

Policy discontinuity and duration outcomes

IFS Working Paper W18/10

Gerard J. van den Berg

Antoine Bozio

Mónica Costa Dias

Policy Discontinuity and Duration Outcomes

Gerard J. van den Berg^{*}

Antoine Bozio[†]

Mónica Costa Dias[‡]

March 2018

Abstract. Causal effects of a policy change on the hazard rates of a duration outcome variable are not identified from a comparison of spells before and after the policy change when there is unobserved heterogeneity in the effects and no model structure is imposed. We develop a discontinuity approach that overcomes this by considering spells that include the moment of the policy change and by exploiting variation in the moment at which different cohorts are exposed to the policy change. We prove identification of average treatment effects on hazard rates without model structure. We estimate these effects by kernel hazard regression. We use the introduction of the NDYP program for young unemployed individuals in the UK to estimate average program participation effects on the exit rate to work as well as anticipation effects.

^{*}University of Bristol, IFAU-Uppsala, IZA, ZEW and CEPR.

[†]Paris School of Economics and Institute for Fiscal Studies.

[‡]Institute for Fiscal Studies, University of Porto and IZA. Contact: monica_d@ifs.org.uk.

Keywords: policy evaluation, hazard rate, identification, causality, regression discontinuity, selectivity, kernel hazard estimation, local linear regression, average treatment effect, job search assistance, youth unemployment. *JEL codes:* J64, C14, C25.

Acknowledgements: We thank the Editor Rosa Matzkin, four anonymous Referees, Richard Blundell, Martin Nybom, Ina Drepper, Tatyana Krivobokova, Enno Mammen, David Margolis, and participants of conferences in Oberwolfach, Bielefeld, Nyborg, Mannheim and Westminster, an Invited Lecture at an Econometric Society European Meeting, and Keynote Lectures at an EALE Conference and an IZA/Worldbank Conference on Policy Evaluation, for useful comments. We gratefully acknowledge financial support from the ESRC (grant numbers RES-000-23-0797, ES/I02574X/1 and ES00116) and the Humboldt Stiftung.

1 Introduction

Most important life events and choices have dynamic consequences and are partly motivated by dynamic considerations. For instance, jobseekers may find it increasingly hard to move into work the longer they remain unemployed, possibly because of some associated stigma reducing the job-offer rate, decreasing job-searching skills or the depreciation of working skills. In turn, these effects may induce the unemployed to revise their job-acceptance criteria over time. Naturally, these dynamic patterns permeate into public policies and call for time-dependencies in successful interventions. In other words, the optimal design of dynamic policies depends crucially on understanding *when to intervene*. For example, should job-search assistance or other active labor market program target the newly unemployed, or should it focus on those who failed to find a job without support? The tradeoff facing the policymaker is that between dead-weight loss in supporting jobseekers who would have found a job anyway, and failing to reach those who would have benefited from assistance being provided early in the unemployment spell.

Despite its relevance, empirical researchers have struggled to uncover measures of policy impact that inform the dynamic design of policies when the object of interest is a duration variable (e.g., unemployment or employment duration, the duration until recovery from a disease) or, more specifically, its hazard rate. To deal with the possibility that policy exposure affects the distribution of unobserved characteristics conditional on survival, inference on hazard rates traditionally relies on rather strong assumptions. Most prominently, it is assumed that the duration dependence effect and the effects of the observed and unobserved explanatory variables on the individual hazard rate are proportional. In addition, independence between observed and unobserved individual characteristics is assumed (see e.g. Meyer, 1996, and Abbring and Van den Berg, 2005).¹

In this paper we show how to fruitfully use policy regime changes or discontinuities to identify and estimate the impact of treatments on duration outcomes without invoking semi-parametric assumptions. Our object of interest is the average effect of exposure to a policy at some duration t_0 on the probability of leaving a given state either instantaneously on the hazard rate, or on the conditional survival probability over some period $[t_0, t_1]$ (for $t_1 > t_0$).

The tradition in the empirical regression-discontinuity literature applied to duration outcomes is to compare spells starting after the reform (treated) with those starting before the reform right-censored at the time of the reform (controls). We show that such practice cannot

¹We discuss these issues in detail in Subsection 3.2 of this paper.

avoid the drawbacks discussed above and, at any post-reform moment, can only be informative about effects on durations shorter than the time since the reform. The latter is both inconvenient and makes it increasingly likely that the post-reform outcomes are affected by other post-reform events.

Our strategy deals with these issues by exploiting *ongoing* spells at the time of the reform. Specifically, we prove identification of an average causal treatment effect on the hazard rate of the duration distribution in the presence of unobserved heterogeneity, in a fully non-parametric setting *without* imposing a (mixed) proportional hazard model structure and without making a “random effects” assumption (i.e., without imposing independence of observed explanatory variables from unobserved heterogeneity). We obtain similar results for effects on conditional survival probabilities.

The basic insight stems from recognizing that the policy change is an exogenous time-varying binary explanatory variable whose discontinuity point varies independently across spells that started before the time of the reform. By comparing survivors who share a given elapsed duration t_0 *at the moment of the reform* to survivors at the same elapsed duration t_0 in an earlier cohort, we effectively compare two cohorts where the dynamic selection of individuals with favorable unobserved characteristics is the same up to t_0 . This means that a cross-cohort comparison of outcomes conditional on survival up to t_0 identifies average causal effects and is not contaminated by selection effects.

Our identification results naturally suggest an empirical implementation. When the outcome of interest is the hazard rate, the inputs for the policy evaluation are provided by estimates of observed hazard rates, i.e., hazard rates as a function of the elapsed duration and observed covariates. In general, such observed hazard rates are selective averages of individual hazard rates; in the paper we show how to carefully combine different observed hazard rates to obtain the average causal effect of interest.

These results are novel and general: our analysis demonstrates that the observed hazards are informative about average policy effects on individual hazard rates, in the presence of unobserved heterogeneity and without model structure. They also carry the perhaps counterintuitive insight that models exploiting policy changes as a time varying covariate may produce valid estimates of the average policy effect even if wrongly assuming away unobserved heterogeneity (and hence imposing homogeneous treatment effects).²

This paper connects to a number of branches of the evaluation literature. The first of these is the literature on “regression discontinuity” methods for policy evaluation (see for example

²These implications as well as empirical examples where they are relevant are discussed in Subsection 4.2.

Hahn, Todd and Van der Klaauw, 2001, Porter, 2003, and Frölich, 2007, for econometric contributions in non-parametric settings). Here, a nuanced difference results from our strategy of combining data from different cohorts to estimate treated and counterfactual hazard rates. It implies that we can drop the typical regression discontinuity assumption of continuity of the hazard rate in the absence of treatment at the duration of interest. This feature is particularly important in applications involving estimates of treatment effects at various durations, as continuity of the hazard function cannot generally be guaranteed everywhere.³ If continuity can be assumed, the alternative before-after comparison around the policy discontinuity point is also valid. In that case, the advantage that the latter has over the “cohort comparison” approach is that it implicitly controls for dynamic selection by relying on within-group comparisons; by contrast, the “cohort comparison” approach relies on matching to select the most appropriate cohort. However, the within-cohort comparison limits the information that can be retrieved: it can adequately estimate the immediate effects of a reform on the hazard rates, but does not handle well the identification of effects on conditional survival probabilities as it would entail contrasting survival probabilities at different elapsed durations, hence for different subgroups within the same cohort.

This paper also relates to the literature on treatment evaluation using “dynamic matching”, where treatment assignment can occur at many possible elapsed durations in the state of interest. This literature considers survivors at an elapsed duration t_0 and, amongst them, compares the treated at t_0 to the not-yet-treated at t_0 . The treatment status of the survivors at t_0 is assumed to be conditionally independent of the potential outcomes after t_0 , conditional on a set of covariates X . This is the identifying conditional independence assumption (CIA). The recent literature takes into account that the not-yet-treated at t_0 may be treated later, but in general it is silent on the dynamic selection before t_0 . Vikström (2014) provides an overview of matching estimators for average effects of a treatment at t_0 on the conditional survival distribution on (t_0, ∞) . Crépon et al. (2009) show that the underlying assumptions for identification are essentially the same as in our case, namely “conditional independence” and “no anticipation” (see Section 3 below). The matching estimator is then similar to our estimator for average effects on conditional survival probabilities. However, our analysis provides a foundation for the CIA, by relating it to events in the duration interval from zero up to t_0 . The analysis carries an important caveat for the application of dynamic matching estimators, namely that the CIA is unlikely to be satisfied if the treatment and comparison groups have

³Indeed, it is a typical characteristic of policies to generate discontinuous duration-dependencies, a feature which may undermine continuity of the hazard rate.

had systematically different event histories between entering the state of interest (say, entry into unemployment) and the moment of treatment t_0 , even if they have the same personal characteristics and the same labor market history before entry.

The outline of the paper is as follows. Section 2 of this paper provides an extensive discussion of the general policy relevance of this approach. Section 3 introduces the notation and discusses the parameters of interest, assumptions and identification results. Section 4 discusses non-parametric estimators of average causal effects on the hazard rates. These are based on the Müller and Wang (1994) boundary kernel hazard estimation method with data-adaptive local bandwidths, and on local linear kernel smoothing along the lines of Wang (2005). We also discuss estimation for conditional survival probabilities. In each case, the empirical setting typically allows for a choice between a range of cohorts that may serve as the comparison group of non-treated. We develop a “matching” procedure to select the most appropriate cohort. In Section 5 we apply our methodology to look at the impact of the New Deal for Young People (NDYP) – an active labor market policy designed to help the young unemployed getting back to work – on duration outcomes. In addition to providing effects of the job search assistance treatment, the application allows us to evaluate the effect of the policy announcement on unemployed individuals before they receive job search assistance. Effectively, our method enables inference on the anticipation of future job search assistance. Section 6 concludes.

2 Policy relevance and applications

This section discusses the general policy relevance of our approach. We describe a range of settings where the approach can be fruitfully applied, and we discuss the plausibility of the underlying assumptions.

Our method exploits exogenous variation from a policy discontinuity to evaluate the impact of such policy on duration outcomes, in the tradition of Difference-in-Differences and Regression Discontinuity Design methods. A prototype empirical setting is that of a reform in active labor market policies targeting the existing unemployed. The NDYP reform that we study in Section 5 is one such reform. It provided job-search assistance and a menu of other treatments to the young unemployed.⁴ Eligibility was limited to those whose elapsed unemployment duration reached an integer multiple of six months,⁵ and participation was compulsory for this group.

⁴The NDYP has been the subject of other evaluations studies; see e.g., Blundell et al. (2004), De Giorgi (2005), and Costa Dias, Ichimura and Van den Berg (2008).

⁵More precisely, eligibility was assessed on the duration of the claiming spell for Job-Seekers Allowance

The NDYP may have affected those enrolling in the program but also those as yet ineligible who foresee gaining eligibility in the future; the latter may change their behavior in an attempt to influence the likelihood of reaching the moment they can enroll.

In this setting, our method can be used to estimate various causal effects of enrollment in the NDYP for the first group of treated – i.e., young unemployed for 6 months, or multiples of 6 months, *at the time of the reform*. Specifically, we can estimate the instantaneous average effect on the exit rates from unemployment and the average effect on the probability of exiting unemployment over some time interval after enrollment, both by elapsed unemployment duration.⁶ Given the high salience of treatment among the treated from the onset of the reform – participants are summoned to the regular 6 months interview at the employment office and personally informed of their enrollment in the program – we expect that the treatment effects estimated for the first group of treated are informative, more generally, about the effects of enrolling in the NDYP. These are policy-relevant parameters because they provide information about how the impact of treatment changes with the elapsed duration at the time of the treatment *for the specific population of survivors at each duration*; in particular they can support the policy decision of *when to provide treatment*.

Clearly, the parameters identified by our method can only replicate the effect of treatment on later cohorts if the composition of the treated remains stable over time. Our method can be used to say something about the latter within the NDYP setting. Specifically, we can estimate the average effect of the reform on those at risk of future treatment by focusing on the unemployed at elapsed durations other than multiples of 6 months at the time of the reform – the effect of anticipating future treatment. We notice, however, that salience at the time of the reform may be weaker among the population of not-yet-treated than among the treated; salience may also change over time as knowledge about the new program becomes more widespread. Together these two features imply that anticipation effects measured at the time of the reform may suffer from attenuation as compared to what they may be later on. But then the finding of even small anticipation effects suggests that these may be non-negligible. This would be relevant information for researchers considering the feasibility of using spells

(JSA, the unemployment benefit in the UK). Eligibility to JSA is universal among the unemployed aged 18 and older, and not time limited. Entitlement is means-tested after the first 6 months, but the means test has no bite among the unemployed young people. Moreover, entitlement is gained upon visiting the employment office. Hence, we use unemployment duration and duration of the claiming spell interchangeably although we always refer to the duration of the claiming spell. The data we use later on, when assessing the impact of the NDYP reform, is from the administrative claiming records. More details can be found in Section 5.

⁶Note that the compulsory nature of the NDYP supports the identification of average treatment effects.

that become treated some time after the reform – estimation of treatment effects using such spells is likely to be confounded by selection bias due to anticipation, despite the mandatory nature of the treatment.

There is a myriad of other relevant policy contexts where our approach can bring important insights. One is that of designing and evaluating a dynamic unemployment insurance. Recent literature has addressed the questions of whether eligibility to unemployment benefits should be time-limited or whether entitlement should change with unemployment duration.⁷ Reforms to the time schedule of the unemployment benefits are salient to those on benefits, who are likely to respond quickly. With such a reform (as in Lalive, 2008) our method is well suited to provide fast and reliable information about its potential effects. Similarly, our approach can be used to assess the effect of pensions on time to retirement, the effect of maternity benefits on parental leave, or the impact of employment subsidies on unemployment durations. Examples are the Canadian Self-Sufficiency Program or the UK Working Families Tax Credits.⁸ The introduction (or reform) of these programs provides critical variation that can be used to estimate their impacts on the duration of out-of-work spells.

Without going into details until the next section, it is clear that our approach requires a number of assumptions to hold true. Most of these assumptions are common in dynamic evaluation settings, although they are often only tacitly made. First, we require a conditional independence assumption (CIA). To explain the CIA we adopt, note that we allow for systematic unobserved heterogeneity across subjects, which affect the individual potential-outcome hazard rates at all durations. The observed and systematic unobserved determinants do not capture all random variation in potential outcomes, by virtue of additional idiosyncratic random shocks. This setting is common in econometric duration models (see e.g. Van den Berg, 2001, Abbring and Van den Berg, 2003 and Abbring and Heckman, 2007, for detailed discussions). The CIA assumes independence between treatment assignment and potential outcomes conditional on both the observed *and* the systematic unobserved determinants. This CIA is weaker than usually encountered in the evaluation literature, as it does not imply a CIA conditional on the observed covariates only.

As a second key assumption, we require absence of anticipation of the event that we aim to evaluate. That is, before the event takes place, behavior does not depend on the residual time until the future event. To avoid confusion, note that the event can be a treatment or the introduction of a new policy but also a restructuring of an old program or the arrival of new

⁷See e.g., Kolsrud et al. (2016) and Stantcheva (2017).

