Institute for Fiscal Studies

Policy discontinuity and duration outcomes

IFS Working Paper W13/27

Gerard van den Berg
Antoine Bozio
Monica Costa Dias

Policy Discontinuity and Duration Outcomes

Gerard J. van den Berg*, Antoine Bozio† and Mónica Costa Dias‡

September 24, 2013

Abstract

A comparison of hazard rates of duration outcomes before and after policy changes is hampered by non-identification if there is unobserved heterogeneity in the effects and no model structure is imposed. We develop a discontinuity approach that overcomes this by exploiting variation in the moment at which different cohorts are exposed to the policy change, i.e. by considering spells crossing the policy change. We prove identification of average treatment effects on hazard rates without model structure. We estimate these

*University of Mannheim, VU University Amsterdam, IFAU-Uppsala, IZA, CEPR, and IFS. Address: Department of Economics, University of Mannheim, L7, 3–5, 68131 Mannheim, Germany.

†Paris School of Economics and Institute for Fiscal Studies.

‡Institute for Fiscal Studies, CEF-UP at the 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.

Acknowledgements: We thank Enno Mammen, David Margolis, Tatyana Krivobokova, Ina Drepper, and participants at conferences in Oberwolfach, Bielefeld, Nyborg, Mannheim and Westminster, and at an Invited Lecture in the Econometric Society European Meeting 2008 in Milan, the Keynote Lecture at the EALE 2008 Conference in Amsterdam, and the Keynote Lecture at the IZA/Worldbank Conference on Policy Evaluation 2009 in Bonn, for useful comments. We gratefully acknowledge financial support from the ESRC.

effects by local linear kernel hazard regression. We use the introduction of the NDYP program for young unemployed individuals to estimate average program participation effects on the exit rate to work.

1 Introduction

Regression discontinuity (or discontinuity design) is often used to evaluate policy effects. In case of a policy change at a point of time τ^*, the idea is that a comparison of observed individual outcomes just before and just after τ^* may provide an estimate of the mean causal effect of the policy change on the individual outcome.

Empirical researchers have struggled to apply this methodology in studies where the outcome of interest is a duration variable, like unemployment duration or the duration until recovery of a disease. The typical approach considers two cohorts of individuals flowing into the state of interest before and after τ^*, in order to compare spells that begin before and after the policy change.[1] But the spells that begin before τ^* do not all end before τ^*, implying that the corresponding duration outcomes are affected by both policy regimes. One could eliminate spells ending after τ^* from the pre-reform spells, following the line of reason that then these are fully observed under the old policy regime. Generally, however, a positive fraction of the earlier spells will be right-censored at τ^*. If the pre-policy-change data (or, to be short, the pre-policy data) are from a cohort flowing in at $\tau_0 < \tau^*$, then one could restrict attention to the truncated duration distribution on $(0, \tau^* - \tau_0)$ as the outcome of interest. More in general, one may right-censor the pre-policy spells at the moment of the policy change. This is what empirical studies typically have done.

These approaches give rise to two problems. First notice that splitting the data into pre- and post-policy spells effectively translates the policy regime into an explanatory variable that is constant within a given spell. It follows from the literature on duration models that any effects on the individual hazard are only identified under rather strong semi-parametric assumptions. Most prominently, (i) proportionality of the duration dependence effect and the effect of the explanatory variable and unobserved explanatory variables on the hazard rate, implying a causal policy effect that is constant across individuals, and (ii) independence

[1] Notice that with a single cohort of individuals flowing in at say $\tau_0 < \tau^*$, the effect of the policy change cannot be distinguished from the duration dependence of the hazard at and after $\tau^* - \tau_0$.

between the observed and unobserved individual characteristics (see e.g. Meyer, 1996, and Abbring and Van den Berg, 2005; we discuss this in detail in section 2.4.1 of this paper.) This is problematic because we are primarily interested in features of individual hazard rates, and because such semi-parametric assumptions may be unappealing.

The second problem is practical. Let τ_0 denote the date at which the first pre-policy spells start. The smaller $\tau^* - \tau_0$, the less interesting the studied outcome is, whereas the larger $\tau^* - \tau_0$, the longer one has to wait before post-policy data become available that enable a full comparison. If one is interested in the effect on the hazard rate after two years of unemployment duration then one would have to wait for two years after the policy change before an estimate can be made.

In this paper we demonstrate that, in fact, ongoing spells at the moment of the policy change can be fruitfully used to identify and estimate causal parameters of interest. 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. independence of observed explanatory variables from unobserved heterogeneity). We obtain the same type of results for survival probabilities conditional on survival up to a given duration. The basic insight follows from the fact that the policy change is an exogenous time-varying binary explanatory variable whose discontinuity point varies independently across spells that started before τ^*. By comparing survivors who share a given elapsed duration t *at the moment of the policy change* to survivors at the same elapsed duration t in an earlier cohort, we effectively compare two cohorts exposed to the same dynamic selection of individuals with favorable unobserved characteristics up to t. So the two cohorts are identical in terms of their unobserved composition at t. This means that a cross-cohort comparison of outcomes conditional on survival up to t identifies average causal effects and is not contaminated by selection effects.

The identification results naturally lead to an empirical implementation. In case the hazard rate is the outcome of interest, this requires estimates of observed hazard rates. In

general, these are selective averages of individual hazard rates, but by carefully combining different observed hazard rates we obtain the average causal effect of interest.

This is a novel result. As noted above, in models where the policy regime is a time-invariant element of X, the observed hazards are uninformative on the average policy effect on the individual hazard rates if one does not impose some untestable model structure, unless one assumes absence of unobserved heterogeneity. In our approach, however, the observed hazards are informative on average policy effects on individual hazard rates, in the presence of unobserved heterogeneity, and without model structure. In particular, this implies that effects that may have been estimated in models where the policy regime is an element of X and under the assumption of no unobserved heterogeneity (and therefore under the assumption of homogeneous effects) are also valid in the presence of unobserved heterogeneity.

We show that the observed hazard rates can be estimated by using non-parametric kernel-type estimation methods. The estimation of the hazard rate at the moment of the policy change involves estimation at the boundary of the relevant duration interval. Standard kernel estimators are heavily biased at such boundaries. We deal with this by using the Müller and Wang (1994) boundary kernel hazard estimation method with data-adaptive local bandwidths. In addition, we use local linear kernel smoothing, along the lines of Wang (2005). We also perform discrete-time analyses with time-aggregated data. The first method in particular has been used in demography and biostatistics but is not well known in econometrics. This is why we explain it in some detail in the paper.

We also consider estimation of average causal effects on conditional survival probabilities, that is, the average effect of being exposed to the policy from duration t_0 onwards on the probability of leaving the state of interest before some duration $t_1 > t_0$. This requires estimates of the corresponding observed probabilities or two cohorts: one that reaches duration t_0 at calender time τ^* and one that enters the state of interest before $\tau^* - t_1$ and hence reaches duration t_0 before τ^*. Here, as well as with estimation of effects on hazard rates, one typically has a choice between a range of cohorts that may serve as the comparison group of non-treated on $[t_0, t_1)$. We develop a "matching" procedure to select the most appropriate cohort.

At least three branches of literature are connected to the present paper. First, our estimation approach is related to the "regression discontinuity" approach for treatment effects and policy evaluation (see Hahn, Todd and Van der Klaauw, 2001, Porter, 2003, and Frölich, 2007, for important recent econometric contributions in a non-parametric setting). One difference is that right-censoring is an essential feature of duration data, which our estimators need to be able to handle. The second difference is that we estimate hazard rates instead of densities, and the empirical hazard rates are not independent over different durations. Another difference is that the hazard estimates that we combine to estimate effects are taken from samples from different cohorts. This does not require that these hazard rates have any determinant in common. As such, we do not assume that the counterfactual hazard rate in the absence of a policy change is continuous as a function of the elapsed duration t. If we do assume continuity of this hazard rate, then we can attempt to make a before-after comparison around the discontinuity point in a given cohort. A before-after comparison has the advantage that we do not need to assume absence of selective cohort differences, although as noted above we could deal with the latter by matching the most appropriate cohort.

The second relevant branch of literature concerns the literature on treatment evaluation using "dynamic matching", where the assignment process is such that treatments can occur at any possible elapsed duration in the state of interest. Typically, this literature considers survivors at a given elapsed duration t_0 and compares individuals whose treatment is observed to start at t_0 to the survivors at t_0 who have not been treated yet at t_0. The treatment status among these individuals 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 literature takes into account that those who have not yet been treated at t_0 may be treated later, but in general it is silent on the dynamic selection before t_0. Fredriksson and Johansson (2008) develop a matching estimator for average effects of treatment at t_0 on the conditional survival distribution on (t_0, ∞). Crépon et al. (2009) show that the underlying assumptions are essentially the same as in our case, namely "conditional independence" and "no anticipation" (see Section 2 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 had systematically different event histories between zero (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 into the state of interest. If the treated are from a region that is equivalent to the comparison region except for an idiosyncratic temporary business cycle shock at say $0.5t_0$, then the composition in terms of unobservables at t_0 is systematically different between treatment and comparison groups, and hence the CIA fails.

Thirdly, there is a literature on identification of duration models with unobserved heterogeneity V and time-varying explanatory variables $X(t)$. In particular, Brinch (2007) shows that certain types of time-varying explanatory variables enable full identification of a generalized Mixed Proportional Hazard (MPH) model in which t and $X(t)$ may interact in the individual hazard rate. However, this requires that the covariates are independent of V and that V acts multiplicatively on the individual hazard rate, effectively ruling out cross-individual heterogeneity in the covariate effects. We do not need to assume either of these for our results. We discuss the connection to this literature in more detail below.

We apply our novel methodological approach to estimate the average effect of participation in the New Deal for Young People (NDYP) program for young unemployed in the UK on the individual transition rate from unemployment to work. All young unemployed individuals enter the program upon reaching 6 months of unemployment. From that moment until 4 months later, they receive intensive job search assistance. This program was implemented on April 1, 1998. Among those unemployed at the implementation date, only those whose elapsed unemployment duration was a multiple of 6 months were allowed in. If the elapsed duration was not a multiple of 6 months, then in principle the individual was only allowed in at the first moment that his or her elapsed duration would equal a multiple of 6 months.

This scheme allows for identification and non-parametric estimation of some additional

treatment effects. From the implementation date onwards, the policy and its enrollment rules are known to the unemployed. This means that individuals who are unemployed for say 1 month at this date know that if they stay unemployed for another 5 months then they will be treated. Our approach can then be used to identify a meaningful average effect of knowing that one will be treated after 5 months. These are effects of anticipation by the individual of the moment at which he or she will be treated. We show that the analysis of effects on hazard rates and conditional exit probabilities provides insights that are not obtained when studying effects on unconditional survival probabilities.

The NDYP has been evaluated before, in a range of studies (see e.g. Blundell et al., 2004, De Giorgi, 2005, Costa Dias, Ichimura and Van den Berg, 2008). In the empirical section we address differences with this literature in terms of methods and results.

The outline of the paper is as follows. In Section 2 we introduce the duration variable and the policy change, and we consider which average causal effects are identified under various assumptions concerning the available data. Section 3 deals with non-parametric kernel-type estimation. Section 4 contains the empirical application. Section 5 concludes.

2 Duration distributions, policy changes, and identification

2.1 Notation

We consider a population of agents or individuals flowing into a state of interest, which we label interchangeably the 'whole population', 'population of interest' or 'inflow population'. We are interested in the durations that these individuals subsequently spend in that state. In particular, we want to measure the causal effect of a single binary treatment that is either assigned to commence at some time in $\mathbb{R}_+ := [0, \infty)$ after entering the state or not assigned at all. As an example, the treatment may be the exposure of an unemployed individual to

intensive job search assistance by his or her case worker. The duration of interest is the unemployment duration, and the treatment may start after having spent a certain amount of time in unemployment. 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 values in some subset of $\mathbb{R}_+ \cup \{\infty\}$ that we denote by $\mathcal{A}$. Here, the point ∞ represents the no-treatment case. To each treatment $s \in \mathcal{A}$ corresponds a random variable $T(s)$ that represents the potential outcome duration if treatment happens at duration s.[2] For each individual we define a vector $\{T(s),\, s \in \mathcal{A}\}$ of potential outcomes under each possible treatment status s. For ease of exposition we assume that each $T(s)$ is a continuous random variable.

