

Long-Run Effects of Dynamically Assigned Treatments: a New Methodology and an Evaluation of Training Effects on Earnings

Gerard J. van den Berg*

Johan Vikström†

September 2021

Abstract

We propose and implement a new method to estimate treatment effects in settings where individuals need to be in a certain state (e.g. unemployment) to be eligible for a treatment, treatments may commence at different points in time, and the outcome of interest is realized after the individual left the initial state. An example concerns the effect of training on earnings in subsequent employment. Any evaluation needs to take into account that some of those who are not trained at a certain time in unemployment will leave unemployment before training while others will be trained later. We are interested in effects of the treatment at a certain elapsed duration compared to “no treatment at any subsequent duration”. We prove identification under unconfoundedness and propose inverse probability weighting estimators. A key feature is that weights given to outcome observations of non-treated depend on the remaining time in the initial state. We study effects of a training program for unemployed workers in Sweden. Estimates are positive and sizeable, exceeding those obtained with common static methods. This calls for a reappraisal of training as a tool to bring unemployed back to work.

*University of Bristol, University of Groningen, University Medical Center Groningen, IFAU, IZA, ZEW, CEPR and CESifo.

†IFAU and UCLS Uppsala University; johan.vikstrom@ifau.uu.se

Keywords: Treatment effects; dynamic treatment evaluation; program evaluation; unconfoundedness; unemployment; wage.

JEL classification: C14, J3.

Acknowledgements: We thank two anonymous Reviewers, Bart Cockx, Bruno Crépon, Bernd Fitzenberger, Anders Forslund, Peter Fredriksson, Martin Lundin, Jeffrey Smith, participants at seminars at Humboldt University Berlin, Ghent University, IFAU, Uppsala Center for Labor Studies, the Swedish Ministry of Finance, the Swedish Public Employment Service, and conference participants at IFAU/IZA, SOLE/EALE, ESWC, RES, Cafe Aarhus and COMPIE conferences, for useful suggestions. Vikström acknowledges support from FORTE.

1 Introduction

Consider a setting where eligibility for a treatment requires the individual to be in a spell in a certain state. In this initial state the date of entry into the treatment may vary, for example because of capacity constraints on the number of treatment slots. The duration of being in the initial state may vary as well, for example due to market frictions. Some of those who do not enter the treatment at a certain point in time will therefore leave the initial state before they are ever treated, while others will be treated later but before they leave the initial state.

We are interested in outcome measures that are realized after the individual has left the initial state. Compared to outcomes capturing the length of stay in the initial state, the former are less sensitive to short-run institutional features and market imperfections. As a leading example, consider long-run annual earnings if the treatment is participation in a training program during unemployment. Longer-run outcomes may better reflect the highest attainable earnings levels. In addition to this, long-run outcomes are relevant for cost-benefits analyses. (See Section 2 for a list of other examples.)

With a sufficiently long time horizon, the long-run earnings of the treated at any elapsed duration t are directly observable. The challenge is that the non-treated individuals at t cannot be straightforwardly used to construct an average counterfactual outcome for the treated. Under unconfoundedness we can adjust for differences between the two groups. However, only those who leave the initial state before treatment provide information on no-treatment potential outcomes, and in the time periods after t some of the non-treated will receive the treatment before they leave unemployment. Individuals with a small probability of receiving the treatment in subsequent periods may remain non-treated longer. Moreover, for a given period-by-period treatment probability, the individuals with a high rate of leaving unemployment are more likely to remain non-treated. Both features (small probability of treatment in subsequent periods and high exit rate) may influence long-run earnings. To infer the causal effect of training, these features need to be controlled for.

We develop a novel approach to resolve this. Besides unconfoundedness we do not impose structure on the assignment process. We establish identification and show that our setting differs from those in other bodies of work. As it turns out, the weights given to observations among those who are not treated after t do not only depend on observed confounders but also on the remaining time spent in unemployment. Simulation results confirm that an Inverse Probability Weighting (IPW) estimator is asymptotically un-

biased and that the proposed bootstrap inference works well. We generalize the results in several ways, e.g. by allowing for selection on time-varying covariates.

In the empirical analysis we estimate the long-run earnings effects of the flagship training program for unemployed workers in Sweden. Earlier studies concluded that the program does not raise post-unemployment earnings (see the survey in Forslund and Vikström, 2011), which motivated policy makers to cut down the program volume. Extensive administrative register data cover over ten years after entry into unemployment. The estimated long-run earnings effects are persistently significantly positive and sizeable, exceeding those obtained by using traditional methods.¹ This calls the view that such programs are rather ineffective into question, and suggests that there is scope for a reappraisal of training.

2 Model

2.1 Potential outcomes framework and observational rule

We let the time clock start at the moment of entry into the initial (or eligibility) state. The duration until the start of the treatment and the duration spent in the initial state are denoted by T_s and T_u , respectively, each with possible values $1, 2, 3, \dots$. T_s is only observed if $T_s \leq T_u$. If the individual leaves the initial state before T_s is realized then eligibility is lost and T_s will not be known. Throughout most of the paper we consider effects on a long-run (or final) outcome Y that is realized after individuals left the initial state and that is observable for all individuals. Thus, in the baseline setting we always observe T_u and Y whereas T_s is only realized and observed if it does not exceed T_u .

In the leading example, the treatment is participation in a training program, the initial state is unemployment, and the final outcome is annual earnings many years later. However, the setting applies to many other cases. In modern economies, individuals in conditions deemed unfavorable often qualify for subsidized treatments that are intended to kick-start the transition out of the unfavorable state and support the individual on an upward-sloping trajectory afterwards. Cost considerations may cause the time

¹Essentially, the static methods ignore that many non-treated individuals have short unemployment durations, and those individuals tend to have more favorable personal characteristics. The ranking of estimated effect sizes across methods is not universal across applications but depends on covariate values and on covariate effects on the conditional probabilities of leaving the initial state at different elapsed durations.

until entry into treatment to vary randomly at the individual level and this may lead to a situation where some individuals move out of unfavorable conditions before any treatment. Concrete examples include integration programs for immigrants, general education programs for homemakers and cognitive-behavioral therapies for depressed individuals. The latter can be generalized to the usage of a particular medication while having a certain illness. In that case, the eligibility period is the duration of the illness during which the medication is prescribed. In demographic applications, fertility is an obvious example of an eligibility state. Fertility treatments or adoption opportunities may only be allowed before an adult has a child. Note that even if a policy or treatment is only intended to reduce the time in a certain state, a comprehensive policy evaluation should consider longer-run effects as these influence the ultimate costs and benefits.