⁸See e.g., Card and Hyslop (2009), Blundell and Hoynes (2004), and Brewer et al. (2006).

information regarding a policy or a future treatment. In the latter case the information arrival may lead to anticipation of the treatment. If we aim to evaluate the effect of the information arrival itself then the moment of the information arrival should not be anticipated.

Finally, we require that the treated and control cohorts at their respective moment of entry into the state of interest have an identical composition in terms of unobserved characteristics. In our NDYP illustration, this means that the composition of the unemployed at entry into unemployment must be the same for the cohort that is treated at a certain elapsed duration t_0 as for the cohort that entered unemployment earlier and hence is exposed to the treatment at an elapsed duration exceeding t_0 . This assumption requires careful consideration, since business cycles, seasonality and institutional features may cause systematic variation across cohorts. It depends on the context whether the assumption is plausible. Often one can use observed covariates to control for major differences. Also, often, it can be argued that cohorts that entered unemployment in close succession do not have a substantially different composition in terms of unobservables. In practice one has the choice between a range of potential control cohorts. In Section 4 we develop a procedure to discard those potential control cohorts that have different compositions in terms of unobservables. The latter include cohorts with different season- or cycle-specific compositions.

A characteristic feature of our methodology is its focus on the immediate or short-run effects of a policy reform. In practical terms, this is valuable because it allows us to build a counterfactual that rules out confounding effects resulting from dynamic selection, other post-reform events or business cycle effects while relying on minimal assumptions. We therefore focus on the short-run effects of the NDYP in our empirical application.

However, the term “short-run” deserves some qualification here. It means that we can only estimate treatment effects for the cohorts that are first exposed to the new policy, but it does not necessarily restrict the evaluation horizon to very brief interventions or very short durations when the goal is to assess the impact of the policy on the conditional survival rates. Indeed, a control group can be used to build the counterfactual over a potentially long spell interval in which it is not exposed to the reform or does not react to it.⁹

Our method is arguably less valuable in contexts where the reform takes time to bear its

⁹If the treatment only takes place after a certain time in the state of interest, then, depending on the specific design of the policy, even control cohorts that enter the state of interest closely after the exposed cohort may offer scope for medium-span comparisons. For instance, eligibility to the NDYP is restricted up to 6 months after the reform for the cohorts completing a 6-month duration in unemployment just before the reform. The earlier part of this period, over which it may be safe to rule out anticipation given that treatment is in the far horizon, can be used to extend the evaluation horizon.

effect. However, it can still be informative in circumstances where the absence of detectable effects of the reform on the hazard rates is indicative of no differential dynamic selection – for instance, where it can be assumed that the effect of treatment is non-negative or non-positive for all (the LATE monotonicity assumption of Imbens and Angrist, 1994). In such a case, non-zero heterogeneous reform effects cannot possibly offset and our method may be used to gauge the time it takes for the treatment to have an effect. Specifically, for a single exposed cohort we may estimate effects as if the reform only starts to kick in at points in time after the actual reform date for that cohort. As time (and hence the duration) elapses, this produces a sequence of estimates until an effect is found. One should be cautious with this approach, however. Even if the causal effect is monotonous, it can not be ruled out that intricate interactions between the treatment effect and unobserved heterogeneity at points in time after the actual reform date cause the dynamic selection to cancel out the treatment effect at such points in time.

3 Duration distributions, policy changes, and identification

3.1 Notation and assumptions

We consider a population of agents or individuals flowing into a state of interest, and we are interested in the durations that these individuals subsequently spend in that state. In particular, we are interested in the causal effect of a single “treatment” that is either assigned to commence at some time $s \in \mathbb{R}_+ := [0, \infty)$ after entering the state or is not assigned at all. We can cast this in the standard potential outcome framework by recognizing that the dynamically assigned binary treatment can be reinterpreted as a set of mutually exclusive treatments indexed by $\mathbb{R}_+ \cup \{\infty\}$ which we denote by \mathcal{A} . Here, the point ∞ represents the no-treatment case. To each treatment $s \in \mathcal{A}$ corresponds a random variable $T(s) \geq 0$, the potential outcome duration in the case that we would intervene and assign treatment s . For ease of exposition we assume that each $T(s)$ for given s is a random variable that is *continuously* distributed.

This framework may look more general than a framework for the evaluation of a single policy reform or a binary “reform exposure” indicator. When comparing outcomes in two mutually exclusive policy regimes, a framework with two mutually exclusive treatment statuses may suffice. However, the treatment at the elapsed duration s can be interpreted as representing the exposure to a reform occurring at the individual elapsed duration s . In the stock of individuals

in the state of interest at the moment of the policy reform, the elapsed duration from the moment of inflow until the moment of exposure to the reform will be dispersed. We therefore do not restrict the number of elements in \mathcal{A} at this stage.

Causal inference is concerned with contrasting potential outcomes corresponding to different treatments. Specifically, we are interested in differences between the distributions of $T(s)$ and $T(s')$ corresponding to treatments $s, s' \in \mathcal{A}$. These differences are called treatment effects. In social sciences, the exit rate or hazard rate of a duration distribution is the most interesting feature of this distribution, as it is directly related to the agent's behavior and his information set and circumstances conditional on survival into the state of interest (see Van den Berg, 2001).¹⁰ Therefore we focus on average effects of the treatments on the individual exit rate out of the state of interest and the individual conditional exit probabilities out of this state.

For arbitrary s , let the distribution function of $T(s)$ be denoted by $F_{T(s)}$. This is a function of the time t since inflow into the state of interest. The corresponding "integrated hazard" $\Theta_{T(s)}(t)$ is defined by $\Theta_{T(s)}(t) := -\log(1 - F_{T(s)}(t))$. We assume that $\Theta_{T(s)}(t)$ has a continuous first-derivative on $(0, \infty)$ except for a finite number of points where it is right-continuous. The hazard rate of $T(s)$ denoted by $\theta_{T(s)}$ can then be formally introduced as the right-derivative of the integrated hazard with respect to t . We assume that the hazard rates satisfy regularity conditions that guarantee existence of all expressions below.

The individual treatment effect of interest is

$$\theta_{T(s')}(t) - \theta_{T(s)}(t) \tag{1}$$

for $t \geq 0$ and for $s', s \in \mathcal{A}$. This is the additive effect on the hazard rate at t of replacing one treatment s by another treatment s' , as a function of t . In the case of a policy reform, this is the additive effect on the hazard rate at t of exposure to the reform at elapsed duration s' instead of at the elapsed duration s .

In addition, we consider the treatment effect on the probability of surviving up to t conditional on survival up to t_0 ,

$$\frac{1 - F_{T(s')}(t)}{1 - F_{T(s')}(t_0)} - \frac{1 - F_{T(s)}(t)}{1 - F_{T(s)}(t_0)} \tag{2}$$

for $t \geq t_0 \geq 0$ and $s', s \in \mathcal{A}$. At $t_0 = 0$ this captures the effect on the unconditional survival function. We also consider the multiplicative or relative treatment effect on the hazard rate at

¹⁰With T continuous, the hazard rate at elapsed duration t is defined as $\theta(t) = \lim_{dt \downarrow 0} \Pr(T \in [t, t + dt] | T \geq t) / dt$.

t ,

$$\frac{\theta_{T(s')}(t)}{\theta_{T(s)}(t)} \tag{3}$$

for all $t \geq 0$ and $s', s \in \mathcal{A}$.

Because the treatments are mutually exclusive, we can never observe potential outcomes corresponding to different treatments simultaneously. Treatments are assigned according to a random variable S with support \mathcal{A} . The actual outcome is $T := T(S)$; all other potential outcomes are counterfactual. Here, we may simply take S to denote the elapsed duration at the moment at which the agent is exposed to the reform.

We allow agents to be *ex ante* heterogeneous in terms of observed characteristics X and unobserved characteristics V . The latter are systematic in that they may affect the individual potential-outcome hazard rates at all durations. Both X and V may be exogenously time-varying, but for ease of exposition we abstract from this. For the same reason, we take V to be a continuous random variable.

As noted in Section 2, we take it that X and V do not capture all random variation in potential outcomes. This means that the distributions of $T(s)|X, V$ are non-degenerate, which in turn allows us to define corresponding hazard rates conditional on X and V . The difference between the distributions and individual drawings from it represent idiosyncratic shocks and capture residual unexplained variation in individual potential-outcome durations. Their distribution is independent of all other model determinants. See Van den Berg (2001), Abbring and Van den Berg (2003) and Abbring and Heckman (2007) for a detailed discussion.

The hazard rate, integrated hazard and the distribution function of $T(s)$ can be defined for individuals with given characteristics (X, V) . We denote these by $\theta_{T(s)}(t | X, V)$, $\Theta_{T(s)}(t | X, V)$ and $F_{T(s)}(t | X, V)$, respectively.¹¹ The survival function is $\bar{F}_{T(s)}(t | X, V) = 1 - F_{T(s)}(t | X, V)$. The individual treatment effects defined above can be defined accordingly as functions of X and V .

Inference is based on a random sample of agents from the population. For each of these we observe the duration outcome T and the observed covariates X . If the treatment S captures the exposure to a policy reform then S is effectively observable to the researcher for all agents (but not necessarily to the agents themselves; see Assumption 2 below). We allow for random right-censoring of T .¹²

¹¹Thus, $\theta_{T(s)}(t | X, V)$ denotes the hazard of $T(s)|X, V$. Arguably, the most accurate notation for this would be $\theta_{T(s)|X, V}(t)$ as it uses the subscript to refer to the random variable of the hazard rate considered. To facilitate readability we adopt the former option. This also applies to Θ and F .

¹²This is usually referred to as “simple random right-censoring”. Extensions to more general forms of inde-

We assume that treatment assignment is randomized conditional on covariates X, V , and also that treatment assignment is independent of V given X ,

Assumption 1 (Assignment). $S \perp\!\!\!\perp \{T(s)\} \mid (X, V)$ and $S \perp\!\!\!\perp V \mid X$.

The first part of this assumption generalizes the usual conditional independence assumption (CIA) in the evaluation literature. The assumption is in line with cases in which a universal policy is rigorously implemented from a specific point in calendar time onwards. In our setting, as we allow for unobserved heterogeneity, the CIA is conditional on observed *and* unobserved explanatory variables. Of course V is not observed so we cannot condition on it in data, but the second part of the assumption provides a way out by imposing that the actual treatment S is unrelated to V conditional on X . This second part of Assumption 1 captures the idea that different cohorts flowing into the state of interest at different points of time have identical compositions of $V|X$. After all, different cohorts correspond to different values of S . Note that a randomized experiment with an instantaneous binary treatment status (i.e. $\mathcal{A} = \{0, \infty\}$) also satisfies the assumption. As shown in Abbring and Van den Berg (2003, 2005), settings in which the assumption that $S \perp\!\!\!\perp V \mid X$ is relaxed require a semi-parametric model framework in order to be able to point-identify objects of interest. However, to some extent, the data may be informative on the violation of that assumption (see Subsection 4.2 and Section 5 below).

Notice that Assumption 1 implies that $S \perp\!\!\!\perp \{T(s)\} \mid X$. The latter is assumed from the outset in the dynamic matching literature (see e.g. Crépon et al., 2009).¹³ Our Assumption 1 provides a foundation for the dynamic matching CIA, by relating it to dynamic selection in the duration interval from zero up to the moment of treatment. We need Assumption 1 to meaningfully compare treated and control cohorts.

Throughout much of the paper, we assume that there is no anticipation by agents of the reform or event that we aim to evaluate. With this we mean that, before a reform takes place, agents' behavior does not depend on the time remaining until the future reform. We formalize this by assuming that current integrated potential hazards do not depend on the moment of future treatment exposure,

Assumption 2 (No anticipation). *For all $s \in (0, \infty)$ and for all $t \leq s$ and all X, V ,*
 $\Theta_{T(s)}(t|X, V) = \Theta_{T(\infty)}(t|X, V)$

pendent censoring and filtering are straightforward (see Andersen et al., 1993, and Fleming and Harrington, 1991).

¹³In the unrealistic special case where V is degenerate, $\Theta_{T(s)}$ can be estimated using standard hazard regression techniques (see e.g. Fleming and Harrington, 1991).

(See Abbring and Van den Berg, 2003, for a detailed discussion.) Here, $\Theta_{T(\infty)}$ is the integrated hazard of the potential duration corresponding to never enrolling in treatment. If agents receive information about a future treatment or reform then this may lead to anticipation of that treatment or reform. For example, caseworkers may privately inform unemployed workers about the date of onset of a future reform, and this may induce the workers to intensify their job search, causing them to leave unemployment before the reform. This violates the above assumption and complicates inference. However, if the moment of the arrival of information is observable to the researcher then one may evaluate the effect of the information arrival (and hence the effect of anticipation of the future treatment or reform) itself. In that case the moment of the information arrival should not be anticipated. Section 5 contains an empirical illustration of this.

3.2 Spells from the steady states before and after the policy change

In this subsection we consider empirical inference if the data collection leads to two samples: one in which $\Pr(S = 0) = 1$ and one in which $\Pr(S = \infty) = 1$. In the context of policy reform evaluation, these samples originate from two sub-populations. One sample is drawn from the inflow into the state of interest after the introduction of a policy, whereas the other sample is drawn from the inflow into the state of interest infinitely¹⁴ long before the introduction of the policy. Figure 1 depicts this setting in a Lexis diagram, where τ denotes calendar time and τ^* denotes the moment at which the reform is implemented. Each diagonal line represents a single cohort. Notice that we tacitly assume that the reform is universal in the sense that all spells starting after τ^* are exposed to the new policy.

The main purpose of the present subsection is to demonstrate that this sampling scheme has limited value for inference on the causal effects of interest. Furthermore, the subsection motivates the study of an alternative sampling scheme and inferential approach in the subsequent subsection.

Note that in the current dichotomous setting, S is observable by the agent from the onset,

¹⁴Or, at least sufficiently long before the reform to observe outcomes in a sufficiently large duration interval. In this case, the outcomes are right-censored at the moment of the reform. Alternatively, one may think of the sample with $\Pr(S = 0) = 1$ as a sample of fully treated agents and the other sample as a sample of controls. Provided that no ambiguity arises, we use the terms “pre-reform policy”, “pre-policy”, and “control” interchangeably. The same applies to “post-reform policy”, “post-policy” and “treatment”, and the same also applies to “moment of the policy change” “reform” and “introduction of the policy”. A more explicit discussion is provided in Subsection 4.2.




Figure 1: *Lexis diagram: “Before” sample and “after” sample based on cohorts flowing into the state of of interest before or after the policy reform, respectively.*

and Assumption 2 is void. Assumption 1 implies that the treatment assignment upon inflow into the state of interest is not selective, conditional on X . In particular, the distribution of characteristics $V|X$ at inflow is the same in each policy regime.

As in most evaluation literature, we study inference on averages of individual treatment effects as that expressed in expression (1). We thus need to average

$$\theta_{T(0)}(t|X, V) - \theta_{T(\infty)}(t|X, V)$$

over the distribution of $V|X$ in the relevant sub-population.

