

Cautionary Tales on Synthetic Controls in Survival Analyses

Alicia Curth*

University of Cambridge

AMC253@CAM.AC.UK

Hoifung Poon

Microsoft Research

HOIFUNG@MICROSOFT.COM

Aditya Nori

Microsoft Research

ADITYA.NORI@MICROSOFT.COM

Javier González

Microsoft Research

GONZALEZ.JAVIER@MICROSOFT.COM

Editors: Francesco Locatello and Vanessa Didelez

Abstract

Synthetic control (SC) methods have gained rapid popularity in economics recently, where they have been applied in the context of inferring the effects of treatments on standard continuous outcomes assuming linear input-output relations. In medical applications, conversely, *survival* outcomes are often of primary interest, a setup in which both commonly assumed data-generating processes (DGPs) and target parameters are different. In this paper, we therefore investigate whether and when SCs could serve as an alternative to matching methods in survival analyses. We find that, because SCs rely on a linearity assumption, they will generally be biased for the true expected survival time in commonly assumed survival DGPs – even when taking into account the possibility of linearity on another scale as in accelerated failure time models. Additionally, we find that, because SC units follow distributions with lower variance than real control units, summaries of their distributions, such as survival curves, will be biased for the parameters of interest in many survival analyses. Nonetheless, we also highlight that using SCs can still improve upon matching whenever the biases described above are outweighed by extrapolation biases exhibited by imperfect matches, and investigate the use of regularization to trade off the shortcomings of both approaches.

Keywords: Synthetic control methods, Matching, Survival analysis, Synthetic control group

1. Introduction

The availability of a suitable control group for evaluating the effectiveness of a treatment, policy or other intervention is the backbone of empirical causal inference. Consider the common scenario in which an analyst has access to a small sample of treated instances – e.g. a single state that experienced a policy shock in economic applications, or a small group that received a novel treatment in an early stage clinical trial – as well as to a much larger observational sample of potential controls that did not experience the intervention of interest (Rosenbaum and Rubin, 1985b). A popular approach to use such observational samples to create a control group *tailored to the treated group at hand* has been to rely on *matching* (Rubin, 1973; Stuart, 2010): pairing each treated unit with the closest control unit in terms of (a summary of) its observed characteristics. While an intuitively appealing approach, success requires close matches for all units (Rosenbaum and Rubin, 1985a), which becomes challenging when the dimensionality of characteristics grows only moderately large (Stuart, 2010).

* Work done while an Intern with Microsoft Research.

As an alternative to matching, synthetic control (SC) methods (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Abadie, 2021) have gained rapid popularity in the economic policy evaluation literature recently (Athey and Imbens, 2017). These methods, instead of matching to a single real control unit, construct *synthetic* units that are a weighted average of multiple existing control units, and as such may be able to match the pre-treatment characteristics of the treated units much more closely. Importantly, the literature on these methods has, to the best of our knowledge, focused exclusively on continuous outcomes assumed linear in inputs – such as GDP per capita (Abadie and Gardeazabal, 2003; Pinotti, 2015), consumption (Abadie et al., 2010), crime rates (Donohue et al., 2019; Cunningham and Shah, 2018) or stock returns (Acemoglu et al., 2016) – because these are usually of primary interest in economic applications.

The outcome of interest in many medical applications, on the other hand, is a so-called time-to-event or survival outcome (Austin and Stuart, 2015), which measures the time elapsed until occurrence of an event (for example, disease progression, an adverse side-effect, or even death). Such survival outcomes are often assumed to follow different data-generating processes (DGPs) than standard continuous outcomes – instead of linear relationships between characteristics and outcomes directly, relationships are assumed to be more likely on the log-time scale (as in accelerated failure time models (Wei, 1992)) or log-hazard scale (as in proportional hazards models (Cox, 1972)). Further, the target parameters of interest in survival analyses are often also different: instead of only considering differences in expected average survival times due to treatment, other summaries of the survival distribution, relating to the difference of survival curves or ratios of hazards, are of primary interest in many medical studies (Clark et al., 2003).

Contributions. In this paper, we investigate whether and when synthetic controls could serve as an alternative to matching methods in applications where time-to-event outcomes are of interest; to the best of our knowledge, this is the first work to consider this question. Through our theoretical investigation, we discover multiple obstacles that prohibit the desirable properties that SCs exhibit for standard outcomes to carry over smoothly to the survival context. In particular, we come across two main obstacles: we find that (i) in commonly assumed survival DGPs, SCs will generally be biased for the true expected survival time (even when taking into account the possibility of linearity on another scale in accelerated failure time models) [Section 4] and (ii) even if they were unbiased for the mean survival time, SCs follow distributions with lower variance than real control units – which is why other summaries of their distribution, such as survival curves, will usually be biased for the true parameters [Section 5]. Nonetheless, we also show that using SCs can still improve upon the use of matching whenever the biases incurred above are outweighed by extrapolation biases exhibited by imperfect matches. Finally, we therefore investigate the use of regularization schemes trading off the shortcomings of both approaches [Section 6] and illustrate how some of the theoretical arguments translate to real data [Section 7].

Remark: Note that the goal of this paper is not to promote the use of any specific method, but rather to start a discussion on the general applicability of SC-inspired methods in survival analyses due to their importance in medicine. As such, we present first analyses and illustrations of problems arising in this context, with the hope of encouraging future methodological work in this space.

2. Background: Synthetic Control in the standard outcome setting

Assume a single individual with pre-treatment characteristics $X_* \in \mathbb{R}^d$ is exposed to a treatment $A_* = 1$, and we observe their outcome $Y_*(1) \in \mathbb{R}$; this is the so-called treated unit. We also have

access to characteristics X_j and outcomes $Y_j(0) \in \mathbb{R}$ of m individuals that were not exposed to treatment ($A_j = 0$); this is the pool of control units at our disposal. The aim of synthetic control methods is to infer the effect of the treatment by comparing the treated outcome $Y_*(1)$ to that of a *synthetic control unit* \hat{Y}_*^{SC} for the treated unit – capturing what would have happened to the treated unit had they not been exposed to treatment, i.e. what would have been their *potential* outcome $Y_*(0)$. Such synthetic controls (SCs) are constructed from the outcomes from the control group as a weighted average

$$\hat{Y}_*^{SC} = \sum_{j=1}^m w_j Y_j(0) \quad (1)$$

where $0 \leq w_j \leq 1$ and $\sum_{j=1}^m w_j = 1$, with the goal that this synthetic control unit approximates the expected outcome without treatment, $\mu_0(x) = \mathbb{E}[Y(0)|X = x]$, well, i.e. that $\mathbb{E}[\hat{Y}_*^{SC}] \approx \mu_0(X_*)$.