We now introduce potential outcomes, following Rubin (1974). Assignment of an individual to a treatment starting date t_s leads to a potential time in the initial state $T_u(t_s)$ and to a potential outcome $Y(t_s)$. Thus, $T_u(\infty)$ and $Y(\infty)$ capture the potential duration and the potential final outcome if the individual is assigned to be “never treated”. In practice, infinity may be replaced by an upper bound on the eligibility interval length.

This setup does allow for indirect effects from t_s via $T_u(t_s)$ to $Y(t_s)$.² We do not attempt to separate these out. Thus we examine the over-all comprehensive effect on Y regardless of the pathway. This is why the potential outcomes are indexed only by t_s and not by t_u .

We take time to be discrete, and assume that, within each time period t with $T_s \geq t$, the binary event governing whether $T_s = t$ vs. $T_s > t$ is realized before the binary event governing whether $T_u = t$ vs. $T_u > t$.

We adopt a “no anticipation” assumption to rule out that a future treatment assignment date affects relevant current events. Specifically, for any given individual,

Assumption 1 (No anticipation).

$$P(T_u(t') = t) = P(T_u(t'') = t), \quad \forall t < \min(t', t''), \text{ and}$$

$$Y(t') = Y(t''), \quad \forall T_u(t') = T_u(t'') < \min(t', t'').$$

²In the leading example, the direct effect may capture that training improves the participants’ human capital and ensuing labor earnings Y , while the indirect effect may capture that training speeds up exit out of unemployment and employers use the realized unemployment duration as a signal of worker quality and hence pay higher earnings Y to trained workers.

The first part rules out that the timing of a future treatment affects the current probability of leaving the initial state (Abbring and Van den Berg, 2003). The second part rules out that the potential outcomes $Y(t_s)$ are affected by behavior before t_s that is driven by the future treatment date t_s , even if this behavior is of no consequence for the moment of leaving the eligibility state before t_s . In particular, in the eligibility state, it is ruled out that private knowledge of the future moment of treatment affects the long-run outcome even in cases where time in the eligibility state is not affected.³ Taken together, the assumption ensures that individual outcomes do not vary with the moment of treatment as long as the treatment is not realized.

For derivations in later sections it is useful to point out that the first part of the assumption implies that $P(T_u(t_s) = t) = P(T_u(\infty) = t)$, $\forall t < t_s$. As we shall see, this allows us at times to replace the conditioning on survival of $T_u(t_s)$ at elapsed durations up to t_s by conditioning on survival of $T_u(\infty)$ at such elapsed durations.

By Assumption 1 we have the following observational rule for Y

$$Y = \sum_{t=1}^{T_u} [\mathbf{I}(T_s = t)Y(t)] + \mathbf{I}(T_s > T_u)Y(\infty), \quad (1)$$

where $\mathbf{I}()$ is the indicator function. The first part of the right-hand side states that if the individual is treated before leaving the initial state, then we observe the long-run potential outcome corresponding to the actually observed time to treatment. The second part states that if an individual exits the initial state without having been treated then the observed outcome is the outcome corresponding to the assignment “never treated”.

Accordingly, the observed outcomes at the individual level can be expressed as:

$$T_u, T_s \cdot \mathbf{I}(T_s \leq T_u), Y(T_s) \cdot \mathbf{I}(T_s \leq T_u), Y(\infty) \cdot \mathbf{I}(T_s > T_u).$$

2.2 The dynamic evaluation challenge

Our prime object of interest is the average treatment effect of t_s on Y among those who are actually treated at t_s

$$\text{ATET}(t_s) = E(Y(t_s) - Y(\infty) \mid T_s = t_s, T_u(t_s) \geq t_s), \quad (2)$$

³The second part of Assumption 1 only considers $T_u(t') = T_u(t'')$ because if $T_u(t') \neq T_u(t'')$ then a difference in the long-run outcomes $Y(t')$ and $Y(t'')$ may be ascribed to an effect of T_u on Y .

This contrasts a treatment start at t_s with the assignment to be “never treated”.⁴

The *first* component of the expression for $\text{ATET}(t_s)$ is identified from a random sample of observed outcomes Y of those treated at time T_s :

$$E(Y(t_s)|T_s = t_s, T_u(t_s) \geq t_s) = E(Y|T_s = t_s, T_u \geq t_s). \quad (3)$$

The evaluation challenge originates from the *second* component of $\text{ATET}(t_s)$, that is, the counterfactual outcome $E(Y(\infty)|T_s = t_s, T_u(t_s) \geq t_s)$. The observational rule in (1) illustrates the main issues. First of all, it is clear that only non-treated individuals, i.e. individuals with $T_s > T_u$, can be informative on the counterfactual outcome under never treatment. Individuals who are not treated at t_s but are treated later are not informative on the counterfactual outcome under never treatment. After all, since they are ultimately treated, their observed final outcome Y corresponds to a potential outcome under a treatment.

Now, the potential control group of non-treated, defined by $T_s > T_u \geq t_s$, in general constitutes a selective sample. It depends on the survival times T_u (i.e., on an intermediate outcome which may affect the outcome of interest), since the probability of treatment enrollment by construction increases with the time in the initial state. Thus, for a given period-by-period treatment probability, individuals with longer unemployment durations are less likely to remain non-treated. This arises purely because of the dynamic nature of the treatment assignment. To phrase this differently, individuals with relatively short durations after a treatment at t_s are overrepresented in the group of non-treated. Thus, even if we assume that treatment assignment is unconfounded, treatment assignments in subsequent time periods cause the group of non-treated to be systematically different.

⁴The focus on ATET is motivated by the policy relevance of ATET and is in line with the prominence of ATET in the economic evaluation literature on which we build. Below we discuss the aggregation of $\text{ATET}(t_s)$ into ATET. An alternative object of interest is the average treatment effect of t_s on Y among all those who, if they were assigned to t_s , would still be in the initial state at that time t_s . We refer to this as $\text{ATE}(t_s) := E(Y(t_s) - Y(\infty) | T_u(t_s) \geq t_s)$ and discuss results on this in the online Appendix B.2.

3 A dynamic evaluation approach

3.1 Sequential unconfoundedness

We consider time-invariant covariates (observed individual characteristics) X and later extend to cases with time-varying covariates. Our key identifying assumption is that in each period the treatment assignment is random among the not-yet-treated survivors, conditional on X . This can be called a sequential unconfoundedness assumption. Formally, let the binary indicator $P_t := 1$ iff $T_s = t$. Then:^{5,6}

Assumption 2 (Sequential unconfoundedness). For all t ,

$$P_t \perp Y(\infty) \mid X, T_s \geq t, T_u(\infty) \geq t.$$

The above assumption is implied by a dynamic version of Rosenbaum and Rubin (1983)'s ignorable assignment or unconfoundedness assumption,

Assumption 3 (Unconfoundedness).

$$T_s \perp Y(\infty) \mid X.$$

Imbens (2000), in his study on multi-valued treatments under unconfoundedness, weakens Rosenbaum and Rubin (1983)'s assumption to pairwise independence of the treatment, at any possible treatment value t , from the corresponding potential outcome conditional on covariates. This resembles the relaxation from Assumption 3 to 2. In each approach, this paves the way for the use of separate propensity scores by each t . However, due to the dynamic nature of our framework, we need to condition on being in the initial state and we need to deal with eligibility after t of those not yet treated at t , so that the inferential method developed in Imbens (2000) is not suitable here. Subsection 3.4 contains additional discussion on differences with approaches in the literature.