The dynamic selection of survivors imposes particular challenges in defining the relevant sub-populations over which to aggregate. As is well known, the distribution of $V|X$ among survivors typically differs from the population distribution of $V|X$. Individuals with values of V that give rise to high hazard rates at durations below t are underrepresented among the survivors at t . This implies, first of all, that it is not informative to average over the distribution of $V|X$ in the full population, since in either policy regime the sub-population of survivors at the elapsed duration t is systematically different from the full population.

Moreover, as indicated by Meyer (1996), if the treatment has a causal effect on the duration, then, typically, the distribution of $V|X$ among the survivors at points in time $t > 0$ depends on the treatment, so $V \not\perp\!\!\!\perp S|X, T > t$. In other words, there is no treatment randomization at $t > 0$ despite the randomization ($V \perp\!\!\!\perp S|X$) at $t = 0$. To illustrate this, let f , F , Θ and \bar{F} be generic symbols for a density, a distribution function, an integrated hazard, and a survivor function, with subscripts denoting the corresponding random variable (note that $\bar{F} = 1 - F = e^{-\Theta}$). By

Bayes' rule, there holds that

$$f_V(v|X, T > t, S) = \frac{\bar{F}_T(t|X, S, V)f_V(v|X)}{\int_0^\infty \bar{F}_T(t|X, S, V)dF_V(v|X)}, \quad (4)$$

which typically varies with S .¹⁵

Thus, contrasting the two sub-populations defined by conditioning on the observed $(T \geq t, X, S)$ does not lead to meaningful average treatment effects, because the sub-populations have systematically different compositions of unobserved characteristics V . To proceed, we consider alternative concepts of average treatment effects. These measures average over sub-populations of individuals for whom one or more *counterfactual* duration outcomes exceed t . This follows Abbring and Van den Berg (2005). Specifically, we consider

$$\begin{aligned} & \mathbb{E} \left[\theta_{T(0)}(t|X, V) - \theta_{T(\infty)}(t|X, V) \mid X, T(0) \geq t \right], \\ & \mathbb{E} \left[\theta_{T(0)}(t|X, V) - \theta_{T(\infty)}(t|X, V) \mid X, T(0) \geq t, T(\infty) \geq t \right], \\ & \mathbb{E} \left[\theta_{T(0)}(t|X, V) - \theta_{T(\infty)}(t|X, V) \mid X, T(\infty) \geq t \right] \end{aligned}$$

which can be called the *Average Treatment effect on the Treated Survivors at t* (ATTS($t|X$)), the *Average Treatment effect on the Survivors at t* (ATS($t|X$)), and the *Average Treatment effect on the Non-Treated Survivors at t* (ATNTS($t|X$)). ATTS($t|X$) averages over the distribution of $V|X$ among the survivors at t if the agents are assigned to the “treatment” (i.e., are assigned to $s = 0$, or, in other words, are exposed to the policy introduced by the reform). Under randomization, this is equivalent to averaging over the distribution of V among the treated survivors at t (so with $X, T \geq t, S = 0$). ATNTS($t|X$) is the counterpart of this for assignment to the control group. ATS($t|X$) averages over the distribution of $V|X$ among individuals who survive up to t under both possible treatment regimes. These measures can subsequently be aggregated over some distribution of X . Analogous additive and multiplicative effects can be defined for the conditional survival probabilities and the hazard rate, respectively (recall equations (2) and (3)).¹⁶ Note that in general all measures are properties of sub-populations whose composition depends on the treatment effect in the duration interval $[0, t)$.

¹⁵It is not difficult to construct examples in which the distribution of $V|X$ among the treated survivors at t is first-order stochastically dominated by the distribution of $V|X$ among the non-treated survivors at t , if there is a strong positive interaction between being treated and V in the individual hazard rates $\theta_{T(s)}(t|X, V)$ and if these hazard rates increase in V and in being treated (see Van den Berg, 2001). In such scenarios, the individual hazard rate at t is disproportionately large if both $S = 0$ and V is large, and as a result the treated survivors at t may contain relatively few treated individuals with a high value of V .

¹⁶The ATS($t|X$) version for the multiplicative effect on the hazard rate basically equals the *survivor average*

The above average effects cannot be estimated non-parametrically from the data design of the present subsection. Non-parametric inference produces sample equivalents of $\theta_T(t|X, S = 0)$ and $\theta_T(t|X, S = \infty)$ and of $\bar{F}_T(t|X, S = 0)/\bar{F}_T(t_0|X, S = 0)$ and $\bar{F}_T(t|X, S = \infty)/\bar{F}_T(t_0|X, S = \infty)$. For given t, s, X , individual and observable hazard rates are connected by the following relation (see e.g. Fleming and Harrington, 1991, and Andersen et al., 1993),

$$\theta_T(t|X, S = s) = \mathbb{E}(\theta_T(t | X, S = s, V) | X, T \geq t, S = s). \quad (5)$$

By definition, therefore,

$$\begin{aligned} & \theta_T(t | X, S = 0) - \theta_T(t | X, S = \infty) \\ & \equiv \mathbb{E}[\theta_T(t | X, S = 0, V) | X, T \geq t, S = 0] - \mathbb{E}[\theta_T(t | X, S = \infty, V) | X, T \geq t, S = \infty] \\ & = \mathbb{E}[\theta_T(t | X, S = 0, V) | X, T \geq t, S = 0] - \mathbb{E}[\theta_T(t | X, S = \infty, V) | X, T \geq t, S = 0] \\ & \quad + \mathbb{E}[\theta_T(t | X, S = \infty, V) | X, T \geq t, S = 0] - \mathbb{E}[\theta_T(t | X, S = \infty, V) | X, T \geq t, S = \infty] \end{aligned}$$

which is the sum of two differences. The first difference is the average treatment effect $\text{ATTS}(t|X)$ (for sake of brevity, we refer to the next subsection for the proof of this statement). The second difference is the selection effect due to the fact that at $T = t$, among the survivors at t , those exposed to the post-reform policy and those not exposed have systematically different unobserved characteristics despite the randomization of the regime status at $t = 0$. Since the second term on the right-hand side reflects the selection effect and is unobserved, we conclude that the left-hand side cannot be used to non-parametrically estimate $\text{ATTS}(t|X)$.¹⁷

The results are straightforwardly extended to more general sets of possible treatments \mathcal{A} as long as we only use data on spells within which the treatment status does not change. To identify average treatment effects in the setting of the current subsection, one needs to adopt a semi-parametric model structure like an MPH model, or one needs to assume absence of unobserved heterogeneity.

Admittedly, if interest is in the average additive treatment effect on the *unconditional* survival probability at t , i.e. $\mathbb{E}[\bar{F}_{T(0)}(t|X, V) - \bar{F}_{T(\infty)}(t|X, V)]$ then the above concerns are less relevant. The latter average effect is identified under a randomization assumption such as Assumption 1, from the observed expression $\Pr(T > t|X, S = 0) - \Pr(T > t|X, S = \infty)$.

causal effect of Rubin (2000) in case the latter measure is applied to the duration outcome itself rather than to non-duration outcomes.

¹⁷By analogy to the remarks on equation (4), one can construct examples where $\theta_T(t|X, S = 0) < \theta_T(t|X, S = \infty)$ even if $\theta_{T(0)}(t|X, V) > \theta_{T(\infty)}(t|X, V)$ almost surely for all (t, V, X) .

Inference on the two survivor functions in this expression is straightforward; see e.g. Andersen et al. (1993). We also point out that under the assumption that all individual treatment effects have the same sign across t and V , this sign is identified from the observed distributions or the observed hazard rates at $t = 0$.

3.3 Spells that are ongoing at the moment of the policy change

In this subsection we consider empirical inference if the data collection is based on random samples from cohorts flowing into the state of interest before the introduction of a universal policy at τ^* .¹⁸ Contrary to the previous subsection, we track duration outcomes in these cohorts beyond τ^* . Figure 2 depicts this setting. As in Figure 1, each diagonal line represents a single entry cohort.




Figure 2: *Lexis diagram: “before” sample based on cohorts flowing into the state of of interest before the policy reform, including spells that are ongoing at the moment of the policy change and that are followed beyond that moment.*

We assume that the post-reform policy regime applies to all agents, from calendar time τ^* onwards, including to those who enter the state of interest before τ^* . Inflow at time $\tau_0 \leq \tau^*$

¹⁸“Universal” here means that the new policy regime is compulsory for all individuals currently in the state of interest at τ^* or entering the state of interest after τ^* . We use the terms “universal” and “compulsory” interchangeably with the term “comprehensive”.

leads to $S := \tau^* - \tau_0$. Thus, there is a one-to-one correspondence between the moment of inflow and the duration at which the treatment starts. However, in this setting, S is not observed by the agent until calendar time τ^* , as there is no anticipation of the introduction of the new policy program (Assumption 2). We rule out that the distributions of $T(s) \mid (X, V)$ are discontinuous at $T(s) = s$ (though of course the hazard rates may be discontinuous there).

Assumption 1 again implies that the treatment assignment upon inflow into the state of interest is not selective, conditional on X . In fact, as we shall see, we only require Assumption 1 for the cohorts flowing in before τ^* . The assumption's implication that the distribution of characteristics $V \mid X$ at inflow is constant over calendar time is therefore only required for inflow dates before τ^* . This is attractive because the effect of a policy reform on the decision to enter the state of interest may vary with unobserved individual characteristics.

Comparing agents who flow out before τ^* to those who flow in after τ^* is hampered by the same problems as in the previous subsection. However, we can now also examine the effect at duration $\tau^* - \tau_0$ of a treatment that starts at duration S , as compared to the case where at duration $\tau^* - \tau_0$ no treatment is assigned yet. To this purpose, we may define the average additive treatment effects on hazard rates at durations t by analogy to those in the previous subsection, contrasting treatment assignment at time s' with treatment assignment at time $s > s'$. In particular,

$$\begin{aligned} \text{ATTS}(s', s, t \mid X) &:= \mathbb{E} \left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X, T(s') \geq t \right] && \text{with } s' \leq t, s; \\ \text{ATNTS}(s', s, t \mid X) &:= \mathbb{E} \left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X, T(s) \geq t \right] && \text{with } s' \leq t, s. \end{aligned}$$

The following proposition is the key to the main result of this paper.

Proposition 1. *Consider a cohort flowing in at calendar time $\tau_0 < \tau^*$ and a cohort flowing in at $\tau_1 < \tau_0$. Let $t_i := \tau^* - \tau_i$. Under Assumptions 1 and 2, $[V \mid T \geq t_0, X, S = t_0]$ and $[V \mid T \geq t_0, X, S = t_1]$ have the same distribution, namely the distribution of $[V \mid T(s) \geq t_0, X]$ with $s \geq t_0$. This distribution does not vary with s for all $s \geq t_0$.*

Proof: Note that $\tau_1 < \tau_0$ implies that $t_0 < t_1$. Let \Pr be a general symbol for a density as well as a probability. By Bayes' law, the density $\Pr(V \mid T \geq t_0, X, S = t_i)$ (with $i = 0, 1$) can be written as

$$\frac{\Pr(T \geq t_0 \mid V, X, S = t_i) \Pr(V \mid X, S = t_i)}{\Pr(T \geq t_0 \mid X, S = t_i)}$$

(see e.g. Lancaster, 1990). In this expression, $\Pr(T \geq t_0 \mid V, X, S = t_i)$ equals $\Pr(T(t_i) \geq t_0 \mid V, X)$ due to the randomized assignment assumption (Assumption 1: $S \perp\!\!\!\perp \{T(s)\} \mid (X, V)$). Moreover,

$\Pr(V|X, S = t_i)$ equals $\Pr(V|X)$ due to the second part of Assumption 1 ($S \perp\!\!\!\perp V \mid X$).¹⁹ This implies that the density $\Pr(V|T \geq t_0, X, S = t_i)$ as a function of V is proportional to $\Pr(T(t_i) \geq t_0|V, X)\Pr(V|X)$ which is proportional to $\Pr(V|T(t_i) \geq t_0, X)$.

Next, we show that $\Pr(V|T(s) \geq t_0, X)$ is the same for every $s \geq t_0$ including $s = t_1$. By analogy to the first part of the proof, the second part of Assumption 1 implies that the density $\Pr(V|T(s) \geq t_0, X)$ as a function of V is proportional to $\Pr(T(s) \geq t_0|V, X)\Pr(V|X)$. The term $\Pr(T(s) \geq t_0|V, X)$ can be expressed as $\exp(-\Theta_{T(s)}(t_0|X, V))$. By virtue of Assumption 2, this equals $\exp(-\Theta_{T(t_0)}(t_0|X, V))$ since $s \geq t_0$. This implies that the density $\Pr(V|T(s) \geq t_0, X)$ as a function of V is proportional to $\Pr(T(t_0) \geq t_0|V, X)\Pr(V|X)$, where the latter is proportional to $\Pr(V|T(t_0) \geq t_0, X)$. Thus, $\Pr(V|T(s) \geq t_0, X)$ is the same for every $s \geq t_0$. \square

The significance of this proposition is that it demonstrates that the sub-population of individuals who are observed to be treated at the elapsed duration t_0 and the sub-population of survivors at t_0 who will be treated at a higher elapsed duration have the same composition. In other words, $V \perp\!\!\!\perp S \mid (T \geq t_0, X, S \geq t_0)$. Clearly, it is crucial that the sub-populations come from populations that are identical to each other at their moment of entry into the state of interest. Moreover, it is crucial that individuals do not act on the future moment of treatment, because then their hazard rates (and consequently the dynamic selection) would already differ before t_0 . Under these two assumptions, the dynamic selection between the moment of entry and the elapsed duration t_0 proceeds identically in both populations, so the resulting sub-populations at t_0 have an identical distribution of unobserved characteristics.

We now apply this to the identification of average treatment effects. This gives the main methodological result of the paper. Recall that $t_i := \tau^* - \tau_i$. From a cohort flowing in at $\tau_i < \tau^*$, we observe the distribution of $[T|X, S = t_i]$. This entails observation of the conditional duration distribution of $[T|T \geq t_0, X, S = t_i]$ and the hazard rate $\theta_T(t_0|X, S = t_i)$ evaluated at $t_i = t_0$.

Proposition 2. *Consider the introduction of a compulsory policy at a given point in time. Suppose we have duration data from cohorts that flow into the state of interest (where they may eventually be exposed to the policy) before this point of time. Under Assumptions 1 and 2, the average treatment effects on the individual hazard rate $ATTS(t_0, t_1, t_0|X)$ and $ATNTS(t_0, t_1, t_0|X)$ are non-parametrically identified and equal the observable $\theta_T(t_0|X, S = t_0) - \theta_T(t_0|X, S = t_1)$ with $t_1 > t_0$. These do not depend on t_1 as long as t_1 exceeds t_0 .*

¹⁹In the setting of this subsection, the assumptions entail that the policy or treatment status is randomized among the stock of subjects in the state of interest, given X . See e.g. Ridder (1984) for an extensive discussion and Heckman and Singer (1984) for a discussion of sampling frames.