The canonical synthetic control method (Abadie and Gardeazabal, 2003; Abadie et al., 2010; Abadie, 2021) aims to reconstruct the pre-treatment characteristics of the treated units perfectly, i.e. to choose $w = (w_1, \dots, w_m)$ such that $\sum_{j=1}^m w_j X_j = X_*$. As this solution is not always part of the convex hull of observed controls (or even feasible at all), they instead solve

$$\arg \min_w \|X_* - \sum_{j=1}^m w_j X_j\|^2 \text{ subject to } 0 \leq w_j \leq 1, \sum_{j=1}^m w_j = 1 \quad (2)$$

Throughout, we will refer to synthetic controls that fulfill $\sum_{j=1}^m w_j X_j = X_*$ exactly as *perfect*.

Underlying assumptions. Under a selection on observables assumption ensuring $Y(0) \perp\!\!\!\perp A|X$ and an overlap assumption ensuring that $X|A = 1$ is contained in the support of $X|A = 0$, the expected untreated outcome is identified and can be estimated using estimators of the form (1), as discussed in e.g. Kellogg et al. (2021)¹. Then, the SC approach of Eq. (2) is well-motivated in settings where $\mu_0(x)$ is *linear in x* , because then² we have (for perfect synthetic control units) that

$$\mathbb{E}[\hat{Y}_*^{SC}] = \sum_{j=1}^m w_j \mathbb{E}[Y_j(0)|X_j] = \sum_{j=1}^m w_j \mu_0(X_j) = \mu_0(\sum_{j=1}^m w_j X_j) = \mu_0(X_*) \quad (3)$$

Clearly, Eq. (3) holds for linear models $Y_i(0) = X_i\beta + \epsilon$ with mean-zero independent errors ϵ . Further, Abadie et al. (2010) show that when X includes pre-treatment outcomes, Eq. (3) also holds for linear factor models allowing for some shared unobserved factors in the DGP³.

Alternative: Nearest neighbor matching. A matching estimator can also be written in the form of Eq. (1), but assigns *all* weight to a single unit j^* satisfying $\arg \min_{j \in [m]} \|X_* - X_j\|$ – i.e. the control unit that is X_* 's nearest neighbor (NN) in terms of observed characteristics (sometimes also coarsened summaries thereof, e.g. the propensity score (Rosenbaum and Rubin, 1985b)). This outputs $\hat{Y}_*^{match} = Y_{j^*}(0) = \sum_{j=1}^m w_j Y_j(0)$ with $w_{j^*} = 1$ and $w_j = 0$ for $j \neq j^*$. Unless there is a perfect match among the control units, we will have that $\|X_* - X_{j^*}\| > 0$ and therefore matching estimators will have to *extrapolate*. In the case of a linear model for a continuous outcome, they would thus incur bias $(X_* - X_{j^*})\beta$ directly proportional to the distance to the match in input space.

3. Problem setup: Survival outcomes and analyses

To consider how SC methods could be applied in the survival context in the remainder of the paper, we first introduce the problem setup of interest. Survival analyses are concerned with modelling

1. Recent work (Shi et al., 2021, 2022; Zeitler et al., 2023) also investigates other identification assumptions for SCs.
2. Shi et al. (2022) show that linearity assumptions can sometimes be circumvented by instead assuming a more fine-grained model, which, as we discuss in detail in Appendix A, is not applicable in the survival setting we study.
3. The fact that the inclusion of pre-treatment outcomes could correct for some forms of unobserved confounding in this way may be a major contributing factor to the popularity of SC methods in economics.

time-to-event outcomes – this could be the time T elapsed until death, disease progression or another adverse event occurs. In addition to estimating standard quantities like the expected potential time $T(A)$ until an event occurs under assignment of different treatments A and given different patient characteristics X (i.e. estimating $\mu_a(x) = \mathbb{E}[T(a)|X = x]$), survival analyses are often interested in modelling quantities capturing more of the *distribution* of survival times (e.g. contrasts of the survival functions $S_a(t; x) = \mathbb{P}(T(a) > t|X = x)$ or hazard functions $h_a(t; x) = \mathbb{P}(T(a) = t|T(a) \geq t, X = x)$) to assess the effects of treatments in this context.

Assumptions. As in the standard setup discussed above, we assume we have access to a control sample $\{X_j, T_j(0)\}_{j=1}^m$ (having received treatment $A_j = 0$), and rely on the overlap assumption and an unconfoundedness assumption (ensuring here that $T(0) \perp\!\!\!\perp A|X$). Further, we restrict our focus to control time-to-event outcomes generated from commonly considered accelerated failure time (AFT) models which assume event times are generated from models of the form

$$\log T_i(0) = X_i\beta + \epsilon_i \tag{4}$$

where the ϵ_i are i.i.d. random variables generated from some common distribution (Wei, 1992) that then determines the shape of survival functions. A log-normal AFT has $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$; other popular distributions for ϵ_i include the logistic distribution and extreme-value distributions, which encompass also the Weibull survival model as a special case.


Remark: Censored outcomes. In survival analyses, observations are often censored (i.e. lost to follow-up) before any event can be observed. Then, we would only observe $T^- = \min(T, C)$ the smaller of event time T and censoring time C . For ease of exposition, we disregard censoring in the discussion below. Under the standard, commonly adopted, censoring-at-random (CAR) assumption (van der Laan and Robins, 2003), which requires that censoring and event times are independent when conditioning on observed covariates X and treatment A (or equivalently, that censoring time C is not correlated with the noise term ϵ in Eq. (4)), the discussion below also directly applies to analyses with censored data, where SC units can then be constructed from uncensored controls only⁴.

Remark: Absence of repeated outcomes. Unlike economic applications where outcomes are often recorded in repeated panel data structures, outcomes in survival analyses are often terminal (e.g. time to death). Therefore, X would usually not include pre-treatment outcomes in the survival context, and hence SCs cannot account for unobserved factors in the same way they are sometimes assumed to in economic panel models. Note, however, that much of the recent methodological literature on SC methods (Abadie and L’hour, 2021; Kellogg et al., 2021) considers a general setup identical to ours, differing only in what kind of information could (not) be included in X in spirit.

4. Obstacles for Synthetic Controls in Survival Analyses (Part 1): Nonlinearity and biases in estimating expected survival time

In this section, we begin our investigation by considering which biases arise in the canonical synthetic control problem of estimating expected outcomes for a single treated unit X_* as a consequence

4. To see why this is true, let $\delta = \mathbf{1}\{C < T\}$ be a binary indicator for whether an individual’s survival time was observed; if $\delta = 0$ their time was censored. Then, we can construct synthetic control units from only uncensored control units and write them as $\hat{T}_*^{SC, \delta=1} = \sum_{j=1}^m w_j^\delta T_j^-$ where $w_j^\delta = 0$ for all $\delta_j = 0$. Note that the CAR assumption implies that $\mathbb{E}[T(0)|X = x, \delta = 0] = \mathbb{E}[T(0)|X = x, \delta = 1] = \mathbb{E}[T(0)|X = x] = \mu_0(X)$, thus we will have that $E[\hat{T}_*^{SC, \delta=1}] = \sum_{j: \delta_j=1} w_j^\delta \mathbb{E}[T(0)|X = X_j, \delta = 1] = \sum_{j: \delta_j=1} w_j^\delta \mu_0(X_j)$. Thus, also in the presence of CAR-censoring, the bias of our synthetic control estimate is determined by the quality of the synthetic match (i.e. the distance of $\sum_{j: \delta_j=1} w_j^\delta X_j$ from X_*) and the linearity of $\mu_0(x)$, as in the absence of censoring.