Causal inference is concerned with contrasting potential outcomes corresponding to different treatments. Specifically, we are interested in the differences between the distributions of $T(s)$ and $T(s')$ corresponding to treatments $s, s' \in \mathcal{A}$. 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). Our empirical application in Section 4 provides an example of the insights that can be learned from the study of outcomes that condition on survival. In sum, we focus on the effect 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 an arbitrary s, let the distribution function of the potential outcome $T(s)$ be denoted by $F_{T(s)}$. This is a function of the time since inflow into the state of interest, t. 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 $(0, \infty)$ possibly 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.

[2]We do not need to specify the length of the time span during which a treatment takes place, as long as it does not vary given s.

We allow agents to be heterogeneous in terms of observed characteristics, X, and unobserved characteristics, V. These may be exogenously time-varying, but for ease of exposition we abstract from this. For simplicity, we take V to be a continuous random variable. The hazard rate, integrated hazard and the distribution function of $T(s)$ can be defined for individuals with characteristics (X, V). We denote these by $\theta_{T(s)}(t \mid X, V)$, $\Theta_{T(s)}(t \mid X, V)$ and $F_{T(s)}(t \mid X, V)$, respectively. The survival function is $\overline{F}_s(t \mid X, V) = 1 - F_{T(s)}(t \mid X, V)$.

Because the treatments are mutually exclusive, we can never observe potential outcomes corresponding to different treatments simultaneously. Treatments are assigned according to the random variable S with support $\mathcal{A}$. The actual observed outcome is $T = T(S)$ for individuals assigned to treatment S; all other potential outcomes are counterfactual. Then, S simply denotes the elapsed duration at the moment when the agent enrolls in the program. The hazard, integrated hazard, distribution and survival functions of the observed outcome T are denoted by θ, Θ, F and $\overline{F}$, respectively.

2.2 Treatment effects

We are interested in measuring the differences between the duration distributions of $T(s)$ and $T(s')$ corresponding to treatments $s, s' \in \mathcal{A}$. These are summarized in so-called treatment effects. The individual additive effect of replacing treatment s with treatment s' for someone with characteristics (X, V) is,

$$\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \tag{1}$$

for $t \geq 0$ and for $s', s \in \mathcal{A}$. This is the additive effect on the hazard rate at t.

For the same individual, we also consider the relative treatment effect on the probability of surviving up to t conditional on survival up to t_0

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

for $t \geq t_0 \geq 0$ and, as before, $s', s \in \mathcal{A}$. At $t_0 = 0$ this captures the effect on the unconditional

survival function. And finally, the relative effect on the hazard rate at t is

$$\frac{\theta_{T(s')}(t \mid X, V)}{\theta_{T(s)}(t \mid X, V)} \tag{3}$$

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

Since we allow for heterogeneity across agents, it is natural to focus on inference of averages of individual treatment effects like (1), as quantities of interest. The averages are taken over the distribution of $V \mid X$ in the relevant sub-population. In the sequel, we define average additive effects on the hazard rates $\theta_{T(s)}(t \mid X, V)$ while assuming the latter satisfy the usual regularity conditions that guarantee existence of the expressions below. Analogous additive and relative effects can be defined for the conditional survival probabilities and the hazard rate, respectively.

We define the additive Average Treatment Effect (ATE) on the hazard rate at duration t of replacing treatment s with treatment s' conditional on observed characteristics X as

$$\text{ATE}\left(t; s', s \mid X\right) \;=\; \mathbb{E}\left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X\right]. \tag{4}$$

This involves aggregation over the whole population. However, the hazard at some duration $t > 0$ concerns a sub-population of survivors at t, which is typically different from the population at inflow. As shown in the next section, it may also depend on the treatment status. Thus, and instead, one would like to take the average over V among survivors at t. Following Abbring and Van den Berg (2005), we propose the following average treatment effects on the individual hazard rate:

$$\begin{aligned}
\text{ATTS}\left(t; s', s \mid X\right) &= \mathbb{E}\left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X, S = s', T(s') \geq t\right] && (5)\\
\text{ATNTS}\left(t; s', s \mid X\right) &= \mathbb{E}\left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X, S = s, T(s) \geq t\right] && (6)\\
\text{ATS}_{\widetilde{s}}\left(t; s', s \mid X\right) &= \mathbb{E}\left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X, T(\widetilde{s}) \geq t\right] && (7)\\
\text{ATS}\left(t; s', s \mid X\right) &= \mathbb{E}\left[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V) \mid X, T(s') \geq t, T(s) \geq t\right] && (8)
\end{aligned}$$

which can be called the *Average Treatment effect on the Treated Survivors at t* (ATTS), the *Average Treatment effect on the Non-Treated Survivors at t* (ATNTS), the *Average Treatment*

effect on the Survivors at t if treatment had been $\widetilde{s}$ ($\text{ATS}_{\widetilde{s}}$) and the *Average Treatment effect on the Survivors at t* (ATS). $\text{ATTS}\,(t; s', s \mid X)$ (ATNTS) averages over the distribution of $(V \mid X)$ among the survivors at t among those assigned to treatment s' (s). $\text{ATS}_{\widetilde{s}}\,(t; s', s \mid X)$ averages over the distribution of $(V \mid X)$ among the survivors at t had treatment been $\widetilde{s}$. In particular, $\widetilde{s}$ can assume the value s' (or s), in which case the population of interest is that of survivors at t had treatment been s' (respectively s). $\text{ATS}_{s'}$ (ATS_s) differs from ATTS (ATNTS) because it is unconditional on treatment assignment. And finally, $\text{ATS}\,(t; s', s \mid X)$ averages over the distribution of $(V \mid X)$ among survivors at t under both treatments, s and s'.

2.3 Assumptions

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. Generally, the treatment S is observed iff $S \leq T$. However, note that if $\mathcal{A} = \{0, s_0\}$ for some $s_0 \in (0, \infty]$ then S is effectively always observed.[3]

We assume that treatment assignment is randomized conditional on covariates (X, V), and also that treatment assignment is randomized over V given X,

Assumption 1 (Assignment)**.** $S \perp\!\!\!\perp \{T(s),\, s \in \mathcal{A}\} \mid (X, V) \quad \textit{and} \quad S \perp\!\!\!\perp V \mid X.$

The significance of this assumption is better understood if we interpret (X, V) as the information available to the individual at inflow. Without loss of generality, and for any treatment status $S = s \in \mathcal{A}$, we can write $T(s) = \overline{T}(s; X, V) + \epsilon_s$ where $\overline{T}(s; X, V)$ is the expected potential outcome at inflow ($\overline{T}(s; X, V) = \text{E}\,[T(s) \mid X, V]$) and ϵ_s represents all unpredictable (at the time of treatment assignment) variation in this outcome. Clearly, the distribution of ϵ_s may depend on (X, V). But the first condition in assumption 1 imposes that, once condi-

[3]We also allow for random right-censoring of T. This is usually referred to as "simple random right-censoring". Extensions to more general forms of independent censoring and filtering are straightforward (see Andersen et al., 1993, and Fleming and Harrington, 1991).

tioning on (X, V), the joint distribution of $\{\epsilon_s,\ s \in \mathcal{A}\}$, is unrelated to treatment assignment. The second half of assumption 1 states that assignment is random at inflow conditional on observed characteristics X.

Assumption 1 is equivalent to $S \perp\!\!\!\perp (V,\ \{T(s),\ s \in \mathcal{A}\}) \mid X$ and implies that $S \perp\!\!\!\perp \{T(s),\ s \in \mathcal{A}\} \mid X$.[4] The latter is assumed from the outset in the dynamic matching literature (see e.g. Fredriksson and Johansson, 2008, and Crépon et al., 2009).[5]

Although assumption 1 ensures that assignment is random at inflow $(S \perp\!\!\!\perp V \mid X)$, it does not preclude selective behavior after that. In general, the distribution of $V \mid X$ among survivors at duration $t > 0$ differs from the distribution of V among the inflow population. Moreover, if the treatment has a causal effect on the duration, then the distribution of V among the survivors at points in time $t > 0$ typically depends on the treatment, so $V \not\perp\!\!\!\perp S \mid X, T > t$.[6] In other words, there is no treatment randomization at $t > 0$ despite the randomization $(V \perp\!\!\!\perp S \mid X)$ at $t = 0$.

[4]To see the former, we first show that assumption 1 implies $S \perp\!\!\!\perp (V,\ W) \mid X$ where $W = \{T(s),\ s \in \mathcal{A}\}$. Let f be the general symbol for density. Then

$$\begin{aligned} f(V, W \mid X, S) &= f(W \mid X, S, V)\, f(V \mid X, S) \\ &= f(W \mid X, V)\, f(V \mid X) \\ &= f(V, W \mid X) \end{aligned}$$

where the second equality results from the application of both conditions in assumption 1.
The reverse implication, that $S \perp\!\!\!\perp (V,\ W) \mid X$ implies assumption 1, can be shown as follows

$$f(W \mid X, S, V) = \frac{f(W, V \mid X, S)}{f(V \mid X, S)} = \frac{f(W, V \mid X)}{f(V \mid X, S)} = \frac{f(W, V \mid X)}{f(V \mid X)} = f(W \mid X, V)$$

where the second equality is a direct application of the independence condition and the third equality results from

$$f(V \mid X, S) = \int f(V, W \mid X, S)\, dW = \int f(V, W \mid X)\, dW = f(V \mid X).$$

proving that both conditions in assumption 1 hold.

[5]Notice that in the unrealistic case where V is degenerate, $\Theta_{T(s)}$ can be estimated using standard hazard regression techniques (see *e.g.* Fleming and Harrington, 1991).

[6]See Meyer (1996).

For some empirical designs, however, differential (by treatment status) selective behavior after inflow but prior to treatment can be ruled out. This amounts to assume that there is no anticipation by agents of the moment of future treatment. With this we mean that agents do not have private information about the moment of realization of a future treatment (or that they do not act on such information). We formalize this condition by assuming that current potential integrated hazards do not depend on the moment of future treatment enrollment,

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

See Abbring and Van den Berg (2003) for a detailed discussion. Recall that $\Theta_{T(\infty)}$ is the integrated hazard of the potential duration corresponding to never enrolling in treatment.

The consequences of selective behavior after inflow are reflected on the distribution of unobserved variables V at some duration $t > 0$. To see this, let g and G represent the density and cumulative distribution functions of V, respectively. It holds that

$$\begin{aligned} g\left(v \mid X, S = s, T \geq t\right) &= \frac{\overline{F}\left(t \mid X, S = s, V = v\right) g\left(v \mid X, S = s\right)}{\int_0^\infty \overline{F}\left(t \mid X, S = s, V\right) dG\left(V \mid X, S = s\right)} \\ &= \frac{\overline{F}_{T(s)}\left(t \mid X, V = v\right) g\left(v \mid X\right)}{\int_0^\infty \overline{F}_{T(s)}\left(t \mid X, V\right) dG\left(V \mid X\right)} \\ &= \frac{\exp\left\{-\Theta_{T(s)}\left(t \mid X, V = v\right)\right\} g\left(v \mid X\right)}{\int_0^\infty \exp\left\{-\Theta_{T(s)}\left(t \mid X, V\right)\right\} dG\left(V \mid X\right)} \end{aligned}$$

where the second equality follows from the randomization assumption 1. Clearly, the distribution of V among survivors is generally not independent of duration t or of the treatment status S. In fact, it is easy to construct examples in which the distribution of V among the treated survivors at t is first-order stochastically dominated by the distribution of V among the non-treated survivors at t.[7] Under assumption 2, however, selection behavior before treatment can

[7]For simplicity, suppose there is only one possible treatment, $S = s_0$, meaning that $\mathcal{A} = \{s_0, \infty\}$. If treatment has a positive impact in the individual hazard rate (so that $\theta_{s_0}(t|X, V) > \theta_{T(\infty)}(t|X, V)$ for $t > s_0$) and $\theta_S(t|X, V)$ is increasing in V, and more strongly so after treatment if $S = s_0$, then the individual hazard rate at $t > s_0$ is larger if both $S = s_0$ and V is large. As a result, the treated survivors at t may contain relatively few treated individuals with a high value of V.

be ruled-out, and assignment to treatment $S = s$ is randomized among the 'not-yet-treated' survivors at s (this is shown below, see Proposition 1).