We first present the proof and then discuss the relevance of the result.

Proof:

$$\begin{aligned}
& \theta_T(t_0|X, S = t_0) - \theta_T(t_0|X, S = t_1) \\
&= \mathbb{E}[\theta_T(t_0|X, V, S = t_0) \mid X, T \geq t_0, S = t_0] - \mathbb{E}[\theta_T(t_0|X, V, S = t_1) \mid X, T \geq t_0, S = t_1] \\
&= \mathbb{E}[\theta_{T(t_0)}(t_0|X, V) \mid X, T(t_0) \geq t_0] - \mathbb{E}[\theta_{T(t_1)}(t_0|X, V) \mid X, T(t_0) \geq t_0]
\end{aligned}$$

The first equality in the above expression follows from applying equation (5) to replace each hazard rate in the first line. By Proposition 1, the distributions over which the expectations are taken in the second line are the same for any $t_1 \geq t_0$ and are equal to the distribution of $[V|T(s) \geq t_0, X]$. This explains the second equality. As a result,

$$\begin{aligned}
& \theta_T(t_0|X, S = t_0) - \theta_T(t_0|X, S = t_1) \\
&= \mathbb{E}[\theta_{T(t_0)}(t_0|X, V) - \theta_{T(t_1)}(t_0|X, V) \mid X, T(t_0) \geq t_0] \\
&= \text{ATTS}(t_0, t_1, t_0|X)
\end{aligned}$$

By substituting into the second-to-last expression that the distributions of $[V|T(t_0) \geq t_0, X]$ and $[V|T(t_1) \geq t_0, X]$ are identical, it also follows that $\text{ATTS}(t_0, t_1, t_0|X)$ equals $\text{ATNTS}(t_0, t_1, t_0|X)$. Moreover, in this second-to-last expression, changing the value of t_1 does not have an effect on the value of the expression as long as $t_1 > t_0$, because of Assumption 2. \square .

The $\text{ATTS}(t_0, t_1, t_0|X)$ and $\text{ATNTS}(t_0, t_1, t_0|X)$ capture the instantaneous causal effect of exposure to the policy (i.e., the instantaneous causal effect of the treatment) at elapsed durations t_0 , compared to when the assigned moment of exposure takes place at a higher duration t_1 . It follows that these measures are identified without any functional-form restriction on the individual hazard rates and without the need to assume independence of the unobserved explanatory variables V from the observed covariates X . From the above proof it is also clear that the results extend to settings where X and/or V are not constant over time, provided that Assumptions 1 and 2 about the assignment process and the absence of anticipation are accordingly modified.

Figure 3 visualizes the underlying idea of the proposition. In each cohort, the dynamic selection between the moment of entry and the elapsed duration t_0 proceeds identically. Therefore, the resulting sub-populations at t_0 have an identical distribution of unobserved characteristics. As a result, any observed difference in the hazard rates at the elapsed duration t_0 must be a causal effect of the policy change.




Figure 3: *Identification based on two cohorts*

Since $ATTS(t_0, t_1, t_0|X)$ and $ATNTS(t_0, t_1, t_0|X)$ are equal and do not depend on t_1 as long as $t_1 > t_0$, we use the short-hand notation $ATS(t_0|X)$ to represent the average instantaneous effect of the policy reform on the survivors with elapsed duration t_0 at the moment of the reform. The effect is measured in deviation from the hazard rate among sub-populations who attained the elapsed duration t_0 strictly before the reform.

The sub-population over which the average $ATS(t_0|X)$ is taken varies with t_0 . This is because the composition of the sub-population changes due to dynamic selection as the elapsed duration t_0 increases. As a result, without further assumptions, it is not possible to combine the average treatment effects for different t_0 in order to estimate how the average effect on the hazard changes over time for a given (sub-)population. Dynamic matching estimators share this limitation (see Crépon et al., 2009).

Under Assumptions 1 and 2, average treatment effects on conditional survival *probabilities* are non-parametrically identified as well. In this case, average effects on treated survivors are defined as follows,

$$\text{ATTS}(s', s, t|X) := \mathbb{E} \left[\begin{array}{l} \Pr(T(s') > t + a | T(s') \geq t, X, V) - \\ \Pr(T(s) > t + a | T(s) \geq t, X, V) \mid X, T(s') \geq t \end{array} \right] \\ \text{with } s' \leq s \text{ and } a > 0.$$

These are identified from their empirical counterpart if $t \leq s'$. For example, take $t = s' = t_0$ and $a = 1$ and $s > t_0 + 1$. The average effect of exposure at t_0 on the probability of exiting before $t_0 + 1$, as compared to when the exposure commences after $t_0 + 1$, equals the observable $\Pr(T > t_0 + 1 | T \geq t_0, X, S = t_0) - \Pr(T > t_0 + 1 | T \geq t_0, X, S = t_0 + 2)$, where instead of $t_0 + 2$ any other number exceeding $t_0 + 1$ can be substituted. Indeed, the observable expression can be replaced by $\Pr(T > t_0 + 1 | T \geq t_0, X, S = t_0) - \Pr(T > t_0 + 1 | T \geq t_0, X, S \geq t_0 + 1)$. Clearly, such results carry over to discrete-time settings (see below).

In Appendix 1, we consider identification of average *multiplicative* effects on individual hazard rates. This requires the additional assumption that the unobserved individual characteristics V affect all counterfactual hazard rates in the same proportional way. In other words, the individual multiplicative effects on the hazard at t are homogeneous across individuals with different V (but not necessarily across X or over time; furthermore, X and V need not be independent). The identification results in the appendix are related to identification results for duration models with unobserved heterogeneity and time-varying explanatory variables in Honoré (1991) and Brinch (2007).

We end this subsection with a brief discussion of the identification of other interesting average treatment effects. Clearly, one cannot hope to identify a full model, that is, the unknown functions $\theta_{T(s)}(t|X, V)$ for all s and the distribution of $V|X$. Now consider average treatment effects on the individual hazard rate $\text{ATTS}(s', s, t|X)$ and $\text{ATNTS}(s', s, t|X)$ if s' is strictly smaller than t and s . In such cases, inference is subject to the same problem as in Subsection 3.2: the dynamic selection between s' and t causes the sub-population with $S = s'$ among the survivors at t to be systematically different from the sub-population with $S = s$ among the survivors at t . This also implies that without functional form assumptions we can not identify accumulation effects of a prolonged exposure to the treatment.

4 Non-parametric estimation

4.1 Boundary kernel hazard estimation

From Subsection 3.3, the identification of average causal effects of the policy change on the individual hazard rates is based on the comparison of observable hazard rates from different entry cohorts into the state of interest. Each observable hazard rate is trivially identified from the corresponding cohort-specific duration data. It is therefore natural to non-parametrically estimate these hazard rates.

Specifically, we are interested in $\theta_T(t_0|X, S = t_0)$ and $\theta_T(t_0|X, S = t_1)$ for some $t_1 > t_0$. In Subsection 4.2 below we consider alternative estimators based on $\lim_{t \uparrow t_0} \theta_T(t|X, S = t_0)$ and $\theta_T(t_0|X, S \geq t_1)$ for some $t_1 > t_0$. In every case, the relevant estimate concerns the hazard at the boundary t_0 . There is no reason to assume a connection between the shape of the individual hazard rate before the policy change at t_0 and the shape after t_0 , so estimation of the hazard rate at one side of the boundary only uses outcomes from that particular side of the boundary. Standard non-parametric hazard estimators are heavily biased at the boundary point. We therefore apply methods that deal with this. Specifically, we use boundary kernel hazard estimators and local linear kernel smoothing estimators.²⁰ See Andersen et al. (1993) for an introduction to these estimators.

In the remainder of this subsection we discuss the second-order boundary kernel hazard estimator of Müller and Wang (1994) in some detail. We use this estimator in the empirical analysis in Section 5. For expositional convenience we restrict attention to hazard estimation at t_0 , and we transform the truncated duration distribution $T|T \geq t_0, X, S$ to the left such that our ultimate interest is in the hazard rate at the boundary 0 when evaluating it from above. Similarly, in the current subsection, we may suppress S in the notation. Since in the empirical analysis we subsume X into V , we do not consider observed explanatory variables X in the current subsection either.²¹

Consider a random sample of n subjects, where the duration outcomes can be independently right-censored. Let T_i denote the minimum of the actual duration outcome and the

²⁰Most of the literature on the non-parametric estimation of hazard rates imposes strong smoothness conditions on the true underlying hazard rate as a function of t and the explanatory variables (in our case, S and X), and the explanatory variables are often assumed to be continuous. In cases where smoothness is absent at a boundary of the support, the hazard rate is often only evaluated at interior points.

²¹If X is exogenously time-varying on $(0, t_0)$ in a specific way across cohorts (e.g., as in a linear time trend), then this may cause the “treatment” and “control” hazard rates to have common determinants, but we do not pursue this here.

censoring outcome for subject i ($i = 1, \dots, n$). Note that this notation deviates from the notation where T denotes the actual duration outcome of interest. Furthermore, let δ_i be a binary variable equalling 1 iff the actual duration outcome is realized before the censoring outcome. Let $(T_{(i)}, \delta_{(i)})$ be the ordered sample with respect to the T_i (so $T_{(1)} \leq T_{(2)} \leq \dots \leq T_{(n)}$).

We assume that the true hazard rate is twice continuously differentiable in an interval A starting at 0. To explain the kernel estimator, consider first the case in which the bandwidth b is global. We distinguish between the boundary region $B = \{t : 0 \leq t < b\}$ and an interior region I which is adjacent to B (we need not discuss the right boundary of A here). In I , the kernel hazard estimator is the standard Ramlau-Hansen kernel hazard estimator,²²

$$\tilde{\theta}(t) = \frac{1}{b} \sum_{i=1}^n K\left(\frac{t - T_{(i)}}{b}\right) \frac{\delta_{(i)}}{n - i + 1}$$

where K is taken to be the Epanechnikov kernel,

$$K(z) = \frac{3}{4}(1 - z^2) \quad \text{for } |z| \leq 1 \quad (6)$$

and $K(z) = 0$ elsewhere, and where b is understood to decrease with n , as explained below.

In B , the above estimator needs to be modified to take account of the bias at the boundary. After all, with the above estimator, it is typically of asymptotic order $O(b)$. In B , the kernel function K is taken to depend on the distance to the left boundary 0, so then K has two arguments, say q and z , where q is the relative distance t/b to the left boundary, and z , as above, attains values $(t - T_{(i)})/b$. Specifically,

$$K(q, z) = \frac{12}{(1 + q)^4} (z + 1) [z(1 - 2q) + (3q^2 - 2q + 1)/2]$$

where $q \in [0, 1]$ and $z \in [-1, q]$. The latter implies that the support of the boundary kernel does not extend beyond the left boundary. Müller and Wang (1994) plot $K(q, z)$ as a function of z for various values of q . As expected $K(1, z)$ is again the Epanechnikov kernel. As q decreases, the kernel becomes more and more skewed, and the weight assigned to values close to the boundary increases strongly. There is a positive probability that the resulting $\tilde{\theta}(0)$ is negative, in which case it is replaced by zero.

²²This smoothes the increments of the Nelson-Aalen estimator $\Lambda_n(t)$ of the integrated hazard based on a random sample of n subjects,

$$\Lambda_n(t) = \sum_{i: T_{(i)} \leq t} \frac{\delta_{(i)}}{n - i + 1}$$

The boundary correction establishes a reduction of the bias. At the same time, the variance of the estimator increases, because the number of observations used to estimate the hazard close to 0 becomes smaller. A further variance reduction can be achieved by choosing a larger bandwidth close to 0 than elsewhere. Müller and Wang (1994) therefore propose to use local bandwidths $b(t)$. In that case, b is replaced by $b(t)$. As functions of n , the local bandwidths $b_n(t)$ are assumed to satisfy the usual conditions (somewhat loosely, $b_n(t) \rightarrow 0, nb_n(t) \rightarrow \infty$). Optimal local bandwidths are such that $nb_n^5(t)$ converges to a number smaller than infinity, so $b_n(t) \sim n^{-\frac{1}{5}}$. The asymptotic behavior of the estimator is not fundamentally different from usual. The convergence rate is $n^{-\frac{2}{5}}$. Optimal global or local bandwidths can be consistently estimated by a data-adaptive procedure, along with the estimates of interest (see Müller and Wang, 1994). In Appendix 2 we present the algorithm, slightly modified in response to our experiences regarding the performance of the estimator.

Asymptotic normality allows for the estimation of a confidence interval for $\theta(0)$. Following the line of reasoning in e.g. Härdle (1994) and Härdle et al. (2004), one could ignore the asymptotic bias term to obtain an approximate 95% confidence interval (see Müller et al., 2004, for an application of the idea of omitting the asymptotic bias in the related case of boundary kernel density estimation). Conceptually, it is not difficult to include the asymptotic bias term in the confidence interval, but in practice this involves non-parametric estimation of the second derivative of the hazard at 0. An alternative that we follow in the empirical application below is to use bootstrapping to obtain confidence intervals.²³

Müller and Wang (1994), Hess et al. (1999) and Jiang and Doksum (2003) provide Monte Carlo simulation results for the estimator. They conclude that it has an excellent performance in samples sizes n as small as 50 to 250. Hess et al. (1999) compare the performance to that of other kernel estimators. They show that the other estimators perform worse, in particular at the left boundary, and they demonstrate that both the boundary correction and the data-adaptive local bandwidth are important in this respect.

Instead of boundary kernel approaches, one may use local linear smoothing (or local linear fitting, or locally weighted least squares) as a non-parametric approach to deal with estimation at a boundary. Wang (2005) gives an intuitive overview of local linear hazard rate estimation, while Nielsen and Tanggaard (2001), Jiang and Doksum (2003) and Bagkavos and Patil (2008)

²³This is proposed by Müller et al. (2004). See e.g. John and Jawad (2010) for an empirical application of non-parametric bootstrapping of standard errors of boundary kernel hazard estimates. Note that despite the local nature of our estimator, the size of treated and control cohorts increases with sample size and the estimator is linear. In such circumstances, bootstrapping is expected to provide correct standard errors.

provide details.²⁴ The asymptotic properties of the estimator are qualitatively identical to those of the boundary kernel hazard estimator. Jiang and Doksum (2003) compare both methods with data-adaptive local bandwidths, in some Monte-Carlo simulation experiments. Both methods give similar results and both perform very well at the boundary, where their relative ranking depends on the shape of the true hazard rate.

The results of this subsection can be straightforwardly applied for inference on the difference of two independently estimated hazard rates. Appendix 1 discusses inference of the ratio of two independently estimated hazard rates.

4.2 Implementation issues

We consider a number of dimensions in which the econometric inference can be improved or modified.