(a) Interpolation bias is high on the non-linear part of the exponential. (b) Interpolation bias is negligible on the approximately linear part of the exponential.

Figure 1: **Stylized illustration of interpolation and extrapolation biases** using a noise-free AFT model ($T = \exp(X)$).

of switching from simple linear DGPs often assumed for outcomes in economics to time-to-event outcomes generated from AFT models. We then consider issues relating to the distribution of synthetic controls and targeting of other parameters of interest in Section 5.

4.1. Interpolation Biases of Synthetic Control in AFT models

The AFT assumption implies that $T_i = \exp(X_i\beta + \epsilon)$ – thus, even for perfect synthetic controls there will likely be substantial bias as $\mu_0(x)$ is not linear in x . Why is that? In the context of standard continuous outcomes, Kellogg et al. (2021) noted that the bias of any estimators of the form Eq. (1) can always be written as

$$\mu_0(X_*) - \sum_{j=1}^m w_j \mu_0(X_j) = \underbrace{[\mu_0(X_*) - \mu_0(\sum_{j=1}^m w_j X_j)]}_{\text{Extrapolation Bias}} + \underbrace{[\mu_0(\sum_{j=1}^m w_j X_j) - \sum_{j=1}^m w_j \mu_0(X_j)]}_{\text{Interpolation Bias}}$$

Further, Kellogg et al. (2021) highlight that synthetic control estimators minimize extrapolation biases (as $X_* \approx \sum_{j=1}^m w_j X_j$), but might suffer from interpolation biases if $\mu_0(\sum_{j=1}^m w_j X_j) \neq \sum_{j=1}^m w_j \mu_0(X_j)$ – i.e. if $\mu_0(x)$ is nonlinear in x . A matching estimator, conversely, which assigns all weight to the single closest control unit j^* and thus uses $w_{j^*} = 1$, will suffer from some extrapolation biases whenever there is no perfect match but has no interpolation bias by construction.

Can we quantify this interpolation bias of perfect synthetic controls in AFT models? Assume the simplest setting where there is no noise, i.e. all $\epsilon_i = 0$. Then, $\mathbb{E}[T_*(0)] - \mathbb{E}[\hat{T}_*^{SC}] = \exp(X_*\beta) - \sum_{j=1}^m w_j \exp(X_j\beta) \leq 0$ and synthetic controls will *overestimate* the expected outcome; this bias will be substantial if the $X_j\beta > 1$. This possibility for interpolation bias is also illustrated in Fig. 1, where we see that the relative bias incurred due to interpolation versus extrapolation indeed depends on *where* on the exponential the treated unit and the control units are located.

4.2. Does constructing synthetic controls on the log-time scale overcome this bias?

It is obvious that the source of interpolation bias here is a result of the AFT assumption: T is *not* linear in X – but $\log(T)$ *is*! It is thus a natural next step to consider whether we can exploit this

assumption, and apply synthetic control weights to log-time instead. Indeed, for perfect synthetic control weights we have that

$$\widehat{\log(T_*)}^{SC} = \sum_{j=1}^m w_j \log(T_j(0)) = \sum_{j=1}^m w_j (X_j \beta + \epsilon_j) = X_* \beta + \sum_{j=1}^m w_j \epsilon_j \quad (5)$$

and further, because the ϵ_j are i.i.d. by assumption and $\sum_{j=1}^m w_j = 1$

$$\mathbb{E}[\widehat{\log(T_*)}^{SC}] = X_* \beta + \sum_{j=1}^m w_j \mathbb{E}[\epsilon_j] = X_* \beta + \mathbb{E}[\epsilon_*] \sum_{j=1}^m w_j = X_* \beta + \mathbb{E}[\epsilon_*] = \mathbb{E}[\log(T_*(0))] \quad (6)$$

Thus, under AFT assumptions, SCs *can* be used to obtain unbiased estimates of *log-survival times*.

Is this all we needed to make SCs fit for use in survival analyses? Unfortunately, no – unless log-time is the outcome of interest! If, however, the quantity of interest is the time on the original scale, we might be tempted to simply consider $\hat{T}_*^{log-SC} = \exp(\widehat{\log(T_*)}^{SC}) = \exp(\sum_{j=1}^m w_j \log(T_j))$. Unfortunately, this will not necessarily lead to unbiased estimates of $T_*(0)$ because for two random variables A and B , $\mathbb{E}[A] = \mathbb{E}[B]$ does not generally imply $\mathbb{E}[f(A)] = \mathbb{E}[f(B)]$ for nonlinear functions f . In particular, if we were to consider a simple log-normal AFT with $\log(T_i(0)) \sim \mathcal{N}(\mu_i, \sigma^2)$, then $\mathbb{E}[T_i(0)] = \exp(\mu_i + \frac{\sigma^2}{2})$ – thus matching the expected value μ_i of log-time alone clearly is not enough. Instead, more of the distribution – in this case also the variance – of true control units would need to be matched. Below, we thus investigate the distribution of synthetic controls further and return to the question of the bias of \hat{T}_*^{log-SC} in Section 5.2.

5. Obstacles for Synthetic Controls in Survival Analyses (Part 2): Distributional mismatches between synthetic & real controls and their consequences

Next, we consider the distribution of synthetic control units beyond just their mean. In Section 5.1, we begin by highlighting that synthetic control units will generally have substantially lower variance than real units. In Section 5.2, we show what this means for the bias of synthetic controls constructed on log-scale (as discussed in Section 4.2 above). In Section 5.3 we then discuss how differences in distribution between synthetic and real units will bias the estimation of survival curves.