(Sequential) unconfoundedness is plausible for some applications but not for others. In the case of training programs for unemployed workers, the argument has been

⁵For the ATE(t_s) discussed in online Appendix B.2 we need conditional independence of P_t from both $Y(t)$ and $Y(\infty)$.

⁶At first sight the special role of $T_u(\infty)$ in the conditioning set may seem puzzling. However, by virtue of Assumption 1, for any treatment assignment date $t_s > t$, the conditioning on $T_u(t_s) \geq t$ is equivalent to the conditioning on $T_u(\infty) \geq t$. Indeed, we may replace $T_u(\infty) \geq t$ at no cost by $T_u \geq t$ although this expresses the assumption in terms of an intermediate outcome.

made that caseworkers determine the assignment and that the register data used in the analysis includes the information used by the caseworkers when deciding on the assignment. In general, unconfoundedness seems more likely if the data include measures of individual motivation, non-cognitive abilities and other personality traits that may affect treatment assignment as well as outcomes, or if the data contain markers of those traits, such as variables describing past choices and outcomes. However, depending on the context and the data, it is possible that caseworkers may select individuals based on characteristics of the latter that we do not observe and cannot capture by other observed indicators, so caseworker selection alone is not sufficient to justify the unconfoundedness assumption. Similarly, workers may aim to schedule their entry into training while taking account of outcome determinants that we do not observe.⁷

As usual with evaluation methods, we impose a “common support” condition: for all t and X , $\Pr(P_t = 1 | T_s \geq t, T_u \geq t, X) < 1$. The usual SUTVA condition also needs to hold, ruling out various types of interference between the units in the sample (see, e.g., Wooldridge, 2010). For expositional convenience, we do not refer to these two conditions in most of the exposition.

3.2 The gist of the approach in a simple setting

This subsection describes the gist of our method in a simplified setting with two periods in which the treatment can take place whereas no treatments can occur thereafter. We consider the effect of treatment in the first period, $ATET(1) = E(Y(1) - Y(\infty) | T_s = 1)$.⁸

Note that $E(Y(1)|T_s = 1) = E(Y|T_s = 1)$ follows directly from the observational rule. We now show that $E(Y(\infty)|T_s = 1)$ is identified as well. If we condition on X then, by Assumption 2, the treated and non-treated are comparable:

$$E(Y(\infty)|T_s = 1, X) = E(Y(\infty)|T_s > 1, X). \tag{4}$$

⁷Note that our assumptions do not rule out that time in the initial state and the long-run outcome are affected by shared unobserved confounders.

⁸Note that for the effect of a treatment in the first period there is no need to condition on survival.

By the law of iterated expectations

$$\begin{aligned}
E(Y(\infty)|T_s > 1, X) = & \tag{5} \\
& \Pr(T_u = 1|T_s > 1, X)E(Y(\infty)|T_s > 1, T_u = 1, X) + \\
& \Pr(T_u > 1|T_s > 1, X)E(Y(\infty)|T_s > 1, T_u > 1, X).
\end{aligned}$$

That is, the counterfactual outcome under never treatment is decomposed into average outcomes for individuals with $T_u = 1$ and with $T_u > 1$. Note that probabilities $\Pr(T_u = 1|T_s > 1, X)$ and $\Pr(T_u > 1|T_s > 1, X)$ are observed.

Next, the group with $T_u = 1$ in (5) consists of non-treated individuals who exit directly in period 1. For this group, the observational rule in (1) and Assumption 1 give

$$E(Y(\infty)|T_s > 1, T_u = 1, X) = E(Y|T_s > 1, T_u = 1, X). \tag{6}$$

This follows since for non-treated individuals with $T_s > T_u$, the observed long-run outcome Y equals the non-treated long-run potential outcome $Y(\infty)$.

For the group with $T_u > 1$, in (5), i.e., those who survive at least one additional time period, we have by Assumption 2:

$$E(Y(\infty)|T_s > 1, T_u > 1, X) = E(Y(\infty)|T_s > 2, T_u > 1, X), \tag{7}$$

i.e. among the non-treated survivors, conditional on X , those who become treated in period 2 are comparable to those who remain non-treated. Here we use that Assumption 1 justifies the replacement of $T_u(\infty) > 1$ by $T_u > 1$ when invoking Assumption 2 (recall the previous subsection). Next, since there are no treatments after period 2 and under no-anticipation, those with $T_s > 2$ remain non-treated, so that for this group the observed long-run outcome Y equals the non-treated long-run potential outcome, $Y(\infty)$:

$$E(Y(\infty)|T_s > 2, T_u > 1, X) = E(Y|T_s > 2, T_u > 1, X). \tag{8}$$

Then, by (4)-(8) we have

$$\begin{aligned}
E(Y(\infty)|T_s = 1, X) = & \tag{9} \\
& \Pr(T_u = 1|T_s > 1, X)E(Y|T_s > 1, T_u = 1, X) + \\
& \Pr(T_u > 1|T_s > 1, X)E(Y|T_s > 2, T_u > 1, X).
\end{aligned}$$

That is, the long-run outcomes for non-treated units with $T_u = 1$ and $T_u > 1$ are weighted in a specific way, based on the exit probabilities $\Pr(T_u = 1|T_s > 1, X)$, and this allows us to recover the long-run outcome under no treatment. In particular, this deals with the fact that some of the non-treated become treated in the second period.

3.3 Identification in the general case

In Appendix A of the paper we generalize the above derivations to the general case by proving the following theorem,

Theorem 1 (ATET). If Assumptions 1 and 2 hold then

$$\begin{aligned}
\text{ATET}(t_s) = & E(Y|T_s = t_s, T_u \geq t_s) - \\
& E_{X|T_s=t_s, T_u \geq t_s} \left[\sum_{k=t_s}^{T_u^{\max}} h(k, X) \left[\prod_{m=t_s}^{k-1} [1 - h(m, X)] \right] E(Y|T_s > k, T_u = k, X) \right],
\end{aligned}$$

where

$$h(t, X) = \Pr(T_u = t|T_s > t, T_u \geq t, X)$$

and where T_u^{\max} is maximum possible time in the initial state, and we apply the notation $\prod_{m=t_s}^{t_s-1} \dots = 1$. In sum, the non-treated group (consisting of individuals leaving the initial state before becoming treated or actually never would have become treated) can be used to identify the counterfactual outcome for those treated after a certain elapsed duration. Theorem 1 shows that this is achieved by giving individuals leaving the initial state a differential weight depending on the duration in that state.

