Curtis Huffman* and Edwin van Gameren

Covariate Balancing Inverse Probability Weights for Time-Varying Continuous Interventions

<https://doi.org/10.1515/jci-2017-0002>

Received January 13, 2017; revised January 24, 2018; accepted March 19, 2018

Abstract: In this paper we present a continuous extension for longitudinal analysis settings of the recently proposed Covariate Balancing Propensity Score (CBPS) methodology. While extensions of the CBPS methodology to both marginal structural models and general treatment regimes have been proposed, these extensions have been kept separately. We propose to bring them together using the generalized method of moments to estimate inverse probability weights such that after weighting the association between time-varying covariates and the treatment is minimized. A simulation analysis confirms the correlation-breaking performance of the proposed technique. As an empirical application we look at the impact the gradual roll-out of Seguro Popular, a universal health insurance program, has had on the resources available for the provision of healthcare services in Mexico.

Keywords: causal inference, covariate balancing propensity score, inverse probability weights, sequential ignorability, time-varying continuous interventions

1 Introduction

Propensity Score (PS) methods can be applied to a wide range of different research settings, and have become a popular tool for making causal inference. As shown by Robins [1], PS methods can even be used to evaluate programs that evolve over long periods of time, reacting to changes in the environment that themselves help to bring about. That class of statistical models came to be known as Marginal Structural Models (MSMs) [2–4]. Despite its theoretical appeal, MSMs as a tool for making causal inference from longitudinal data have yet to extend to general treatment regimes, that is, from binary to (potentially) continuous treatment intensities. Surprisingly, while several extensions of propensity score methods to general treatment regimes have been proposed [5, 2, 6–8] and keep gaining popularity among applied researchers [9], these methodologies haven't been coupled with the MSM framework.

Perhaps behind this gap between methodological and applied research is the practical difficulty of estimating the inverse probability weights that are required for MSMs to generate a pseudo-sample from which to estimate the causal quantities of interest. The correct model is generally unknown and its misspecification can actually increase bias, even if the selection on observables assumption holds. In contrast with cross-sectional PS methods, in MSMs the effect of model misspecification is exacerbated, for it propagates across time since the algorithm provided by Robins et al. [2] to estimate the weights in these models involves the product of propensity scores estimated for each time period.

Since, in theory at least, inverse probability weighting renders the intervention independent of time-varying covariates, usual practice involves iteratively checking if weighting on the estimated probability meets this expected property – revising the model until said correlation is minimized. The downside of this

*Corresponding author: Curtis Huffman, Programa Universitario de Estudios del Desarrollo, Coordinación de Humanidades, Universidad Nacional Autónoma de México, Ciudad Universitaria, Coyoacán, Mexico City, Mexico, e-mail: chuffman@colmex.mx
Edwin van Gameren, Center for Economic Studies, El Colegio de México A.C., Mexico City, Mexico, e-mail: egameren@colmex.mx

algorithm is that it can lead to ad hoc specifications of the probability models which may not agree with the preferred theoretical specification.

As a way to get around this practical difficulty surrounding MSMs, Imai and Ratkovic [10] came up with the idea of formally dovetailing this correlation minimizing property of the weights into the model estimation method itself, eliminating the need to manually and iteratively check the treatment assignment model. They accomplished this by making use of the Generalized Method of Moments (GMM), adding to the model's score vector – that is, the gradient of the log-likelihood – the correlation minimizing property as an extra moment condition. They have called this new approach the Covariate Balancing Propensity Score (CBPS) [11] and it can easily be implemented through publicly available open-source software [12]. Simulation and empirical studies suggest that the CBPS improves the empirical performance of MSMs by making the treatment assignment model depend less heavily on assumptions which are often not met in practice. However, so far the CBPS extension to MSMs has been confined to binary time-dependent treatments.

Here we propose a further extension to the CBPS methodology that can be used by researchers facing continuous time-dependent treatments. Our proposal builds on the recent generalization of the CBPS methodology to general treatment regimes [13] taking it to the longitudinal setting, thus bridging the gap between two extensions of the original CBPS estimator.

In Section 2 we briefly review MSMs and their assumptions. Section 3 we describe the proposed extension to the CBPS methodology. In Section 4 we present our simulation study. In addition, we present an empirical application in Section 5 to show the performance of the proposed methodology with real data. Finally, we offer concluding comments in Section 6.

2 Marginal structural models

Assume a program that evolves over long periods of time intervening with different intensity across recipients at every stage of their implementation, reacting to changes in the environment which the program itself contributed to bring about; that is to say, reacting to its own previous results in a feedback loop. In this context, variables that can both affect future exposure to the program and be affected by past exposure may confound the impact of the program in a way that no cross-sectional impact evaluation method can sort out [4]. In this Section we briefly review the MSM framework [2–4] that has been developed to analyze the impact of such programs in case of binary treatments.

Denote recipient i 's observed k -dimensional vector of time-varying covariates at stage t of the program as $\mathbf{X}_{it}$, with $i = 1, \dots, N$ and $t = 1, \dots, T$. At every stage t of the program, recipients are observed receiving the exposure level a_{it} to the benefits of the program; that is, one possible realization of the exposure variable A_{it} , with $a_{i0} = 0$. Collecting all the observed exposures to the program for a given recipient, from its launch up to stage t , gives us the history of exposure $\underline{a}_{it} = (a_{i1}, \dots, a_{it})$. Independent, identically distributed realizations of $(Y_{it}, \underline{A}_{it}, \underline{\mathbf{X}}_{it})$ are observed for each recipient at every stage of the program, where Y_{it} is some outcome variable of interest. Let $\underline{\mathbf{X}}_{it}$ and $\underline{\mathbf{x}}_{it}$ be similarly defined for a covariate history, where $\mathbf{X}_{it}$ is the most recent set of variables that could affect A_{it} , and are also affected by past exposure history, $\underline{A}_{it-1}$.

Following Robins' adaptation of the potential outcome framework [14, 15] to longitudinal settings [16], we use $Y_{it}^{\underline{a}_{it}}$ to represent the potential outcome for recipient i measured at time t associated to the exposure history $\underline{a}$ up to t , which we assume does not depend on the exposure histories of other recipients. This is often referred to in the literature as the stable unit treatment value assumption. Naturally, any recipient of the program exhibits only one of these potential outcomes at any given stage t ; that is, the one associated to its own particular exposure history up to that point in time $Y_{it} = Y_{it}^{\underline{a}_{it}}$, the consistency assumption. All other, unobserved, potential outcomes are said to be counterfactual.

As with other PS methods, MSMs are based on the assumption of no unmeasured confounding, in this case at each time period. This assumption, also referred to as sequential ignorability [4] or conditional interchangeability [3], states that the exposure level at stage t is statistically independent (ignorable) of the

potential outcomes, conditional on the covariate and exposure histories up to that point. Using previous notation, sequential ignorability says that for any exposure history up to t , $a_t, Y_{it}^{a_t} \perp A_{it} \mid \underline{A}_{it-1} = \underline{a}_{t-1}, \underline{X}_{it} = \underline{x}_t$. Additionally, also typical in PS methods, MSMs also assume that at any stage t , there is no covariate history $\underline{x}_t$ and past exposure $\underline{a}_{t-1}$ such that all recipients with such histories are certain to receive the identical exposure a_t . That is, each exposure history must have some positive probability of occurring. This last assumption can be formally written as, if $f(\underline{A}_{it-1}, \underline{X}_{it}) \neq 0$ then $f(\underline{A}_{it} \mid \underline{A}_{it-1}, \underline{X}_{it}) > 0$, where f stands for the probability density function (pdf).