5.1. Aside: On the variance of synthetic controls

Clearly, Eq. (6) can be generalized to any model in which an invertible transformation of f of the random variable Y_i follows a linear model $f(Y_i) = X_i \beta + \epsilon_i$ with independent error terms ϵ_i with equal expected value $\mathbb{E}[\epsilon_i] = \mu_\epsilon$: In any such model, for a perfect synthetic control unit

$$\mathbb{E}[\widehat{f(Y_*)}^{SC}] = \mathbb{E}[\sum_{j=1}^m w_j f(Y_j)] = X_* \beta + \mu_\epsilon \sum_{j=1}^m w_j = \mathbb{E}[f(Y_*)] \quad (7)$$

Thus, the expected value can always be matched on the f -scale. What about the variance of such perfect synthetic controls? As synthetic controls are simply weighted averages, it is easy to see that, for independently generated homoskedastic error-terms with variance σ_ϵ^2 ,

$$\text{Var}(\widehat{f(Y_*)}^{SC}) = \text{Var}(\sum_{j=1}^m w_j \epsilon_j) = \sigma_\epsilon^2 \sum_{j=1}^m w_j^2 \quad (8)$$

Further, because synthetic control weights are restricted to the convex hull of the data (thus all $0 \leq w_j \leq 1$ and $\sum_{j=1}^m w_j = 1$), we have that $\min(\sum_{j=1}^m w_j^2) = \frac{1}{m}$ and $\max(\sum_{j=1}^m w_j^2) = 1$,




Figure 2: **Illustration of the effect of error variance on bias of different estimators** for a log-normal AFT with $\log(T) = X + \epsilon$ and $\epsilon \sim \mathcal{N}(0, \sigma)$.

which are attained at $w_j = \frac{1}{m}$ (the sample average) and $w_j = 1\{j = k\}$ (for k the index of a single control unit), respectively. Thus, $\text{Var}(\widehat{f(Y_*)}^{SC}) \leq \text{Var}(f(Y_*))$, and the variance of synthetic control units $\widehat{f(Y_*)}^{SC}$ is strictly smaller than that of real outcomes $f(Y_*)$ whenever there is more than one contributor (which would generally be the case unless there is a perfect match for X_* in the control sample). Note that this observation holds generally for synthetic controls, including the standard synthetic control setting where f is the identity function.

5.2. Bias of log-time synthetic controls for the expected survival time in log-normal AFT models (ct'd from Section 4.2)

For the special case of log-normal AFT models, we can use Eq. (8) to deduce the bias of perfect synthetic controls $\hat{T}_*^{log-SC} = \exp(\widehat{\log(T_*)}^{SC})$ constructed on log-scale. In particular, due to the properties of the normal distribution, we know that $\widehat{\log(T_*)}^{SC} \sim \mathcal{N}(\mu_*, \sigma_\epsilon^2 \sum_{j=1}^m w_j^2)$, where $\mu_* = \mathbb{E}[\log(T_*)] = X_*\beta + \mathbb{E}[\epsilon]$. Therefore, using properties of log-normals, $\mathbb{E}[\hat{T}_*^{log-SC}] = \exp(\mu_* + \frac{\sigma_\epsilon^2}{2} \sum_{j=1}^m w_j^2)$, which is a biased estimator for $T_*(0)$ whenever there is more than one contributor because then $\sum_{j=1}^m w_j^2 < 1$ while the true expected value of $T_*(0)$ is $\mathbb{E}[T_*(0)] = \exp(\mu_* + \frac{\sigma_\epsilon^2}{2})$.⁵

The estimator $\hat{T}_*^{log-SC} = \exp(\widehat{\log(T_*)}^{SC})$ will thus generally *underestimate* the expected (mean) survival time, and this bias will be more severe a) the larger the error variance is relative to μ_* (i.e. the lower the signal-to-noise ratio) and b) the more contributors there are (this is because $\sum_{j=1}^m w_j^2$ becomes smaller as the w_j become more uniform and *less sparse*). We illustrate the bias induced by variance-mismatches for a simple log-normal AFT in Fig. 2.

In Fig. 2, we also observe that the relative performance of a nearest neighbor match and the standard synthetic control is unaffected by the error variance, while the relative performance of synthetic control constructed at log-time scale greatly deteriorates as the variance increases. Further insight on this special case can therefore be gained by considering explicit expressions for the bias:

5. Note that, if time is assumed log-normal and σ_ϵ can be estimated or is known, \hat{T}_*^{log-SC} could therefore in principle be debiased by multiplying it by a factor $\frac{\exp(\sigma_\epsilon^2/2)}{\exp(\sigma_\epsilon^2/2 \sum_{j=1}^m w_j^2)}$. We do not investigate this avenue further because usability of this specific debiased estimator in practice requires both i) availability of a perfect synthetic control unit (as otherwise the mean μ_* is also not matched) and ii) *log-normality* of the true time (as different AFT models would require different correction factors depending on the underlying relationship between the distributions of log-time and time), which is a very strong assumption and unlikely to hold in many applications.

$$\mathbb{E}[T_*(0)] - \mathbb{E}[\widehat{\exp(\log(T_*)^{SC})}] = \underbrace{\left(\exp\left(\frac{\sigma^2}{2}\right) - \exp\left(\frac{\sigma^2 \sum_{j=1}^m w_j^2}{2}\right) \right)}_{\text{always } \geq 0} \exp(X_*\beta) \quad (9)$$

$$\mathbb{E}[T_*(0)] - \mathbb{E}[\hat{T}_*^{SC}] = \exp\left(\frac{\sigma^2}{2}\right) \underbrace{\left(\exp(X_*\beta) - \sum_{j=1}^m w_j \exp(X_j\beta) \right)}_{\text{always } \leq 0 \text{ as } \exp(\cdot) \text{ is convex}} \quad (10)$$

$$\mathbb{E}[T_*(0)] - \mathbb{E}[\hat{T}_*^{match}] = \exp\left(\frac{\sigma^2}{2}\right) \underbrace{\left(\exp(X_*\beta) - \exp(X_{match}\beta) \right)}_{\leq \geq 0 \text{ depending on exponents}} \quad (11)$$

Thus, for log-normal AFTs, a standard synthetic control estimator will *overestimate* the expected time while a synthetic control constructed at log-time scale will *underestimate* the expected time. Further, the relative performance of standard synthetic control and matching is determined by the distance of the individual selected control units from the treated unit X_* , while the bias of a (perfect) synthetic control constructed at log-time scale is entirely driven by differences in variance due to multiple contributors. More generally, regardless of whether $T(0)$ actually follows any AFT model, we always have $\hat{T}_*^{log-SC} = \exp(\sum_{j=1}^m w_j \log(T_j(0))) \leq \exp(\log \sum_{j=1}^m w_j (T_j(0))) = \hat{T}_*^{SC}$ because of concavity of the log. That is, a log-time synthetic control will always output an estimate of time that is lower than the standard synthetic control estimate.

Remark: Over- vs underestimation. Comparing standard and log-time synthetic controls, a natural question that arises is thus whether one would prefer an estimator that overestimates or one that underestimates expected survival time. While neither is ideal, the answer may be application-dependent: if the observational bias one is trying to correct for is that the control sample overall has lower outcomes already (e.g. because less healthy individuals are less likely to receive a promising yet invasive treatment), using an estimator that may systematically underestimate time will likely exacerbate this bias and may therefore be less desirable (and vice versa).