ATET(t_s) can be aggregated over the distribution of T_s among those with a realized treatment (so with $T_u \geq T_s$) in order to identify an over-all average effect of the program for the population of treated over all possible t_s . Specifically, $\text{ATET} = E_{T_s}[E(Y(t_s) - Y(\infty) | T_s = t_s, T_u(t_s) \geq t_s)]$ where the outer expectation is taken over the distribution

of $T_s|T_u \geq T_s$.

3.4 Relation to approaches in the literature

In biostatistics, Robins and co-authors have developed a methodology for dynamic evaluation based on sequential unconfoundedness (see Robins and Hernán, 2009, and Hernán and Robins, 2020, for extensive overviews). This aims to compare dynamic *regimes* g covering a time period with length K and characterized by a sequence of *treatments* or *exposures* $a(t)$ that may vary over time t . In a given regime, the value of $a(t)$ may depend on treatments in previous periods and on current and previous values of observed time-varying covariates $L(t)$. The latter may depend on its previous values and on previous exposures as well as on unobserved confounders. In common applications, $a(t)$ denotes the dosage of a medication provided at time t and $L(t)$ is a biomarker capturing the current health status of the patient. The key objects of interest are the ATE that compare potential outcomes Y_g across two regimes g' and g'' , that is, $\mathbb{E}(Y_{g'}) - \mathbb{E}(Y_{g''})$. The ATET is not considered because there is no single regime that is promoted as the benchmark regime over and above other treatment regimes. An assumption similar to common support assumptions states that for every possible history of covariates and past treatments, each current treatment exposure value has a positive probability density of being attained in the data. A natural interpretation is that, in real-life data, for any given history up to t , some patients will receive a medication at t and others will not.

We consider the non-parametric setting closest to ours, i.e., with continuous outcomes Y that are realized and observed for every subject at time K (Robins and Hernán, 2009, Chapter 23). In this setting we may define regimes g by the first point in time t_s at which a treatment is given. However, this does not lend itself to inclusion of an eligibility state. Among individuals with a history stipulating that they leave an eligibility state at $t_s - 1$ and never entered treatment before or at $t_s - 1$, we can rule out with probability one that they will begin the treatment at t_s , which violates their version of the common support assumption. Indeed, their ATE is not identified in our approach. The literature by Robins and co-authors does not discuss eligibility states. In medical applications, a setting where a treatment can be switched on or off in any period is often more reasonable.⁹

⁹To further push for formal equivalence, one may introduce ad hoc restrictions on the set of possible regimes g . For this, one may include the duration in the eligibility state T_u in the time-varying covari-

Approaches by Robins and co-authors have been adopted in economics by Lechner (2009) and Lechner and Miquel (2010). Here, the treatment in a certain period may depend on current values of intermediate outcomes, and covariates and treatments are essentially observed for all units in all periods. In their innovative application to active labor market policies, Lechner and Wiehler (2013) compare different sequences of program participation (exposure) and unemployment before eligibility loss comes into play, and they examine if the timing of the exposure within the sequence matters. This approach warrants groups with different types of sequences to have a similar total sequence length, which is pointed out and discussed carefully by Lechner and Wiehler (2013). In our terminology, this involves conditioning on remaining in the eligibility state for a certain amount of time, which is not compatible with our setting. The Lechner and Wiehler (2013) approach is attractive when evaluating sequences of different types of treatments if the main interest is in the comparison of sequences before eligibility loss comes into play.¹⁰

3.5 Estimation

The identification results suggest estimators of the ATETs. In the online Appendix B.1 it is shown that if Assumptions 1 and 2 hold, then an unbiased estimator of $\text{ATET}(t_s)$ is:

$$\widehat{\text{ATET}}(t_s) = \frac{1}{\pi(t_s)N_{t_s}} \sum_{i \in T_{s,i}=t_s, T_{u,i} \geq t_s} Y_i - \frac{1}{\sum_{i \in T_{s,i} > T_{u,i} \geq t_s} w^{t_s}(T_{u,i}, X_i)} \sum_{i \in T_{s,i} > T_{u,i} \geq t_s} w^{t_s}(T_{u,i}, X_i) Y_i, \quad (10)$$

ates $L(t)$ and impose that the only possible regimes are those where the treatment cannot begin if T_u has already been realized. However, this subsequently requires a number of non-trivial changes in the model framework and in the end leads to an alternative framework (details available upon request). Robins and co-authors have not considered settings where the set of *all* possible regimes is restricted in such a way. Again, this is reflected in a discrepancy in terms of the ATE objects of interest. Even if we take $t_s'' = \infty$, the ATE as defined by Robins and co-authors is smaller in absolute value than ours because the former includes the (zero) contributions by individuals who leave the eligibility state before treatment in each of the two regimes considered.

¹⁰As an alternative approach, Lechner (1999, 2002) and Lechner et al. (2011) developed the “hypothetical treatment durations” method of inference, where hypothetical treatment dates are generated for each non-treated individual, and the actual and hypothetical treatment dates are used as covariates in the propensity score.

where N_t is the number of non-treated survivors at the beginning of t and $\pi(t) = \Pr(T_s = t | T_s \geq t, T_u \geq t)$ while the weights $w^{t_s}(\cdot, \cdot)$ are specified as

$$\begin{aligned} w^{t_s}(t_u, X) &= \frac{p(t_s, X)}{1 - p(t_s, X)} \cdot \frac{1}{\prod_{m=t_s+1}^{t_u} [1 - p(m, X)]} \\ p(t, X) &= \Pr(T_s = t | T_s \geq t, T_u \geq t, X). \end{aligned} \tag{11}$$

These weights $w^{t_s}(\cdot, \cdot)$ are replaced by estimated weights containing estimates of the treatment probabilities (propensity scores) $p(\cdot, \cdot)$ which also estimate $\pi(\cdot)$. This follows the ideas of Horovitz and Thompson (1952) and weighs the outcome responses of the treated and non-treated toward the target population.¹¹

Note that in the estimator, only non-treated individuals provide information about the counterfactual outcome under never treatment for those treated at t_s . The term $p(t_s, X_i)/(1 - p(t_s, X_i))$ in the weights follows IPW estimators in the static evaluation literature (e.g., Wooldridge, 2010), and adjusts for covariate differences between the treated t_s and those still waiting for treatment at t_s . The second term in the weights corrects for the fact that some non-treated at t_s start treatment at $t_s + 1, \dots$, with individuals with a high treatment probability and/or a high likelihood of remaining in the initial state for a long time being less likely to remain non-treated. This is reflected in the weights depending on both the observed characteristics and the time in the initial state. After estimation of the ATET(t_s), the over-all average effect ATET can be estimated using the ATET expression as identified in Subsection 3.3.