Under the above assumptions (sequential ignorability, consistency and positivity), Robins [1] showed that the conditional probability density of exposure history can be used to consistently estimate the impact of such dynamic interventions. In this generalization of Inverse Probability of Treatment Weights (IPTW), the probability of observing a particular exposure history is expressed as the product of the respective probability density at every stage, $\prod_{t=1}^T f(A_t \mid \underline{A}_{t-1}, \underline{X}_t)$. However, using this probability directly as weights when fitting the outcome model leads to highly variable estimates and researchers generally follow Cole and Hernán [17] and use stabilized weights of the form

$$SW_t(\underline{A}_t, \underline{X}_t) = \prod_{t'=0}^t \frac{f(A_{t'})}{f(A_{t'} \mid \underline{A}_{t'-1}, \underline{X}_{t'})} \quad (1)$$

where, for convenience, we define $SW_0 = 1$.

In nonrandomized evaluations, irrespective of the nature of the intervention whose impact one might be trying to estimate – whether it's best coded as binary, multivalued or continuous –, these weights are unknown, and have to be estimated. Usually, a parametric model is used to estimate the probability density of exposure given the past at each time period. When the intervention comes as a single action that either occurs or not at every stage of the program (i.e. coded as a binary variable), a common approach is to estimate this probability density at every stage with a standard logistic regression (logit) model. In the case of continuous interventions, Gaussian models (with constant variance) are often assumed, i.e., $A_t \mid \underline{A}_{t-1}, \underline{X}_t \sim N(h(\underline{A}_{t-1}, \underline{X}_t; \boldsymbol{\gamma}), \sigma^2)$, which can be estimated by the GMM.¹ If there is reason to believe that h is a linear, additive function of exposure and covariates histories, and parameters $\boldsymbol{\theta}_t = (\boldsymbol{\alpha}, \boldsymbol{\beta}, \sigma^2)$, researchers may regress the current level of exposure on all past exposure and covariates,² e.g.,

$$f_{\boldsymbol{\theta}_t}(A_t \mid \underline{A}_{t-1}, \underline{X}_t) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2\sigma^2} [A_t - (\underline{A}_{t-1}^\top \boldsymbol{\alpha} + \underline{X}_t^\top \boldsymbol{\beta})]^2\right\}. \quad (2)$$

The numerator of the stabilized weights is typically estimated using the empirical distribution of the exposure. If the assumptions behind equation 2 hold, and further assuming the marginal distribution of A_t to be Gaussian with mean zero (possibly after the pertinent transformation), the stabilizing weight is given by³

$$SW_{it}(\underline{A}_t, \underline{X}_t) = \prod_{t'=0}^t \exp\left[\frac{1}{2\sigma^2} \{-2A_{it'}(\underline{A}_{it'-1}^\top \boldsymbol{\alpha} + \underline{X}_{it'}^\top \boldsymbol{\beta}) + (\underline{A}_{it'-1}^\top \boldsymbol{\alpha} + \underline{X}_{it'}^\top \boldsymbol{\beta})^2\}\right]. \quad (3)$$

Once these stabilized weights are estimated, all the researchers have to do to is use them in running whatever model of the outcome, as function of treatment history alone, they would have used in the case of a randomized evaluation.⁴

¹ Note that there is no need to assume cross-time homoscedasticity or even a Gaussian distribution. More in general, one can estimate $f(A_{t'} \mid \underline{A}_{t'-1}, \underline{X}_{t'})$ as the exposure assignment model pdf evaluated at the observed exposure and covariates histories for the distribution and function h that best suits the program under analysis. However, some assumption about the distribution family of the exposure is needed here if one is to estimate the relevant density. We thank an anonymous referee for bringing this to our attention.

² The Least Square Estimator can also be viewed as a special case of the GMM estimator.

³ We have dropped the sub-index of the parameters to simplify the notation. Equations are derived in detail in the Appendix.

⁴ For $t = 1$, these weights are exactly the same as those proposed by Fong et al. [13] for continuous treatments in cross sectional settings.

3 Correlation breaking IPTW

As with other PS methods, Robins' stabilized weights break the link between the exposure history and the time-varying covariates, guaranteeing the exposure history is unconfounded. That is, at any given stage of the program, we expect the histories of both exposure and the time-varying covariates to be uncorrelated in the weighted data [4, 18, 19].

Formally, under our time-varying setting, we can see this correlation-breaking property for all t as

$$E [SW_{it} (\underline{A}_{it}, \underline{X}_{it}) \underline{A}_{it} \underline{X}_{it}] = E (\underline{A}_{it}) E (\underline{X}_{it}). \quad (4)$$

Note that if we center both the exposure variable and the covariates, equation 4 equals zero. We propose to use this weighted cross moment as an extra moment condition in a GMM framework, which for $t = 1$ is the same Fong et al. [13] used in the cross-sectional setting.

Including equation 4 among the score conditions – the gradient of the log-likelihood with respect to θ – as moments, allows to account for both properties we would expect the stabilized weights satisfy in a single estimator: predict the exposure among recipients and break the link between the exposure and the time-varying covariates.

3.1 Estimation

There are many ways to specify the moment conditions for the GMM estimator depending on the assumptions one is willing to make in estimating the probability models and how to include the correlation-breaking condition, each with different computational burden. Assuming $A_t \mid \underline{A}_{t-1}, \underline{X}_t$ independent and identically distributed Gaussian (normal) across time periods, and making use of the fact that it also follows from the correlation-breaking property that current exposure is uncorrelated with the past, we have the following moment conditions,

$$g_{\theta} (\underline{A}_t, \underline{X}_t) = \begin{pmatrix} \frac{1}{\sigma^2} (A_{it} - (\underline{A}_{it-1}^{\top} \alpha + \underline{X}_{it}^{\top} \beta)) (\underline{A}_{it-1}, \underline{X}_{it}) \\ -\frac{1}{2\sigma^2} \left\{ 1 - \frac{1}{\sigma^2} [A_{it} - (\underline{A}_{it-1}^{\top} \alpha + \underline{X}_{it}^{\top} \beta)]^2 \right\} \\ \prod_{t'=0}^t \exp \left[\frac{1}{2\sigma^2} \left\{ -2A_{it'} (\underline{A}_{it'-1}^{\top} \alpha + \underline{X}_{it'}^{\top} \beta) + (\underline{A}_{it'-1}^{\top} \alpha + \underline{X}_{it'}^{\top} \beta)^2 \right\} \right] A_{it} (\underline{A}_{it-1}, \underline{X}_{it}) \end{pmatrix}. \quad (5)$$

Where the first two entries correspond to the first-order-conditions of the log-likelihood of the exposure assignment model (eq. 2) with respect to the model parameters (α, β) and σ^2 , respectively, and the third entry captures the correlation-breaking condition (eq. 4), after centering both the exposure and the covariates.

The GMM estimator of θ is obtained minimizing

$$\hat{\theta}_{CBSW} = \operatorname{argmin}_{\theta} \bar{g}_{\theta} (\underline{A}_t, \underline{X}_t)^{\top} \sum_{\theta} (\underline{A}_t, \underline{X}_t)^{-1} \bar{g}_{\theta} (\underline{A}_t, \underline{X}_t), \quad (6)$$

where $\bar{g}_{\theta}$ is the mean over i of g_{θ} , and $\sum_{\theta} (\underline{A}_t, \underline{X}_t)$ is the mean over i of the 3×3 covariance matrix.⁵

Even though assuming an independent correlation structure is important in keeping the score conditions and the covariance matrix as simple as possible, its role can be attenuated depending on how the exposure and covariates histories enter the exposure assignment model.⁶ Also, in practice, researchers faced with time

⁵ Equations are derived in detail in the Appendix.