5.3. Bias in survival curve estimation

Finally, we now consider biases in estimating other summaries of the survival distribution for a control sample – e.g. marginal survival curves $S_0(t) = \mathbb{P}(T_*(0) > t)$ (marginalised over the distribution of covariates X_* among the treated). To do so, we move beyond finding matches for single units and consider creation of an entire synthetic control *group* by creating individual synthetic control units for the outcomes $T_{k*}(0)$ of an entire cohort of n treated units for which we have access to $\{X_{k*}, T_{k*}(1)\}_{k=1}^n$. This setup is motivated by, for example, the problem of constructing a virtual (external) control arm for a one-armed clinical trial (Thorlund et al., 2020; Mishra-Kalyani et al., 2022), which is often tackled either by matching or by weighting observed control units in downstream analyses. Here, we are therefore interested in understanding whether one could create synthetic units to be used in such downstream analyses instead.

Exaggerated concentration around the mean survival time. To illustrate challenges with this approach even in the simplest of settings, we now assume that time indeed follows a simple linear model $T_i(0) = X_i\beta + \epsilon_i$ so that even standard synthetic control is unbiased for the expected (mean) survival time. However, when the goal is to estimate entire survival functions from the constructed

cohorts, the full implied distribution of survival times – and not just their means – matters. In particular, because – as we have shown in Section 5.1 – the variance of synthetic control units is smaller than that of real units, the synthetic survival times will be *too concentrated around their mean*, which means that survival curves constructed from such synthetic cohorts will likely both i) underestimate the probability of early events (i.e. underestimate $P(T_*(0) < t_1)$ for t_1 small) and ii) underestimate the probability of surviving for long times (i.e. underestimate $P(T_*(0) > t_2)$ for t_2 large).

We illustrate this for a very simple example in Fig. 3. We revisit the stylized setup from the previous figures with $x_* = 2$ and two control units $x_1 = 1.5$ and $x_2 = 3$ but let survival time be linear $T(0) = 3 + x + \epsilon$, $\epsilon \sim \mathcal{N}(0, 2^2)$ to illustrate that issues arise even in this setting. We generate 1000 examples at each datapoint, and can thus construct the survival curves from 1000 mutually independent nearest neighbor matches and synthetic control units, respectively. Fig. 3 clearly illustrates the anticipated issues with the survival curve of the synthetic control cohort (orange, dotted): while the mean of the true survival times (blue, solid) is matched well, the estimated survival curve does not perform well for low and high survival times. The survival curve created through matching (green, dashed), on the other hand, is shifted to the left due to imperfect matching but captures the shape of the true survival curve well.




Figure 3: **Stylized illustration of biases in survival curve estimation.**

Dependence between synthetic control units. A final issue that arises once we move from individual synthetic control units to constructing a full control group to be used in downstream analyses is that – if some synthetic units share donors – they are no longer independent. In particular, for any two synthetic control units we have

$$\text{Cov}(\hat{Y}_{k*}^{SC}, \hat{Y}_{l*}^{SC}) = \text{Cov}(\sum_{j=1}^m w_{jk} Y_j(0), \sum_{j=1}^m w_{jl} Y_j(0)) = \sigma^2 \sum_{j=1}^m w_{jk} w_{jl} \quad (12)$$

Thus, the larger $\sum_{j=1}^m w_{jk} w_{jl}$, the higher the correlation between synthetic control units, which will also lead to *underestimation* of the spread of the survival distribution (as survival times will be artificially close across units). Note that even matching estimators – which have only a single contributor j that gets weight $w_{jk} = 1$ (and hence do not underestimate the individual variance) – can encounter this issue if matches are chosen with replacement (Stuart, 2010).

6. Towards Overcoming Obstacles: Penalizing deviations from real distributions

In the preceding sections, we highlighted that the use of synthetic controls for survival analyses can lead to bias because their distribution will likely differ from real control units – depending on the behaviour of the weight vectors $\mathbf{w}_k = (w_{1k}, \dots, w_{mk})$ for each synthetic unit k . Motivated by the discussions above, it is therefore a natural next step to consider whether one can trade off the quality of the synthetic match with implied deviations from the true distribution by *penalizing* the \mathbf{w}_k .

Here, we consider doing so by trading off the fit in input space (i.e. extrapolation bias) with bias incurred due to *artificially low variance of synthetic control units* by *maximising* $\|\mathbf{w}_*\|^2$ (or equivalently, minimizing $-\|\mathbf{w}_*\|^2$) which maximises the variance of the synthetic control outcome. For a single synthetic unit, this gives rise to the following penalized objective:

$$\mathbf{w}_* = \arg \min_{w: 0 \leq w_i \leq 1, \sum_i w_i = 1} \|X_* - \sum_{j=1}^m w_j X_j\|^2 - \lambda_{var} \|\mathbf{w}\|^2 \quad (13)$$

Note that – because $\|\mathbf{w}_*\|^2$ is maximised when all weight is given to *a single donor unit* – this objective essentially interpolates between a perfect synthetic control (when $\lambda_{var} = 0$) and a nearest neighbor matching estimator (when $\lambda_{var} \rightarrow \infty$) at its extremes⁶. Because $\|\mathbf{w}_*\|^2$ grows as weight is assigned less uniformly, this objective may encourage *more sparsity* in the weights at intermediate values – which, as a side effect, thus also leads to *less interpolation bias*. While we do not investigate this further here, we note that if the goal was to also correct for (i.e. minimize) the covariance between units, Eq. (12) suggests that this can be achieved by minimizing $\mathbf{w}_k^T \mathbf{w}_l$ for $l \neq k$, i.e. adding a term $+\lambda_{cov} \sum_{k=1}^n \sum_{l>k} \mathbf{w}_k^T \mathbf{w}_l$, which will reduce overlap between donor units.

Remark: Initialisation. Whenever the size of the donor pool is large relative to the number of features, the minimizer of the synthetic control objective will not necessarily be unique. In our experiments (see Fig. 7), we found that initialising \mathbf{w} at the nearest neighbor solution leads to much better performance than initialising it with random donors; this initialisation scheme also naturally encourages solutions with large $\|\mathbf{w}_*\|^2$.

Remark: Hyperparameter choice. Hyperparameters in synthetic control methods are usually chosen by cross-validation on i) pre-intervention outcomes of the treated or ii) post-intervention outcomes of the untreated (Abadie and L’hour, 2021). In our context, where no pre-treatment outcomes are available, only option ii) is feasible. We would thus suggest choosing hyperparameters through performing cross-validation *among the controls only* by constructing synthetic control outcomes for held-out out controls which can then be compared to their real outcomes. The used validation metric should depend on the goal of the downstream analysis: this could either be prediction metrics (e.g. the mean absolute error as in our experiments) if getting the individual survival times right is of greater importance, or metrics capturing the realism of the distribution of survival times (e.g. a Kolmogorov-Smirnov statistic as in our experiments) if that aligns better with the final goal of the analysis.

