

IZA DP No. 8450

Policy Discontinuity and Duration Outcomes

Gerard J. van den Berg
Antoine Bozio
Mónica Costa Dias

August 2014

Policy Discontinuity and Duration Outcomes

Gerard J. van den Berg

*University of Mannheim, IFAU-Uppsala, VU University Amsterdam,
IZA, CEPR, CEMMAP and IFS*

Antoine Bozio

*Paris School of Economics
and Institute for Fiscal Studies*

Mónica Costa Dias

*Institute for Fiscal Studies,
CEF-UP at the University of Porto and IZA*

Discussion Paper No. 8450
August 2014

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Policy Discontinuity and Duration Outcomes^{*}

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

JEL Classification: J64, C14, C25

Keywords: policy evaluation, hazard rate, identification, causality, regression discontinuity, selectivity, kernel hazard estimation, local linear regression, average treatment effect, job search assistance, youth unemployment

Corresponding author:

Gerard J. van den Berg
Department of Economics
University of Mannheim
L7, 3–5
68131 Mannheim
Germany
E-mail: gjvdberg@uni-mannheim.de

^{*} We thank Richard Blundell, Ina Drepper, Tatyana Krivobokova, Enno Mammen, David Margolis, and participants of conferences in Oberwolfach, Bielefeld, Nyborg, Mannheim and Westminster, an Invited Lecture at an Econometric Society European Meeting, and Keynote Lectures at an EALE Conference and an IZA/World Bank Conference on Policy Evaluation, for useful comments. We gratefully acknowledge financial support from the ESRC. Van den Berg is Alexander von Humboldt Professor of Econometrics and Empirical Economics and thanks the Humboldt Stiftung for financial support.

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 the unemployment duration or the duration until recovery of a disease) or, more specifically, where the object of interest is the hazard rate of the duration outcome. Consider the use of data on individuals flowing into the state of interest before τ^* as well as individuals flowing in after τ^* , in order to compare spells that occur before τ^* to spells that occur after τ^* .¹ A first practical problem with such an approach is that spells that start before the policy change, i.e. spells that are informative on the outcome distribution before the policy change, do not all end before the policy change. The corresponding duration outcomes are then affected by both policy regimes. This may be dealt with by right-censoring the pre-policy-change spells (or, to be short, the pre-policy data) at τ^* . Such an approach is common in empirical studies.

However, this approach in which the data are split into pre-policy spells and post-policy spells has some more fundamental problems. To explain these, notice that the approach effectively translates the policy regime into an explanatory variable that is constant within any given spell. From the literature on duration models it follows that effects on the individual hazard rate are only identified under rather strong semi-parametric assumptions. Most prominently, it is assumed that the duration dependence effect and the effects of the observed and unobserved explanatory variables on the individual hazard rate are proportional. This implies a causal policy effect that is constant across individuals. In addition, independence between observed and unobserved individual characteristics is assumed (see e.g. Meyer, 1996, and Abbring and Van den Berg, 2005; we discuss this in detail in Subsection 2.2 of this paper.) Such semi-parametric assumptions may be unappealing.

Yet another problem with the above approach is that policy evaluation requires a certain waiting time before the post-policy data can be observed. For example, if one is interested in the effect on the hazard rate after two years of

¹Notice 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 $\tau^* - \tau_0$.

unemployment duration then one has to wait for two years after the policy change before an estimate can be made. As the time between the policy reform and the observation of post-policy outcomes increases, it becomes increasingly likely that the post-policy outcomes are affected by events after the reform, complicating the evaluation. Clearly, this is at odds with the spirit of the “regression discontinuity” approach in which observations are used that are close to the point in time at which some policy exposure changes.

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 where the dynamic selection of individuals with favorable unobserved characteristics is the same 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 suggest an empirical implementation. If the hazard rate is the outcome of interest, this requires estimates of observed hazard rates, meaning hazard rates as a function of the elapsed duration and observed covariates. In general, observed hazard rates are selective averages of individual hazard rates, but by carefully combining different observed hazard rates we obtain the average causal effect of interest.