⁶ The assumption leads naturally to the score conditions corresponding to a Pooled OLS estimation, across i and t , of the treatment assignment model, which reduces significantly the computational cost of the estimation method; otherwise, the number of parameters grows arithmetically with every point in time, along with computing time. The validity of these assumptions can be tested using standard tests such as the Kolgomorv-Smirnov. One way to assess the plausibility of the cross-time homoscedastic assumption is fitting panel-data linear models, by using Feasible Generalized Least Squares, with and without heteroscedasticity across time (instead of panels). Since iterated FGLS with time-level heteroscedasticity produces maximum-likelihood parameter estimates, it is rather straightforward to do a Likelihood Ratio Test with the number of time periods minus one degrees of freedom. Or, equivalently, regress the square of the residuals, rescaled so that their mean is 1, on a set of indicator variables for each – but one – of the time points, where the test statistic is the model's p-value; that is, the F-test of the overall significance. However, we cannot produce a similar test for the proposed method's estimates. In Section 4 we show the resulting performance of the pro-

series data serially correlated of unknown form can preserve the original time series structure resorting to block bootstrap, bootstrapping whole sequences at unit level i . We follow this strategy in Section 5.⁷

In programming the optimization procedure we have followed closely the original CBPS algorithm of Fong et al. [12]. In particular, we have modified their subroutine for Continuous CBPS to accommodate the longitudinal case estimating all the weights at once.

Given the complexity of the objective function in equation 6, for practical purposes we recommend to compare the performance of different optimization algorithms and starting points (use a search net) including the OLS solution, and apply non-derivative-based methods because numerical derivatives can become a serious hurdle for the optimization algorithm – at least for Stata's optimize function. Since the estimation of the covariance matrix of the sample moments requires the coefficient estimates, the optimization procedure can start with the maximum likelihood values of β , and successively update the GMM estimators throughout the optimization. This, of course, increases the computation time, to the point of making it impractical for simulations. The method can be made faster by fixing the estimate of the covariance matrix at the maximum likelihood value. In Section 4, all the simulation are carried out using the Nelder–Mead method, using the OLS solution as starting points, and no updating of the covariance (weight) matrix.

4 Simulation studies

In order to assess the empirical performance of these correlation-breaking stabilized weights, we carry out some simulation studies. In particular, we consider the case of three time periods, i. e., $t = 1, 2, 3$, and use five different sample sizes $n = 300, 600, 900, 1200$ and 1500 , corresponding to 100, 200, 300, 400 and 500 exposure-recipients across our time frame. In our setup, summarized in Figure 1, the treatment-generating process is a function of a vector of exogenous time-varying covariates, $\mathbf{X}_{it}$, and the previous history of exposure, $\mathbf{A}_{it-1}$. The outcome variable, Y , is a function of both entire histories of covariates and exposure.

We use this data generating process (DGP) to examine the performance of the proposed methodology, vis-à-vis the usual Ordinary Least Squares (OLS) approach, when faced with both misspecification of the treatment-assignment model – incorrectly specifying the lag structure, distributional family and cross-time homoscedasticity – and measurement error.

We use a similar simulation setup as Imai and Ratkovic [10] in two scenarios. In the first scenario, the simulation consists of four covariates. For time t , we use the covariates

$$\mathbf{X}_{it} = (Z_{it1} U_{it}, Z_{it2} U_{it}, |Z_{it3} U_{it}|, |Z_{it4} U_{it}|) \quad (7)$$

Figure 1: Data Generating Process in Simulation Study.

posed method under different misspecifications of the treatment assignment model regarding lag structure, measurement error, distributional family and cross-time heteroscedasticity.

⁷ We thank an anonymous referee for bringing this to our attention.

where each Z_{itk} is an *i.i.d.* draw from the standard normal distribution, and U_{it} is constructed as

$$U_{it} = \prod_{t'=1}^t (1 + .4 A_{it'-1}) \quad (8)$$

for $t = 2, 3$ and $U_{it} = 1$ for $t = 1$. The continuous treatment A_{it} is generated by

$$A_{it} = \sum_{t'=1}^t \left\{ \left(\frac{1}{2} \right)^{t'-1} (\alpha \mathbf{X}_{it'} + A_{it'-1}) \right\} + v_i \quad (9)$$

where $\alpha = (.8, -.2, .4, -.6)^\top$ and v_i is a standard normal disturbance – a cross-time homoscedastic normal model. All parameters were selected in an almost random manner looking only to guarantee a minimum dispersion. The continuous outcome Y is generated by a linear combination of the whole exposure and confounders histories such that

$$Y_{it} = 275 + 10 \sum_{t'=1}^t A_{it'} + \delta \sum_{t'=1}^t \mathbf{X}_{it'} + u_i \quad (10)$$

where $\delta = (-23.5, 12.5, 7.3, -17.4)^\top$ and u_i is another normal disturbance with mean zero and standard deviation five. Note that we have generated the outcome variable for every t , making it a function of both entire histories of covariates and exposure.⁸

There are many ways in which researchers can misspecify the treatment assignment model. In our simulations we explore two scenarios. In the first misspecified scenario, we use the following non-linear transformation of the covariates

$$X_{it} = \left(X_{it1}^2, \exp(X_{it2}), \log(X_{it3}), \frac{1}{X_{it4}} \right) \quad (11)$$

as measurement error, as well as incorrectly specify the lag structure of the treatment-assignment model using only the data from the immediately previous period such that

$$A_{it} = \alpha \mathbf{X}_{it} + A_{it-1} + v_i, \quad (12)$$

but maintain the correct outcome model. We report results of 1,000 simulated datasets.

In this first scenario, with every dataset, we fit a Pooled OLS model (GLM) as the treatment assignment model, using correct and incorrect model specifications as discussed above. Then we fit the same models using the proposed correlation-breaking IPTW methodology (CBIPTW) in two ways: first using both score and correlation-breaking conditions (Over-identified Correlation Breaking Inverse Probability of Treatment Weights (CBIPTW)) as described in section 2, and with the correlation-breaking condition alone (Just-identified CBIPTW), ignoring the score conditions. This gives us three different estimates according to the importance placed on both expected properties of the weights. The Generalized Linear Model (GLM) estimates obviate the correlation-breaking property, ignoring the subvector g_3 , while the just-identified CBIPTW pays no attention to predict exposure assignment, ignoring the subvector $(g_1, g_2)^\top$. Over-identified CBIPTW stands in the middle attending to both criteria.⁹ Finally, we regress in a pooled model the outcome variable on the cumulated exposure, $\sum_{t'=1}^t A_{it'}$ using the stabilized MSM weights. The coefficient that results from this last regression is compared with the numerical estimate obtained from a thousand simulations using large datasets ($n = 30,000$) and the true exposure assignment probabilities.

⁸ We haven't included these elements in the DGP depicted in Figure 1 in order not to over-complicate the diagram.

⁹ Note that both the over and just-identified CBIPTW work within the Gaussian probability model framework as described by the product in equation 3. The implied probability structure of the weights limits their flexibility as, at every point in time, the exposure probability models share the same parameters (a consequence of the iid assumption), where the latest weight is the product of all previous exposure probabilities (i. e., the probability of exposure history). Dropping these assumptions would allow for more flexibility to the weights, at the cost of losing the structure of the probability model. The limit case would be the one in which the weights are estimated without any reference to a probability model via Lagrange multipliers [13].