7. Empirical illustrations

Finally, we illustrate some of the theoretical discussion presented above on real data. We use the Rotterdam breast cancer dataset (Foekens et al., 2000; Royston and Altman, 2013)⁷, which contains 2982 female breast cancer patients of the Rotterdam tumour bank. It provides information on patient characteristics (age, menopausal status, tumor size, tumor differentiation grades, number of positive lymph nodes, progesterone receptors and estrogen receptors), information on treatment (whether hormone therapy and/or chemotherapy was received) and information on a survival outcome (the time from primary surgery to death or censoring). We plot the marginal survival curves by treatment in Fig. 4. Throughout, we focus on survival within the first 10 years to ensure sufficient samples with follow-up.




Figure 4: **Kaplan-Meier Survival Curves by Treatment.**

6. This behaviour is similar to the penalized objectives of Kellogg et al. (2021) and Abadie and L’hour (2021), which are derived with the explicit goal of interpolating between nearest neighbor and synthetic control estimator and therefore use a penalty term $+\lambda \sum_{j=1}^m w_j \|X_* - X_j\|^2$ instead of our term $-\lambda_{var} \|\mathbf{w}\|^2$ that is motivated from concerns about variance.

7. Retrieved from the R package `survival` (Therneau, 2023).

7.1. Experiment 1: Evaluating relative performance through biased resampling of real data

In this section, we use the data of the $n = 2091$ patients that received neither treatment to evaluate the performance of the different approaches. Inspired by negative control outcome methods (Lipsitch et al., 2010; Shi et al., 2020), we create biased subsets (“treatment groups”) of this data and evaluate methods based on their ability to reduce bias that we induced ourselves in a setting where we know there is *no effect at all* (all patients actually received the same treatment!). In particular, as we detail further in Appendix B, we sample a “target group” of size $\approx .05n$ with *higher* expected survival time (estimated using a Cox model) than the average patient. While unadjusted comparisons will thus find better survival in this target group compared to the remaining “controls”, correctly adjusting for patient characteristics should lead to the conclusion that target and control group have the same expected survival. Note that, because the remaining control donor pool is still quite large, *very* close neighbors for most samples in the target group exist in the control pool. To gain more interesting insights on the effects of inter- and extrapolation biases, we therefore also investigate the performance of the different methods when we remove observations within δ_{min} average squared Euclidean distance across all features to any of the samples in the target group.

• **Handling censoring.** Unlike the setup considered in the preceding theoretical sections, there is censoring in the Rotterdam data. While this would not be problematic if the CAR assumption was fulfilled, it seems that censoring is actually *not* at random here, because longer times are naturally more likely to be censored as follow-up is limited (indeed, only 50-70% of individuals experience an event by 120 months). Constructing matched survival curves from only uncensored individuals, be it through matching or through synthetic control, therefore leads to downward bias as we can see in Fig. 5. Instead, we therefore apply the synthetic control weights to both the censored times $T_j^- = \min(T, C)$ and the event indicators $E_j = \mathbf{1}\{T < C\}$ and heuristically censor synthetic control units with $\sum_{j=1}^m w_j E_j < 0.5$. In Fig. 5 we observe that this clearly improves upon using uncensored individuals only, so we use this approach throughout the section.

• **The shapes of survival curves.**

In Fig. 6 we illustrate the effect of the different approaches when used to create marginal survival curves. We make multiple interesting observations: First, when there are very close matches in the control data ($\delta_{min} = 0$, left panel), all methods essentially perform the same – this because there are nearest neighbor matches that are simultaneously an (almost) perfect synthetic control unit. When not ($\delta_{min} = 0.1$, right panel), we observe that, as expected, standard SC tends to overestimate time more often while log-time SC tends to underestimate. Also, both versions of SC result in survival curve shapes that are qualitatively less close to the target distribution in shape than matching – as expected, both SC distributions underestimate the occurrence of early events.




Figure 5: **Illustrating censoring bias.** Kaplan-Meier curves for $\delta_{min} = 0.1$.




Figure 6: **Illustrating differences in shapes of (Kaplan-Meier) control survival curves.**

• **The effect of initialisation.** In Fig. 7, we illustrate that this effect is exacerbated when we change the initialisation scheme to initialise SC with random weights: in this case, because solutions are not unique, SC tends to converge to a solution with many more contributors, so that survival times are indeed much too concentrated around their mean as expected. Throughout, we therefore initialise SCs from the nearest neighbor (NN) match.




Figure 7: **SC Initialization effects.** Kaplan-Meier curves for $\delta_{min} = 0.1$.

• **Trade-off between performance in terms of time prediction and survival curve estimation.** In Fig. 8, we now compare SC and matching more quantitatively in terms of (a) a metric for how close the survival curve is to the target survival curve (we use the Kolmogorov-Smirnov (KS) statistic for this) and (b) a metric capturing how close the individual predictions of survival time are (we use the mean absolute error (MAE) of predicting the Restricted Mean Survival Time (RMST) for this). We find that, as expected, NN matching does better at mimicking the survival curve shape – but SC can do better at predicting the individual times. We also find that log-time SC does substantially worse than standard SC here, which might be due to any of the following reasons: i) underestimation of times exacerbates the existing observational bias here, ii) the signal-to-noise ratio might be low and iii) the true survival time might not actually follow an AFT model.

• **Effects of regularization.** Finally, we investigate whether incorporating weight penalty λ_{var} can help with the observation above. In Fig. 8, we find that using small regularization penalties indeed leads to control cohorts with properties between the extremes of NN matching and SC: they perform better in terms of distribution metrics than SC, and better in terms of prediction metrics than matching – but not as good as the best at either. This indicates that the best strategy to use in practice depends on the objective: If getting the individual time predictions right is of most importance and there are no close matches in the data, then one may want to rely on synthetic control methods directly. If only the marginal distribution is of interest, then using matching alone may suffice. If, however, a practitioner is interested in trading off the two objectives, then the use of the regularization penalty can allow to find a place on the Pareto frontier between the two extremes.

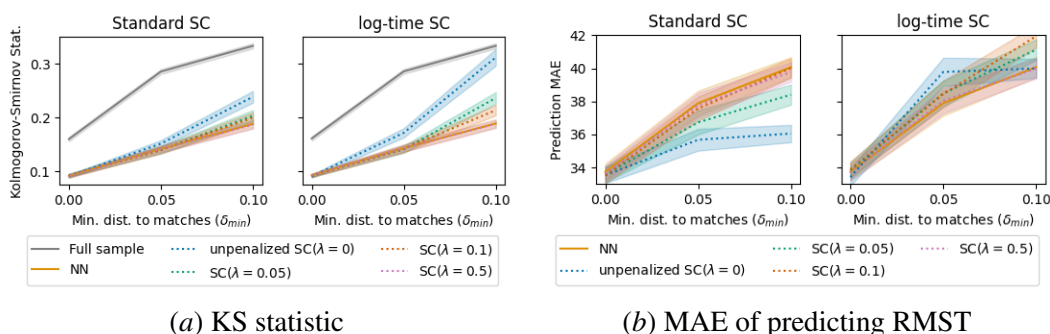


Figure 8: **Performance of different methods by δ_{min} .** Averaged over 20 experiments, 2SEs shaded.