These results are novel. As noted above, in models where the policy regime is a time-invariant covariate, 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. This leads to the insight that models in which the pol-

icy regime is a time-varying covariate and in which unobserved heterogeneity is absent (so that the effects are assumed to be homogeneous) give rise to policy effect estimates that may also be valid in the presence of unobserved heterogeneity. In this case, the fact that in linear regression models orthogonal omitted variables can be subsumed into the residual term carries over to duration analysis. In the paper we mention examples of empirical studies in which such models are estimated. It follows that the estimates, obtained under the assumption of no unobserved heterogeneity, may also be valid without that assumption.

The observed hazard rates can be non-parametrically estimated by using kernel hazard estimation methods. 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 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 two non-parametric methods have been used in demography and biostatistics but they have not yet been widely used in econometrics.

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, for the cohort for which t_0 is reached at calendar time τ^* and for a cohort that entered 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 connected 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 econometric contributions in a non-parametric setting). Obvious differences are (a) that right-censoring is an essential feature of duration data, which our estimators need to be able to handle, and (b) that we estimate hazard rates instead of densities. 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 everywhere 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 recent 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 or unobserved heterogeneity before t_0 . Vikström (2014) discusses matching estimators for average effects of a treatment at t_0 on the conditional survival distribution on (t_0, ∞) (see also Fredriksson and Johansson, 2008, and Crépon et al., 2009). 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. For example, if the treated are from a region that is equivalent to the comparison region except for an idiosyncratic temporary business cycle shock at say $t_0/2$, then the composition in terms of unobservables at t_0 is systematically different between treatment and comparison groups, and hence the CIA at t_0 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 vari-

ables 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 enroll in a job search assistance program upon reaching 6 months of unemployment. This program was implemented on April 1, 1998.

Among those unemployed at the implementation date, only those whose elapsed unemployment duration was an integer 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 equaled a multiple of 6 months. This scheme allows for identification and non-parametric estimation of some additional causal 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 4 months at this date know that if they stay unemployed for another 2 months then they will receive intensive job search assistance. Our approach can then be used to identify a meaningful average causal effect of knowing at an elapsed duration of 4 months that one will be treated 2 months later. These are effects of anticipation by the individual of the future job search assistance treatment. This illustrates that the analysis of effects on hazard rates and conditional exit probabilities provides insights that cannot be 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, and 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. We also discuss how the method can be used to evaluate the effect of the arrival of information. For example, a new policy may be announced that affects individuals in the state

of interest once their spell duration reaches a certain length. Our method then enables inference on anticipation effects. The method can also be used to evaluate treatments occurring at some point in calendar time that varies across individuals. Section 4 contains the empirical application. Section 5 concludes.

2 Duration distributions, policy changes, and identification

2.1 Notation and assumptions

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

This framework may look more general than a framework for the evaluation of a single policy reform or a binary “reform exposure” indicator as sketched in Section 1. When comparing outcomes in two mutually exclusive policy regimes, a framework with two mutually exclusive treatment statuses may suffice. However, the treatment at the elapsed duration s can be interpreted as representing the exposure to a reform occurring at the individual elapsed duration s . In the stock of individuals in the state of interest at the moment of the policy reform, the elapsed duration from the moment of inflow until the moment of exposure to the reform will be dispersed. We therefore do not restrict the number of elements in \mathcal{A} at this stage.

Causal inference is concerned with contrasting potential outcomes corresponding to different treatments. Specifically, we are interested in differences between the distributions of $T(s)$ and $T(s')$ corresponding to treatments $s, s' \in \mathcal{A}$. These differences are called treatment effects. In social sciences, the exit rate or hazard rate of a duration distribution is the most interesting feature of this distribution, as it is directly related to the agent’s behavior and his information set and cir-

cumstances conditional on survival into the state of interest (see Van den Berg, 2001).² Therefore we focus on average effects of the treatments on the individual exit rate out of the state of interest and the individual conditional exit probabilities out of this state.

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

The individual treatment effect of interest is

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

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

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

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

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

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

for all $t \geq 0$ and for all $s', s \in \mathcal{A}$. Below we also consider alternative treatment effects.