Table 1: Performance of Probability Weights Estimation Methods (Normal Homoscedastic Treatment-Assignment DGP).

Estimator	$n = 300$		$n = 600$		$n = 900$		$n = 1200$		$n = 1500$	
	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
Correct Specification										
True	1.19	3.79	0.94	3.41	0.70	3.38	0.83	3.31	0.47	3.09
GLM	1.19	3.61	0.92	3.05	0.70	3.23	0.77	3.26	0.46	2.96
CBIPW Over-identified	1.44	3.94	1.26	3.38	1.11	3.01	0.97	2.79	0.99	2.90
CBIPW Just-identified	1.47	3.93	1.07	3.33	1.07	3.06	1.16	3.34	0.98	2.85
Misspecification (measurement and lag structure)										
GLM	4.87	8.66	5.07	23.20	3.38	15.04	3.86	18.70	3.58	24.41
CBIPW Over-identified	3.89	6.42	4.49	8.87	3.65	7.27	3.73	7.35	3.77	8.66
CBIPW Just-identified	4.16	6.04	4.48	6.63	3.84	6.72	3.97	6.39	3.99	7.63

Table 2: Performance of Probability Weights Estimation Methods (Gamma Heteroscedastic Treatment-Assignment DGP).

Estimator	$n = 300$		$n = 600$		$n = 900$		$n = 1200$		$n = 1500$	
	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE	Bias	RMSE
Correct Specification										
True	1.84	4.97	1.52	4.23	1.27	4.09	1.04	4.15	1.00	4.07
Misspecification (distribution family and cross-time homoscedasticity)										
GLM	1.75	9.83	2.39	18.16	3.04	23.03	4.19	21.32	5.61	27.47
CBIPW Over-identified	2.56	4.84	2.53	4.85	2.47	4.76	2.50	4.38	2.53	4.62
CBIPW Just-identified	2.59	4.84	2.51	4.75	2.43	4.68	2.50	4.40	2.57	4.63

Table 1 presents the results for our five sample sizes. The columns show the bias and Root Mean Squared Error (RMSE) from the pooled weighted linear regression of the outcome on the cumulated exposure using the stabilized weights. The results are in line with those of previous Covariate Balance Propensity Scores for MSMs [10]. When the exposure assignment model is correctly specified, all methods have small bias and small RMSE. In contrast, when the model is misspecified, the bias and RMSEs are large and even grow with the sample size. However, the RMSE of the CBIPW estimates grow at a much smaller rate, thereby outperforming OLS.

In the second misspecified scenario, we use almost the same DGP, only this time v_i is a Gamma disturbance with different variance at every time-point: .5, 1 and 2 for t_1 , t_2 and t_3 respectively. In this case, the misspecification consists in fitting a normal cross-time homoscedastic model for the treatment A_{it} , with no measurement error and the correct lag structure.

Table 2 shows the result of this second misspecified scenario, fitting a Gaussian cross-time homoscedastic treatment-assignment model, just as in the first scenario, but this time to a Gamma, cross-time heteroscedastic DGP. Here too we see the RMSEs of the CBIPW grow at a much smaller rate, outperforming again the traditional OLS approach, this time also in bias.

What the statistical evidence shows is that the proposed method effectively improves the data fitting processes narrowing the estimator's dispersion, albeit at the expense of bias, resulting in more precise estimates than GLM. In a way, the extra balancing conditions work as a muffling effect that prevents the OLS estimates (that only consider the score conditions) from drifting away whenever the model is misspecified.

In short, Tables 1 and 2 tell a simple story for time-varying continuous interventions: if researchers are lucky enough to know the true DGP, and it happens to be Normal, they are probably better off with OLS. If, on the other hand, researchers have reasonable doubts with respect to the origin of their data, they are probably better off with CBIPW, given that it is less affected by (small) departures from the assumptions underlying the traditional OLS approach.

Together, what these two sets of simulations show is that the balancing conditions counterweight quite general misspecifications of the treatment-assignment probability model at roughly the same bias – or less – and quickly stabilizing an otherwise exploding RMSE.

5 Empirical application

As an empirical application of the proposed method, we estimate the impact of the Mexican universal health insurance program, Seguro Popular (SP), on key variables associated with the provision of healthcare services by the Secretaría de Salud – Ministry of Health – (SSA). We analyze data from the federal records on infrastructure and human resources employed by the SSA,¹⁰ and the coverage records of the program at a yearly frequency from 2001, the year prior the introduction of the program, to 2013. The information is aggregated to the Sanitary Jurisdiction (SJ) level, the most basic administrative units of SSA in charge of the operation of healthcare services and its programs. Time-varying covariates include doctors' offices, staffed and non-staffed hospital beds, physicians and nurses, both with and without day-to-day contact with patients, all per 1,000 people outside the social security network.¹¹ In order to make the common support assumption more credible and the estimates less model dependent, 39 out of 233 SJs are dropped from the analysis.¹²

The estimations of the treatment exposure model fit a single probability model, regressing the yearly increment of SP coverage (per 1,000 people outside the social security network) on the most recent set of covariates and past (cumulative) SP coverage.¹³ Looking at the degree of covariate balance achieved by the CBIPTW (Fig. 2, Appendix Table A.1), we see that it is improved in all cases, as medians are closer to zero, and that is best achieved using the just-identified approach, both on location and dispersion. It is worth noting that, in this case, fitting only the probability model (GLM) does not break the correlation between the exposure to treatment and three covariates: doctors' offices, physicians and nurses in day-to-day contact with patients.

Using weighted linear regressions to estimate the impact of (cumulated) SP coverage on the change in the (relative) numbers of doctors' offices, physicians and nurses (per 1,000 people outside the social security

10 Available from the National Health Information System's web site: <http://www.dgis.salud.gob.mx/contenidos/sinais/estadisticas.html>

11 As in all observational studies, the covariates should allow to approximate as closely as possible the characteristics that explain the units' treatment level as well as their outcome (its reaction to the intervention). Thus, covariate selection should aim to make the ignorability assumption more plausible. Even though there is, in principle, a risk of going too far following those who believe that "there is no reason to avoid adjustment for a [pretreatment] variable" [20, p. 76], this seems rather unlikely in most circumstances. However, if researchers find themselves forced to discriminate between covariates due to computational burden, empirical studies [21] suggest to privilege those with the strongest relation with the outcome. [22]. Indeed, sometimes is quite hard to inform model selection in practice, and this only gets more complicated in a dynamic context. There are, however, new tools that have come to replace earlier approaches, such as the stepwise regression, that required some amount of eyeballing the right set of covariates (see Ratkovic and Tingley [23]).