7.2. Experiment 2: Exploratory analyses of different treatment groups in real data

Plotting marginal survival in the real data as in Fig. 4 appears to – counter-intuitively – show that, on average, individuals who received a hormonal treatment or chemotherapy have shorter survival than those patients that received neither. We now investigate whether this unexpected observation can

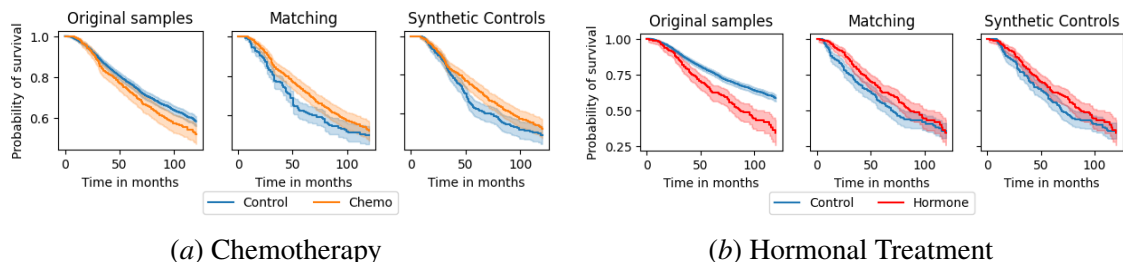


Figure 9: **Exploratory results adjusting for biases in real data.** Kaplan-Meier survival curves for the original treated patients and control cohorts created using different methods.

be *explained away* by adjusting for observed differences between the two cohorts through creating real and synthetic matches in Fig. 9. For both treatments, we observe that creating matched cohorts through either NN-matching or standard SCs (we drop log-time SC here due to consistently worse performance in Section 7.1) leads to qualitatively the same conclusion and indeed results in reversal of the originally counterintuitive finding: once we create a control group more aligned with the treated group, treated individuals indeed *no longer* seem to live shorter – to the contrary, there now appears to be some evidence that treatment increases the expected survival time. This shift in survival curves once patient characteristics are adjusted for is entirely expected – only node positive patients actually received either therapy in the data (Foekens et al., 2000) – and the evidence for treatment benefit, once groups are more matched in terms of patient characteristics, is in line with the original findings of the study where therapy indeed had a positive effect on survival (Royston and Altman, 2013).

8. Conclusion

In this paper, we investigated the possibility of using synthetic control methods in survival analyses, and discovered multiple biases arising uniquely in this context. We showed that these are rooted in differences in both (i) commonly assumed DGPs and target parameters of interest in survival analyses relative to standard settings in economics, and (ii) differences in distributions of real and synthetic control units. We also highlighted that regularization schemes interpolating between matching and synthetic control methods may constitute a promising avenue to control some of these biases.

We hope that future methodological work will build upon this initial evidence to develop more sophisticated approaches to overcome the challenges investigated in this paper. Further, we put our focus on a specific family of survival models (AFT models) in the theoretical part of this paper. It would therefore be another interesting avenue for future research to investigate whether our analyses can be extended to the proportional hazards family of survival models. Finally, we chose to focus on comparing synthetic controls to matching, but did not investigate the possibility of imputing missing counterfactual outcomes in *other* ways, e.g. using model-based approaches. Additionally comparing to model-based imputation of control survival times, using e.g. AFT- or Cox-regression models fitted using the control data, will add the potential for model misspecification as another dimension to the comparison, and may thus lead to further interesting tradeoffs (in addition to trading off interpolation and extrapolation biases as in the current comparison). Investigating the (dis)advantages of synthetic control methods relative to other possible approaches in the survival context could thus also be a fruitful direction to explore in future work.

References

- Alberto Abadie. Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2):391–425, 2021.
- Alberto Abadie and Javier Gardeazabal. The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132, 2003.
- Alberto Abadie and Jeremy L’hour. A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association*, 116(536):1817–1834, 2021.
- Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American statistical Association*, 105(490):493–505, 2010.
- Daron Acemoglu, Simon Johnson, Amir Kermani, James Kwak, and Todd Mitton. The value of connections in turbulent times: Evidence from the united states. *Journal of Financial Economics*, 121(2):368–391, 2016.
- Susan Athey and Guido W Imbens. The state of applied econometrics: Causality and policy evaluation. *Journal of Economic perspectives*, 31(2):3–32, 2017.
- Peter C Austin and Elizabeth A Stuart. Optimal full matching for survival outcomes: a method that merits more widespread use. *Statistics in medicine*, 34(30):3949–3967, 2015.
- Taane G Clark, Michael J Bradburn, Sharon B Love, and Douglas G Altman. Survival analysis part i: basic concepts and first analyses. *British journal of cancer*, 89(2):232–238, 2003.
- David R Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.
- Scott Cunningham and Manisha Shah. Decriminalizing indoor prostitution: Implications for sexual violence and public health. *The Review of Economic Studies*, 85(3):1683–1715, 2018.
- John J Donohue, Abhay Aneja, and Kyle D Weber. Right-to-carry laws and violent crime: A comprehensive assessment using panel data and a state-level synthetic control analysis. *Journal of Empirical Legal Studies*, 16(2):198–247, 2019.
- John A Foekens, Harry A Peters, Maxime P Look, Henk Portengen, Manfred Schmitt, Michael D Kramer, Nils Bruenner, Fritz Jaenicke, Marion E Meijer-van Gelder, Sonja C Henzen-Logmans, et al. The urokinase system of plasminogen activation and prognosis in 2780 breast cancer patients. *Cancer research*, 60(3):636–643, 2000.
- Maxwell Kellogg, Magne Mogstad, Guillaume A Pouliot, and Alexander Torgovitsky. Combining matching and synthetic control to tradeoff biases from extrapolation and interpolation. *Journal of the American statistical association*, 116(536):1804–1816, 2021.
- Marc Lipsitch, Eric Tchetgen Tchetgen, and Ted Cohen. Negative controls: a tool for detecting confounding and bias in observational studies. *Epidemiology (Cambridge, Mass.)*, 21(3):383, 2010.