Assumptions 1 and 2 are in line with cases in which a comprehensive policy is rigorously implemented from a specific point in calendar time onwards. This is further discussed in the next section. Another example is a randomized experiment with an instantaneous binary treatment status (i.e. $\mathcal{A} = \{0, \infty\}$). As shown in Abbring and Van den Berg (2003, 2005), relaxation of assumption 1 implies that a semi-parametric model framework needs to be adopted in order to be able to identify objects of interest.

2.4 Inference

In this section we consider inference under two distinct empirical designs that are commonly available and generally satisfy assumptions 1 and 2. Both explore a policy change, the first to compare the steady states before and after the reform, the second to explore variation in the spell duration at treatment assignment for spells starting at different times prior to the reform. We show that non-parametric identification of the impact of treatment on the duration outcomes of survivors can only be ensured for the latter.

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

We consider empirical inference if the data collection leads to two samples: one in which $\Pr(S = 0) = 1$ and one in which $\Pr(S = \infty) = 1$. These can be thought of as samples from the inflow into the state of interest: the sample with $\Pr(S = 0) = 1$ is taken after the introduction of a comprehensive policy that consists of an immediate treatment for all who enter the state of interest, whereas the other sample is taken infinitely long before the introduction of the policy. One may also think of the sample with $\Pr(S = 0) = 1$ as a sample of fully treated agents and the other sample as a sample from the comparison population. The purpose of this subsection is to demonstrate that this sampling scheme, which has been widely used in the

empirical literature, has limited value for inference of effects of interest.

Note that, in this setting, S is observed by the agent from the beginning, and assumption 2 is void. Assumption 1 states that the treatment assignment upon inflow into the state of interest is not selective, conditional on X. In particular, it requires that the distribution of characteristics $V \mid X$ at inflow is the same under both policy regimes.

Average treatment effects can be defined for the comparison of treated ($S = 0$) and non-treated ($S = \infty$), in a analogous way to equations (4)-(8). So, for instance, the average effect on the treated survivors at t is

$$\text{ATTS}(t; 0, \infty \mid X) \;=\; \mathbb{E}\left[\, \theta_{T(0)}\left(t \mid X, V\right) - \theta_{T(\infty)}\left(t \mid X, V\right) \;\middle|\; X, S = 0, T(0) \geq t \,\right].$$

Notice that the $\text{ATS}\left(t; 0, \infty \mid X\right)$ for the relative change in the hazard rate basically equals the *survivor average causal effect* of Rubin (2000) in case the latter would be applied to the duration outcome itself rather than to effects on non-duration outcomes.

The measures of interest that we introduced cannot be estimated non-parametrically from this data design. Non-parametric inference produces the sample counterparts of the following quantities

$$\theta\left(t \mid X, S = 0\right), \;\; \theta\left(t \mid X, S = \infty\right), \;\; \frac{\overline{F}\left(t \mid X, S = 0\right)}{\overline{F}\left(t_0 \mid X, S = 0\right)} \;\text{ and }\; \frac{\overline{F}\left(t \mid X, S = \infty\right)}{\overline{F}\left(t_0 \mid X, S = \infty\right)}$$

where $\theta\left(t \mid X, S = s\right)$ and $\overline{F}\left(t \mid X, S = s\right)$ are, respectively, the observed hazard rate and the survival function at duration t among those with observed characteristics X assigned to treatment s. The individual and observed hazard rates for any triple (t, s, X) are related by

$$\begin{aligned} \theta\left(t \mid X, S = s\right) \;&=\; \mathbb{E}\left[\, \theta\left(t \mid X, S = s, V\right) \;\middle|\; X, S = s, T \geq t \,\right] \\ &=\; \mathbb{E}\left[\, \theta_{T(s)}\left(t \mid X, V\right) \;\middle|\; X, S = s, T(s) \geq t \,\right] \end{aligned} \tag{9}$$

where $\theta_{T(s)}\left(t \mid X, V\right)$ is the individual hazard rate for potential outcome $T(s)$ at duration t among those with treatment status s, and the expectations are taken over the distribution of V conditional on survival up to t. The second equality is a result of the first condition in

assumption 1. Consequently,

$$
\begin{aligned}
&\theta\left(t \mid X, S=0\right) - \theta\left(t \mid X, S=\infty\right) \\
&= \mathbb{E}\big[\theta_{T(0)}\left(t \mid X, V\right) \;\big|\; X, S=0, T(0) \geq t\big] - \mathbb{E}\big[\theta_{T(\infty)}\left(t \mid X, V\right) \;\big|\; X, S=\infty, T(\infty) \geq t\big] \\
&= \Big\{\mathbb{E}\big[\theta_{T(0)}\left(t \mid X, V\right) \;\big|\; X, S=0, T(0) \geq t\big] - \mathbb{E}\big[\theta_{T(\infty)}\left(t \mid X, V\right) \;\big|\; X, S=0, T(0) \geq t\big]\Big\} \\
&\quad + \Big\{\mathbb{E}\big[\theta_{T(\infty)}\left(t \mid X, V\right) \;\big|\; X, S=0, T(0) \geq t\big] - \mathbb{E}\big[\theta_{T(\infty)}\left(t \mid X, V\right) \;\big|\; X, S=\infty, T(\infty) \geq t\big]\Big\}
\end{aligned}
$$

Thus, the difference between the observed hazard rates is the sum of two terms. The first term (in line 3 of the above expression) is the average treatment effect $\mathrm{ATTS}\left(t; 0, \infty \mid X\right)$. The second term (line 4) is the selection effect that, among the survivors at t, the treatment and comparison groups have systematically different unobserved characteristics at $T = t$ despite the randomization at $t = 0$. A similar decomposition applies to the other objects of interest. Clearly, the left-hand side does not capture any meaningful treatment effect because the second term on the right-hand side reflects the selection effect and is unobserved.

In fact, one can construct examples where $\theta(t \mid X, S=0) < \theta(t \mid X, S=\infty)$ even if

$$\theta_{T(0)}(t \mid X, V) > \theta_{T(\infty)}(t \mid X, V) \quad \text{almost surely for all } t, V, X.$$

by analogy to the analysis in Van den Berg (2001) of observed covariate effects in MPH models.

The results are straightforwardly extended to more general $\mathcal{A}$ as long as we only use data on spells in which the treatment status does not change.[8] To identify average treatment effects

[8] We should emphasize that it is possible to identify other treatment effect measures, notably the average additive treatment effect on the unconditional survival probability at t, i.e. $\mathbb{E}\left[\overline{F}_0(t \mid X, V) - \overline{F}_\infty(t \mid X, V) \mid X\right]$. This equals $\overline{F}_0(t \mid X) - \overline{F}_\infty(t \mid X)$. With randomization, as in assumption 1, this in turn is equal to $\overline{F}(t \mid X, S=0) - \overline{F}(t \mid X, S=\infty)$. The two survivor functions in the latter can be estimated straightforwardly taking account of right-censoring. For example, in the absence of X, non-parametric survival estimators like Kaplan-Meier estimators can be used (see e.g. Andersen et al., 1993, and Fleming and Harrington, 1991). One can derive uniform confidence bounds on the potential duration distributions and tests of hypotheses like $\overline{F}_0 = \overline{F}_\infty$ (see Andersen et al., 1993). One may also obtain point-wise results for isolated survival probabilities, e.g. to assess the effect of training on the probabilities of staying unemployed for 6 or 12 months. Furthermore, 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 $t = 0$.

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.

2.4.2 Spells that are interrupted by the policy change

In this subsection we consider empirical inference if the data collection is based on random samples from the cohorts flowing into the state of interest before the unexpected introduction of a comprehensive policy. Let τ denote calendar time, and let τ^* denote the moment at which the policy is implemented. We assume that we follow samples from the inflows at calendar times $\tau < \tau^*$ at least until and including the moment τ^*.

We assume that the policy applies to all agents from calendar time τ^* onwards, including to those who enter the state of interest before τ^*. As a result, each agent has a positive probability of being exposed to the policy. Inflow at time $\tau_0 < \tau^*$ 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 known until calendar time τ^* as there is no anticipation of the introduction of the policy program and, thus, of the future moment of treatment (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 this empirical setting, assumption 1 implies that the distribution of characteristics V in the inflow sample is constant over calendar time.

Comparing agents who flow in before τ^* to those who flow in after τ^* is hampered by the problems discussed in the previous subsection. However, we can now also focus on the effect at duration t of a treatment that starts at duration $s' = \tau^* - \tau_0$, as compared to the case where at duration s' no treatment is assigned yet. We consider the average treatment effects as defined in equations (4)-(8). For instance, the ATTS for treatment s' at duration $t > s'$ is

$$\text{ATTS}\,(t; s', s \mid X) \;=\; \mathbb{E}\big[\theta_{T(s')}(t \mid X, V) - \theta_{T(s)}(t \mid X, V)\big| X, S = s', T(s') \geq t\big] \qquad \text{for } s' \leq t, s.$$

The following proposition is the key to the main results of the paper.

Proposition 1. *Consider two cohorts flowing in at τ_0 and τ_1 such that $\tau_1 < \tau_0 < \tau^*$. 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 any $s \geq t_0$.*

Proof: Take any $s \geq t_0$. As before, let g denote the density function of V and $\overline{F}$ and $\overline{F}_{T(s)}$ the survival functions for the observed and potential outcomes, T and $T(s)$ respectively. The density of V conditional on $(T \geq t_0, X, S = s)$ can be written as

$$g\left(V \mid T \geq t_0, X, S = s\right) = \frac{\overline{F}\left(t_0 \mid V, X, S = s\right) g\left(V \mid X, S = s\right)}{\overline{F}\left(t_0 \mid X, S = s\right)}$$

In this expression,

$$\begin{aligned} \overline{F}\left(t_0 \mid V, X, S = s\right) &= \overline{F}_{T(s)}\left(t_0 \mid V, X\right) \\ g\left(V \mid X, S = s\right) &= g\left(V \mid X\right) \end{aligned}$$

where the first and second equalities are implied by, respectively, the first and second conditions in assumption 1 (namely $S \perp\!\!\!\perp \{T(s),\, s \in \mathcal{A}\} \mid (X, V)$ and $S \perp\!\!\!\perp V \mid X$). This means that the density $g\left(V \mid T \geq t_0, X, S = s\right)$ as a function of V is proportional to $\overline{F}_{T(s)}\left(t_0 \mid V, X\right) g\left(V \mid X\right)$, which, in turn, is proportional to $g\left(V \mid T(s) \geq t_0, X\right)$. In particular, this holds when S is t_0 or t_1 since $t_1 > t_0$.

Next, we show that $g\left(V \mid T(s) \geq t_0, X\right)$ does not vary with s for all $s \geq t_0$, and including $s = t_1$. We notice that

$$\begin{aligned} \overline{F}_{T(s)}\left(t_0 \mid V, X\right) &= \exp\left\{-\Theta_{T(s)}(t_0 \mid X, V)\right\} \\ &= \exp\left\{-\Theta_{T(t_0)}(t_0 \mid X, V)\right\} \end{aligned}$$

where the second equality is ensured by assumption 2 for all $s \geq t_0$. This implies that the density $g\left(V \mid T(s) \geq t_0, X\right)$ as a function of V is proportional to $\overline{F}_{T(t_0)}\left(t_0 \mid V, X\right) g\left(V \mid X\right)$,

where the latter is proportional to $g\left(V \mid T(t_0) \geq t_0, X\right)$. Thus, $g\left(V \mid T(s) \geq t_0, X\right)$ is the same for every $s \geq t_0$, and particularly for $s = t_0$ and $s = t_1$. $\square$

The significance of this proposition is that it shows that there is no selection problem if we compare the sub-population of individuals who are observed to be treated at the elapsed duration t_0 to the sub-population of survivors at t_0 who will be treated at a higher elapsed duration, in the sense that these sub-populations 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 develops equally in both populations, and the resulting sub-populations at t_0 are equal.