12 Behind common support considerations lies the concern for making comparisons between comparable units. The same logic applies in dynamic interventions, only this time comparisons are made across histories. The farther away recipient i covariate-history is from the rest, the less comparable it is since its particular history of $\mathbf{X}_i$ sets it apart from the rest of the recipients with values-in-time that no other possesses. King and Zeng [24, 25] have proposed the use of Gower's [26, 27] metric as the basis of this assessment in a cross-sectional context. The Gower distance between two points x_i and x_j is defined as the average absolute distance between the elements of the two points, divided by the range of the data:

$$G_{ij}^2 = \frac{1}{K} \sum_{k=1}^K \frac{|x_{ik} - x_{jk}|}{r_k}, \quad (13)$$

where the range is $r_k = \max(X_k) - \min(X_k)$ and the $\min$ and $\max$ functions return the smallest and largest elements, respectively, in the set including the k th element of the covariates X . We can see G^2 as the distance between two points expressed as a proportion of the distance across the data. Every analysis requires a cut-off value to distinguish those recipients too far away to be considered comparable to the rest. A natural possibility is to use the mean distance between the recipient under scrutiny and the rest of the observations. In this fashion a mean value of G_{ij}^2 over all j s equal to 0.5 means that the average distance between recipient i and all other recipients requires to travel the equivalent distance as 50 % of the way across the data set. Given the progressive way the weights are estimated, we have imposed restrictions on the data sequentially, keeping only those recipients whose covariates $\mathbf{X}_i$ are deemed "nearby" at every stage t . We kept only those SJs whose average distance to the rest was less than 0.5 at every stage. Another possibility to assess the distance between the covariate-histories would be through *latent class trajectory analysis*, something a lot more computationally intensive than Gower distance.

13 Although we do not explicitly model serial correlation, we try to preserve the original time series structure by bootstrapping at the Sanitary Jurisdiction level (block bootstrap).

Figure 2: Covariate-Balance for Time-varying SP coverage on Sanitary Jurisdictions.

Source: Own elaboration based on data from the National Health Information System (SINAIS) and the administrative records of Seguro Popular.

The bottom and top of each box are the first and third quartiles, and the band inside is the median. Whiskers represent the smallest observation still within 1.5 Inter Quartile Range of the lower quartile, and the largest observation still within 1.5 Inter Quartile Range of the upper quartile. Outside values are not shown.

* In day-to-day contact with patients. ** Without day-to-day contact with patients

Coefficients result from pooled linear regressions of each covariate on the yearly increment of SP coverage (affiliates as proportion of the population outside the social security network (IMSS, ISSSTE, Pemex, the Ministry of Defense or the Navy)), weighted accordingly to the pseudo-sample in question. UW corresponds to the unweighted 194 sanitary jurisdictions used for the analysis; GLM to the sample weighted using OLS and an empty model as stabilizer; CBIPW-OVER to the sample weighted using over-identified correlation-breaking inverse-probability weights; and CBIPW-JUST to the sample weighted using Just-identified CBIPW. Estimates are based on 1000 replications; bootstrapping takes into account the matching algorithm.

network) providing clinical care, we can see (Fig. 3, Appendix Table A.2) that in all probability the effect of SP is positive in all three outcome variables, and grows bigger the less correlated are the covariates with the treatment in the pseudo-sample.

Analysis by Bosch and Campos-Vzquez [28] on the impact of SP on the expansion of healthcare services ignored both the duration of the roll-out process and the different levels of treatment intensities observed during that process, applying a differences-in-differences approach where treatment was dichotomized (municipalities were considered as ‘treated’ once more than 10 individuals had been affiliated to SP). Huffman and van Gameren [19] improve on that by both considering a continuous treatment indicator (as in this paper) as well as a better-suited unit of analysis. While our results are somewhat similar to earlier findings, to the extent that positive impacts of SP on healthcare services is reported, our results – slightly bigger than Huffman and van Gameren [19] – should be considered more precise as we use an estimator both less model dependent and that outperforms the balance achieved through GLM.

Figure 2: (continued).

It is important to note that, in empirical applications, the least comparable the units are, that is to say, the least reasonable the common support assumption is, the more relevant the CBIPW become. This is because the CBIPW automatically compensate for the instability in the probability model generated by the lack of common support, something King and Zeng [24] identified as the risk of extreme counterfactuals. In this sense the CBIPW may also contribute to expand the validity of the results, allowing the analysis of larger data sets that otherwise would be unreasonable.

6 Conclusions

We have shown that the idea behind the Covariate Balancing Propensity Score [11] is flexible enough to cover also longitudinal settings with general treatment regimes. Here we provide an estimation method for Marginal Structural Models with continuous treatments whose main advantage is that it directly includes the correlation-breaking property of Robins' [1] stabilized weights in the estimation procedure. This new tool avoids the manual process of checking the correlation amongst the current exposure to treatment and the covariate histories in the weighted data and then respecifying. This new extension to the CBPS may help clear the way for those interested in causal analysis in these complex settings.¹⁴

Since the balance conditions do not require making any modeling assumptions in themselves, but only that the weights balance the covariates – something we know to be true regardless of the true DGP –, it can be argued that minimizing the correlation between covariates can be done more efficiently than with already available methods, such as the non-parametric extension of the CBPS to general treatment regimes [12] or

¹⁴ Stata code is available upon request from the authors.

Figure 3: Average Treatment Effects of SP coverage on Sanitary Jurisdictions.

Source: Own elaboration based on data from the National Health Information System (SINAIS) and the administrative records of Seguro Popular.

The bottom and top of each box are the first and third quartiles, and the band inside is the median. Whiskers represent the smallest observation still within 1.5 Inter Quartile Range of the lower quartile, and the largest observation still within 1.5 Inter Quartile Range of the upper quartile. Outside values are not shown.

* In day-to-day contact with patients

Coefficients result from pooled linear regressions (2328 obs.) of each outcome variable on SP coverage (affiliates as proportion of the population outside the social security network), weighted accordingly to the pseudo-sample in question. UW corresponds to the unweighted 194 sanitary jurisdictions used for the analysis; GLM to the sample weighted using OLS and an empty model as stabilizer; CBIPW-OVER to the sample weighted using over-identified correlation-breaking inverse-probability weights; and CBPS-EXACT to the sample weighted using Just-identified CBIPW. Estimates are based on 1000 replications; bootstrapping takes into account the matching algorithm.

Genetic Matching [29]. This point grows in importance as treatment assignment models gain in complexity.¹⁵ However, unlike these other methods, the methodology we propose here provides the researcher also with a probability model that outperforms the traditional GLM approach from which it is possible to start building bridges between MSMs and the decomposition methods literature, something that remains a pending task for future research.

Acknowledgment: Both authors are grateful for the comments received during various presentations of early versions, while the observations of two anonymous referees have importantly contributed to the final version. Special thanks are due to Dr Christian Fong for very helpful, clarifying comments on the CBPS Package.

Funding: A preliminary version of the paper formed part of the first author's PhD thesis, financially supported by a Consejo Nacional de Ciencia y Tecnología scholarship (CVU 173055), under the supervision of the second

¹⁵ We thank an anonymous referee for bringing this to our attention.

author. In the final phase the first author received support from the UNAM. Programa de Becas Posdoctorales en la UNAM, Becario del Programa Universitario de Estudios del Desarrollo.

Appendix I. Tables

Table A.1: Covariate-Balance for Time-varying SP coverage on Sanitary Jurisdictions^a.