- PS Mishra-Kalyani, L Amiri Kordestani, DR Rivera, H Singh, A Ibrahim, RA DeClaro, Y Shen, S Tang, R Sridhara, PG Kluetz, et al. External control arms in oncology: current use and future directions. *Annals of Oncology*, 33(4):376–383, 2022.
- Paolo Pinotti. The economic costs of organised crime: Evidence from southern italy. *The Economic Journal*, 125(586):F203–F232, 2015.
- Paul R Rosenbaum and Donald B Rubin. The bias due to incomplete matching. *Biometrics*, pages 103–116, 1985a.
- Paul R Rosenbaum and Donald B Rubin. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *The American Statistician*, 39(1):33–38, 1985b.
- Patrick Royston and Douglas G Altman. External validation of a cox prognostic model: principles and methods. *BMC medical research methodology*, 13:1–15, 2013.
- Donald B Rubin. Matching to remove bias in observational studies. *Biometrics*, pages 159–183, 1973.
- Claudia Shi, Dhanya Sridhar, Vishal Misra, and David Blei. On the assumptions of synthetic control methods. In *International Conference on Artificial Intelligence and Statistics*, pages 7163–7175. PMLR, 2022.
- Xu Shi, Wang Miao, and Eric Tchetgen Tchetgen. A selective review of negative control methods in epidemiology. *Current epidemiology reports*, 7:190–202, 2020.
- Xu Shi, Wang Miao, Mengtong Hu, and Eric Tchetgen Tchetgen. Theory for identification and inference with synthetic controls: a proximal causal inference framework. *arXiv preprint arXiv:2108.13935*, 2021.
- Elizabeth A Stuart. Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(1):1, 2010.
- Terry M Therneau. *A Package for Survival Analysis in R*, 2023. URL <https://CRAN.R-project.org/package=survival>. R package version 3.5-5.
- Kristian Thorlund, Louis Dron, Jay JH Park, and Edward J Mills. Synthetic and external controls in clinical trials—a primer for researchers. *Clinical epidemiology*, pages 457–467, 2020.
- Mark J van der Laan and James M Robins. *Unified methods for censored longitudinal data and causality*. Springer, 2003.
- Lee-Jen Wei. The accelerated failure time model: a useful alternative to the cox regression model in survival analysis. *Statistics in medicine*, 11(14-15):1871–1879, 1992.
- Jakob Zeitler, Athanasios Vlontzos, and Ciarán Mark Gilligan-Lee. Non-parametric identifiability and sensitivity analysis of synthetic control models. In *Conference on Causal Learning and Reasoning*, pages 850–865. PMLR, 2023.

Appendix A. On the necessity of linearity assumptions for synthetic control methods

In this section, we briefly discuss whether linearity assumptions on the outcome-generating process, as in Eq. (3), can sometimes be circumvented while retaining the ability of synthetic control methods to succeed. Indeed, Shi et al. (2022) show that linearity assumptions on the outcome-generating process are not strictly necessary for synthetic control methods *if one makes a different assumption* on the DGP (which, as we discuss at the end of this section, is not applicable in our survival setting).

Shi et al. (2022) make use of a *more fine-grained model assumption* than is usually the case in the synthetic control literature. Where classical synthetic control settings usually just assume the goal to estimate $\mu_{j,t}$ for units j and times t , Shi et al. (2022) make the additional assumption that every macro-unit j is an *average* of individual units l , i.e. that $\mu_{j,t} = E_l[\mu_{l,j,t}]$. In this case, they show that “linearity is a consequence of the fact that expectation is a linear operator” (Shi et al., 2022, p.2) – i.e. linearity is an automatic feature of the DGP of the *aggregate units* even if the DGP of the individual units is not linear. They conclude that this is why synthetic control applied to aggregate units can work even if the DGP of the individual units are nonlinear. Conversely, they show that even within this setting, if the aggregation across individuals is *not linear* (e.g. if using the group median instead of the group mean to construct aggregate units from the individual units), synthetic controls will be biased. Further, they also discuss that, if one wanted to include covariates beyond lagged outcomes, the “use of auxiliary covariates requires us to assume that the underlying data generating process is linear” (Shi et al., 2022, p. 7). In summary, Shi et al. (2022) thus show that one may sometimes be able to remove the (parametric) linearity assumption usually used in synthetic control papers by instead assuming a fine-grained data-generating model that *linearly aggregates unobserved individual units into macro-units* by averaging them. That is, some form of assumed linearity appears to remain crucial for successful application of synthetic control methods.

Why the fine-grained model is less applicable in survival analyses. The additional assumption on the decomposition of aggregate units into individuals made in the fine-grained model of Shi et al. (2022) can be very reasonable within the economic applications considered in most of the original synthetic control literature: in such applications, one is often analyzing macro-level data from e.g. states or countries, which are indeed usually reported in aggregated form across counties, cities or even individuals. In our medical setting, however, the fine-grained model is not appropriate, as the final outcomes we consider are patient-level survival outcomes which are by construction already at an individual level; no further disaggregation is possible. Thus, the arguments of Shi et al. (2022), which rely on linear aggregation across individual units in the generation of observed outcomes, do not apply to our setting and can hence not help in mitigating the biases arising due to non-linearity present in survival DGPs.

Appendix B. Experimental details

Sampling scheme. For the experiments in Section 7.1, the target and control groups are constructed from the real data through the following biased sampling scheme: We first fit a cox proportional hazards model using all covariates on the full $n = 2091$ patients, and use it to predict the expected median survival time $T_i^{med}(X_i)$ of each patient, which we normalize across the sample to give one value $Z_i = \text{normalize}(T_i^{med}(X_i))$ per patient. Second, we then split the n patients randomly into two samples: with probability .1 they become part of the target pool \mathcal{I}_t , and with probability .9 they become part of the control pool \mathcal{I}_c . Third, from the target pool \mathcal{I}_t , we then create

a biased target sample with higher expected survival times by sampling patients into the sample \mathcal{S}_t with probability $p_i = \text{expit}(3 * Z_i)$. Finally, when $\delta_{min} > 0$, we then remove all close neighbors of patients in \mathcal{S}_t from \mathcal{I}_c to create the control sample \mathcal{S}_c , where the squared distance for each patient j in the control sample for each patient k_* in the target sample is calculated from the normalized d-dimensional covariate vectors as $\delta_{jk_*} = \frac{1}{d} \sum_{d=1}^D (X_j^d - X_{k_*}^d)^2$.

Metrics. The Komolgorov-Smirnov statistic measures the distance between two CDFs $F(x)$ and $G(x)$ as $\sup_x |F(x) - G(x)|$. We approximate this from the target samples and the different control groups by first computing the empirical survival functions \hat{S}^t and \hat{S}^c using the Kaplan-Meier estimator, and then computing $\max_{T \in \mathcal{S}_t} |\hat{S}^t(T) - \hat{S}^c(T)|$.

The restricted mean survival time (RMST) relative to time $t_{end} = 120$ months is defined as $T_i^{RMST} = \min(T_i, t_{end})$. To evaluate individual predictions, we compute the mean absolute error between predicted and observed RMST for each unit in the target sample, where we can only use individuals i for which both observed and predicted time are not censored before t_{end} :

$$MAE = \frac{\sum_{i \in \mathcal{S}_t} \mathbf{1}\{i : (E_j = 1 | T_j^- > t_{end}) \& (\hat{E}_j = 1 | \hat{T}_j^- > t_{end})\} |T_j^{RMST} - \hat{T}_j^{RMST}|}{\sum_{i \in \mathcal{S}_t} \mathbf{1}\{i : (E_j = 1 | T_j^- > t_{end}) \& (\hat{E}_j = 1 | \hat{T}_j^- > t_{end})\}} \quad (14)$$