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

²With T continuous, the hazard rate at elapsed duration t is defined as $\theta(t) = \lim_{dt \downarrow 0} \Pr(T \in [t, t + dt) | T \geq t) / dt$.

simply take S to denote the elapsed duration at the moment at which the agent is exposed to the reform.

We allow agents to be ex ante heterogeneous in terms of observed characteristics X and unobserved characteristics V . These characteristics may be exogenously time-varying, but for ease of exposition we abstract from this. For the same reason, we take V to be a continuous random variable.

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 $\theta_{T(s)}(t | X, V)$, $\Theta_{T(s)}(t | X, V)$ and $F_{T(s)}(t | X, V)$, respectively. The survival function is $\bar{F}_s(t | X, V) = 1 - F_{T(s)}(t | X, V)$. The individual treatment effects defined above can be defined accordingly as functions of X and V .

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

As we shall see, the assumption is in line with cases in which a comprehensive policy is rigorously implemented from a specific point in calendar time onwards. 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), settings in which the assumption is relaxed require a semi-parametric model framework in order to be able to identify objects of interest.

Notice that the assumption implies that $S \perp\!\!\!\perp \{T(s)\} | X$. The latter is assumed from the outset in the dynamic matching literature (see e.g. Crépon et al., 2009).⁴

Throughout much of the paper, we assume that there is no anticipation by agents of the moment of future reforms. With this we mean that agents do not have private information on the moment of realization of a future reform (or that they do not act on such information). We formalize this by assuming that current potential integrated hazards do not depend on the moment of future treatment exposure,

³This 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).

⁴In the unrealistic special case where V is degenerate, $\Theta_{T(s)}$ can be estimated using standard hazard regression techniques (see e.g. Fleming and Harrington, 1991).

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

(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.

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

In this subsection we consider empirical inference if the data collection leads to two samples: one in which $\Pr(S = 0) = 1$ and one in which $\Pr(S = \infty) = 1$. In the context of policy reform evaluation, these samples originate from two sub-populations. One sample is drawn from the inflow into the state of interest after the introduction of a policy, whereas the other sample is drawn from the inflow into the state of interest infinitely⁵ long before the introduction of the policy. Figure 1 depicts this setting in a Lexis diagram, where τ denotes calendar time and τ^* denotes the moment at which the reform is implemented. Each diagonal line represents a single cohort. Notice that we tacitly assume that the reform is comprehensive.




Figure 1. “Before” sample and “after” sample.

⁵Or, at least sufficiently long before the reform to observe outcomes in a sufficiently large duration interval. In this case, the outcomes are right-censored at the moment of the reform. Alternatively, one may think of the sample with $\Pr(S = 0) = 1$ as a sample of fully treated agents and the other sample as a sample of controls. Provided that no ambiguity arises, we use the terms “pre-reform policy”, “pre-policy”, and “control” interchangeably. The same applies to “post-reform policy”, “post-policy” and “treatment”, and the same also applies to “moment of the policy change” “reform” and “introduction of the policy”. A more explicit discussion is provided in Subsection 3.2.

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

Note that in the current dichotomous setting, S is observable by the agent from the onset, and Assumption 2 is void. One may say that in this case, anticipation is perfect. Assumption 1 implies that the treatment assignment upon inflow into the state of interest is not selective, conditional on X . In particular, the distribution of characteristics $V|X$ at inflow is the same in each policy regime.

Since we allow for unobserved heterogeneity across agents, it is natural to focus the inference on averages of individual treatment effects like (1) as quantities of interest. We thus need to average

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

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

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

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

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

which typically varies with S .⁶

⁶In fact it is not difficult to construct examples in which the distribution of $V|X$ among the treated survivors at t is first-order stochastically dominated by the distribution of $V|X$ among the non-treated survivors at t , if there is a strong positive interaction between being treated