Variable ^b	Sample ^c	Coef. ^d	Bootstrap Std. Err.	z	p > z	Normal-based [95 % Conf. Interval]	
Doctor's offices	UW	0.10	0.03	3.75	0.00	0.05	0.15
	GLM	0.05	0.02	2.89	0.00	0.02	0.09
	CBIPW-OVER	0.01	0.06	0.20	0.84	-0.10	0.13
	CBIPW-EXACT	0.01	0.02	0.47	0.64	-0.03	0.05
Staffed hospital beds	UW	0.00	0.04	0.02	0.98	-0.07	0.07
	GLM	0.01	0.02	0.46	0.65	-0.03	0.05
	CBIPW-OVER	0.02	0.05	0.43	0.67	-0.07	0.11
	CBIPW-EXACT	0.01	0.02	0.40	0.69	-0.04	0.05
Non-staffed hospital beds	UW	0.24	0.07	3.61	0.00	0.11	0.37
	GLM	-0.01	0.03	-0.32	0.75	-0.06	0.04
	CBIPW-OVER	-0.08	0.12	-0.67	0.50	-0.31	0.15
	CBIPW-EXACT	0.01	0.02	0.40	0.69	-0.04	0.06
Physicians with day-to-day contact with patients	UW	0.21	0.06	3.54	0.00	0.09	0.33
	GLM	0.13	0.04	3.66	0.00	0.06	0.21
	CBIPW-OVER	0.11	0.08	1.27	0.21	-0.06	0.27
	CBIPW-EXACT	0.02	0.04	0.55	0.58	-0.05	0.09
Physicians without day-to-day contact with patients	UW	-0.01	0.01	-1.37	0.17	-0.02	0.00
	GLM	0.00	0.00	-0.57	0.57	-0.01	0.01
	CBIPW-OVER	0.00	0.01	0.06	0.95	-0.03	0.03
	CBIPW-EXACT	0.00	0.00	0.25	0.80	-0.01	0.01
Nurses with day-to-day contact with patients	UW	0.25	0.09	2.88	0.00	0.08	0.41
	GLM	0.14	0.05	2.95	0.00	0.05	0.24
	CBIPW-OVER	0.11	0.13	0.82	0.41	-0.15	0.36
	CBIPW-EXACT	0.03	0.04	0.70	0.49	-0.06	0.12
Nurses without day-to-day contact with patients	UW	-0.02	0.01	-1.93	0.05	-0.04	0.00
	GLM	0.00	0.01	-0.61	0.54	-0.01	0.01
	CBIPW-OVER	0.00	0.02	-0.12	0.91	-0.04	0.04
	CBIPW-EXACT	0.00	0.01	0.37	0.71	-0.01	0.01

Source: Own elaboration based on data from the Sistema Nacional de Información en Salud (National Health Information System) and Seguro Popular administrative records

Estimates based on 1000 replications. Bootstrapping takes into account the matching algorithm.

^aSP coverage refers to affiliates as proportion of the population without traditional (work-related) health insurance (IMSS, ISSSTE, Pemex, the Ministry of Defense or the Navy).

^bPer thousand population without traditional health insurance.

^cUW corresponds to the unweighted 194 sanitary jurisdictions with the smallest mean distance to the rest of the observations in the data; GLM to the sample weighted using OLS and an empty model as stabilizer; CBIPW-OVER to the sample weighted using over-identified correlation-breaking inverse-probability weights, which gives equal importance to both correctly predict treatment-assignment and balancing covariates and CBPS-EXACT to the sample weighted using Just-identified CBIPW, which privileges covariate balance over the probabilistic model.

^dCoefficients result from a pooled regressions (2328 obs.) of each covariate on the yearly increment of SP coverage.

Table A.2: Average Treatment Effect of SP coverage on Sanitary Jurisdictions^a.

Outcome Variable ^b	Sample ^c	Coef. ^d	Bootstrap Std. Err.	z	p > z	Normal-based [95 % Conf. Interval]	
Doctor's offices	UW	0.20	0.01	16.90	0.00	0.18	0.23
	GLM	0.19	0.01	17.37	0.00	0.16	0.21
	CBIPW-OVER	0.19	0.04	5.31	0.00	0.12	0.26
	CBIPW-EXACT	0.19	0.01	16.03	0.00	0.17	0.22
Physicians with day-to-day contact with patients	UW	0.53	0.02	22.42	0.00	0.49	0.58
	GLM	0.51	0.02	21.50	0.00	0.46	0.56
	CBIPW-OVER	0.52	0.05	10.36	0.00	0.42	0.61
	CBIPW-EXACT	0.51	0.03	15.23	0.00	0.45	0.58
Nurses with day-to-day contact with patients	UW	0.73	0.04	19.00	0.00	0.66	0.81
	GLM	0.71	0.03	20.81	0.00	0.65	0.78
	CBIPW-OVER	0.72	0.05	15.46	0.00	0.63	0.82
	CBIPW-EXACT	0.72	0.04	18.09	0.00	0.64	0.80

Source: Own elaboration based on data from the Sistema Nacional de Información en Salud (National Health Information System) and Seguro Popular administrative records

Estimates based on 1000 replications. Bootstrapping takes into account the matching algorithm.

^aSP coverage refers to affiliates as proportion of the population without traditional (work-related) health insurance (IMSS, ISSSTE, Pemex, the Ministry of Defense or the Navy).

^bIncrement on the relative number (per thousand population without traditional health insurance) with respect to year prior the implementation of the program.

^cUW corresponds to the unweighted 194 sanitary jurisdictions with the smallest mean distance to the rest of the observations in the data; GLM to the sample weighted using OLS and an empty model as stabilizer; CBIPW-OVER to the sample weighted using over-identified correlation-breaking inverse-probability weights, which gives equal importance to both correctly predict treatment-assignment and balancing covariates and CBIPW-EXACT to the sample weighted using Just-identified CBIPW, which privileges covariate balance over the probabilistic model.

^dCoefficients result from pooled regressions (2328 obs.) of each outcome variable on (cumulated) SP coverage.

Appendix II. Derivation of Formulas

Equation 3 follows from plugging the Gaussian distribution,

$$\begin{aligned}
 SW_{it}(\underline{A}_t, \underline{X}_t) &= \prod_{t'=0}^t \frac{f_{\sigma^2}(\underline{A}_{it'})}{f_{\theta}(\underline{A}_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \\
 &= \prod_{t'=0}^t \frac{\frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2\sigma^2} \underline{A}_t^2\right\}}{\frac{1}{\sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2\sigma^2} [\underline{A}_t - (\underline{A}_{t-1}^{\top} \underline{\alpha} + \underline{X}_t^{\top} \underline{\beta})]^2\right\}} \\
 &= \prod_{t'=0}^t \exp\left\{\frac{1}{2\sigma^2} ([\underline{A}_t - (\underline{A}_{t-1}^{\top} \underline{\alpha} + \underline{X}_t^{\top} \underline{\beta})]^2 - \underline{A}_t^2)\right\} \\
 &= \prod_{t'=0}^t \exp\left\{\frac{1}{2\sigma^2} (\underline{A}_t^2 - 2\underline{A}_t (\underline{A}_{t-1}^{\top} \underline{\alpha} + \underline{X}_t^{\top} \underline{\beta}) + (\underline{A}_{t-1}^{\top} \underline{\alpha} + \underline{X}_t^{\top} \underline{\beta})^2 - \underline{A}_t^2)\right\} \\
 &= \prod_{t'=0}^t \exp\left[\frac{1}{2\sigma^2} \{-2\underline{A}_{it'} (\underline{A}_{it'-1}^{\top} \underline{\alpha} + \underline{X}_{it'}^{\top} \underline{\beta}) + (\underline{A}_{it'-1}^{\top} \underline{\alpha} + \underline{X}_{it'}^{\top} \underline{\beta})^2\}\right].
 \end{aligned}$$