Bootstrapping is one way to obtain standard errors. A simulation study in the online Appendix C shows that the bias of our estimator is virtually zero in all simulations, that the bootstrap estimator for the standard errors has the correct size, and that the estimated standard error decreases by roughly 50% when the sample size is increased by a factor four (suggesting that the estimator is \sqrt{N} -convergent). As expected, a static IPW estimator is biased, and the bias is increasing in the share of treated.

¹¹It is well known that IPW estimation may be sensitive to extreme values of the propensity scores, since single observations may receive a high weight (see e.g., Frölich 2004, Huber et al. 2013, Busso et al. 2014). One way to overcome this problem is trimming. The three-step trimming approach by Huber et al. (2013) can be applied to our estimator; note that in our case the weights are a function of several propensity scores.

4 Extensions

4.1 Time-varying covariates

This section discusses extensions that are often empirically relevant. For sake of brevity, some details and derivations are relegated to the online Appendix B. In the first extension, we allow for selection on *time-varying* observed covariates. This allows the start of the treatment to depend on characteristics that change during the spell in the initial state. We use the notation X_{t^-} for the observed covariates at t , where t^- indicates that X is measured at least slightly before t . This rules out effects of the treatment in period t on X . The sequential unconfoundedness assumption now is:

Assumption 4 (Sequential unconfoundedness with time-varying covariates). For all t ,

$$P_t \perp Y(\infty) \mid X_{t^-}, T_s \geq t, T_u(\infty) \geq t.$$

By analogy to Assumption 1, we also require a “no-anticipation” assumption regarding future changes in the time-varying covariates. Similarly, future treatments and outcomes are not allowed to affect current covariates.

To identify the counterfactual outcome $E(Y(\infty) \mid T_s = t_s, T_u(t_s) \geq t_s)$, we can use similar reasoning as in Section 3. The main difference is that we now first average over $X_{t_s}^-$, then over $X_{t_s+1}^-$ and so on. We obtain

Theorem 2 (ATET with time-varying covariates). If Assumption 4 and no-anticipation hold then

$$\begin{aligned} \text{ATET}(t_s) &= E(Y \mid T_s = t_s, T_u \geq t_s) - \\ & [E_{X_{t_s}^- \mid T_s = t_s, T_u \geq t_s} [h(X_{t_s}^-) E(Y \mid T_s > t_s, T_u = t_s, X_{t_s}^-)] + \\ & [1 - h(X_{t_s}^-)] E_{X_{t_s+1}^- \mid T_s > t_s, T_u > t_s, X_{t_s}^-} [h(X_{t_s+1}^-) E(Y \mid T_s > t_s + 1, T_u = t_s + 1, X_{t_s+1}^-)] + \\ & [1 - h(X_{t_s+1}^-)] E_{X_{t_s+2}^- \mid T_s > t_s+1, T_u > t_s+1, X_{t_s+1}^-} [h(X_{t_s+2}^-) E(Y \mid T_s > t_s+2, T_u = t_s+2, X_{t_s+2}^-)] + \dots + \\ & [1 - h(X_{T_u^{\max}-1}^-)] E_{X_{T_u^{\max}}^- \mid T_s > T_u^{\max}-1, T_u > T_u^{\max}-1, X_{T_u^{\max}-1}^-} [h(X_{T_u^{\max}}^-) \times \\ & E(Y \mid T_s > T_u^{\max}, T_u = T_u^{\max}, X_{T_u^{\max}}^-)] \dots]]. \end{aligned}$$

where

$$h(X_t^-) = \Pr(T_u = t | T_s > t, T_u \geq t, X_t^-).$$

Proof See online Appendix B.3.

Corresponding estimators are also discussed in online Appendix B.3.

4.2 Right-censored durations

We now allow for right-censoring of the durations T_u in the initial state. Let T_c denote the time until the unit is right-censored. We assume that right-censoring can only take place in the initial state. At any point in time t in the initial state, we take censoring to be realized before any other event at t .^{12,13} If $T_c = t$, any treatment realizations at t and after t are unobserved. Also, the long-run outcome is assumed to be unobserved if the unit is right-censored. This introduces another selection problem because the durations of certain types of individuals may be censored at a higher rate. Since both treated and non-treated durations can be censored, this problem occurs for both non-treated and treated units.

Let the binary indicator $C_t := 1$ iff $T_c = t$. As usual with right-censoring, we assume unconfoundedness of the censoring process,

Assumption 5 (Right-censored process). For all t ,

$$C_t \perp Y(t), Y(\infty) \mid X, T_s \geq t, T_c \geq t, T_u(\infty) \geq t.$$

Using similar reasoning as above we can handle the selection due to both the treatment assignment and the right-censoring processes. We arrive at Theorem 3 (with results on estimators in online Appendix B.4),

Theorem 3 (ATET with right-censored durations). If Assumptions 1, 2 and 5 hold

¹²That is, the binary event governing whether $T_c = t$ vs. $T_c > t$ is realized before the binary events governing whether $T_s = t$ vs. $T_s > t$ and $T_u = t$ vs. $T_u > t$.

¹³Since the realization of the right-censoring outcome is the first event within each period, technically we now have the following average treatment effect of t_s on Y among those who are actually treated at t_s : $\text{ATET}(t_s) = E(Y(t_s) - Y(\infty) \mid T_s = t_s, T_c > t_s, T_u(t_s) \geq t_s)$.

then

$$\begin{aligned} \text{ATE}(t_s) = & E_{X|T_s=t_s, T_c>t_s, T_u \geq t_s} \left[\sum_{k=t_s}^{T_u^{\max}} h_{c1}(t, X, t_s) \left[\prod_{m=t_s}^{k-1} [1 - h_{c1}(t, X, t_s)] \right] E(Y|T_s = t_s, T_c > k, T_u = k, X) \right] - \\ & E_{X|T_s=t_s, T_c>t_s, T_u \geq t_s} \left[\sum_{k=t_s}^{T_u^{\max}} h_c(k, X) \left[\prod_{m=t_s}^{k-1} [1 - h_c(m, X)] \right] E(Y|T_s > k, T_c > k, T_u = k, X) \right]. \end{aligned}$$

where

$$\begin{aligned} h_{c1}(t, X, t_s) &= \Pr(T_u = t \mid T_s = t_s, T_c > t, T_u \geq t, X). \\ h_c(t, X) &= \Pr(T_u = t \mid T_s > t, T_c > t, T_u \geq t, X). \end{aligned}$$

Proof See online Appendix B.4.