(i) The “comparison” cohort(s). In Subsection 4.1, we used a boundary-corrected estimator for the observed hazard $\theta_T(t_0|X, S = t_1)$ at t_0 in the cohort that is eventually exposed to the reform at some elapsed duration $t_1 > t_0$. Instead, one may use a standard kernel (or local linear or local constant) hazard estimator, if one is prepared to assume that this hazard is smooth in an interval around t_0 , since then the estimation concerns the interior of an interval on which the hazard is smooth. Whether this assumption makes sense depends on the setting at hand. At certain elapsed durations t_0 of interest, the eligibility to other policy measures may change, causing the individual hazard rates $\theta_{T(s)}(t|X, V)$ to be discontinuous at $t = t_0$ for all s . The application in Section 5 is a case in point. To rule out that this affects the estimated effects, one needs to resort to boundary correction methods.

Analogously, one may examine the left-hand limit of the observed $\theta_T(t|X, S = t_0)$ at $t = t_0$ in order to estimate the “control” hazard, but this also requires the assumption that there are no other sources of discontinuities at t_0 .

Note that one may widen the “control group” and increase the precision of the estimates of interest, by estimating $\theta_T(t_0|X, t_2 > S \geq t_1)$ with $t_0 < t_1 < t_2 \leq \infty$, instead of $\theta_T(t_0|X, S = t_1)$. This does come at a price, namely that Assumption 1, ruling out the absence of cohort effects, needs to be extended to multiple comparison cohorts flowing in at or before $\tau^* - t_1$. Recall that we require unobserved cohort effects to be absent, since otherwise $S \perp\!\!\!\perp V|X$ so that Assumption 1 is violated. Observable cohort indicators may be included in X , but note that

²⁴Local linear estimation of hazard rates is related to fixed design non-parametric regression.

in non-parametric analysis any addition to X adds to the curse of dimensionality.²⁵

Instead of enlarging the “control group”, one may use the availability of multiple potential comparison cohorts in order to select the most similar cohort (or set of cohorts) among the cohorts flowing in before $\tau^* - t_0$. We do not observe the distribution of $V|X$ in a cohort, but we observe outcomes that are informative on it, namely the duration distribution on the duration interval $[0, \tau^*)$ in the corresponding cohort. As a selection mechanism, one may match on the survival probability in the cohort at duration τ^* , or, even stronger, on the shape of the duration distribution in the cohort on the duration interval $[0, \tau^*)$. The more similar this shape, the more similar the composition of survivors at the duration τ^* . At the extreme, if alternative cohorts are equally adequate, this provides overidentifying information that can be used to test the validity of our framework.

If one comparison cohort is to be selected, then one may consider a cohort that flowed in only marginally earlier than the “treated” cohort, following the line of thought that unobserved changes of the entry composition of the cohorts are a smooth function of the moment of entry. However, such a choice of t_1 being almost equal to t_0 has a practical disadvantage. To see this, notice that $\theta_T(t|X, S = t_1)$ may display a discontinuity at t_1 , so the value $\theta_T(t_0|X, S = t_1)$ at the elapsed duration $t_0 < t_1$ can only be estimated from observed realized durations in an interval to the right of t_0 that does not stretch beyond t_1 . Spells in the comparison cohort with durations exceeding t_1 should be treated as right-censored at t_1 . Consequently, the measure of realized duration outcomes that is informative on $\theta_T(t_0|X, S = t_1)$ is very small if t_1 is only marginally larger than t_0 .

(ii) Observed covariates. Including many elements in X raises a curse of dimensionality in the non-parametric estimation. One may therefore choose to treat the observed covariates X as unobservables and hence subsume them into V . This involves a strengthening of Assumption 1, in the sense that it requires $S \perp\!\!\!\perp X$. The latter can be empirically verified by examining the composition of the cohorts used to estimate the objects of interest. If $S \perp\!\!\!\perp X$ is satisfied then treating X as unobservables in the estimation of the objects of interest does not involve a strengthening of Assumption 1. In practice one may therefore verify that $S \perp\!\!\!\perp X$ and, if this holds, proceed by ignoring X in the duration analysis.²⁶

With discrete X , non-parametric inference would typically lead to separate estimations for

²⁵It may be an interesting topic for further research to examine under which conditions the presence of multiple comparison cohorts enables the identification of anticipation effects that violate Assumption 2.

²⁶Such a pre-test affects the precision of the inference on the effect of interest.

each value of X . This would also allow for the selection of the most similar control cohort for each value of X separately. To aggregate the estimated average effects over X , one may average the estimated effects given X over the relevant distribution of X .²⁷

Alternatively, one can use standard matching techniques to balance observed variables across the “treated” and control cohorts. For high-dimensional problems, or when some of the matching variables are continuous, semi-parametric methods such as propensity score matching can be used to limit the dimensionality of the problem.

(iii) Discrete time. Now let us reconsider the continuous nature of the duration variable. Sometimes a continuous-time analysis may be unfeasible. For example, the data may be time-aggregated in the sense that events are recorded in time intervals (e.g. unemployment duration is collected in months even though individuals may enter employment on any given workday). Alternatively, duration outcomes may be discrete due to institutional constraints (e.g. in certain occupations a job can only start on the first day of a month).

Accordingly, we distinguish between two frameworks. In one, the model is in continuous-time and the duration outcomes are in discrete time. In the other, both are in discrete time. In the first framework, the results of Section 3 apply but we cannot estimate hazard rates. However, we can estimate conditional survival probabilities and their differences, as outlined in Section 3. In general, results obtained in this framework can be viewed as approximations of those for hazard rates obtained in a genuine continuous-time framework. Because of the ease with which survival probability outcomes can be estimated, this approach may be useful from a practical point of view. However, this may come at the cost of ignoring short-term fluctuations of duration determinants such as seasonal effects. As for the second framework, the analysis of Section 3 is straightforwardly modified to such settings by working with a genuine discrete-time framework. We examine this empirically in Section 5 below.

(iv) Reduced-form model estimation. Empirical duration analysis often specifies parameterized models to estimate the objects of interest. In a setting with a compulsory reform affecting ongoing spells, an obvious choice would be to estimate a Proportional Hazard (PH) model for the distribution of $T|X, S$, using all available cohorts, with S represented by way of a simple time-varying covariate $I(t \geq \tau^* - \tau_0)$,

$$\theta_T(t|X, \tau_0) = \lambda(t) \exp(X'\beta + \alpha I(t \geq \tau^* - \tau_0))$$

²⁷With time-varying X , the approaches of this Subsection 4.2.(ii) also apply.

where α is the parameter of interest. Such a model does not explicitly include unobserved heterogeneity. However, one may simply interpret it as a parsimonious description of patterns in the data averaged over unobservables, i.e. as a simple representation of a distribution that is generated by an underlying model for individual hazard rates with unobserved heterogeneity. In that case, the question is whether the parameter α still captures a causal effect. It follows from our results that to some extent it does. Specifically, if we abstract from the restrictiveness of the PH model structure and potential misspecification, then α in this model captures the average causal policy effect at $\tau^* - \tau_0$ on the hazard rate. So, the averaged policy effect is estimated correctly when using a model that ignores unobserved heterogeneity even if in reality there is unobserved heterogeneity. It follows that the estimate obtained under the assumption of no unobserved heterogeneity is in fact controlling for unobserved heterogeneity. This is in marked contrast to single-spell duration analysis with time-invariant covariates, where ignoring unobserved heterogeneity typically leads to biased covariate effect estimates (see Van den Berg, 2001, for an overview). Obviously, it is essential that the analysis uses data that include spells that are ongoing at the moment of the policy change. Note that the model does not require continuity of the hazard function in t .

Hall and Hartman (2010) provide an example of a study in which a PH model is estimated using spells interrupted by a policy change. Specifically, they estimate a PH model for the transition rate from unemployment into sickness absence as a function of the sickness benefit policy regime, using unemployment spells that cover a date at which a policy regime change was implemented. They find that a reduced cap for sickness benefits lowers the transition rate to sickness absence by about 35% in the treated population. In their study, they also estimate MPH-type model extensions that allow for unobserved heterogeneity as proportional fixed effects in the individual hazard rates, exploiting the fact that the data contain multiple unemployment spells for many subjects. Interestingly, they find that the estimated policy effect is virtually identical to that in their main analysis, suggesting that, indeed, this coefficient is estimated correctly even when ignoring unobserved heterogeneity.

(v) Dynamic treatment evaluation. The results of Section 3 can be applied to dynamic treatment evaluation settings. In such settings, the exposure to a treatment is not necessarily due to some institutional change at a fixed point in time. Rather, different individuals in the same cohort are exposed to a treatment at different elapsed durations, where the treatment may affect the individual hazard from the moment of exposure onwards. Typically, S is only observed if $S \leq T$. It may be interesting to reassess Assumptions 1 and 2 in such settings.

Abbring and Van den Berg (2003) demonstrate that Assumptions 1 and 2 are fundamental in the following sense: any causal model is observationally equivalent to a model in which these two assumptions are satisfied. If one of the assumptions is relaxed then point-identification requires additional structure.

5 Empirical application

5.1 The New Deal for Young People (NDYP): policy regime and treatment

The NDYP was a welfare-to-work program introduced in the UK in the late 1990s. It targeted the young unemployed, aged 18 to 24, who have claimed unemployment benefits (UB, known as Job Seekers' Allowance in the UK) for at least 6 months. Participation was compulsory upon reaching 6 months in the claimant count, and refusal to participate could be punished with a temporary benefits withdrawal. Since entitlement to UB in the UK is neither time-limited nor dependent on past working history, and eligibility is constrained only by a means-test that most young people pass, the NDYP was effectively targeted at all young long-term unemployed (long-term being defined as 6 months or more). Thus, and for simplicity, we use "unemployed" to signify those in the UB claiming count.

After enrollment,²⁸ the treatment was split into three stages: a first period of up to 4 months of intensive job search assistance called the Gateway, with fortnightly meetings with a personal adviser; next, for those unable to find work in that period, a period with four alternative programs including subsidized employment, full-time education or training, work experience in the voluntary sector and work experience in an environment-focused organization;²⁹ and a final period of intensive job search assistance called the Follow Through.³⁰

The NDYP treated millions of people before being replaced by another program in 2009, the Flexible New Deal. For instance, 172,000 new participants entered the NDYP in 2006 alone, and the average number of participants at any month during that year was 93,000. The

²⁸Throughout the section we use "enrollment" to denote actual mandatory participation in the job search assistance program and subsequent programs. As we shall see, actual participation may start strictly later than the moment at which the NDYP policy was introduced.

²⁹These options would last for up to 6 months (or 12 months in the case of education).

³⁰Repeated participation in the four alternative treatments could be arranged if perceived beneficial. More details on the program can be found in White and Knight (2002), Podivinsky and McVicar (2002), Blundell et al. (2004), Van Reenen (2004), or Dorsett (2006).

per-year expenditure of the NDYP during the 2000s was in the order of GBP 200 million, excluding administrative costs (DWP (Department for Work and Pensions), 2006), but a large proportion of this cost is UB payments that would be due independently of the program.

5.2 The introduction of the policy

The NDYP was released nationwide on April 1, 1998; this is calendar time τ^* . The existing stock of those who were unemployed for at least 6 months at τ^* was gradually moved into the program. At τ^* , only those whose elapsed unemployment duration was an integer multiple of 6 months were enrolled. Enrollment took place at the job-focused interviews the unemployed individuals attend every 6 months, provided they are aged 25 at the first interview after the NDYP was implemented. If the elapsed duration was not an integer multiple of 6 months, then the individual was enrolled at the moment that his or her elapsed duration attained an integer multiple of 6 months, provided that he or she was not yet 25 years old at that point in time. In the empirical analysis we do not exploit the age eligibility criterion except for robustness checks.

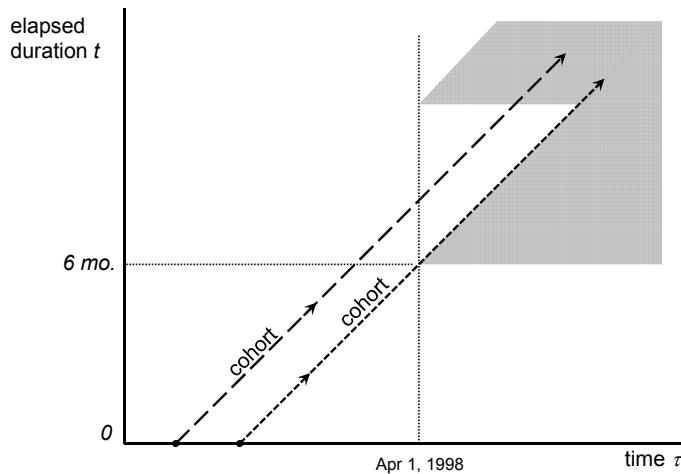


Figure 4: Introduction of the NDYP

Notes: diagonal lines depict cohorts. Each cohort moves along its own line as t and τ increase in tune. In the grey area, individuals are enrolled in the job search assistance program.

Figure 4 illustrates the enrollment scheme in the years around τ^* . This scheme is somewhat

more complicated than that described in Subsection 3.3, but it supports the identification and non-parametric estimation of average causal effects of the arrival of information about the new policy for all elapsed durations $t \geq 0$. For elapsed durations that are integer multiples of 6 months, this is the effect of enrolling in the program; for other elapsed durations, it is the effect of anticipating potential participation at the time of the next interview. To avoid misunderstandings, note that the latter is not in conflict with Assumption 2, which is still required. Specifically, individuals are not allowed to anticipate the moment at which the information arrives, which is the time of the reform, τ^* .³¹ Notice that it is not possible to identify effects of the actual participation in the NDYP at 6 months among those who are unemployed for less than 6 months at τ^* , since such individuals may act upon the information about the future treatment throughout the time interval between τ^* and the moment that $t = 6$.

Two minor institutional features constitute deviations from the above description of the introduction of the policy. First, individuals with $t < 6$ at τ^* can try to apply for early enrollment, especially if they are disadvantaged (disabled, former convict or lacking basic skills). Such applications were rare at the initial stages of the program. Secondly, the NDYP was introduced in a few small pilot areas on January 1, 1998, i.e., three months before the national rollout. We use the data from these areas and shift calendar time with 3 months when combining these data with the data from the rest of the country. Since the pilot study did not receive massive attention before April 1, and the evaluation of the pilot was not completed on April 1, we feel that the risk of information spillovers is small.

5.3 Data

The empirical analysis is based on the JUVOS longitudinal dataset, a 5% random sample of the register data of all UB claiming spells. JUVOS records the entire claiming histories of sampled individuals since 1982. Information includes the start and ending dates of each claiming spell as well as the destination upon leaving (since 1996, but plagued with missing values), and a small number of demographic variables such as age, gender, marital status, geographic location, previous occupation and sought occupation. JUVOS contains no information about what happens while off benefits and the information on destination upon leaving is plagued with missing values (over 30% of spells). Therefore, we will focus on “all exits from the claimant count” as the outcome of interest, irrespective of destination.

³¹Blundell et al. (2004) study anticipation of the reform (and hence of the moment of the corresponding information arrival) on April 1, 1998, by exploiting spatial and age discontinuities. No significant anticipatory effects are found.

The estimation sample is formed of men who were aged between 20 to 24 upon reaching 6 months in unemployment. We discard observations for younger individuals to avoid having to deal with education decisions.