Equation 4 follows from plugging the definition of Robins' stabilized weights,

$$E[SW_{it}(\underline{A}_{it}, \underline{X}_{it}) \underline{A}_{it} \underline{X}_{it}] = E\left[\prod_{t'=0}^t \frac{f(\underline{A}_{it'})}{f(\underline{A}_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \underline{A}_{it} \underline{X}_{it}\right]$$

$$\begin{aligned}
&= \int \left\{ \int \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'})} A_{it'} dF(\underline{A}_{it} | \underline{\mathbf{X}}_{it}) \right\} \underline{\mathbf{X}}_{it} dF(\underline{\mathbf{X}}_{it}) \\
&= E(\underline{A}_{it}) E(\underline{\mathbf{X}}_{it}),
\end{aligned}$$

where the last equality follows from assuming that current variables can only be affected by the past and the definition of conditional distribution, i. e.,

$$\begin{aligned}
f(\underline{A}_{it} | \underline{\mathbf{X}}_{it}) &= f(A_{it} | \underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) f(A_{it-1} | \underline{A}_{it-2}, \underline{\mathbf{X}}_{it}) \cdots f(A_{i1} | \underline{\mathbf{X}}_{i1}) \\
&= f(A_{it} | \underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) f(A_{it-1} | \underline{A}_{it-2}, \underline{\mathbf{X}}_{it-1}) \cdots f(A_{i1} | \underline{\mathbf{X}}_{i1}) \\
&= \prod_{t'=0}^t f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'}).
\end{aligned}$$

The first two lines of the score vector $\mathbf{g}_\theta$, equation 5, correspond to the first order conditions of the log-likelihood

$$l(\boldsymbol{\alpha}, \boldsymbol{\beta} | \underline{A}_t, \underline{\mathbf{X}}_t) = \frac{1}{2} \sum_{i=1}^N \left\{ \log(2\pi\sigma^2) + \frac{1}{\sigma^2} [A_{it} - (\underline{A}_{it-1}^\top \boldsymbol{\alpha} + \underline{\mathbf{X}}_{it}^\top \boldsymbol{\beta})]^2 \right\}. \quad (14)$$

In the third line we have conveniently rearranged the histories of exposure and covariates in the correlation-breaking condition for symmetry's sake, the proof of which follows the same general outline of equation 3.

The entries of the 3×3 covariance matrix $\sum_\theta(\underline{A}_t, \underline{\mathbf{X}}_t)$ are

$$\sum_\theta[1, 1] = \frac{1}{\sigma^2} (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it})^\top, \quad (15)$$

$$\sum_\theta[1, 2] = \sum_\theta[1, 2] = 0, \quad (16)$$

$$\sum_\theta[1, 3] = \sum_\theta[3, 1] = \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'})} (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it})^\top, \quad (17)$$

$$\sum_\theta[2, 2] = \frac{1}{2\sigma^4}, \quad (18)$$

$$\sum_\theta[2, 3] = \sum_\theta[3, 2] = \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'})} \left\{ \frac{-(\underline{A}_{it-1}^\top \boldsymbol{\alpha} + \underline{\mathbf{X}}_{it}^\top \boldsymbol{\beta})}{\sigma^2} \right\} (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}), \quad (19)$$

$$\begin{aligned}
\sum_\theta[3, 3] &= \prod_{t'=0}^{t-1} \frac{f(A_{it'})^2}{f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'})^2} \left\{ \sigma^2 + (\underline{A}_{it-1}^\top \boldsymbol{\alpha} + \underline{\mathbf{X}}_{it}^\top \boldsymbol{\beta})^2 \right\} \\
&\quad \exp \left\{ \frac{(\underline{A}_{it-1}^\top \boldsymbol{\alpha} + \underline{\mathbf{X}}_{it}^\top \boldsymbol{\beta})^2}{\sigma^2} \right\} (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it})^\top.
\end{aligned} \quad (20)$$

The upper left 2×2 submatrix of $\sum_\theta(\underline{A}_t, \underline{\mathbf{X}}_t)$ is well-known since corresponds to the covariance of the score conditions of the normal distribution. The rest of the elements are obtained by working out the respective covariances involving the third term of the score vector $\mathbf{g}_\theta$ as follows.

$$\begin{aligned}
\sum_\theta[1, 3] &= \sum_\theta[3, 1] \\
&= E[g_{13} g_{31} | (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it})] \\
&= E \left[\frac{1}{\sigma^2} (A_{it} - (\underline{A}_{it-1}^\top \boldsymbol{\alpha} + \underline{\mathbf{X}}_{it}^\top \boldsymbol{\beta})) (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'})} A_{it'} (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) | (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) \right] \\
&= \frac{1}{\sigma^2} \left\{ \int_{-\infty}^{\infty} (A_{it}^2 - A_{it} (\underline{A}_{it-1}^\top \boldsymbol{\alpha} + \underline{\mathbf{X}}_{it}^\top \boldsymbol{\beta})) \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{\mathbf{X}}_{it'})} f(A_{it} | \underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) dA_{it} \right\} \\
&\quad (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it}) (\underline{A}_{it-1}, \underline{\mathbf{X}}_{it})^\top
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{\sigma^2} \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \left\{ \int_{-\infty}^{\infty} (A_{it}^2 - A_{it} (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)) \frac{f(A_{it})}{f(A_{it} | \underline{A}_{it-1}, \underline{X}_{it})} f(A_{it} | \underline{A}_{it-1}, \underline{X}_{it}) dA_{it} \right\} \\
&\quad (\underline{A}_{it-1}, \underline{X}_{it}) (\underline{A}_{it-1}, \underline{X}_{it})^\top \\
&= \frac{1}{\sigma^2} \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \left\{ E(A_{it}^2) - E(A_{it}) (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta) \right\} (\underline{A}_{it-1}, \underline{X}_{it}) (\underline{A}_{it-1}, \underline{X}_{it})^\top \\
&= \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} (\underline{A}_{it-1}, \underline{X}_{it}) (\underline{A}_{it-1}, \underline{X}_{it})^\top,
\end{aligned}$$

where the last equality follows from assuming the marginal distribution of A_t to be Gaussian with mean zero, and thus $E(A_{it}) = 0$ and $E(A_{it}^2) = \sigma^2$.

$$\begin{aligned}
\sum_{\theta} [2, 3] &= \sum_{\theta} [3, 2] \\
&= E[g_2 g_3 | (\underline{A}_{it-1}, \underline{X}_{it})] \\
&= E \left[-\frac{1}{2\sigma^2} \left\{ 1 - \frac{1}{\sigma^2} [A_{it} - (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)]^2 \right\} \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} A_{it} (\underline{A}_{it-1}, \underline{X}_{it}) | (\underline{A}_{it-1}, \underline{X}_{it}) \right] \\
&= E \left[-\frac{1}{2\sigma^2} \left\{ 1 - \frac{1}{\sigma^2} [A_{it}^2 - 2A_{it} (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta) + (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)^2] \right\} \right. \\
&\quad \left. \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} A_{it} (\underline{A}_{it-1}, \underline{X}_{it}) | (\underline{A}_{it-1}, \underline{X}_{it}) \right] \\
&= E \left[\frac{1}{2\sigma^4} \{ A_{it}^2 - 2A_{it} (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta) + (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)^2 - \sigma^2 \} \right. \\
&\quad \left. \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} A_{it} (\underline{A}_{it-1}, \underline{X}_{it}) | (\underline{A}_{it-1}, \underline{X}_{it}) \right] \\
&= \frac{1}{2\sigma^4} \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \\
&\quad \left\{ \int_{-\infty}^{\infty} (A_{it}^3 - 2A_{it}^2 (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta) + A_{it} (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)^2 - A_{it} \sigma^2) \frac{f(A_{it})}{f(A_{it} | \underline{A}_{it-1}, \underline{X}_{it})} f(A_{it} | \underline{A}_{it-1}, \underline{X}_{it}) dA_{it} \right\} \\
&\quad (\underline{A}_{it-1}, \underline{X}_{it}) \\
&= \frac{1}{2\sigma^4} \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \left\{ E(A_{it}^3) - 2E(A_{it}^2) (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta) + E(A_{it}) [(\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)^2 - \sigma^2] \right\} \\
&\quad (\underline{A}_{it-1}, \underline{X}_{it}) \\
&= \frac{1}{2\sigma^4} \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \left\{ -2\sigma^2 (\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta) \right\} (\underline{A}_{it-1}, \underline{X}_{it}) \\
&= \prod_{t'=0}^{t-1} \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} \left\{ -\frac{(\underline{A}_{it-1}^\top \alpha + \underline{X}_{it}^\top \beta)}{\sigma^2} \right\} (\underline{A}_{it-1}, \underline{X}_{it}),
\end{aligned}$$