4.3 Short-run outcomes

So far we have considered a long-run outcome realized after the units left the initial state. We now consider effects on shorter-run outcomes. Define Y_t as the observed outcome at the point in time t (relative to the moment of entry into the initial state). The corresponding potential outcomes are denoted by $Y_t(t_s)$. Here, the object of interest is the average effect of treatment at t_s on the outcome at points in time $t_s + \tau$ (i.e., τ periods after the start of the treatment) among the treated at t_s :

$$\text{ATE}(t_s, \tau) = E(Y_{t_s+\tau}(t_s) - Y_{t_s+\tau}(\infty) | T_s = t_s, T_u(t_s) \geq t_s). \quad (12)$$

In our leading example, this may be the effect on earnings at a certain number of periods (e.g., quarters) after the start of the training.

The observation rule is: $Y_t = \sum_{m=1}^t [\mathbf{I}(T_s = m) Y_t(m)] + \mathbf{I}(T_s > \min(T_u, t)) Y_t(\infty)$. Thus, if the individual is treated before t then we observe the potential outcome corresponding to the actually observed moment at which the treatment starts, but if the individual left the initial state before t without having been treated or if the individual is treated after t , then the observed outcome at t is the outcome corresponding to the assignment “never treated”. Note that Y_t may be realized before leaving the initial state, i.e. before T_u is realized. Hence, we need some additional assumptions regarding

no-anticipation. Without going into detail, we argue in online Appendix B.5 that under suitable assumptions, $\text{ATET}(t_s, \tau)$ is identified by

$$\begin{aligned} \text{ATET}(t_s, \tau) &= E(Y_{t_s+\tau} | T_s = t_s, T_u \geq t_s) - \\ &E_{X|T_s=t_s, T_u \geq t_s} \left[\sum_{k=t_s}^{t_s+\tau} h(k, X) \left[\prod_{m=t_s}^{k-1} [1 - h(m, X)] \right] E(Y | T_s > k, T_u = k, X) - \right. \\ &\left. \left[\prod_{m=t_s}^{t_s+\tau} [1 - h(m, X)] \right] E(Y_{t_s+\tau} | T_s > t_s + \tau, T_u > t_s + \tau, X) \right], \end{aligned}$$

where

$$h(t, X) = \Pr(T_u = t | T_s > t, T_u \geq t, X).$$

Estimators are discussed in online Appendix B.5.

5 Long-run effects of Swedish training programs

The Swedish training program called AMU aims to improve the skills of unemployed workers. The training courses are directed towards the upgrading of skills that are in short supply or that are expected to be in short supply. While the training program is for unemployed workers, it is important to know if training affects earnings in the long run. For participants, the long-run earnings under no treatment are unobserved, and under unconfoundedness the treated and the not-yet treated at t are comparable. However, since training may start at any point during the unemployment spell, some non-treated receive treatment in subsequent periods. This creates the dynamic evaluation problem studied in our paper.

Several register-based datasets are used in the analysis. From the Swedish employment office registers, we observe day-by-day information on unemployment status and participation in training as well as personal characteristics and information on the unemployment history. Additional background characteristics are obtained from an annual population register.¹⁴ The population register also includes information on labor

¹⁴Jointly, these registers include gender, age, marital status, number of children by age bracket, level of education, region of residence, UI entitlement, job search area, and a detailed unemployment and labor earnings history. See Table D.1 in online Appendix D for variables and sample statistics for the treated and non-treated.

earnings (all cash compensation paid by employers, consumer price adjusted) for any calendar year.

For the estimation we aggregate the daily data into monthly intervals. We sample all unemployment spells that start in the period 1995–1998, and focus on unemployed individuals aged 25–55 at the time, with no immediate previous unemployment record.¹⁵ We consider the effect of the first training program during the first unemployment spell that the individual experiences. All propensity scores are estimated using logistic regression models and standard errors are obtained by bootstrapping (500 replications).¹⁶ The analysis sample includes 792,580 unemployment spells of which 57,033 (7.2%) include a participation in the training program (Table D.1).

The evaluation approach rests on the unconfoundedness assumption. In the AMU program, the timing of entry into training is primarily determined by the caseworker (see Richardson and Van den Berg, 2013, for details on the assignment process). The administrative information that caseworkers have at their disposal concerning clients corresponds to the variables in our register data. Unobserved personality traits may be captured by features of the individual’s past labor market history. At the individual level, residual randomness in the start date of treatments is created by scheduling vacation in courses due to demand- and supply-driven capacity constraints.¹⁷ Our selection of covariates relies on the evidence in previous studies that have examined the relevance of various types of covariates in propensity scores for active labor market program participation. Notably, Heckman et al. (1998, 1999), Lechner and Wunsch (2013) and Biewen et al. (2014) stress that it is important to control for previous unemployment, lagged earnings, and local labor market conditions. We also control for various socio-economic characteristics.

The main ATET estimation results are presented in Table 1. In particular, Column 2 reports the key results using our dynamic estimator. We find sizable positive long-run ATET effects on yearly labor earnings. For instance, five years after the year of entry into the training program, the estimated average long-run effect is 6470 SEK which, in

¹⁵Specifically, no unemployment within 180 days before the unemployment spell. The age restrictions are imposed because the benefits entitlement rules and active labor market policy programs were different for persons aged below 25 or above 55 during the period studied here.

¹⁶We explore common support restrictions using a variant of the three-step approach in Huber et al. (2013). We set an upper bound on the weight given to the outcome of any given individual (1%) but obtain similar results as in the unrestricted case.

¹⁷Effects of anticipation of T_s on unemployment exit are unlikely to be relevant, as the assignment and the actual start date are very close in time and the latter is not determined a long time in advance; see Richardson and Van den Berg (2013) for details.

that year, is 4.9% of the average annual earnings among those who were treated. The effect is highly persistent as we further increase the long-run time horizon. Nine years after the year of entry into training, the estimated effect equals 5900 SEK.¹⁸ All the average effects are significant at the 1% level for any number of years after entry into training.

We may contrast this to results based on a static IPW approach which uses the following estimator,

$$\delta_{Static\ IPW} = \frac{1}{n_1} \sum_{i=1}^N W_i Y_i - \frac{1}{n_0} \sum_{i=1}^N \left(\frac{1}{n_0} \sum_{i=1}^N \frac{\hat{e}(x_i)(1 - W_i)}{1 - \hat{e}(x_i)} \right)^{-1} \frac{\hat{e}(x_i)(1 - W_i)Y_i}{1 - \hat{e}(x_i)}, \quad (13)$$

where W_i is a binary treatment indicator for individual i and $\hat{e}(X_i)$ is the estimated score for the probability that $W_i = 1$. Furthermore, n_1 and n_0 are the numbers of treated and non-treated observations, respectively, and $N = n_0 + n_1$. In this static setting, we take $W = 1$ if training is obtained within the first 12 months in unemployment.