5.4 The choice of treatment and comparison groups

We estimate the $ATS(t_0)$ for t_0 equal to 6 months (precisely 182 days). The estimation relies on comparing the survivors among the cohort attaining the elapsed duration t_0 at τ^* (which we call the treatment group, or the treated) with a similar sample of survivors from earlier cohorts (the comparison groups). Following the discussion in Subsection 4.2, we do not condition on covariates X .

The continuous-time framework must be reconciled with the requirement of a positive sample size. In practice, we need samples of cohorts flowing into unemployment within time intervals rather than at two singular points in time. To proceed we consider full monthly cohorts. For instance, the treated sample includes all spells starting in October 1997 (or July 1997 in pilot areas), lasting for at least 6 months. This choice is not innocuous as discussed in Subsection 5.2: those starting a spell towards the end of October 1997 will have been exposed to the new policy for a few weeks before enrolling, and this may lead to biased inference. The bias should be negligible if the anticipatory effects within weeks of the reform are much smaller than the impact of participation. We argue below that the distortion may lead to an under-estimation of $ATS(t_0)$ at 6 months.³²

We define comparison groups in an analogous way, selecting individuals reaching 182 days in unemployment over an entire calendar month prior to April 1998. As candidate groups we consider the cohorts flowing in during June 1997 (pilot areas) and September 1997 (non-pilot areas), or May 1997 and August 1997, or July 1996 and October 1996, or the combination of June and September 1997 with July and October 1996. We include data on both pilot and non-pilot regions in all that follows but designate each cohort by the month of inflow in non-pilot areas for simplicity. Following the discussion in Subsection 4.2, different candidate groups are assessed based on two outcomes: the distribution of T on days 1 to 181, and the distribution of observed characteristics X among survivors at 182 days.

Figure 5 shows the survival functions for the treatment and comparison groups in the pre-treatment period, up to 181 days into unemployment. For the combined cohort, the matching

³²Clearly, it is preferable to apply an estimator in which the inflow time interval shrinks as the sample size increases, such that observations from cohorts close to the inflow date of interest are given more weight. Given our modest sample sizes we do not pursue such an approach, and we leave this as a topic for further research.

is so close that the curve is hardly distinguishable from the curve for the treatment group. The survival function for the September 1997 cohort diverges from that for the treatment group during the December/January period but quickly returns to match it over the final 2 months of the interval. For our purposes, what matters is whether treatment and comparison groups are similar at the time of enrollment. We cannot reject such hypothesis for the September 1997 cohort. The August 1997 cohort curve also converges towards the treatment cohort curve in the last month before enrollment, but the match is not as close as for the September 1997 cohort. The exception to this pattern is the October 1996 cohort. The survival function for this cohort is systematically above that for the treatment group over the entire interval, suggesting that aggregate conditions in the market changed during the year.




Figure 5: *Treatment vs. comparison groups – empirical survival functions between 0 and 6 months after inflow*

Further evidence of the similarities between the “treated” and control cohorts can be gathered from contrasting their hazard rates in the pre-treatment period. Figure 6 does so for durations 90 to 181 days. The differences between the hazard rates are small and mostly not statistically significant at conventional levels. In just one case, when contrasting the treated with the August 1997 cohort, is there evidence of seasonal effects differentially affecting the hazard rates, but this effect is mild even in this case. We conclude that the evidence on dynamic selection supports the use of two cohorts, September 1997 and October 1996, for constructing the counterfactual.

Table 1 compares the empirical distributions of observed covariates among the survivors

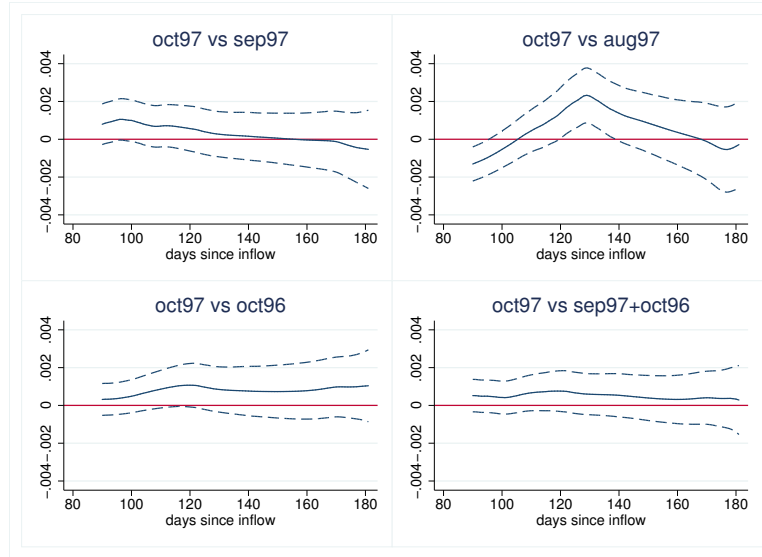


Figure 6: *Treatment vs. comparison groups – differences in the hazard rates before the policy reform, at durations 90 to 181 days*

Note: Dashed lines represent 95% confidence intervals.

in the treatment and comparison groups. The September 1997 cohort displays no discernible differences to the treatment group (column 1). The combined cohort does not perform as well, with systematic differences in the history of unemployment up to three years prior to inflow (column 4). Based on these results, we select the September 1997 cohort as the comparison group and discard the alternative candidates. It turns out, however, that our results are robust to the choice of the comparison group (estimates available upon request). In total, the sample size of individuals completing 182 days in unemployment during March and April 1998 while aged 20 to 24 is 902. This is almost equally split between the treatment (April 1998) and comparison (March 1998) groups.

Individuals in the September 1997 comparison cohort only become entitled to treatment at 12 months as they are past the 6 months threshold at the time of the reform. The arrival of information about the new program may affect their behavior and confound estimates of the ATS at 6 months. This source of bias can be eliminated by right-censoring spells in the comparison group when they cross the time of the reform, τ^* . However, any anticipating response is likely to be negligible as eligibility is far into the future. We checked the sensitivity of the results, and they are robust to right-censoring (estimates available upon request).

Table 1: Treatment vs. comparison groups – p -values for Hotelling statistics comparing the distribution of covariates conditional on survival up to 181 days of unemployment.

	Comparison cohort			
	September 97 (1)	August 97 (2)	October 96 (3)	Sep97 + Oct96 (4)
Marital status	0.997	0.643	0.114	0.509
Age	0.307	0.299	0.916	0.942
Region	0.276	0.095	0.112	0.083
Occupation	0.767	0.575	0.302	0.532
Time U in the past	0.363	0.846	0.021	0.046
U spells in the past	0.801	0.454	0.000	0.006
Zero U spells in the past	0.353	0.747	0.020	0.164
All covariates	0.761	0.517	0.001	0.085
No. observations	456	368	557	1013

Notes: The treatment group is the October 1997 cohort. The variables in rows 5 to 7 describe the UB claiming history in the 3 years preceding inflow into current unemployment spells. Numbers in bold highlight statistically significant differences in the distribution of the covariate, at the 5% level.

We also estimate anticipatory effects induced by the arrival of information about the reform among those approaching enrollment. These are interesting per se, and are also informative about the accuracy of estimates of the impact of program participation that ignore anticipation, by exposing the extent to which anticipation alters the composition of the treatment group prior to participation. The population we consider is again that of individuals who will be aged 20 to 24 when reaching 6 months of unemployment. We may estimate the anticipatory effects of future enrollment evaluated at each duration t_0 shorter than 6 months (182 days). Now, “treatment” means exposure to information about future enrollment while “comparison” means the absence of such exposure. Thus, the treatment and comparison groups are now defined in reference to whether they are exposed to the information arrival at the reform date. For a given $t_0 < 6$ months, the treatment group now consists of individuals who reach t_0 during April 1998. In accordance to the earlier procedure to choose comparison groups, we select individuals who flowed in one month before $\tau^* - t_0$ and who hence reach the duration $t_0 < 6$ during March 1998.

Figure 7 displays the survival functions up to t_0 for the treatment and comparison groups for t_0 equal to 2, 3, 4 and 5 months. There are some signs of seasonal differential selection during December/January. The survival functions for later cohorts, crossing December/January earlier in their spells (panels B and C), diverge throughout the duration interval $(0, t_0)$, especially at the end of the period, when approaching April (treatment) or March (comparison). Post December/January cohorts (panel A) are unaffected by conditions in those months and exhibit very similar survival functions. Earlier cohorts (panel D) are also affected but return quickly to a common path. The latter finding echoes the observed patterns for the October and September cohorts in Figure 5.

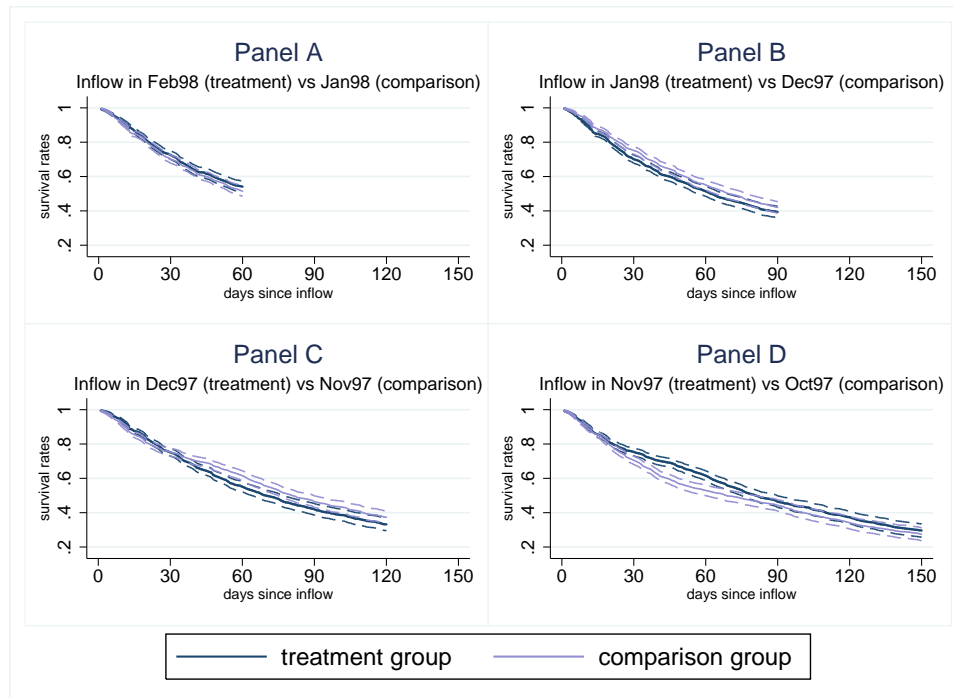


Figure 7: *Empirical survival functions for cohorts reaching durations of 2 (panel A) to 5 (panel D) months in April 1998 (treatment group) or March 1998 (comparison group).*

Note: Dashed lines represent 95% confidence intervals. Treatment and comparison here refer to exposure to information about future enrollment.

Table 2 compares the empirical distributions of observed covariates among the survivors in the treatment and comparison groups for different t_0 . Column 2 shows that the December 1997 and January 1998 cohorts are compositionally different upon having reached 3 months in unemployment. For earlier cohorts, the absence of statistically significant differences further

supports their comparability (columns 3 and 4). In the light of these findings, our analysis of anticipatory effects focuses on durations from 4 to 5 months. These are the most relevant comparisons as anticipatory effects will be larger at dates closer to enrollment in the NDYP.

Table 2: Treatment vs. comparison groups – p -values for Hotelling statistics comparing the distribution of covariates conditional on survival up to 2 to 5 months of unemployment

Treatment group	Month of inflow			
	Feb 98	Jan 98	Dec 97	Nov 97
Comparison group	Jan 98	Dec 97	Nov 97	Oct 97
Elapsed duration t_0	2 months	3 months	4 months	5 months
	(1)	(2)	(3)	(4)
Marital status	0.471	0.339	0.790	0.656
Age	0.120	0.263	0.366	0.318
Region	0.425	0.304	0.671	0.858
Occupation	0.338	0.234	0.410	0.603
Time U in the past	0.188	0.015	0.439	0.921
U spells in the past	0.303	0.021	0.242	0.387
Zero U spells in the past	0.626	0.167	0.271	0.589
All covariates	0.222	0.037	0.599	0.893

Notes: Row 1 (2) details the inflow date of the treatment (comparison) group for the evaluation of the effect at the elapsed duration in row 3. Treatment and comparison here refer to exposure to information about future enrollment. The variables in rows 8 to 10 describe the UB claiming history in the 3 years preceding inflow into current unemployment spells. Numbers in bold highlight statistically significant differences in the distribution of the covariate, at the 5% level.

In contrast to our earlier discussion on the estimation of enrollment effects, right-censoring at times shortly after calendar time τ^* is expected to be important here, as the comparison group will itself be subject to the information arrival on April 1st, 1998. Moreover, the treatment group will enroll into job search assistance upon 6 months of unemployment, with, potentially, causal effects on their hazard rate from that moment onwards. We examine these two issues below.

5.5 The effect of enrollment in the NDYP

We estimate effects in discrete and in continuous time. Estimates in discrete time capture effects on aggregate monthly conditional transition probabilities while estimates in continuous time do the same for daily hazard rates. Both sets of estimates are based on the same treated and comparison samples.

Table 3 presents the main results for the discrete-time setting. The figure in column (1) is the estimated ATS at 6 months for the average causal effect of enrolling in the NDYP program on the conditional probability of leaving unemployment within one month of enrolment. We find a substantial impact of 4.5 percentage points (pp). This estimate is significantly positive at the 5% level. The corresponding relative increase in the conditional probability is about 35% and is in line with the results in Blundell et al. (2004) based on a difference-in-differences approach.

Columns (2) and (3) of the table display placebo effects. The figure in column (2) is the ATS estimate at 6 months for the non-existing reform of April 1997; the entry in column (3) is the ATS at 6 months for the older non-eligible group of 25-29 year olds. If our method is appropriate, these estimates should be statistically insignificant. Alternatively, they may reflect the presence of seasonal effects, macro effects or other differential changes affecting the composition of contrasting groups at inflow or their dynamic selection prior to reaching 6 months in the claimant count. As it turns out, neither of the two estimates is statistically significant. This result does not invalidate our approach.

Table 3: Non-parametric discrete-time estimation of the average causal effect $ATS(t_0)$ of enrolling into the NDYP program at the elapsed unemployment duration of $t_0 = 6$ months on the conditional probability of leaving unemployment within one month.

	Treatment effect	Placebo effects	
	20-24 years	20-24 years	25-29 years
age	April 1, 1998	April 1, 1997	April 1, 1998
τ^*	(1)	(2)	(3)
Estimate	.045	.014	-.009
Standard error	(.023)	(.022)	.021
No. individuals	911	1118	1365

Note: Estimates in bold are statistically significant at the 5% level.

Figure 8 displays the continuous-time counterparts of the ATS estimates in column (1) of the table. It shows both the additive and the multiplicative average effects, together with 95% confidence intervals based on the analytic asymptotic variance without bias correction.³³ Although $t_0 = 182$ days is the minimum elapsed unemployment duration for enrollment in the NDYP, it is conceivable that program participation requires a positive amount of time to act and exert any effect due to short delays in the timing of the 6-month interview (when the information about the treatment is shared with the treated). Therefore, we plot estimates at elapsed durations from 182 to 212 days. A zero effect in the early days after 182 suggests no differential dynamic selection between the treatment and comparison groups at this stage, allowing the ATS to be identified later on.