We now apply this result to the identification of average treatment effects. These are the main methodological results of the paper. We first enunciate the result and proof and then discuss the meaning and relevance of the identified measures.

Recall that $t_i = \tau^* - \tau_i$. From a cohort flowing in at $\tau_i < \tau^*$, we observe the distribution of $(T \mid X, S = t_i)$. This entails observation of the conditional duration distribution of $(T \mid T \geq t_0, X, S = t_i)$ and the hazard rate $\theta(t_0|X, S = t_i)$ evaluated at t_0.

Proposition 2. *Consider the introduction of a comprehensive policy at time τ^*. Suppose we have duration data from cohorts that flow in before τ^*. Under assumptions 1 and 2, the average treatment effects on the individual hazard rate at duration t_0 of treatment at t_0 as compared to treatment at $t_1 > t_0$, ATTS$(t_0; t_0, t_1)$, ATNTS$(t_0; t_0, t_1)$ and ATS$(t_0; t_0, t_1)$, are non-parametrically identified and equal $\theta(t_0|X, S = t_0) - \theta(t_0|X, S = t_1)$. These do not depend on t_1 as long as t_1 exceeds t_0.*

Proof: Contrasting the hazard rates for the observed durations at t_0 yields

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

The first equality follows from the application of equation (9) to each term of the left-hand side of the first line. Then, by Proposition 1, the distributions over which the expectations are taken in the second term (line 2) are the same for any treatment $t_1 \geq t_0$, explaining the second equality. The third equality follows from the first condition in assumption 1, just like in expression (9). Thus the difference in observed hazard rates identifies the ATTS.

Since the distributions of $(V \mid X, S = t_0, T \geq t_0)$ and $(V \mid X, S = t_1, T \geq t_0)$ are identical (Proposition 1), it also follows that $\text{ATTS}\left(t_0; t_0, t_1 \mid X\right)$ equals $\text{ATNTS}\left(t_0; t_0, t_1 \mid X\right)$. Moreover, since both former distributions are the same as the distribution of $(V \mid X, T(s) \geq t_0)$ for any $s \geq t_0$ (again shown in Proposition 1), it also follows that it equals the distribution of $(V \mid X, T(s) \geq t_0, T(s') \geq t_0)$ for any $s, s' \geq t_0$. In particular, it is the same as the distribution of $(V \mid X, T(t_0) \geq t_0, T(t_1) \geq t_0)$ for $t_1 > t_0$, implying that the ATS is also equal to the ATTS and the ATNTS.

Finally, assumption 2 ensures that changing the value of t_1 does not affect the value of the treatment effect as long as $t_1 > t_0$. □

The $\text{ATTS}\left(t_0; t_0, t_1\right)$ and $\text{ATNTS}\left(t_0; t_0, t_1\right)$ 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. These measures are identified without any functional-form restriction on the individual hazard rates and without the need to assume independence of observed and unobserved explanatory variables. 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.

Since $\text{ATTS}\,(t_0; t_0, t_1)$, $\text{ATNTS}\,(t_0; t_0, t_1)$ and $\text{ATS}\,(t_0; t_0, t_1)$ are all equal and do not depend on t_1 as long as $t_1 > t_0$, we may denote them by a short-hand measure $\text{ATS}(t_0)$ giving the average instantaneous effect of the treatment at t_0 on the survivors at t_0.

The sub-population over which the average is taken depends on 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, 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 have the same problem (see Crépon et al., 2009).

Under assumptions 1 and 2, average treatment effects on the individual conditional survival probabilities are also non-parametrically identified. In this case, we define average effects such as the ATTS as

$$\mathbb{E}\big[\overline{F}_{T(s')}\,(t+a \mid X, V) \;-\; \overline{F}_{T(s)}\,(t+a \mid X, V) \;\big|\, X, S = s', T(s') \geqslant t\big]$$

with $s' \leq s$ and $a > 0$. It follows that these are identified if $t \leq s'$ for the empirical design we are considering under assumptions 1 and 2. In particular, take $t = s' = t_0$, $a = 1$ and $s > t_0$. The average causal effect on survivors of starting the treatment at t_0 on the probability of surviving up to $t_0 + 1$, as compared to when the treatment starts sometime after t_0, is identified from

$$\overline{F}\,(t_0 + 1 \mid T \geq t_0, X, S = t_0) \;-\; \overline{F}\,(t_0 + 1 \mid T \geq t_0, X, S > t_0)\,.$$

Notice that the counterfactual may include cases where treatment happens at a later stage but before the outcome is realised, in which case the treatment effect parameter identifies the impact of being treated at t_0 versus not being treated at t_0 (but possibly being treated later) among the 'not-yet-treated' survivors at t_0 (see Sianesi, 2004, for a discussion of this parameter). Under the current assumptions, however, it is also possible to identify the causal effect on survivors of treatment at t_0 versus no treatment up to the time when the outcome is

measured. For instance, taking again $t = s' = t_0$ and $a = 1$, one can choose the counterfactual $s > t_0 + 1$ to identify the impact on survivors of treatment at t_0 versus no treatment before $t_0 + 1$ on the probability of surviving up to $t_0 + 1$:

$$\overline{F}\left(t_0 + 1 \mid T \geq t_0, X, S = t_0\right) - \overline{F}\left(t_0 + 1 \mid T \geq t_0, X, S > t_0 + 1\right).$$

Clearly, these results enable applications in discrete-time settings as well (see below).

Now consider the average relative effect on the individual hazard rates. By analogy to the proof of Proposition 2, it follows that

$$\begin{aligned}
\frac{\theta\left(t_0 \mid X, S = t_0\right)}{\theta\left(t_0 \mid X, S = t_1\right)} &= \frac{\mathbb{E}\left[\theta\left(t_0 \mid X, V, S = t_0\right) \middle| X, S = t_0, T \geq t_0\right]}{\mathbb{E}\left[\theta\left(t_0 \mid X, V, S = t_1\right) \middle| X, S = t_1, T \geq t_0\right]} \\
&= \frac{\mathbb{E}\left[\theta_{T(t_0)}\left(t_0 \mid X, V\right) \middle| X, S = t_0, T(t_0) \geq t_0\right]}{\mathbb{E}\left[\theta_{T(t_1)}\left(t_0 \mid X, V\right) \middle| X, S = t_0, T(t_0) \geq t_0\right]} \\
&= \frac{\mathbb{E}\left[\theta_{T(t_0)}\left(t_0 \mid X, V\right) \middle| X, T(s) \geq t_0\right]}{\mathbb{E}\left[\theta_{T(t_1)}\left(t_0 \mid X, V\right) \middle| X, T(s) \geq t_0\right]}
\end{aligned} \tag{10}$$

for any $s \geq t_0$ and, in particular, for $s = t_0$. Thus, the ratio of the observable average hazard rates equals the relative effect on the average counterfactual hazard rates (averaged over the same sub-population). This does not necessarily equal the average effect on the ratio. For this we make the additional assumption,

Assumption 3 (Multiplicative unobserved heterogeneity)**.** $\theta_{T(s)}\left(t \mid X, V\right) = \theta^0_{T(s)}\left(t \mid X\right) V$.

Assumption 3 imposes that the individual characteristics V affect the counterfactual hazard rates in the same proportional way. Note that this is weaker than assuming an MPH model for $T(s) \mid X, V$ or $T \mid X, S, V$. First, it does not rule out that t and X and the treatment status interact in the hazard rates of $T(s) \mid X, V$ or $T \mid X, S, V$. 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

$$\frac{\theta^0_{T(s')}\left(t \mid X\right)}{\theta^0_{T(s)}\left(t \mid X\right)}$$

so they are homogeneous across individuals with different V (although not necessarily across X or over time). Indeed, the individual effects at t equal $\text{ATE}(t; s', s \mid X)$, $\text{ATS}(t; s', s \mid X)$, $\text{ATTS}(t; s', s \mid X)$ and $\text{ATNTS}(t; s', s \mid X)$ on the relative magnitude of the hazard rate.

By substituting Assumption 3 into (10) it follows that, for $t_1 > t_0$, $\text{ATS}(t_0; t_0, t_1 \mid X)$ (and thus $\text{ATE}(t_0; t_0, t_1 \mid X)$) on the relative magnitude of the hazard rate is identified by

$$\frac{\theta(t_0 \mid X, S = t_0)}{\theta(t_0 \mid X, S = t_1)}$$

The next proposition summarizes this result.

Proposition 3. *Consider the introduction of a comprehensive policy at time τ^*. Suppose we have duration data from cohorts that flow in before τ^*. Under Assumptions 1 to 3, the relative treatment effect on the individual hazard rate at t_0 is non-parametrically identified and equals*

$$\frac{\theta(t_0 \mid X, S = t_0)}{\theta(t_0 \mid X, S = t_1)}$$

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 cohort with $t_1 > t_0$). 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 a point in time where the covariate value changes for a subset of individuals. This approach is essentially equal to 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 $\widetilde{x}(t)$ 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 $\widetilde{x}(t)$ 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 $\widetilde{x}(t)$ 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 $\widetilde{x}(t)$ 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.

We end this subsection with a discussion of the identification of average treatment effects that are not mentioned in the above propositions. Clearly, one cannot hope to identify a full model, that is, the unknown functions $\theta_{T(s)}(t \mid X, V)$ for all s and the distribution of $V|X$. What about the average treatment effects on the individual hazard rate $\text{ATTS}(t; s', s)$ and $\text{ATNTS}(t; s', s)$ if s' is strictly smaller than t and s? In such cases, inference is subject to the same problem as discussed in section 2.4.1: 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 additional exogenous variation in the treatment duration, and without any functional form assumptions, we cannot identify any accumulation effects of a prolonged exposure to the treatment or delayed effects of a treatment, if the object of interest is the hazard rate. Notice that these shortcomings are averted if the conditional survival probability is the object of interest. In this case one may estimate $\overline{F}_{T(t_0)}(t_0 + a \mid T(t_0) \geq t_0, X) - \overline{F}_{T(t_1)}(t_0 + a \mid T(t_1) \geq t_0, X)$ with $t_1 \geq t_0 + a$, where a captures the length of time needed for accumulation effects or delayed responses to kick in.

3 Hazard rate estimation at a discontinuity using non-parametric smoothing

3.1 Hazard rates of interest

From section 2.4.2, the identification of the average causal effect of the policy change on the individual hazard rates is based on the comparison of observed hazard rates for different entry cohorts into the state of interest. Each of these is separately identified from the corresponding cohort-specific duration data. It is therefore a natural approach to non-parametrically estimate both hazard rates, i.e. $\theta(t_0|X, S = t_0)$ and $\theta(t_0|X, S = t_1)$ for some $t_1 > t_0$.

The estimation of the hazard rate $\theta(t|X, S = t_0)$ at the moment of the policy change, for cohorts flowing into the state of interest at τ_0 such that $t_0 = \tau^* - \tau_0$ (the treatment group), involves estimation at the left boundary t_0 of the relevant duration interval. After all, the shape of the individual hazard rate after the policy change t_0 may not have anything in common with the shape before t_0, so we restrict attention to duration outcomes exceeding t_0. The estimated hazard rate for the treatment group is then contrasted against a similar estimate of the hazard rate at the left boundary t_0 for cohorts reaching such duration before the policy reform at τ^*.