where the last two equalities follow from assuming the marginal distribution of A_t to be Gaussian with mean zero, making its third moment equal zero, $E(A_{it}^3)$, and its second moment equal its variance, $E(A_{it}^2) = \sigma^2$.

$$\begin{aligned}
\sum_{\theta} [3, 3] &= E[g_3 g_3 | (\underline{A}_{it-1}, \underline{X}_{it})] \\
&= E \left[\prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} A_{it} (\underline{A}_{it-1}, \underline{X}_{it}) \prod_{t'=0}^t \frac{f(A_{it'})}{f(A_{it'} | \underline{A}_{it'-1}, \underline{X}_{it'})} A_{it} (\underline{A}_{it-1}, \underline{X}_{it}) | (\underline{A}_{it-1}, \underline{X}_{it}) \right]
\end{aligned}$$

[illegible]

where the last equality follows from recognizing the Gaussian distribution of A_{it} with mean $-(\underline{A}_{it'}^\top \alpha + \underline{X}_{it'}^\top \beta)$ and variance σ^2 , whose second noncentral moment appears enclosed in the first set of brackets.

References

1. Robins JM. Marginal structural models versus structural nested models as tools for causal inference. In: Halloran M, Berry D, editors. *Statistical models in epidemiology, the environment, and clinical trials. The IMA Volumes in Mathematics and its Applications*. vol. 116. Berlin: Springer; 2000. p. 95–133, chap 2.
2. Robins JM, Hernn MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000;11(5):550–60.
3. Robins JM, Hernán MA. Estimation of the causal effects of time-varying exposures. In: Fitzmaurice G, Davidian M, Verbeke G, Molenberghs G, editors. *Longitudinal data analysis, Chapman & Hall/CRC Handbooks of Modern Statistical Methods*, Boca Raton: CRC Press; 2008. p. 553–99, chap 23.
4. Blackwell M. A framework for dynamic causal inference in political science. *Am J Polit Sci*. 2013;57(2):504–20.
5. Imbens GW. The role of the propensity score in estimating dose-response functions. *Biometrika*. 2000;87(3):706–10.
6. Imai K, Van Dyk DA. Causal inference with general treatment regimes. *J Am Stat Assoc*. 2004;99(467):854–66.
7. Hirano K, Imbens GW. The propensity score with continuous treatments. In: Gelman A, Meng XL, editors. *Applied Bayesian modeling and causal inference from incomplete-data perspectives: An essential journey with Donald Rubin's statistical family. Wiley Series in Probability and Statistics*. vol. 226164. New York: John Wiley & Sons, Ltd; 2004. p. 73–84, chap 7.
8. Flores CA, Flores-Lagunes A, Gonzalez A, Neumann TC. Estimating the effects of length of exposure to instruction in a training program: The case of Job Corps. *Rev Econ Stat*. 2012;94(1):153–71. http://www.mitpressjournals.org/doi/abs/10.1162/REST_a_00177.
9. Zhao S, van Dyk DA Imai K. Propensity-score based methods for causal inference in observational studies with fixed non-binary treatments. 2013. <https://imai.princeton.edu/research/files/gpscore.pdf>, mimeo.
10. Imai K, Ratkovic M. Robust estimation of inverse probability weights for marginal structural models. *J Am Stat Assoc*. 2015;110(511):1013–23. <https://doi.org/10.1080/01621459.2014.956872>.
11. Imai K, Ratkovic M. Covariate balancing propensity score. *J R Stat Soc, Ser B, Stat Methodol*. 2014;76(1):243–63.
12. Fong C, Ratkovic M, Hazlett C, Imai K. Cbpps: R package for covariate balancing propensity score. Comprehensive R Archive Network (CRAN). 2015. <https://cran.rproject.org/web/packages/CBPS/CBPS.pdf>.
13. Fong C, Hazlett C, Imai K. Parametric and nonparametric covariate balancing propensity score for general treatment regimes. 2015. <https://pdfs.semanticscholar.org/89d0/fe3d3d9a0270a60e368950c14498bd1c051e.pdf>, mimeo.
14. Splawa-Neyman J, Dabrowska D, Speed T, et al.. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Stat Sci*. 1990;5(4):465–72.
15. Rubin DB. Matching to remove bias in observational studies. *Biometrics*. 1973;29(1):159–83.
16. Robins J. A new approach to causal inference in mortality studies with a sustained exposure period-application to control of the healthy worker survivor effect. *Math Model*. 1986;7(9):1393–512.
17. Cole SR, Hernán MA. Constructing inverse probability weights for marginal structural models. *Am J Epidemiol*. 2008;168(6):656–64.
18. Huffman C. Causal inference with time-varying continuous interventions: Evaluating the Mexican universal health insurance program Seguro Popular. PhD thesis. Centro de Estudios Econmicos, El Colegio de Mxico, A.C. 2016.
19. Huffman C, van Gameren E. Time-varying continuous interventions: The roll-out of a universal health insurance program and its impact on the supply of health care services. 2017. mimeo.
20. Rosenbaum PR. Observational studies. In: *Observational studies*. Berlin: Springer; 2002. p. 1–17.
21. Brookhart MA, Schneeweiss S, Rothman KJ, Glynn RJ, Avorn J, Stürmer T. Variable selection for propensity score models. *Am J Epidemiol*. 2006;163(12):1149–56.
22. Ding P, Miratrix LW. To adjust or not to adjust? sensitivity analysis of m-bias and butterfly-bias. *J Causal Inference*. 2015;3(1):41–57.
23. Ratkovic M, Tingley D. Sparse estimation and uncertainty with application to subgroup analysis. *Polit Anal*. 2017;25(1):1–40. <https://doi.org/10.1017/pan.2016.14>.
24. King G, Zeng L. The dangers of extreme counterfactuals. *Polit Anal*. 2006;14(2):131–59.
25. King G, Zeng L. When can history be our guide? The pitfalls of counterfactual inference. *Int Stud Q*. 2007;51(1):183–210.
26. Gower JC. Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*. 1966;53(3–4):325–38.
27. Gower J. A general coefficient of similarity and some of its properties. *Biometrics*. 1971;27(4):857–71.
28. Bosch M, Campos-Vzquez RM. The trade-offs of social assistance programs in the labor market: The case of the “Seguro Popular” program in Mexico. Documento de Trabajo DT-2010-12. El Colegio de Mxico, A.C. 2010. <http://cee.colmex.mx/documentos/documentos-de-trabajo/2010/dt201012.pdf>.
29. Diamond A, Sekhon JS. Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. *Rev Econ Stat*. 2013;95(3):932–45.