Figure 8: *Non-parametric continuous-time estimates of the average effect of enrolling in the NDYP at elapsed duration of 6 months on the hazard rate of leaving unemployment*

Note: Estimates using Müller and Wang method with optimal local bandwidths. Dashed lines represent 95% confidence intervals.

³³Pointwise confidence intervals, not uniform. With bootstrapping we obtain virtually the same intervals. The estimated optimal local bandwidth for the additive effect at the boundary of 182 days is 80 days with a standard error of 30 days.

The estimated patterns are similar for the additive and multiplicative average effects.³⁴ The focus should be on the impact at the first duration beyond 182 days for which the effect is significant. Any features after that may be due to duration dependence of the treatment effect or to differential dynamic selection, or both. We find significant effects of enrollment only after about a week into the program. At that time, the estimated effect as a function of the elapsed duration jumps rather abruptly to a positive level of about 0.006 per day. This amounts to more than doubling the hazard rate in the absence of NDYP, as can be seen from the figure for the multiplicative effect (right-hand-side panel of Figure 8). The estimated effect then drops to a lower positive level that just misses the 95% significance level, although at this stage we can no longer separate causal and confounding compositional effects. We conclude that among those who enter the new policy regime at 6 months of unemployment duration, the program has a significant and sizeable positive effect on the hazard rate at 6 months.

5.6 The effect of information about future enrollment

Table 4 presents the discrete-time estimates of the additive ATS at 4 months and 5 months (column (1)). Both are negative, but none is significantly different from zero, suggesting that individuals do not react in advance to the prospect of future participation in the NDYP. This suggests that estimates of the impact of participation in the NDYP using methodologies that ignore anticipation are not biased. Yet these are averages over potentially heterogeneous effects, so a zero average does not rule out that some individuals anticipate the job search assistance and, hence, that Assumption 2 is violated. Moreover, the finer continuous-time analysis may reveal other patterns.

The estimates in columns (2) and (3) are placebo effects similar to those in Table 3. They are statistically insignificant, except for 20–24 year olds 5 months after inflow in 1997, when no reform took place. The latter estimate is based on a comparison of those who flowed in in November 1996 to those who flowed in in October 1996. While we cannot exclude the presence of seasonality (or other) effects driving this estimate, it is telling that similar patterns cannot be found for the older group over the same period or for the same age-group over other adjacent months (September to October or from November to December).³⁵

The estimates in Table 4 may be biased if the one-month time interval used in estimating the conditional outflow probability for the comparison group crosses April 1, 1998, when

³⁴Inference for the multiplicative effect warrants an additional assumption (Assumption 3 in the Appendix 1) whereas inference on the additive effect does not.

³⁵Results available from the authors.

Table 4: Non-parametric discrete-time estimation of the average causal effect $ATS(t_0)$ of receiving information at elapsed durations $t_0 = 4$ or 5 months about enrollment at 6 months, on the conditional probability of leaving unemployment within one month.

	Treatment effect	Placebo effects	
	20-24 years	20-24 years	25-29 years
age	April 1, 1998	April 1, 1997	April 1, 1998
τ^*	(1)	(2)	(3)
4 months after inflow	-.015 (.021) 1328	.006 (.022) 1365	-.022 (.020) 1826
5 months after inflow	-.017 (.021) 1098	.057 (.021) 1228	.033 (.020) 1571

Notes: Estimates, standard errors and numbers of observations are in the first, second and third line, respectively. Estimates in bold are statistically significant at the 5% level.

information about the NDYP is released – a bias towards zero if the treated and comparison groups react similarly to the disclose of information. We therefore right-censor spells for the comparison group when they cross April 1, 1998. Likewise, treatment spells are right-censored at 6 months. We show in Subsection 5.5 that the causal effect of participation kicks in at an elapsed duration of 189 days, so we use this as the right-censoring value. We then use the resulting possibly right-censored data to estimate the effects of information arrival at elapsed durations 4 to 5 months in continuous time. Given the high data demands of this procedure, which requires a sufficiently large number of informative spells at each elapsed duration, we restrict the continuous-time analysis to durations 123 to 181 days.

Figure 9 shows the continuous-time estimates of the ATS for the arrival of information on future participation. Notice that the computational burden required to produce Figure 9 is much higher than for Figure 8. Figure 8 is based on the estimation of two non-parametric hazard rates, for two cohorts with one boundary at 182 days. In contrast, estimates in Figure 9 are for moving boundaries (the boundary coinciding with the elapsed duration in the x -axis), and, hence, moving cohorts. That is, Figure 9 represents the differences of two hazard rates at

the boundary t_0 , for $t_0 \in [123, 172]$.

The results in Figure 9 provide evidence of anticipatory behavior, showing a drop in the hazard rates starting after the beginning of the 5th month. Despite the wide 95% confidence intervals towards the end of the period (due to the bias corrections discussed above), the anticipatory effect is statistically significant during the second half of month 5.^{36,37} This evidence of anticipatory behavior is new, as previous NDYP results, like ours in Table 4, did not consider changes in behavior closely before the moment of enrollment.³⁸

The empirical finding on the anticipation of future enrollment in job search assistance has implications for standard evaluation approaches of participation in job search assistance. It is likely that those who postpone job search until after the enrollment into job search assistance at 6 months are on average more work-prone than those who remain unemployed for 6 months in a world without the program. In that case, a comparison of spells with elapsed durations of 6 months, before and after the introduction of the program (and censoring any spells crossing the reform date), would lead to an upward bias in the estimated effect of the job search assistance.

6 Conclusions

This paper merges regression discontinuity methods and duration analysis. We have shown that, in order to study causal policy effects on hazard rates, one may usefully exploit spells crossing the moment of the introduction of the policy, even if the individual hazard rates depend on unobserved covariates. The approach does not need any functional form assumption on the hazard rate or its determinants. This stands in marked contrast to standard duration analysis, which has been plagued by proportionality assumptions on the hazard rate, functional form assumptions on the duration dependence and the unobserved heterogeneity distribution, and a

³⁶Pointwise confidence intervals, not uniform.

³⁷The anticipatory effect was not detected in our discrete-time analysis with a monthly time unit in Table 4, although the signs of the estimates in Column 1 of Table 4 are consistent with the results in Figure 9. The lack of significance of the estimated anticipatory effect in the 5th month in Table 4 is due to time aggregation: it conflates a zero effects in the first half of the 5th month with a negative effect in the second half, hence diluting the over-all effect over the month and making it more difficult to detect.

³⁸Black et al. (2003) provides evidence that many unemployed workers in the U.S. dislike being an unemployment insurance claimant if it involves mandatory participation in programs of job search counselling. This does not seem to apply to the long-term young unemployed in the UK, approaching eligibility to a job-search assistance program. Indeed, it is conceivable that this is the result of a dynamic selection where those who find a job relatively fast anyway leave unemployment very fast and hence are underrepresented in the sub-populations over which our estimated ATS are defined.




Figure 9: *Non-parametric continuous-time estimates of the average effect of the arrival of information about future enrollment at 6 months on the hazard rate of leaving unemployment, by elapsed unemployment duration between 123 and 172 days*

Note: Dashed lines represent 95% confidence intervals.

“random effects” assumption for observed and unobserved covariates. An additional advantage of the new approach is that it enables policy evaluation shortly after the introduction of a new policy.

Our analysis shows that details of how a new policy for subjects in a certain state is implemented have implications for the quality and timing of evaluation studies regarding effects on the hazard rate out of the state. A reform that immediately applies to all subjects in the state of interest alleviates the need for strong identifying assumptions and supports the early production of evaluation results on the hazard rate. Conversely, a reform that applies only to new entrants into the state of interest will have to deal with differential dynamic selection and possibly with differential selection at inflow once the new regime is announced; it will also require waiting for at least t periods before the impact at duration t can be evaluated.

Our approach is also suitable to study the causal effect of the arrival of information on the hazard rate in a certain state. If the information captures the future moment at which the

subject will be exposed to a certain treatment then the approach provides estimates of the anticipatory effect of the treatment without having to rule out unobserved heterogeneity. In our empirical application, one of the effects we study concerns the causal effect of receiving information at elapsed unemployment durations below 6 months about an intensive job search assistance treatment at 6 months, on the hazard rate of leaving unemployment. Using fully non-parametric inference allowing for unobserved heterogeneity, we conclude that anticipatory effects on the hazard rate are present in the weeks before the onset of the treatment. In those weeks, individuals reduce their search effort.

Our study provides some suggestions and implications for existing methods of policy evaluation. First, consider semiparametric estimation of simple models for the observed hazard rate (i.e., without unobserved heterogeneity) in which exposure to the new policy is a time-varying covariate and in which the data include spells crossing the reform date. Such simple models may be regarded as a representation of the distribution of observables that is generated by underlying individual hazard rates with unobserved heterogeneity. In our non-parametric approach, observed hazards are informative on average policy effects on individual hazard rates, in the presence of unobserved heterogeneity, and without any identified model structure. This leads to the insight that the estimated policy exposure coefficient in a simple model can be informative on the causal policy effect. In this sense, estimates obtained under the assumption of no unobserved heterogeneity are also informative without this assumption. This is an improvement over the conventional state of affairs in hazard rate analysis.

Secondly, consider “dynamic matching” approaches. These make a conditional independence assumption (CIA) on the treatment status at some elapsed duration t_0 but they are silent on how this assumption depends on dynamic selection due to unobserved heterogeneity in the interval between inflow and t_0 . Our analysis carries the caveat that such a CIA is unlikely to be satisfied if the treatment and comparison groups have had systematically different event histories between inflow and t_0 even if they have the same personal characteristics and the same labor market history before inflow. Hence, it is useful to ensure that, after propensity score matching, the treatment and comparison groups are identical in terms of (i) the duration distribution between inflow and t_0 , and (ii) the distribution of observed characteristics X among survivors at t_0 . Effectively, satisfaction of these conditions means that one matches on the distribution of unobservable characteristics among survivors as well as on the propensity score.

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.
- Abbring, J.H. and G.J. van den Berg (2005), “Social experiments and instrumental variables with duration outcomes”, Working paper, IFAU, Uppsala.
- Abbring, J.H. and J.J. Heckman (2007): “Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation,” in J.J. Heckman and E. Leamer, editors, *Handbook of Econometrics, Volume VI.b*, North Holland, Amsterdam.
- Andersen, P.K., Ø. Borgan, R.D. Gill, and N. Keiding (1993), *Statistical Models Based on Counting Processes*, Springer, New York.
- Bagkavos, D. and P. Patil (2008), “Local polynomial fitting in failure rate estimation”, *IEEE Transactions on Reliability*, 57, 41–52.
- Black, D.A., J.A. Smith, M.C. Berger and B.J. Noel (2003), “Is the threat of reemployment services more effective than the services themselves? Evidence from random assignment in the UI system”, *American Economic Review*, vol. 93: 1313–1327.
- Blundell, R., M. Costa Dias, C. Meghir, and J. Van Reenen (2004), “Evaluating the employment impact of a mandatory job search program”, *Journal of the European Economic Association*, 2, 569–606.
- Blundell, R. and H. Hoynes (2004). “Has ‘In-Work’ Benefit Reform Helped the Labour Market?,” in R. Blundell, D. Card and R. Freeman (eds), *Seeking a Premier League Economy*, 411-460, Chicago: University of Chicago Press.
- Brewer, M., A. Duncan, A. Shepard and M. Suarez, (2006). “Did Working Families’ Tax Credit work? The impact of in-work support on labour supply in Great Britain,” *Labour Economics*, vol. 13(6): 699-720.
- Brinch, C.N. (2007), “Nonparametric identification of the mixed hazards model with time-varying covariates”, *Econometric Theory*, 23, 349–354.
- Card, D. and R. Hyslop (2005). “Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers,” *Econometrica*, vol. 73(6): 1723-1770.
- Costa Dias, M., H. Ichimura and G.J. van den Berg (2008), “The matching method for treatment evaluation with selective participation and ineligibles”, Working paper, IFS, London.

- Crépon, B., M. Ferracci, G. Jolivet, and G.J. van den Berg (2009), “Active labor market policy effects in a dynamic setting”, *Journal of the European Economic Association*, 7, 595–605.
- De Giorgi, G. (2005), “Long-term effects of a mandatory multistage program: the New Deal for Young People in the UK”, Working paper, IFS, London.
- Dorsett, R. (2006) “The New Deal for Young People: effect on the labour market status of young men”, *Labour Economics*, 13, 405–422.
- DWP (2006), *Departmental Report 2006*, Working paper, Department for Work and Pensions, London.
- Fleming, T.R. and D.P. Harrington (1991), *Counting Processes and Survival Analysis*, Wiley, New York.
- Frölich, M. (2007), “Regression discontinuity design with covariates”, Working paper, IZA Bonn.
- Hahn, J., P. Todd and W. van der Klaauw (2001), “Regression-discontinuity design”, *Econometrica*, 69, 201–209.
- Hall, C. and L. Hartman (2010), “Moral hazard among the sick and unemployed: evidence from a Swedish social insurance reform”, *Empirical Economics*, 39, 27–50.
- Härdle, W. (1994), *Applied Nonparametric Regression*, Cambridge University Press.
- Härdle, W., M. Müller, S. Sperlich and A. Werwatz (2004), *Nonparametric and Semiparametric Models*, Springer, Berlin.
- Heckman, J.J. and B. Singer (1984), “Econometric duration analysis”, *Journal of Econometrics*, 24, 63–132.
- Hess, K.R., D.M. Serachitopol and B.W. Brown (1999), “Hazard function estimators: a simulation study”, *Statistics in Medicine*, 18, 3075–3088.
- Honoré, B.E. (1991), “Identification results for duration models with multiple spells or time-varying covariates”, Working paper, Northwestern University, Evanston.
- Imbens, G. and J. Angrist (1994), “Identification and Estimation of Local Average Treatment Effects”, *Econometrica*, 62(2): 467–475.
- Jiang, J. and K. Doksum (2003), “On local polynomial estimation of hazard rates and their derivatives under random censoring”, in: M. Moore et al., eds., *Mathematical statistics and applications*, Institute of Mathematical Statistics, Beachwood, OH.