Standard kernel hazard estimators are heavily biased at the boundary point. In this section we discuss the application of methods that are designed to handle discontinuities. Specifically, we propose boundary kernel hazard estimators and local linear kernel smoothing estimators. The statistical literature on non-parametric hazard rate estimation in the presence of discontinuities in the hazard rate and discretely jumping explanatory variables is not yet well-developed or well-known.[9] In section 3.2 we discuss a non-parametric estimator that is mostly unknown in the economics field, the boundary kernel hazard estimators of Müller

[9]Most 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.

and Wang (1994). An alternative to the boundary kernel hazard estimator is the local linear smoothing (or local linear fitting, or locally weighted least squares). We do not discuss it here since it is more widely applied (see Wang, 2005, for an intuitive overview in the field of hazard rate estimation, and Nielsen and Tanggaard, 2001, Jiang and Doksum, 2003, and Bagkavos and Patil, 2008, for more details). In the empirical application below, we have applied both methods to reach very similar conclusions.[10]

Note that the non-parametric estimation restricts attention to the truncated duration variables $T \mid S, T \geq t_0$. For expositional convenience and without loss of generality, we transform the left-truncated durations by shifting them to the left, so that our ultimate interest is in the hazard rate at the boundary 0 when evaluating it from above. When discussing the estimators we also suppress S in the notation as all the estimation exercise is conditional on treatment assignment and the same principles apply regardless of treatment assignment. In the first part of this section, and in line with the empirical analysis below, we do not consider observed explanatory variables X.[11]

3.2 Boundary kernel hazard estimators

In the empirical analysis of section 4 we adopt the second-order boundary kernel hazard estimator of Müller and Wang (1994). Since this is still a largely unknown estimator among economists, we discuss it here in some detail.

Consider a random sample of n subjects, where the duration outcomes can be independent right-censored. Let T_i denote the minimum of the duration outcome and the censoring outcome for subject i ($i = 1, \ldots, n$). Note that this notation deviates from the notation where T denotes

[10] See Andersen et al. (1993) for an introduction to these approaches.

[11] If X is exogenously time-varying on $(0, t_0)$ in a systematic way across cohorts, then this may cause the two hazard rates to have common determinants, but we do not pursue this here. Also, in the special case where V is assumed fixed over time and Assumption 3 applies, each $T \mid X, S$ has a survival function that is a Laplace transform of a monotone function of the duration variable, but we do not exploit this restriction in the estimation procedure.

the duration outcome of interest and $T(s)$ denotes the potential outcome when treatment assignment is s. In fact, we abstract from treatment assignment in this section as all estimation is conditional on it. The difference between T and T_i will always be clear by the presence or absence of an index. Furthermore, δ_i is a binary variable equalling 1 iff the duration outcome is realized before the censoring outcome. Let $(T_{(i)}, \delta_{(i)})$ be the ordered sample with respect to the T_i (so $T_{(1)} \leq T_{(2)} \leq \cdots \leq T_{(n)}$).

We assume that the true hazard rate is twice continuously differentiable in an interval A starting at 0, $A = [0, U]$ where $U > 0$ is the right boundary of the interval. 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 the interior region $I = \{t : b \leq t \leq U\}$, which is adjacent to B (we do not discuss estimation on the right boundary of A here, see Müller and Wang, 1994, for details). In I, the kernel hazard estimator is the standard Ramlau-Hansen kernel hazard estimator,[12]

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

where K is taken to be the Epanechnikov kernel,

$$K(z) \;=\; \frac{3}{4}(1-z^2) \quad \text{for} \quad |z| \leq 1 \tag{11}$$

and where b is understood to decrease with n, as explained below. Notice that z in (11) in the standardized difference between the duration of interest, t, and observations $T_{(i)}$, $(t - T_{(i)})/b$.

In B, the above estimator needs to be modified to account for the bias at the boundary, which is typically of asymptotic order $O(b)$ with the above estimator. The kernel function K is modified to depend on the distance to the left boundary (0). So then K has two arguments, say q and z, where q equals the relative distance t/b to the left boundary and z, as above,

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

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

equals $(t - T_{(i)})/b$. Thus $q \in [0, 1)$ and $z \in [-1, q]$ for any point $t \in B$. This ensures that the support of the boundary kernel does not extend beyond the left boundary. The modified Kernel function is

$$K(q, z) = \frac{12}{(1+q)^4}(z+1)\left[z(1-2q) + \frac{3q^2 - 2q + 1}{2}\right] \quad \text{for} \quad q \in [0, 1],\ z \in [-1, q]$$

which simplifies to (11) if $q = 1$ (that is, in I). Müller and Wang (1994) plot $K(q, z)$ as a function of z for various values of q. As mentioned, $K(1, z)$ 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. At the left boundary, q equals zero and the estimator of $\theta(0)$ equals

$$\widetilde{\theta}(0) = \frac{1}{b}\sum_{i=1}^{n} K\left(0, \frac{t - T_{(i)}}{b}\right)\frac{\delta_{(i)}}{n - i + 1}$$

with

$$K(0, z) = 6(z+1)(2z+1)$$

There is a positive probability that $\widetilde{\theta}(0)$ is negative since the kernel function may assume negative values when q is close to 0 and z is small (close to -1). In such cases it must be replaced by zero.

The boundary correction reduces the bias. At the same time, the variance of the estimator increases, because the number of points used to estimate the hazard close to 0 becomes smaller. The above boundary kernel generates a smaller variance than many other boundary kernels, but 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 data-adaptive bandwidths $b(t)$. Their hazard estimator combines the boundary kernel with varying degrees of smoothing at different points in the distribution of duration t,

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

where K_t is defined as

$$K_t(z) = \begin{cases} K(1,z) & \text{if } t \in I \\ K\left(\frac{t}{b(t)}, z\right) & \text{if } t \in B \end{cases}$$

That is, both the kernel function, K_t and the bandwidth $b(t)$ depend on the point t where the estimate is being computed.

The crucial building blocks of the data-adaptive boundary kernel estimator are the local bandwidths. Their optimal choice is the one minimizing the asymptotic mean squared error (MSE), but this solution 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 estimation. Müller and Wang (1994) propose $b_0 = R \Big/ \left(8n_u^{1/5}\right)$ 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 $\widetilde{T} = \left\{\widetilde{t}_1, \ldots, \widetilde{t}_M\right\}$. 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, $\widetilde{T}$ 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 $\widetilde{\widetilde{T}} = \left\{\widetilde{\widetilde{t}}_1, \ldots, \widetilde{\widetilde{t}}_P\right\}$, where $P > M$.
4. Construct an equidistant grid for bandwidth b in $\left[\underline{b}, \overline{b}\right]$, call it $\widetilde{B} = \left\{\widetilde{b}_1, \ldots, \widetilde{b}_L\right\}$. Müller and Wang (1994) propose using $\underline{b} = 2b_0/3$ and $\overline{b} = 4b_0$. For the empirical application discussed in section 4, we found that this interval was too tight as the optimal choice often coincided with its boundaries. We used $\left[\underline{b}, \overline{b}\right] = [b_0/6, 6b_0]$.

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

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

for $p = 1, \ldots P$.

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

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

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

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

$$\widehat{\text{Var}}\left(\widetilde{t}_m, \widetilde{b}_l\right) = \frac{1}{n\widetilde{b}_l}\int_0^R K^2_{\widetilde{t}_m}\left(\frac{\widetilde{t}_m - t}{\widetilde{b}_l}\right)\frac{\widehat{\theta}_0(t)}{\overline{F}_n(t)}\,dt$$

$$\widehat{\text{bias}}\left(\widetilde{t}_m, \widetilde{b}_l\right) = \int_0^R K_{\widetilde{t}_m}\left(\frac{\widetilde{t}_m - t}{\widetilde{b}_l}\right)\widehat{\theta}_0(t)\,dt - \widehat{\theta}_0\left(\widetilde{t}_m\right)$$

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

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

The integrals can be approximated numerically. For a generic function $g(t)$, a simple numerical approximation over a grid $\widetilde{\widetilde{T}}$ including the lower and upper boundaries of the integrating interval (in this case 0 and R) is

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

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 $\widetilde{t}_m$ over the grid $\widetilde{B}$:

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

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

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

Step 5 Estimate the data-adaptive kernel hazard rates for points in $\widetilde{\widetilde{T}}$ using the bandwidths $\widehat{b}\left(\widetilde{\widetilde{t}}_p\right)$ for $p = 1, \ldots, P$

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

As functions of the number of observations n, the optimal bandwidths satisfy the usual conditions (somewhat loosely, $b(t) \to 0, nb(t) \to \infty$ as n increases). The asymptotic behavior of the estimator is not fundamentally different from usual. The optimal bandwidths are such that $nb^5(t)$ converges to a number smaller than infinity, so $b(t) \sim n^{-\frac{1}{5}}$.

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.

Müller and Wang (1994), Hess et al. (1999) and Jiang and Doksum (2003) provide Monte Carlo simulation results for the above boundary kernel hazard estimator. They conclude that

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

3.3 Implementation issues

We consider some alternatives for the above estimators, and some aggregate effect measures.

The "comparison" cohort(s). Consider the estimation of $\theta(t_0 \mid X, S = t_1)$ with some $t_1 > t_0$. One could argue that this involves estimation in the interior of an interval around t_0 on which the hazard is smooth. In that case, standard kernel hazard (or local linear) estimators can be used. However, one may not want to rule out that the individual hazard rates $\theta_{T(s)}(t \mid X, V, S = s)$ are discontinuous at $t = t_0$ even if $S > t_0$. The application in the next section is a case in point. In that case, one needs to resort to boundary correction methods.

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

Note that one may widen the "comparison group" by taking $\theta(t_0 \mid X, S > t_0)$ instead of $\theta(t_0 \mid X, S = t_1)$. This does come at a price, namely that the validity of assumption 1 needs to hold for all cohorts flowing in before $\tau^* - t_0$, thus ruling out cohort effects. Recall that unobserved cohort effects must be absent, because otherwise $S \not\perp\!\!\!\perp V \mid X$ and assumption 1 would be violated. Observed cohort indicators may be included in X, but note that in non-parametric analysis any addition to X adds to the curse of dimensionality.

Instead of enlarging the comparison group, one may use the availability of multiple com-

parison cohorts in order to select the most similar cohort (or set of cohorts) among those flowing in before $\tau^* - t_0$. We do not observe the distribution of $V \mid X$ 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^*.

If one comparison cohort is to be selected, then one may tend to choose a cohort that flowed in only marginally earlier than the "treated" cohort, following the line of thought that any unobserved change of the entry composition of the cohorts is 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, note that $\theta(t \mid X, S = t_1)$ may display a discontinuity at t_1, so the value $\theta(t_0 \mid X, S = t_1)$ at the elapsed duration $t_0 < t_1$ is to be estimated from observed realized durations in an interval to the right of t_0 that should 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 number of realized duration outcomes providing information on $\theta(t_0 \mid X, S = t_1)$ is very small if t_1 is very close to t_0.

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. Notice, however, that this involves a strengthening of Assumptions 1 and 3. Now suppose that $S \perp\!\!\!\perp X$. This can be empirically verified by examining the composition of the cohorts used to estimate the objects of interest. It is not difficult to demonstrate that Assumption 1 and $S \perp\!\!\!\perp X$ jointly imply that $S \perp\!\!\!\perp V$. So in this case, 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. The only remaining disadvantage is that this does not provide estimates by X.

With discrete X, non-parametric inference would typically lead to separate estimations for each value of X. This would also allow for the selection of the most similar comparison 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.

Discrete time. Now let us reconsider the continuous nature of the duration variable. In practice, a continuous-time analysis may sometimes 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 move to work 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 2 apply but we cannot estimate hazard rates. However, we can estimate conditional survival probabilities and their differences, as outlined in Section 2. 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 can be useful from a practical point of view.

As for the second framework, the analysis of Section 2 is straightforwardly modified to such settings by working with a genuine discrete-time framework. This is pursued in section 4 below.

Reduced form model estimation. The identification results in Section 2 are constructive in that they can be translated into estimation methods. In duration analysis it has been common to view identification results as a justification for the estimation of parameterized

reduced-form model specifications, implicitly assuming that identification results that do not rely on functional form assumptions imply that related estimates are also not fundamentally driven by functional-form assumptions (see Van den Berg, 2001). We may follow this approach as well, and specify classes of models to estimate the objects of interest. An obvious choice is to estimate separate PH models by whether $t \gtrless \tau^*$, using all available cohorts, including in each case calender time as a time-varying regressor.