The results based on the static approach are in Column 1. We see striking differences with the results based on the dynamic approach. The static approach suggests large negative effects of training, which would plausibly lead to different policy implications. Since the static approach does not adjust for differences in pre-treatment durations, its control group includes too many short-term unemployed workers. This leads to substantial downward bias, as short-term unemployed workers tend to have more favorable characteristics.¹⁹

6 Conclusions

This paper presents new identification results and proposes new estimators for treatment evaluations based on unconfoundedness assumptions in a dynamic setting. In the application, we estimate the long-run earnings effects of a flagship training program for unemployed workers in Sweden. Traditional static evaluation approaches lead to estimates with a negative sign. However, using our new approach, we find the effects to

¹⁸This uses the estimator of Subsection 3.5. For illustrative purposes we also display estimates for relatively low numbers of years after entry into training, where we apply the same estimator though a non-negligible fraction of individuals has not left unemployment then yet.

¹⁹Average earnings five years after entry into unemployment are about SEK 185,000 for the previously short-term unemployed (less than four months of unemployment) and SEK 129,000 for the previously long-term unemployed (more than 12 months of unemployment).

Table 1: ATET effects of Swedish AMU training on earnings, and comparison of estimators

	Static IPW [1]	Dynamic IPW [2]	Difference [3]:=[2]-[1]
Program entry year + 3	-3.68 (0.45)	8.88 (1.08)	12.56 (1.17)
Program entry year + 4	-4.45 (0.48)	7.65 (0.97)	12.11 (1.08)
Program entry year + 5	-5.63 (0.51)	6.47 (0.89)	12.10 (1.02)
Program entry year + 6	-6.37 (0.54)	5.19 (0.80)	11.56 (0.96)
Program entry year + 7	-7.24 (0.54)	4.91 (0.80)	12.15 (0.96)
Program entry year + 8	-7.44 (0.55)	5.23 (0.87)	12.67 (1.03)
Program entry year + 9	-7.46 (0.58)	5.90 (0.98)	13.36 (1.14)

Note: Outcome: labor earnings in 1000 SEK. “Dynamic IPW” is the dynamic IPW estimator introduced in this paper. “Static IPW” is a standard static IPW estimator with normalized weights. Bootstrapped standard errors (500 replications).

be significantly positive. Taking the dynamic treatment assignment into account thus turns out to be empirically important. Indeed, for reasonable values of the discount rate, our results lead to a reassessment of the net over-all benefits of the program.

This is interesting in the light of the literature and policy debates which seem to have agreed on a consensus that training programs are ineffective (see surveys by e.g. Heckman, LaLonde and Smith, 1999, and Crépon and Van den Berg, 2016). Much of the evidence was based on short-run observed outcomes and/or methodologies that could not appropriately deal with the dynamic assignment of training. Our results support a reappraisal of training.

The analyses in this paper do not explicitly distinguish between the direct causal effect of the treatment and an indirect effect that may run through the length of stay in the initial eligibility state. Often one would expect the former to dominate in the longer run. However, we view it as an interesting topic for further research to develop a formal statistical framework for mediation analysis that distinguishes between the two channels.

References

- Abbring J.H. and G.J. van den Berg (2003), “The non-parametric identification of treatment effects in duration models”, *Econometrica*, 71, 1491–1517.
- Biewen M., B. Fitzenberger, A. Osikominu and M. Paul (2014), “The Effectiveness of Public Sponsored Training Revisited: The Importance of Data and Methodological Choices”, *Journal of Labor Economics* 32(4), 837—897.
- Busso M., J. DiNardo and J. McCrary (2014), ““New Evidence on the Finite Sample Properties of Propensity Score Reweighting and Matching Estimators”, *Review of Economics and Statistics* 96(5), 885—897.
- Crépon, B. and G.J. van den Berg (2016), “Active Labor Market Policies Setting”, *Annual Review of Economics* 8(1), 521–546.
- Forslund A. and J. Vikström (2011) “Arbetsmarknadspolitikens effekter på sysselsättning och arbetslöshet – en översikt (The effects of labor market policies on employment and unemployment – an overview)”, IFAU Report 2011:7.
- Frölich M. (2004), “Finite Sample Properties of Propensity-Score Matching and Weighting Estimators”, *Review of Economics and Statistics* 86, 77–90.
- Heckman J.J, H. Ichimura and P. Todd (1998), “Matching As An Econometric Evaluation Estimator”, *Review of Economic Studies* 65, 261–294.
- Heckman, J.J., R.J. Lalonde and J.A. Smith (1999), “The Economics and Econometrics of Active Labor Market Programs”, in *Handbook of Labor Economics*, Vol. 3, eds. O. Ashenfelter and D. Card, Amsterdam: North-Holland, pp. 1865–2097.
- Hernán, M. and J.M. Robins (2020), *Causal Inference: What If*, CRC Press.
- Horovitz D.G. and D.J. Thompson (1952), “A Generalization of Sampling Without Replacement From a Finite Universe”, *Journal of the American Statistical Association*, 47, 663–685.
- Huber M., M. Lechner, and C. Wunsch (2013), “The Performance of Estimators Based on the Propensity Score”, *Journal of Econometrics* 175, 1–21.
- Imbens, G.W. (2000), “The role of the propensity score in estimating dose-response functions”, *Biometrika*, 87, 706–710.
- Lechner M. (1999), “Earnings and employment effects of continuous off-the-job training in East Germany after unification”, *Journal of Business and Economic Statistics*,