- John, M. and A.F. Jawad (2010), “Assessing the regression to the mean for non-normal populations via kernel estimators”, *North American Journal of Medical Sciences*, 2(7): 288–292.
- Kolsrud, J., C. Landais, P. Nilsson and J. Spinnewijn (2016). “The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden,” mimeo, LSE.
- Lalive, R. (2008). “How do extended benefits affect unemployment duration? A regression discontinuity approach,” *Journal of Econometrics*, 142(2): 785–806.
- Lancaster, T. (1990), *The Econometric Analysis of Transition Data*, Cambridge University Press, Cambridge.
- Meyer, B.D. (1996), “What have we learned from the Illinois reemployment bonus experiment?”, *Journal of Labor Economics*, 14, 26–51.
- Müller, H.G. and J.L. Wang (1994), “Hazard rate estimation under random censoring with varying kernels and bandwidths”, *Biometrics*, 50, 61–76.
- Müller, H.G., J.L. Wang, J.R. Carey, E.P. Caswell-Chen, C. Chen, N. Papadopoulos et al. (2004), “Demographic window to aging in the wild: constructing life tables and estimating survival functions from marked individuals of unknown age”, *Aging Cell*, 3, 125–131.
- Nielsen, J.P. and C. Tanggaard (2001), “Boundary and bias correction in kernel hazard estimation”, *Scandinavian Journal of Statistics*, 28, 675–698.
- Podivinsky, J.M. and D. McVicar (2002), “Unemployment duration before and after New Deal”, Working paper, Northern Ireland Economic Research Centre.
- Porter, J. (2003), “Estimation in the regression discontinuity model”, Working paper, Harvard University.
- Ridder, G. (1984), “The distribution of single-spell duration data”, in: G.R. Neumann and N. Westergård-Nielsen, eds., *Studies in labor market analysis*, Springer Verlag, Berlin.
- Rubin, D. (2000), “Discussion of causal inference without counterfactuals”, *Journal of the American Statistical Association*, 95, 435–438.
- Stantcheva, S. (2017). “Optimal Taxation and Human Capital Policies over the Life Cycle,” *Journal of Political Economy*, 125(6).
- Van den Berg, G.J. (2001), “Duration models: Specification, identification, and multiple durations”, in J.J. Heckman and E. Leamer, editors, *Handbook of Econometrics, Volume V*, North Holland, Amsterdam.

- Van Reenen, J. (2004), “Active labour market policies and the British New Deal for Unemployed Youth in context”, in R. Blundell et al., eds., *Seeking a Premier Economy*, University of Chicago Press.
- Vikström, J. (2014), “IPW estimation and related estimators for evaluation of active labor market policies in a dynamic setting”, Working paper, IFAU, Uppsala.
- Wang, J.L. (2005), “Smoothing hazard rates”, in: P. Armitage and T. Colton, eds., *Encyclopedia of Biostatistics*, Wiley, Chichester.
- White, M. and G. Knight (2002), “Benchmarking the effectiveness of NDYP”, Working paper, National Audit Office, London.

Appendix 1. Average multiplicative effects on individual hazard rates

A1.1 Identification

By analogy to the proof of Proposition 2, it follows that

$$\begin{aligned} \frac{\theta_T(t_0|X, S = t_0)}{\theta_T(t_0|X, S = t_1)} &= \frac{\mathbb{E}[\theta_T(t_0|X, V, S = t_0) \mid X, T \geq t_0, S = t_0]}{\mathbb{E}[\theta_T(t_0|X, V, S = t_1) \mid X, T \geq t_0, S = t_1]} \\ &= \frac{\mathbb{E}[\theta_{T(t_0)}(t_0|X, V) \mid X, T(t_0) \geq t_0]}{\mathbb{E}[\theta_{T(t_1)}(t_0|X, V) \mid X, T(t_0) \geq t_0]} \end{aligned} \quad (7)$$

with $t_1 > t_0$. Thus, the ratio of the observable average hazard rates equals the ratio of the average counterfactual hazard rate (averaged over the same sub-population). This does not necessarily equal an average multiplicative effect (i.e. an average of the ratio). For this we make the additional assumption,

Assumption 3 (Multiplicative unobserved heterogeneity).

$$\theta_{T(s)}(t|X, V) = \theta_{T(s)}^0(t|X)V \quad (8)$$

This imposes that the unobserved individual characteristics V affect the counterfactual hazard rates in the same proportional way. Note that this is weaker than adopting an MPH model framework for $T(s)|X, V$ or $T|X, S, V$. First, it does not rule out that t and X and the treatment status interact in the hazard rates of $T(s)|X, V$ or $T|X, S, V$. And secondly, it does not make the MPH assumption that $V \perp\!\!\!\perp X$. But it does imply that individual treatment effects on the hazard at t can be expressed as $\theta_{T(s')}^0(t|X)/\theta_{T(s)}^0(t|X)$, so they are homogeneous across individuals with different V (but not necessarily across X or over time). Indeed, the individual effects at t equal the average multiplicative effects on the hazard rate given X , as defined by versions of ATTS($s', s, t|X$) and ATNTS($s', s, t|X$).

By substituting Assumption 3 into (7), we obtain that $\theta_T(t_0|X, S = t_0)/\theta_T(t_0|X, S = t_1)$ for $t_1 > t_0$ identifies the average multiplicative effects ATNTS($t_0, t_1, t_0|X$) and thus ATTS($t_0, t_1, t_0|X$). In sum,

Proposition 3. *Consider the introduction of a compulsory policy at a given point of time. Suppose we have duration data from cohorts that flow in before this point of time. Under Assumptions 1, 2, and 3, the multiplicative treatment effect on the individual hazard rate at t_0 given X is non-parametrically identified and equals $\theta_T(t_0|X, S = t_0)/\theta_T(t_0|X, S = t_1)$ with $t_1 > t_0$. This does not depend on t_1 as long as t_1 exceeds t_0 .*

This result can be related to identification results for duration models with unobserved heterogeneity and time-varying explanatory variables. Honoré (1991) considers an MPH model with a time-varying explanatory variable that is equal across individuals at short durations but different for some individuals at high durations (notice that our variable S can be re-expressed like that if we only use one value $t_1 > t_0$). He shows that the MPH model is fully identified without assumptions on the tail of the distribution of V . He identifies the effect of the time-varying covariate on the individual hazard rate by considering the ratio of the observable hazard rates at point in time where the covariate value changes for a subset of individuals. Clearly, this resembles the approach in the proof of Proposition 3. Brinch (2007) considers a hazard rate model where X is absent and S is replaced by a time-varying explanatory variable $\tilde{x}(t)$ that is different across individuals at short durations but equal for some individuals at high durations. His model is more general than an MPH model because t and $\tilde{x}(t)$ may interact in the individual hazard rate, like in our Assumption 3. However, it does not allow for covariates X that are dependent on V , and it requires a monotonicity assumption on the over-all effect of the past path of $\tilde{x}(t)$ on the observed survival probability, which we do not need. Brinch (2007) shows that his model is fully identified. His proof is a mirror-image of the proof of Proposition 3: he exploits variation in the value of $\tilde{x}(t)$ at short durations in order to gather information on the unobserved heterogeneity distribution, whereas we exploit the lack of variation in the dynamic selection up to t_0 in order to gather information on the causal effect of S .

A1.2 Inference

We start out by pointing out that if Assumption 3 applies, $T|X, S$ has a survival function that is a Laplace transform of a monotone function of the duration variable. We do not exploit this restriction in the estimation procedure.

We are interested in estimating the ratio of two hazard rates, based on different independent samples, and each evaluated at the left boundary. In obvious notation, we denote this ratio by

$$r(0) = \frac{\theta_1(0)}{\theta_2(0)}$$

and we denote its estimator by $\tilde{r}(0) := \tilde{\theta}_1(0)/\tilde{\theta}_2(0)$, where $\tilde{\theta}_i(0)$ is the boundary kernel hazard estimator of Section 3 (or, alternatively, a local linear hazard rate estimator). We may distinguish between three different methods to obtain a confidence interval for $r(0)$. All three of these methods are more generally applicable to ratio estimators. First, we may perform bootstrapping simultaneously on both samples. Secondly, we may apply the delta method. If,

following Tu (2007), we again ignore the asymptotic biases, then we obtain that the estimator $\tilde{r}(0)$ has an asymptotically normal distribution with mean $r(0)$ and variance

$$\frac{\text{AVar}(\tilde{\theta}_1(0)) + r^2(0)\text{AVar}(\tilde{\theta}_2(0))}{\theta_2^2(0)}$$

For this, we need to assume that, in obvious notation, the fraction $n_1 b_{1,n}/(n_2 b_{2,n})$ converges to a finite number. The confidence interval follows immediately (see Müller et al., 2004, which also contains an empirical example in the related case of boundary kernel estimation of a ratio of densities). Also, a local bandwidth may be used. The approach can be straightforwardly extended to allow for asymptotic biases (see e.g. Porter, 2003, for the relevant delta method result).

The third approach is to use Fieller type confidence intervals (see Tu, 2007). The basic idea is to make a confidence interval for $\tilde{\theta}_1(0) - r(0)\tilde{\theta}_2(0)$ and to convert this into a confidence interval for $\tilde{r}(0)$. This again requires that $n_1 b_{1,n}/(n_2 b_{2,n})$ converges to a finite number.

Appendix 2. Algorithm for the data-adaptive boundary kernel estimator with local bandwidths

Müller and Wang's (1994) optimal local bandwidths minimize the asymptotic mean squared error (MSE). However, this objective function is impractical since it depends on unknown quantities, like the hazard rates themselves. Instead, the optimal local bandwidths can be consistently estimated by minimizing an estimate of the local mean squared error (see Müller and Wang, 1990 and 1994 for a discussion). The following algorithm details the computational implementation stages of the local data-adaptive kernel hazard estimator:

Step 1 Choose initial value of bandwidth and construct grids

1. The initial value of the bandwidth, b_0 , is to be used as global bandwidth to start off the estimation. Müller and Wang (1994) propose $b_0 = R / \left(8n_u^{1/5}\right)$ if data is available in the time interval $[0, R]$, where n_u is the number of uncensored observations.
2. Construct an equidistant grid for duration variable T in the domain $A = [0, R]$, call it $\tilde{T} = \{\tilde{t}_1, \dots, \tilde{t}_M\}$. Computation time depends crucially on the size of this grid, so one may start with a parsimonious choice of M .
3. If computation time is important and, as a consequence, \tilde{T} is sparse, construct a second, finer, equidistant grid for duration variable T in the domain $A = [0, R]$ to estimate the hazard functions. Call it $\tilde{\tilde{T}} = \{\tilde{\tilde{t}}_1, \dots, \tilde{\tilde{t}}_P\}$, where $P > M$.

4. Construct an equidistant grid for bandwidth b in $[\underline{b}, \bar{b}]$, call it $\tilde{B} = \{\tilde{b}_1, \dots, \tilde{b}_L\}$. Müller and Wang (1994) propose using $\underline{b} = 2b_0/3$ and $\bar{b} = 4b_0$. In the empirical analysis in Section 5, this interval is too tight as the optimal choice often coincides with its boundaries. Therefore we use $[\underline{b}, \bar{b}] = [b_0/6, 6b_0]$.

Step 2 Obtain an initial estimate of the hazard rates in all points of the grid \tilde{T} using the initial global bandwidth b_0 :

$$\hat{\theta}_0(\tilde{t}_p) = \frac{1}{b_0} \sum_{i=1}^n K_{\tilde{t}_p} \left(\frac{\tilde{t}_p - t_{(i)}}{b_0} \right) \frac{\delta_{(i)}}{n - i + 1}$$

for $p = 1, \dots, P$.

Step 3 For each point $\tilde{t}_m \in \tilde{T}$ ($m = 1, \dots, M$), estimate the optimal local bandwidth by minimizing the local MSE:

1. Compute the MSE at \tilde{t}_m for each bandwidth $\tilde{b}_l \in \tilde{B}$ ($l = 1, \dots, L$). This is

$$\text{MSE}(\tilde{t}_m, \tilde{b}_l) = \text{Var}(\tilde{t}_m, \tilde{b}_l) + \text{bias}^2(\tilde{t}_m, \tilde{b}_l)$$

where the $\text{Var}(\tilde{t}_m, \tilde{b}_l)$ and $\text{bias}(\tilde{t}_m, \tilde{b}_l)$ are, respectively, the asymptotic variance and bias of the hazard estimator at duration \tilde{t}_m when using bandwidth \tilde{b}_l . The following are consistent estimators of these two quantities,

$$\begin{aligned} \widehat{\text{Var}}(\tilde{t}_m, \tilde{b}_l) &= \frac{1}{n\tilde{b}_l} \int_0^R K_{\tilde{t}_m}^2 \left(\frac{\tilde{t}_m - t}{\tilde{b}_l} \right) \frac{\hat{\theta}_0(t)}{\bar{F}_n(t)} dt \\ \widehat{\text{bias}}(\tilde{t}_m, \tilde{b}_l) &= \int_0^R K_{\tilde{t}_m} \left(\frac{\tilde{t}_m - t}{\tilde{b}_l} \right) \hat{\theta}_0(t) dt - \hat{\theta}_0(\tilde{t}_m) \end{aligned}$$

where the function \bar{F} is the empirical survival function of the uncensored observations. \bar{F} can be estimated at each grid point \tilde{t}_p as follows:

$$\bar{F}(\tilde{t}_p) = 1 - \frac{1}{n+1} \sum_{i=1}^n \mathbf{1}(t_i \leq \tilde{t}_p, \delta_i = 1).$$

The integrals can be approximated numerically. For a generic function $g(t)$, a simple numerical approximation over a grid \tilde{T} including the lower and upper boundaries

of the integrating interval (in this case 0 and R) is

$$\int_0^R g(t) dt \simeq \frac{R}{P-1} \left\{ \sum_{p=2}^{P-1} g(\tilde{t}_p) + \frac{g(\tilde{t}_1) + g(\tilde{t}_P)}{2} \right\}.$$

An alternative is to estimate the variance and bias by varying t (the integrating variable) over the observations instead of over the grid.

2. Select the bandwidth that minimizes the estimated MSE at point \tilde{t}_m over the grid \tilde{B} :

$$b^*(\tilde{t}_m) = \operatorname{argmin}_{\tilde{b}_l} \left\{ \widehat{\text{MSE}}(\tilde{t}_m, \tilde{b}_l), \tilde{b}_l \in \tilde{B} \right\}.$$

Step 4 Smooth the bandwidths b^* to obtain the bandwidths \hat{b} over the grid on which the hazard rates are to be estimated, \tilde{T} . The optimal data-adaptive local bandwidths (using the initial bandwidth b_0 to smooth the original estimates) are

$$\hat{b}(\tilde{t}_p) = \left[\sum_{m=1}^M K_{\tilde{t}_p} \left(\frac{\tilde{t}_p - \tilde{t}_m}{b_0} \right) \right]^{-1} \sum_{m=1}^M K_{\tilde{t}_p} \left(\frac{\tilde{t}_p - \tilde{t}_m}{b_0} \right) b^*(\tilde{t}_m)$$

Step 5 Estimate the data-adaptive kernel hazard rates for points in \tilde{T} using the bandwidths $\hat{b}(\tilde{t}_p)$ for $p = 1, \dots, P$

$$\hat{\theta}(\tilde{t}_p) = \frac{1}{\hat{b}(\tilde{t}_p)} \sum_{i=1}^n K_{\tilde{t}_p} \left(\frac{\tilde{t}_p - t_{(i)}}{\hat{b}(\tilde{t}_p)} \right) \frac{\delta_{(i)}}{n - i + 1}.$$

See also Hess et al. (1999) for useful details on the implementation of the estimator.