4 Empirical illustration

4.1 The New Deal for Young People

The New Deal was a flagship welfare-to-work program in the UK, first introduced in the early years of the Labour government in the late 1990s. There were a myriad of New Deals for different groups and addressing different employment problems, the largest being the New Deal for the Young People (NDYP). The NDYP was targeted at the young unemployed, aged 18 to 24, who have claimed unemployment benefits (UB, known as Job Seekers' Allowance in the UK) for 6 months. Participation was compulsory upon reaching 6 months in the claimant count, and refusal to participate was sanctioned by a temporary withdrawal from benefits.

Since entitlement to UB is not time-limited nor dependent on past working history in the UK, and eligibility is constrained only by a means-test, the NDYP was effectively targeted at all young long-term unemployed. Thus, and for simplicity, we use 'unemployed' to signify those in the UB claiming count in what follows.

After enrollment, treatment was split into three stages. It comprised a first period of up to 4 months of intensive job search assistance, with fortnight meetings between the participant and a personal adviser. This was called the Gateway. For those still unemployed after the Gateway, the NDYP offered four alternative treatments: subsidized employment, full-time education or training, working on an organization in the voluntary sector and working in an

environment-focused organization. Participation in one of these four options was compulsory for individuals completing 4 months into the NDYP but could be arranged earlier. The options would last for up to 6 months (or 12 months in the case of education), after which those still unemployed would go through another period of intensive job-search assistance. This was called the Follow Through. If perceived beneficial to the worker, repeated participation in the four alternative options could be arranged.[13]

The NDYP treated millions of people before being replaced by another program in 2009, the Flexible New Deal. In 2006, 172 thousand new participants enrolled in the NDYP and the average number of participants at any month during that year was 93 thousand. According to the UK Department for Work and Pensions statistics, the per-year expenditure of the NDYP during the 2000s was in the order of GBP 200 million, excluding administrative costs.[14] However, a large proportion of this concerns UB that would be due independently of the program, for as long as individuals remain unemployed.

The NDYP was first introduced in a few small pilot areas on January 1, 1998, and extended nation-wide on April 1, 1998. During the implementation stage, the existing stock of long-term unemployed was gradually moved into the program. New participants were called as their duration in unemployment reached 6 months or, for the stock of long-term unemployed at the time of the reform, a multiple of 6 months. Enrollment happened during the intensive job-focused interviews regularly scheduled to happen every 6 months during the unemployment spells. Individuals at other durations could apply for early enrollment, but such behavior is unlikely to be prevalent at the initial stages of the program, when information about the NDYP was still limited.

This scheme is somewhat more complicated than the design considered in section 2. However, it allows for the identification and non-parametric estimation of some average causal

[13]More details on the program can be found in White and Knight (2002), Blundell et al. (2004), Van Reenen (2004), or Dorsett (2006).

[14]See DWP (Department for Work and Pensions), 2006, and the DWP website for recent official statistics on the NDYP.

effects of enrolling in the NDYP as well as of anticipating future enrolment. Indeed, from April (or alternatively January in the pilot areas) 1998 onwards, those still ineligible may anticipate future participation and react in advance. Our methodology will allow us to measure the size of such anticipation responses.

Notably, anticipation effects of this type do not conflict with Assumption 2, and thus do not represent a threat to the identification of the parameters of interest. Crucial to this is that both the treated and comparison groups are drawn either at the time of the reform or before it, having no prior information of the new policy. Section 4.3 below discusses the (potentially more serious) consequences of anticipation when the treated and/or comparison groups are drawn sometime after the reform.

4.2 Data

Data is from the JUVOS longitudinal dataset. This is a random sample of the register data on all UB claiming spells. JUVOS contains information on 5% of the UK population, recording 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 (only since 1996), 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 except for the destination upon leaving the UB claimant count, but even this is plagued with missing values. In total, 5.7% of the spells end in 'unknown destination' and almost 25% end in 'failed to attend'. Subsequent transitions are unobserved if they do not involve a claim of UB.

The estimation sample is formed of men aged between 20 to 24 when reaching 6 months in the claimant count. We discard observations for younger individuals to avoid having to deal with education decisions.

4.3 The choice of treated and comparison groups

In terms of our identification strategy, the time of policy change (τ^*) is January 1, 1998 for a small number of pilot regions, and April 1, 1998 for most of the UK. We aim to estimate the ATS(t_0) for t_0 equal to 6 months, or 182 days. Estimation relies on comparing the survivals among the cohort completing duration t_0 at τ^* (the treatment group, or the treated) with a similar sample of survivals from an earlier cohort (the comparison group).

The continuous-time framework must be reconciled with the requirement of a positive sample size. In practice, we need samples of cohorts flowing in unemployment during two time intervals rather than at two singular points in time. So instead of restricting attention to those individuals completing 6 months in the claimant count in a particular day, we consider a full monthly cohort. For instance, the treated sample includes all claiming spells starting in October 1997 in non-pilot areas, and lasting for at least 6 months. Upon completion of this time in unemployment, which happens during April 1998, these individuals will enrol in the NDYP. We also add the sample of claiming spells starting in July 1997 in pilot areas and lasting until January 1998, when the NDYP is locally available.

It is important to realize that expanding the inflow period may not be innocuous. In particular, those starting a claiming spell towards the end of October will have some time to react to the new information becoming available on April 1st, 1998, before gaining eligibility. They may attempt to influence participation by anticipating or postponing their exit from unemployment. This behavior violates Assumption 2, endogenously affecting the composition of the treated group and leading to biased estimates of the impact of treatment. Such bias should be negligible if the anticipatory effect of the new information is much smaller than the effect of actual participation in the NDYP. We show in the next section that the distortion may lead to an under-estimation of ATS(t_0) at 6 months.

We define several comparison groups in an analogous way, to be formed of individuals completing 182 days in the claimant account over an entire calendar month prior to April 1998. Alternative groups were assessed based on two main outcomes: the distribution of T on days 1

to 181 in the claimant count and the distribution of observed characteristics at 182 days. Close resemblance to the treatment group along these two dimensions supports the randomization hypotheses 1 and 2. Candidate groups varied by date of inflow. We considered the cohorts moving into the claimant count during June 1997 (pilot areas) and September 1997 (non-pilot areas), May 1997 and August 1997, July 1996 and October 1996 or the combination of June and September 1997 with July and October 1996. For simplicity, we designate each alternative cohort by the month of inflow in non-pilot areas as these represent a larger proportion of the population. However, data includes data on both pilot and non-pilot regions in all that follows.

Figure 1 displays the survival functions for the treated and comparison groups up to 181 days into unemployment, prior to the release of NDYP. The survival curve for the combination cohort can hardly be distinguished from the treated survival function, so close is the matching. The survival function for the September 1997 diverges from that for the treatment cohort during the December/January period but quickly returns to match it over the last 2 months of the time window. For our analysis, the important aspect is that early selection ensures treatment and comparison groups are similar at the time of enrolment. We cannot reject such hypothesis for the September cohort. Similarly, the August 1997 cohort does converge towards the treatment cohort curve in the last month before enrolment, but matching 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 for the whole duration of the period, suggesting aggregate conditions in the market may have changed through the 1-year interval.

Table 1 compares the empirical distributions of observed variables among treatment and comparison groups. The September 1997 cohort displays no discernible differences to the treatment cohort (column 1 in the table). The combination cohort does not perform as well, with systematic differences on the history of unemployment up to three years prior to inflow (column 4 in the table).

We follow the empirical evidence in favor of the September 1997 cohort and confine the discussion to estimates obtained with this comparison group. Spells in this cohort will not be

Figure 1: Empirical survival functions between 0 and 6 months after inflow for treated and alternative comparison groups

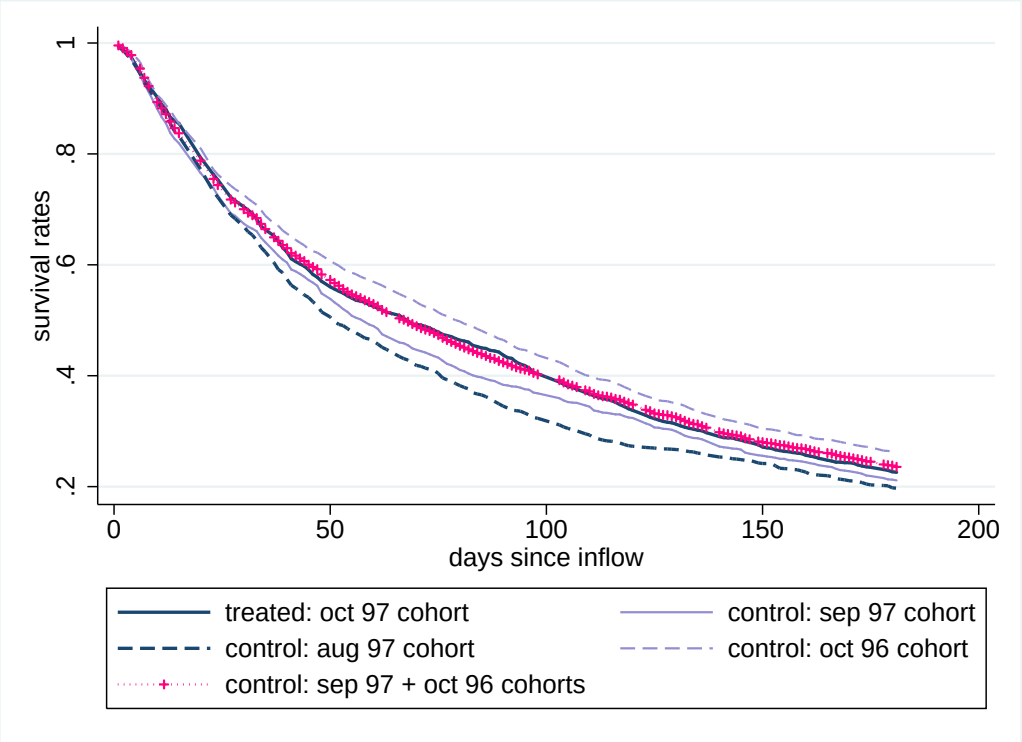

drawn into the NDYP before reaching 12 months of elapsed duration as they are past the 6 months threshold at the time of the reform. But the behavior of this group may be affected earlier by the information becoming available on April 1st, 1998, confounding the treatment effects estimates. This source of bias can be simply eliminated by censoring comparison spells at the time of the reform. In practice, however, one needs to balance the gains and losses from such procedure. In our case, the instantaneous effect of the new information is likely to be small, particularly as the NDYP is in its early days and the prospect of participation among the non-treated is a long distance away. On the contrary, the right-censoring would substantively reduce the information in the comparison cohort sample. Therefore, we chose not to right-censor. But we did check the sensitivity of our results and it turns out that they

Table 1: Treated versus comparison cohorts - p-values for Hotelling statistics comparing distribution of observables characteristics at completion of 181 days in claimant count

	Control cohort			
	September 97	August 97	October 96	Sep97 + Oct96
	(1)	(2)	(3)	(4)
Nr observations	456	368	557	1013
(1) marital status	0.997	0.643	0.114	0.509
(2) age	0.307	0.299	0.916	0.942
(3) region	0.276	0.095	0.112	0.083
(4) occupation	0.767	0.575	0.302	0.532
(5) time U in the past	0.363	0.846	**0.021**	**0.046**
(6) U spells in the past	0.801	0.454	**0.000**	**0.006**
(7) Zero U spells in the past	0.353	0.747	**0.020**	0.164

Notes: Data on men aged 20 to 24 years old 6 months after enrolment. The treatment group is the October 1997 cohort. Variables being compared in rows 5 to 7 describe the claiming history in the 3 years preceding inflow into current unemployment spells. Numbers in bold highlight statistically significant differences in the distribution of observables at 5% level.

are robust to both right-censoring and the choice of the comparison cohort.[15]

By varying the time of entrance in the claimant count, t_0, we can also recover the impact of introducing the NDYP on the hazard rates at different durations. Crucially, the arrival of new information about the possibility of future participation can be used to estimate the anticipatory effects of approaching enrolment. These effects are interesting per se. They are also informative about the reliability of estimates of the impact of program participation by exposing the significance of endogenous selection behavior prior to participation.