- 17, 74–90.
- Lechner M. (2002), “Some practical issues in the evaluation of heterogenous labour market programmes by matching methods”, *Journal of Royal Statistical Society A* 165(1), 59–82.
- Lechner M. (2009), “Sequential Causal Models for the Evaluation of Labor Market Programs”, *Journal of Business & Economic Statistics*, 27:1, 71–83.
- Lechner M. and R. Miquel (2010), “Identification of the effects of dynamic treatments by sequential conditional independence assumptions”, *Empirical Economics*, 39, 111–137.
- Lechner M., R. Miquel and C. Wunsch (2011), “Long-run Effects of Public Sector Sponsored Training in West Germany ”, *Journal of European Economic Association*, 9(4), 742–784.
- Lechner M. and S. Wiehler (2013), “Does the Order and Timing of Active Labour Market Programmes Matter?”, *Oxford Bulletin of Economics and Statistics*, 75(2), 180–212.
- Lechner M. and C. Wunsch (2013), “Sensitivity of matching-based program evaluations to the availability of control variables,” *Labour Economics*, 21, 111–121.
- Richardson K. and G.J. van den Berg (2013), “Duration dependence versus unobserved heterogeneity in treatment effects: Swedish labor market training and the transition rate to employment”, *Journal of Applied Econometrics*, 28, 325–351.
- Robins J.M. and M. Hernán (2009), “M. Estimation of the causal effects of time-varying exposures”, *Advances in Longitudinal Data Analysis*, eds: Fitzmaurice G., Davidian M., Verbeke G., and Molenberghs G., Boca Raton, FL: Chapman Hall, 553–599.
- Rosenbaum, P.R. and D.B. Rubin (1983), “The central role of the propensity score in observational studies for causal effects”, *Biometrika*, 70, 41–55.
- Rubin, D.B. (1974), “Estimating causal effects of treatments in randomized and non-randomized studies”, *Journal of Educational Psychology*, 66, 688–701.
- Wooldridge J.M. (2010). *Econometric Analysis of Cross Section and Panel Data (2nd ed.)*, MIT Press.

Appendix A. Identification in the general case

We first rewrite the counterfactual outcome using Assumption 1,

$$E(Y(\infty) | T_s = t_s, T_u(t_s) \geq t_s) = E(Y(\infty) | T_s = t_s, T_u(\infty) \geq t_s). \quad (\text{A.1})$$

Next, we condition on X and apply Assumption 2 for period t_s :

$$E(Y(\infty) | T_s = t_s, T_u(\infty) \geq t_s, X) = E(Y(\infty) | T_s > t_s, T_u(\infty) \geq t_s, X). \quad (\text{A.2})$$

Then, by the law of iterated expectations

$$E(Y(\infty) | T_s > t_s, T_u(\infty) \geq t_s, X) = \quad (\text{A.3})$$

$$\begin{aligned} & \Pr(T_u = t_s | T_s > t_s, T_u(\infty) \geq t_s, X) E(Y(\infty) | T_s > t_s, T_u(\infty) = t_s, X) + \\ & \Pr(T_u > t_s | T_s > t_s, T_u(\infty) \geq t_s, X) E(Y(\infty) | T_s > t_s, T_u(\infty) > t_s, X). \end{aligned}$$

That is, the counterfactual outcome under “never treatment” can be decomposed into average outcomes for individuals with $T_u = t_s$ and for individuals with $T_u > t_s$. The former, with $T_u = t_s$, in (A.3) are non-treated individuals who leave the initial state directly in period t_s . For this group, the observational rule in (1) and Assumption 1 give

$$E(Y(\infty) | T_s > t_s, T_u(\infty) = t_s, X) = E(Y | T_s > t_s, T_u = t_s, X). \quad (\text{A.4})$$

Also, note that the treatment probabilities $\Pr(T_u = t_s | T_s > t_s, T_u(\infty) \geq t_s, X) = \Pr(T_u = t_s | T_s > t_s, T_u \geq t_s, X)$ and $\Pr(T_u > t_s | T_s > t_s, T_u(\infty) \geq t_s, X) = \Pr(T_u > t_s | T_s > t_s, T_u \geq t_s, X)$ are observed.

For the other group, with $T_u > t_s$, in (A.3), i.e., for those who survive at least one additional time period, we have

$$E(Y(\infty) | T_s > t_s, T_u(\infty) > t_s, X) = E(Y(\infty) | T_s > t_s + 1, T_u(\infty) > t_s, X) =$$

$$E(Y(\infty) | T_s > t_s + 1, T_u(\infty) \geq t_s + 1, X),$$

where the first equality follows from Assumption 2 for period $t_s + 1$, and the second equality from re-writing. Next, using (A.3) by replacing t_s with $t_s + 1$ we have

$$E(Y(\infty) | T_s > t_s + 1, T_u(\infty) \geq t_s + 1, X) = \quad (\text{A.5})$$

$$\begin{aligned} & \Pr(T_u = t_s + 1 | T_s > t_s + 1, T_u(\infty) \geq t_s + 1, X) E(Y(\infty) | T_s > t_s + 1, T_u(\infty) = t_s + 1, X) + \\ & \Pr(T_u > t_s + 1 | T_s > t_s + 1, T_u(\infty) \geq t_s + 1, X) E(Y(\infty) | T_s > t_s + 1, T_u(\infty) > t_s + 1, X). \end{aligned}$$

For the sake of presentation, we introduce some auxiliary notation for the distribution of T_u :

$$h(t, X) = \Pr(T_u = t \mid T_s > t, T_u \geq t, X).$$

Using this notation, (A.3)-(A.5) and (A.4) for period $t_s + 1$ give

$$\begin{aligned} E(Y(\infty) | T_s > t_s, T_u(\infty) \geq t_s, X) &= h(t_s, X) & E(Y | T_s > t_s, T_u = t_s, X) + \\ & [1 - h(t_s, X)] h(t_s + 1, X) & E(Y | T_s > t_s + 1, T_u = t_s + 1, X) + \\ [1 - h(t_s, X)] [1 - h(t_s + 1, X)] & E(Y(\infty) | T_s > t_s + 1, T_u(\infty) > t_s + 1, X). \end{aligned}$$

All parts of this expression are observed except $E(Y(\infty) | T_s > t_s + 1, T_u(\infty) > t_s + 1, X)$. However, for this outcome, we can iteratively use (A.3) and (A.4) for $t_s + 2, \dots, T_u^{\max}$. This gives:

$$E(Y(\infty) | T_s > t_s, T_u(\infty) \geq t_s, X) = \sum_{k=t_s}^{T_u^{\max}} h(k, X) \left[\prod_{m=t_s}^{k-1} [1 - h(m, X)] \right] E(Y | T_s > k, T_u = k, X) \quad (\text{A.6})$$

(where we apply the convention that $\prod_{m=t_s}^{t_s-1} \dots = 1$). Interestingly, this is the expectation of the random outcome variable Y over a discrete-time competing risks duration distribution. In this interpretation, the competing risks are treatment and exit out of the initial state, and only the observed outcomes after the second risk (exit) are used.

Next, note that

$$E(Y(\infty) | T_s = t_s, T_u(t_s) \geq t_s) = E_{X | T_s = t_s, T_u \geq t_s} [E(Y(\infty) | T_s = t_s, T_u(t_s) \geq t_s, X)]. \quad (\text{A.7})$$

Then, $E(Y(\infty) \mid T_s = t_s, T_u(t_s) \geq t_s)$ is identified from (A.1)-(A.7). This gives the second component of $\text{ATET}(t_s)$, and (3) gives the first component of $\text{ATET}(t_s)$. This is summarized in Theorem 1.