOL

Guo et al. (2021) pointed out that due to the dependence of the causal excursion effect on the treatment protocol, modifying the treatment protocol based on the excursion effect can be a challenging task. This is because as the treatment protocol is modified, the causal excursion effect would change accordingly. We very much agree. Below we lay out our current understanding and thoughts on this topic.

Using the treatment regime notation mentioned by Guo et al. (2021) in their § 1, let $Y_{t,\Delta}(G)$ be the potential outcome of $Y_{t,\Delta}$ under a regime G . Let G_b be the regime, i.e., the treatment protocol, in the micro-randomized trial. Consider a new candidate regime, G_c . A natural comparison between the two regimes is the log relative risk,

$$\gamma_t(G_c, G_b) = \log \frac{E\{Y_{t,\Delta}(G_c)\}}{E\{Y_{t,\Delta}(G_b)\}}. \quad (1)$$

For clarity, we only consider the marginal comparison and do not condition on any effect modifiers S_t . Inverse probability weighting can be used to estimate the numerator in (1) using data collected under G_b ; note that to ensure the weights are stable, G_c cannot be too different from G_b . One realistic situation is where G_c represents a small deviation from the treatment protocol G_b . For example, $\gamma_t(G_c, G_b)$ represents an approximation to the derivative

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} \left(\log E[Y_{t,\Delta}\{(1-\epsilon)G_b + \epsilon G_c\}] - \log E\{Y_{t,\Delta}(G_b)\} \right), \quad (2)$$

where $(1-\epsilon)G_b + \epsilon G_c$ is a mixture of the two regimes. Equation (2) represents the relative change in the log expected potential proximal outcome when the regime changes in the direction of G_c . The sign of $\gamma_t(G_c, G_b)$ as well as how it varies over time, plus expert input and qualitative user studies, inform how to alter the current treatment protocol. A caveat is that these are high-dimensional optimization problems and the gradient at any point in a particular direction may be misleading, e.g., it may lead to a saddle point or a local optimum.

Using the above notation, the causal excursion effect defined in the main manuscript can be written as $\gamma_t(G_c, G'_c)$, where G_c is the regime that is the same as the treatment protocol G_b up to $t-1$, then it assigns the future Δ treatments deterministically as $(1, \vec{0})$, and G'_c is the regime that is the same as the treatment protocol G_b up to $t-1$, then it assigns the future Δ treatments deterministically as $(0, \vec{0})$. As a result, $\gamma_t(G_c, G'_c)$ can be viewed as an approximation to

$$\lim_{\epsilon \rightarrow 0^+} \frac{1}{\epsilon} \left(\log E[Y_{t,\Delta}\{(1-\epsilon)G_b + \epsilon G_c\}] - \log E[Y_{t,\Delta}\{(1-\epsilon)G_b + \epsilon G'_c\}] \right),$$

i.e., the difference in the relative change in the log expected proximal outcome when the regime changes in the direction of G_c versus when the regime changes in the direction of G'_c . One recent work that leverages the above idea is Liao et al. (2020), where the authors used effect modification analyses of micro-randomized trial data to form a warm-start policy for a reinforcement learning algorithm in a subsequent study.

4. OTHER POINTS RAISED BY DISCUSSANTS

4.1. Causal excursion effects with different Δ

In the main manuscript, the causal excursion effect is defined for a specific Δ , the time window over which the proximal outcome is measured, and Δ is typically predetermined based on domain knowledge.

For example, in Sense2Stop (Sarker et al., 2017; Dempsey et al., 2020), a micro-randomized trial for developing stress reduction interventions for people who recently quit smoking, the decision point is per minute and Δ is 120 minutes. The particular choice of Δ was based on the research team's prior knowledge of the time duration for the intervention effect.

Kim et al. (2021) in their § 3 propose to model the excursion effect for each Δ separately for a collection of Δ s. Such analysis might be used to inform the duration for an intervention to take effect as well as the duration over which the intervention's effect lasts. We agree that this is a fruitful option for exploratory analysis. However, the warning in (Guo et al., 2021, § 2.2) on history-adjusted marginal structural models indicates that great care must be taken to avoid producing model incoherence. We look forward to seeing the work of Kim et al. (2021) in this direction.

4.2. Extension to observational data; efficiency improvement

Both Guo et al. (2021) and Kim et al. (2021) discussed extending the estimating equation (10) of the main manuscript to observational study settings. The estimator $\hat{\beta}$ solving estimating equation (10) is consistent for the parameter in the marginal excursion effect, β in equation (9), if the propensity score $p_t(H_t)$ can be consistently estimated. However, as Guo et al. (2021) and Kim et al. (2021) pointed out, when $p_t(H_t)$ is incorrectly modelled, $\hat{\beta}$ will generally not be consistent even if $\exp\{g(H_t)^\top \alpha\}$ is a correct model for the mean potential outcome under no-future-treatment excursion, $E\{Y_{t,\Delta}(\bar{A}_{t-1}, 0, \bar{0}) \mid H_t, I_t = 1, A_t = 0\}$. This is due to the fact that the causal excursion effect $\beta_M(S_t) = S_t^\top \beta$ is a marginal model; the fully conditional, on H_t , causal effect is not necessarily $S_t^\top \beta$. Therefore, even when $\exp\{g(H_t)^\top \alpha\}$ is a correct model for $E\{Y_{t,\Delta}(\bar{A}_{t-1}, 0, \bar{0}) \mid H_t, I_t = 1, A_t = 0\}$, the term $Y_{t,\Delta} - \exp\{g(H_t)^\top \alpha + A_t S_t^\top \beta\}$ in the estimating equation (10) does not generally have conditional expectation zero given H_t . Guo et al. (2021) proposed to subtract the projection of $\prod_{j=t+1}^{t+\Delta-1} \mathbb{1}(A_j = 0)/1 - p_j(H_j)$ onto the scores for treatment to both obtain a doubly robust version and to improve efficiency. Fortunately, for the micro-randomized trials one knows the randomization probability, and the one-sided robustness, consistency under misspecified $\exp\{g(H_t)^\top \alpha\}$, is adequate there. We thank the discussants for pointing this out, and we leave the detailed derivation of a doubly robust version of the estimator for the causal excursion effect for future work.