We estimate the anticipatory effects of introducing the NDYP at each duration x shorter than 6 months (or 182 days). The treated and comparison groups are formed of individuals aged 20 to 24 at 6 months into the claimant spell and completing elapsed duration x (below 6

[15]Estimates can be obtained from the authors upon request.

months) during April 1998 and March 1998, respectively. In contrast with our earlier discussion on the effects of participation, right censoring can be key to the identification of the parameters of interest here. On the one hand, the comparison group will itself be exposed to the new information on April 1st, 1998. Since their unemployment duration is itself approaching 6 months, they may as well react in advance to influence participation. On the other hand, the treated group will enrol in the NDYP once they reach 6 months in unemployment, with potential effects on their outflow rates from that time onwards. We deal with this two potential sources of bias as explained below, in Section 4.5.

4.4 Results for the average causal effect of program participation

In what follows, the outcome of interest is "all exits from the claimant count", independently of destination, as exits by destination are plagued by missing information. We estimate the impact of program participation in discrete and continuous time by varying the length of the time unit. Estimates in discrete time measure the effects on aggregate monthly outflows while estimates in continuous time do the same for daily outflows. Both sets of estimates are based on the same treated and comparison samples. In total, the sample size of individuals completing 182 days in the claimant account during March and April 1998 while aged 20 to 24 is 902, almost equally split between the treatment (April 1998) and comparison (March 1998) cohorts.

Table 2 presents the main discrete time estimates. We now focus on the results in row (1) and postpone the discussion of the remaining estimates to the next section. Row (1) compares the effect of the NDYP on survivors at 6 months (column (1)) with intention to treat effects at the same duration (columns (2) and (3)). The estimate in column (1) contrasts the one month outflow rates for the treated (October 1997 cohort, completing 6 months in the claimant count during April 1998) and comparison group just excluded from the program (September 1997 cohort, completing the same duration during March 1998). It suggests the outflows from the claimant count increased by 4.5% in the first month after enrolment, representing a raise in

the probability of leaving of about 35%. It is in line with earlier results.[16] The intention to treat effects contrast similar groups one year earlier (column (2)) or contemporaneous older groups (column (3)). None of the estimates is statistically significant.

Table 2: Discrete time estimates - Effects on the treated and intention to treat on outflows from UI claimant count at different elapsed durations; men only

		Treatment effect	Intention to treat	
		20-24 year olds	20-24 year olds	25-29 year olds
	elapsed duration	1998	1997	1998
		(1)	(2)	(3)
Effect at enrolment into treatment				
(1)	6 months	**.045**	.014	-.009
	-inflow Sep (T) vs Oct (C) -	(.023)	(.022)	.021
		911	1118	1365
Effects in anticipation of enrolment				
(2)	4 months	-.015	.006	-.022
	- inflow Nov (T) vs Dec (C) -	(.021)	(.022)	(.020)
		1328	1365	1826
(3)	5 months	-.017	**.057**	.033
	- inflow Oct (T) vs Nov (C) -	(.021)	(.021)	(.020)
		1098	1228	1571

Notes: T and C stand for treatment and comparison groups, respectively. Rows (1) to (3) display estimates of the effects of the NDYP or intention to treat at different elapsed unemployment durations. Each row presents estimates, standard errors and number of observations in the first, second and third lines, respectively. Values in columns (1) and (2) respect to men aged 20 to 24 years old 6 months after inflow into the claimant count. Treatment effects in column (1) compare ongoing spells at the specified elapsed duration in April 1998 and March 1998. Intention to treat effects in column (2) compare ongoing spells at the specified elapsed duration in April 1997 and March 1997. Column (3) presents intention to treat effects at completion of the specified elapsed durations in April and March 1998 for men aged 25 to 29 years old 6 months after inflow into the claimant count.

Estimates in bold are statistically significant at 5% level.

[16] See Blundell et al., 2004.

Figure 2 displays the continuous time counterpart of the treatment effect in row (1) column(1) of table 2. The plot contrasts the outcomes of the October 1997 and September 1997 cohorts using Müller and Wang estimator with optimal local bandwidths. Treatment effects are displayed both in differences and ratios together with the 95% confidence intervals using the analytic asymptotic variance without bias correction.[17] We notice that although $t_0 = 182$ days is the minimum elapsed duration for job search assistance, it is conceivable that the program requires a minimum length of time to act and exert any effect due to the administrative procedures involved in enrolling individuals and passing on the information about the treatment. Thus, figure 2 shows estimates for elapsed durations 182 to 212 days. Zero effects at the start of the eligibility period support comparisons at higher durations by suggesting that the selection process for those in the treatment and comparison groups remains identical. Once the treatment and comparison hazards diverge, differential changes in composition may undermine future comparisons.[18]

The pattern of results is almost identical whether these are estimated in differences or ratios.[19] We therefore discuss the former only. Clearly, the main interest is in the first set of significant values after 182 days. Any features after that may be due to duration dependence, dynamic selection, or both. We find significant effects of participation in the NDYP only after about a week into the program, when the estimates quickly peak to a positive increase of about 0.6% on the rate at which claimants leave the claimant count. This amounts to more than doubling the hazard rate in the absence of treatment as can be seen from the right hand side graph. The effect then drops to a lower positive level that just misses the 96% significance level given the wide confidence bands. However, at this stage we can no-longer separate causal and

[17] With bootstrapping we obtain virtually the same intervals.

[18] Estimates of treatment effects at durations beyond 182 days may be affected by an additional source of bias as late September entrants in the comparison group may cross April 1, 1998 before their hazard being assessed. As discussed before, reactions in anticipation of future eligibility (6 months into the future) may bias estimates of the hazard rates for the comparison cohort. Correcting the comparison group to account for the arrival of information on April 1, 1998 does not change the results. Robustness to the choice of the comparison groups also suggests this source of bias may be irrelevant in estimation.

[19] Recall that the ratio estimate requires Assumption 3 to hold, whereas the difference estimate does not.

Figure 2: Average impact of the NDYP on those eligible to treatment - October 1997 *versus* September 1997 inflow into UI claimant count

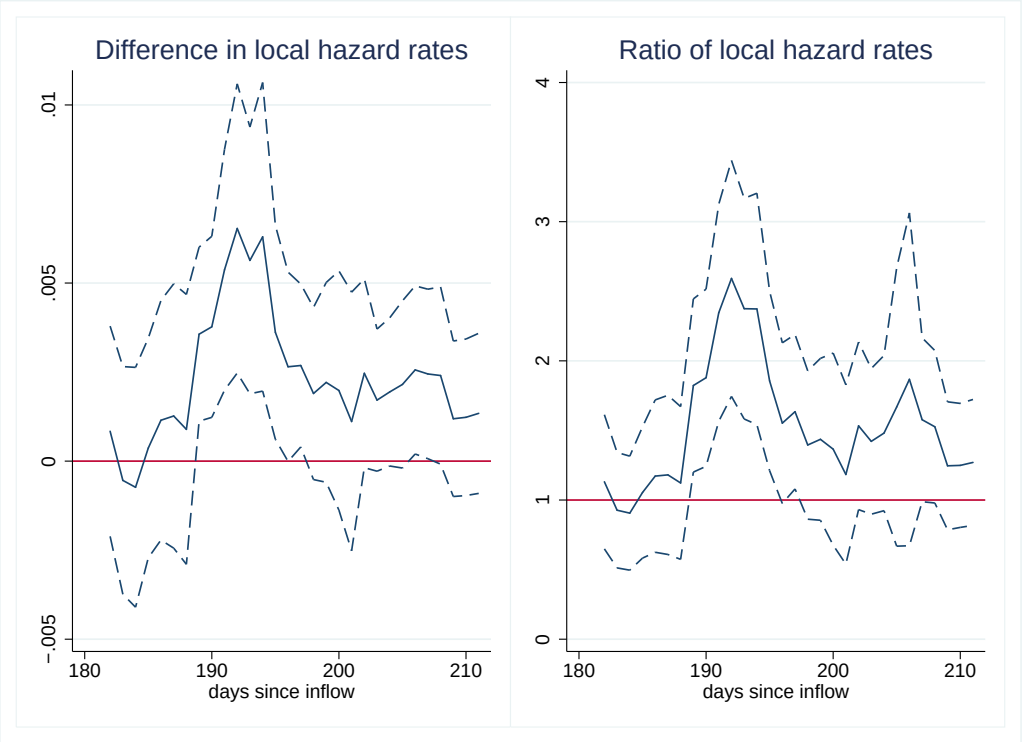

Notes: The treated (comparison) group is formed of spells flowing into unemployment during October (September) 1997 and lasting for 6 months. For both treated and comparison groups, the average optimal local bandwidth is 80 days with a standard deviation of 30 days.

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 important positive average causal effect on the exit rate at 6 months. This is a robust result, valid in the presence of unobserved heterogeneity.

4.5 Results for the average causal effect of receiving information on the future treatment

We now consider spells with elapsed durations shorter than 182 days at the time of the reform to estimate the anticipatory effects of the NDYP. The new comparison groups are formed of individuals completing elapsed duration t_0 during April 1998 (treatment group) or March 1998 (comparison group).

To ensure comparability, we contrast the treated and comparison groups with respect to earlier survival rates and the distribution of observables at elapsed durations of interest. Figure 3 plots the survival functions for 4 sets of entry cohorts. There are some signs of differential selection, apparently associated with the seasonal effects of December/January. For later cohorts, crossing December/January earlier in their spells (panels B and C), the survival functions continue diverging throughout the observed durations and especially at the end of the period, when reaching April (treated) or March (comparison groups). Post December/January cohorts (panel A), unaffected by this seasonal variation, exhibit very close survival functions. Earlier cohorts (panel D) are also affected but return quickly to a common path and look comparable towards the end of the period. The latter result is very similar to the observed patterns for the October and September cohorts.

Table 3 compares the distribution of observed variables for entry cohorts one month apart conditional on reaching elapsed durations 2 to 5 months during March or April 1998. Column 2 shows that the December 1997 and January 1998 cohorts are compositional different upon reaching 3 months in the claimant count. The absence of statistical significant differences in the distribution of observables among earlier cohorts further supports their comparability (columns 3 and 4). On the light of these results, our analysis of anticipatory effects focuses on durations from 4 to 6 months. In the presence of anticipatory effects, this is when they are expected to peak.

Rows 2 and 3 in column 1 of table 2 present the discrete time anticipatory effects at 4 and 5 months of elapsed duration, respectively. Although both negative, none of the estimates

Figure 3: Survival functions for cohorts reaching elapsed durations of 2 (panel A) to 5 (panel D) months during April 1998 (treatment group) *versus* March 1998 (comparison group).

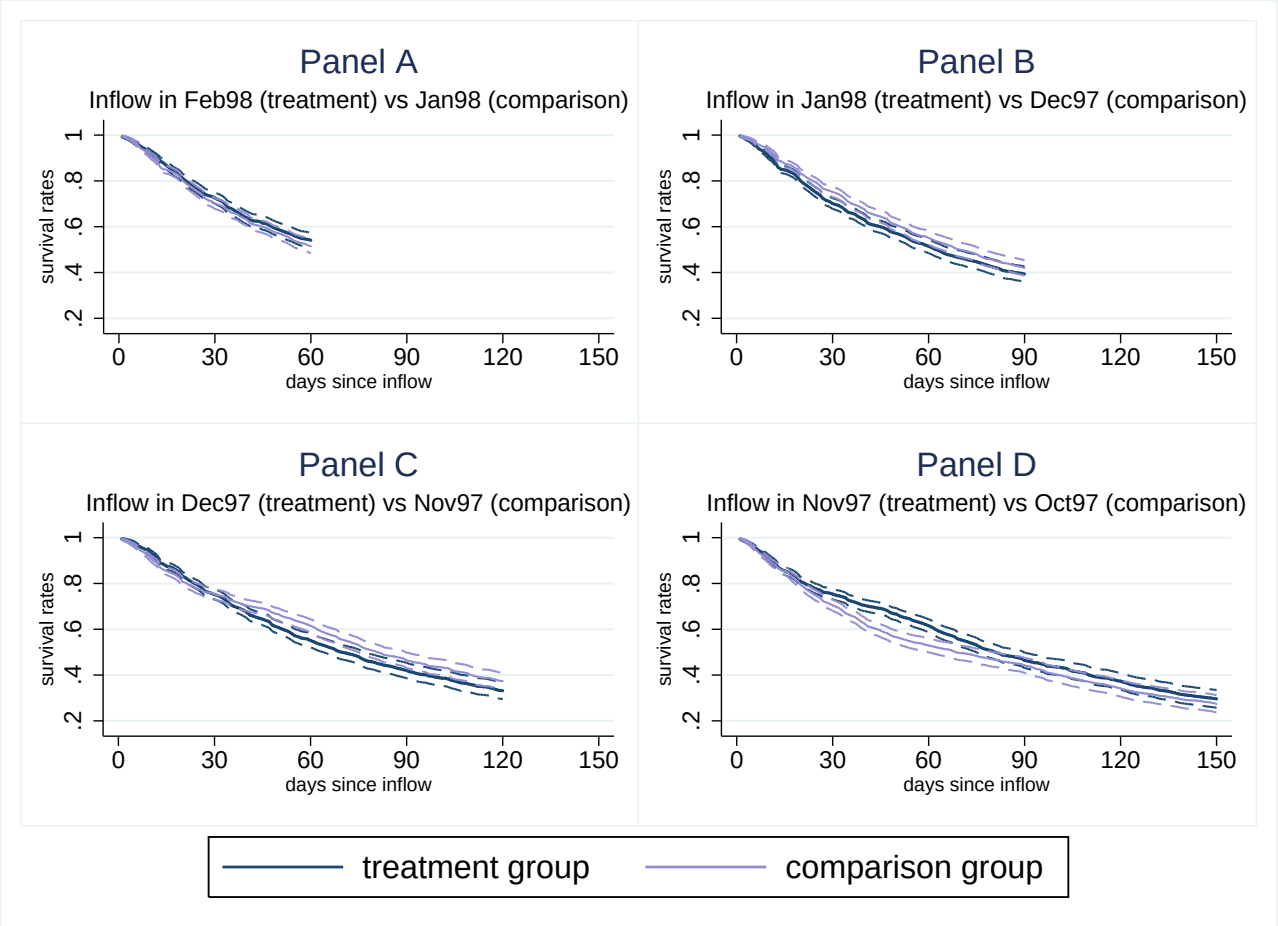

Notes: Each graph title details the month of inflow of treatment and comparison groups constructed to reach 2 (panel A) to 5 (panel D) months of elapsed duration in April (treatment) versus March (comparison) 1998. Men only. Dash lines represent 95% confidence intervals.

is statistically significant. This result is in line with other existing assessments of the NDYP anticipatory effects (Blundell et al., 2004, de Giorgi, 2005). It suggests individuals do not react in advance to the prospect of future treatment. Similar estimates for the effects on intention to treat in columns 2 and 3 are also statistically zero except for the odd case of 20 to 24 years old one year earlier.

Table 3: Treatment versus comparison cohorts - p-values for Hotelling statistics comparing distribution of observables at completion of 2 to 5 months in the UI claimant count

		Month of inflow			
(1)	treatment cohort	Feb 98	Jan 98	Dec 97	Nov 97
(2)	comparison cohort	Jan 98	Dec 97	Nov 97	Oct 97
(3)	elapsed duration	2 months	3 months	4 months	5 months
		(1)	(2)	(3)	(4)
(4)	marital status	0.471	0.339	0.790	0.656
(5)	age	0.120	0.263	0.366	0.318
(6)	region	0.425	0.304	0.671	0.858
(7)	occupation	0.338	0.234	0.410	0.603
(8)	time U in the past	0.188	**0.015**	0.439	0.921
(9)	U spells in the past	0.303	**0.021**	0.242	0.387
(10)	Zero U spells in the past	0.626	0.167	0.271	0.589

Notes: Treatment and comparison groups composed of men in the claimant count conditional on completing the elapsed duration of interest during April 98 (in the treatment group) or March 98 (in the comparison group). Row 1 (2) details the enrolment date for the treatment (comparison) group in the evaluation of the effect at the elapsed duration in row 3. All men aged 20 to 24 years of age 6 months after enrolment. The variables being compared in rows 8 to 10 describe the claiming history in the 3 years preceding inflow into current unemployment spells. Numbers in bold highlight statistically significant differences in the distribution of observables at 5% level.

However, our estimates of the anticipatory effects of treatment may be biased as the one month time window used to estimate outflows crosses April 1, 1998 for the comparison group. Thus, the comparison and treatment groups will be exposed to the same information about the NDYP for part of the evaluation period. In these circumstances, one would expect a bias towards zero if treatment and comparison groups react similarly to the prospect of future treatment.

We deal with this form of bias in the continuous time approach by right censoring spells in the comparison group at the time the NDYP is introduced. There is an additional prob-

lem affecting continuous time estimates resulting from the eligibility rule at 182 days. Local estimates close to that point result from the weighted average of hazard rates in the neigbourhood. Close to the eligibility point, estimates of the anticipatory effects are vulnerable to the inclusion of hazards at durations above 182 days in the estimation procedure. We therefore limit analysis of anticipatory effects to durations below 172 days and condition the choice of the bandwidth to ensure hazard rates affected by the treatment are not considered. In fact, we use our results from the previous section in conditioning the choice of the bandwidth and consider 189 days to be the limit in duration before participation in the NDYP exerts some effects.

Figure 4 shows the continuous time estimates of anticipatory effects. It is evident from the graph that only late in the spell, within 15-20 days before gaining eligibility, does anticipation gain importance. This result would not be noticed in a discrete time analysis. The treatment effects on the hazard rates drop strongly from the beginning of the 5th month. Despite the wide 95% interval bands towards the end of the period (due to the bias corrections discussed above), the effect is statistically significant at high durations.

This result is new, as previous studies looking at anticipatory responses fail to consider changes in behavior very close to the eligibility point.[20] It suggests estimates of the impact of the NDYP using spells reaching eligibility after the program has been running for some time may be upward biased if those in the treatment group are more work-prone than those in the comparison group.

5 Conclusions

We show that, even in the presence of dynamic selection, one can use RD to identify the average impact of a new policy on the hazard rates at different durations in the state of

[20]De Giorgi (2005) estimates $\Pr(T < 6|X,$ inflow after Apr98$) - \Pr(T < 6|X,$ inflow before Oct97$)$. This method is only applicable for unconditional survival probabilities. Blundell et al. (2004) study anticipation *before* April 1998 by exploiting regional and age discontinuities. Both studies find no evidence of anticipation.

Figure 4: Anticipatory effects of the NDYP - difference in hazard rates by elapsed duration prior to completion of 6 months in the claimant count

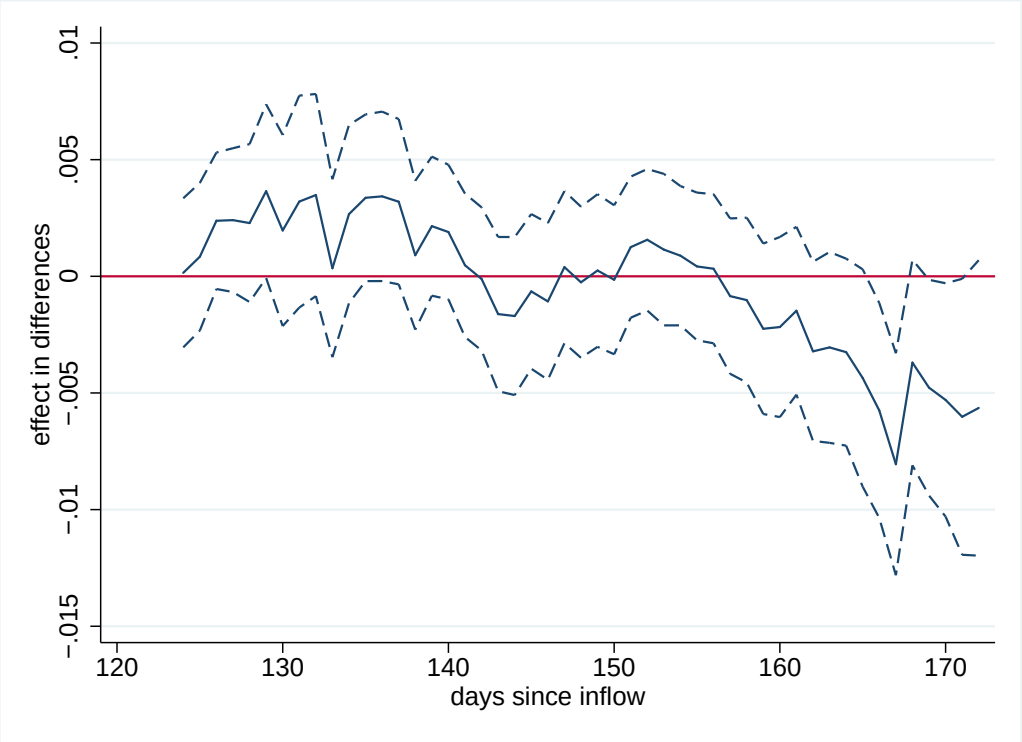

Notes: Effects estimated on (sliding) inflow cohorts 1 month apart, conditional on completing x days in the claimant count in April 1998 (treatment group) versus March 1998 (comparison group), where x is the duration specified in the x-axis. Dash lines represent 95% confidence intervals.

interest. This is done in a completely non-parametric framework, without resorting to the typical proportional hazard specification or exogeneity assumptions. Our results are most useful in the presence of a policy reform when it is possible to explore variation in the time of exposure to the new regime. Identification relies on the comparison of cohorts flowing in the state of interest at different points in time prior to the reform. For a later cohort, survivors at the time of the policy reform are exposed to the new regime when elapsed duration in the state of interest is, say, s, while for an earlier cohort, survivors at duration s still face the old regime. The assumptions required to ensure that the two cohorts are comparable at duration s are generally satisfied in this empirical setting, implying that both cohorts experience the

same dynamic selection up to s.

Our results also show that the implementation details framing the introduction of a new policy have important consequences for the quality and timing of potential evaluation exercises. Specifically, a policy that applies to all cases in the state of interest at the time of the reform alleviates the need for strong identifying assumptions and supports the early production of evaluation results. On the contrary, a policy reform that applies only to the new spells will have to deal with differential dynamic selection, possibly differential selection at inflow once the new regime is announced, and wait for at least t periods before the impact of the new policy can be evaluated on the outcomes at duration t. The main drawback to our approach, though, is that only it only applies in the short run, at the time of the reform. The average treatment effects on the hazard rates cannot be assessed once the policy has matured.

We illustrate the use of our suggested method in the evaluation of the NDYP, using non-parametric estimates of the hazard rates at the boundary as suggested by Müller and Wang (1994). Our results on the impact of the program on the exit rates from unemployment after enrolment are consistent with those of other studies (e.g. Blundell et al., 2004). However, contrary to others, we find statistically significant anticipation effects. These effects are evident when treating time as a continuous variable as they are important only 2 weeks prior to gaining eligibility to the program.

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, IZA, Bonn.

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.

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.

Brinch, C.N. (2007), "Nonparametric identification of the mixed hazards model with time-varying covariates", *Econometric Theory*, 23, 349–354.

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.

Fredriksson, P. and P. Johansson (2008), "Dynamic treatment assignment - the consequences for evaluations using observational data", *Journal of Business and Economic Statistics*, 26, 435–445.

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.

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.

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.

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.

Meyer, B.D. (1996), "What have we learned from the Illinois reemployment bonus experiment?", *Journal of Labor Economics*, 14, 26–51.

Muller, H.G. and Wang, J.L. (1990). "Locally adaptive hazard smoothing", *Probability Theory and Related Fields*, 85, 523–38.

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.

Porter, J. (2003), "Estimation in the regression discontinuity model", Working paper, Harvard University.

Rubin, D. (2000), "Discussion of causal inference without counterfactuals", *Journal of the American Statistical Association*, 95, 435–438.

Tu, D. (2007), "Longitudinal data and survival analysis under-smoothed kernel confidence intervals for the hazard ratio based on censored data", *Biometrical Journal*, 49,474–483.

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.

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.