Of course, the semiparametric locally efficient estimator for the fully conditional, on H_t , effect with $\Delta = 1$, which was derived in § 4 of the main manuscript based on the results in Robins (1994), has the double robustness property in the sense that it is consistent if either the propensity score model is correct or $\exp\{g(H_t)^\top \alpha\}$ is a correct model for $E\{Y_{t,1}(\bar{A}_{t-1}, 0) \mid H_t, I_t = 1, A_t = 0\}$. We look forward to the future work of Kim et al. in extending this method to Poisson distributed outcomes.

4.3. Connection and difference between the causal excursion effect and structural models

As we mentioned in § 2.2 of the main manuscript, the causal excursion effect differs from the causal effects considered in marginal structural models and structural nested mean models, and the marginalization aspect of the causal excursion effect is related to the history-restricted version of marginal structural models. Guo et al. (2021), in their Figs. 1–3, rather nicely illustrate the connection and difference between the causal excursion effect and the structural nested mean models and a variety of marginal structural models, including the classic, history-restricted and history-adjusted marginal structural models. We appreciate Guo et al. (2021) making the differences between these effects so clear.

4.4. Identifiability for other types of excursions

In the definition of the causal excursion effect, we considered the two excursions $(\bar{A}_{t-1}, 1, \bar{0}_{\Delta-1})$ and $(\bar{A}_{t-1}, 0, \bar{0}_{\Delta-1})$. Guo et al. (2021), in their § 1, pointed out that these two are the only static regimes that are identified without further assumptions. Other static regimes will have $a_m = 1$ for some $m > t$, and the average potential outcome under such regimes is generally not identifiable without additional assumptions because of the violation of positivity assumption when $I_m = 0$ with positive probability. We agree with their statement, and appreciate the clarification.

In the Sense2Stop micro-randomized trial that motivated us to consider causal excursion effects with general Δ , other static regimes were not of scientific interest. If a micro-randomized trial is well designed,

regimes such as those that violate the positivity assumption should have been a priori decided to be uninteresting for scientific or ethical reasons. Absent further assumptions, the regimes that are identified by the micro-randomized trial must assign no treatment at time t whenever $I_t = 0$. Note that I_t can be affected by prior treatments. A rather interesting question is the following. Suppose under the treatment protocol of the micro-randomized trial the probability of $I_t = 1$ is low, and suppose the interest is in estimation of the effect of a new regime that is expected to influence I_t so much that I_t will be 1 with high probability. While this new regime might not violate the positivity assumption, it is likely that the estimator of the effect of this new regime will be quite variable. This can be an area in which it is critical to follow the recommendation of Guo et al. (2021) to extend our current estimator to a doubly robust estimator with improved efficiency.

ACKNOWLEDGEMENT

Research reported in this paper was supported by the National Institutes of Health (R01AA23187, P50DA039838, R01DA039901, P41EB028242, U01CA229437, R01HL125440 and UG3DE028723). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

REFERENCES

- DEMPSEY, W., LIAO, P., KUMAR, S. & MURPHY, S. A. (2020). The stratified micro-randomized trial design: Sample size considerations for testing nested causal effects of time-varying treatments. *Ann. Appl. Statist.* **14**, 661–84.
- GUO, F. R., RICHARDSON, T. S. & ROBINS, J. M. (2021). Discussion of ‘Estimating time-varying causal excursion affects in mobile health with binary outcomes’. *Biometrika* **108**, 541–50.
- KIM, S., CHO, H., BANG, D., DE MARCH, D., EL-ZAATARI, H., SHAH, K. S., VALANCIUS, M., ZIKRY, T. M. & KOSOROK, M. R. (2021). Discussion of ‘Estimating time-varying causal excursion affects in mobile health with binary outcomes’. *Biometrika* **108**, 529–33.
- KLASNJA, P., SMITH, S., SEEWALD, N. J., LEE, A., HALL, K., LUERS, B., HEKLER, E. B. & MURPHY, S. A. (2018). Efficacy of contextually tailored suggestions for physical activity: A micro-randomized optimization trial of HeartSteps. *Ann. Behav. Med.* **53**, 573–82.
- LIAO, P., GREENEWALD, K., KLASNJA, P. & MURPHY, S. (2020). Personalized HeartSteps: A reinforcement learning algorithm for optimizing physical activity. *Proc. ACM on Interact. Mobile Wear. Ubiqu. Tech.* **4**, 1–22.
- ROBINS, J. M. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Commun. Statist. Theory Meth.* **23**, 2379–412.
- SARKER, H., HOVSEPIAN, K., CHATTERJEE, S., NAHUM-SHANI, I., MURPHY, S. A., SPRING, B., ERTIN, E., AL’ABSI, M., NAKAJIMA, M. & KUMAR, S. (2017). From markers to interventions: The case of just-in-time stress intervention. In *Mobile Health: Sensors, Analytic Methods, and Applications*, J. M. Rehg, S. A. Murphy & S. Kumar, eds. pp. 411–433. New York: Springer.
- ZHANG, Y. & LABER, E. B. (2021). Discussion of ‘Estimating time-varying causal excursion affects in mobile health with binary outcomes’. *Biometrika* **108**, 535–9.

[Received on 8 June 2021. Editorial decision on 9 June 2021]