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

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

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

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

which can be called the *Average Treatment effect on the Treated Survivors at t* (ATTS($t|X$)), the *Average Treatment effect on the Survivors at t* (ATS($t|X$)), and the *Average Treatment effect on the Non-Treated Survivors at t* (ATNTS($t|X$)). ATTS($t|X$) averages over the distribution of $V|X$ among the survivors at t if the agents are assigned to the “treatment” (i.e., are assigned to $s = 0$, or, in other words, are exposed to the policy introduced by the reform). Under randomization, this is equivalent to averaging over the distribution of V among the treated survivors at t (so with $X, T \geq t, S = 0$). ATNTS($t|X$) is the counterpart of this for assignment to the control group. ATS($t|X$) averages over the distribution of $V|X$ among individuals who survive up to t under both possible treatment regimes. These measures can subsequently be aggregated over some distribution of X . Analogous additive and multiplicative effects can be defined for the conditional survival probabilities and the hazard rate, respectively (recall equations (2) and (3)).⁷ Note that in general all measures are properties of sub-populations whose composition depends on the treatment effect in the duration interval $[0, t]$.

The above measures of interest cannot be estimated non-parametrically from the data design of the present subsection. Non-parametric inference produces sample equivalents of $\theta_T(t|X, S = 0)$ and $\theta_T(t|X, S = \infty)$ and of $\bar{F}_T(t|X, S =$

and V in the individual hazard rates $\theta_{T(s)}(t|X, V)$ and if these hazard rates increase in V and in being treated (see Van den Berg, 2001). In such scenarios, the individual hazard rate at t is disproportionately large if both $S = 0$ and V is large, and as a result the treated survivors at t may contain relatively few treated individuals with a high value of V .

⁷The ATS($t|X$) version for the multiplicative effect on the hazard rate basically equals the *survivor average causal effect* of Rubin (2000) in case the latter measure is applied to the duration outcome itself rather than to non-duration outcomes.

$0)/\bar{F}_T(t_0|X, S = 0)$ and $\bar{F}_T(t|X, S = \infty)/\bar{F}_T(t_0|X, S = \infty)$. For given t, s, X , individual and observable hazard rates are connected by

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

By definition, therefore,

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

which is the sum of two differences. The first difference is the average treatment effect $\text{ATTS}(t|X)$ (for sake of brevity, we refer to the next subsection for the proof of this statement). The second difference is the selection effect due to the fact that at $T = t$, among the survivors at t , those exposed to the post-reform policy and those not exposed have systematically different unobserved characteristics despite the randomization of the regime status at $t = 0$. A similar decomposition applies to the other objects of interest. Since the second term on the right-hand side reflects the selection effect and is unobserved, we conclude that the left-hand side cannot be used to non-parametrically estimate $\text{ATTS}(t|X)$.⁸

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. To identify average treatment effects in the setting of the current subsection, one needs to adopt a semi-parametric model structure like an MPH model, or one needs to assume absence of unobserved heterogeneity.⁹

⁸By analogy to the remarks on equation (4), one can construct examples where $\theta_T(t|X, S = 0) < \theta_T(t|X, S = \infty)$ even if $\theta_{T(0)}(t|X, V) > \theta_{T(\infty)}(t|X, V)$ almost surely for all t, V, X .

⁹The average additive treatment effect on the unconditional survival probability at t , i.e. $\mathbb{E}[\bar{F}_{T(0)}(t|X, V) - \bar{F}_{T(\infty)}(t|X, V)]$, is identified under a randomization assumption such as Assumption 1, from the observable expression $\Pr(T > t|X, S = 0) - \Pr(T > t|X, S = \infty)$. Inference on the two survivor functions in this expression is straightforward; see e.g. Andersen et al. (1993). We also point out that under the assumption that all individual treatment effects have the same sign across t and V , this sign is identified from the observed distributions or the observed hazard rates at $t = 0$.

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

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




Figure 2. “Before” sample including spells that are ongoing at the moment of the policy change and that are followed beyond that moment.

We assume that the post-reform policy regime applies to all agents, from calendar time τ^* onwards, including to those who enter the state of interest before τ^* . Inflow at time $\tau_0 \leq \tau^*$ leads to $S := \tau^* - \tau_0$. Thus, there is a one-to-one correspondence between the moment of inflow and the duration at which the treatment starts. However, in this setting, S is not observed by the agent until calendar time τ^* , as there is no anticipation of the introduction of the new policy program (Assumption 2). We rule out that the distributions of $T(s)|X, V$ are discontinuous at $T(s) = s$ (though of course the hazard rates may be discontinuous there).

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

Comparing agents who flow out before τ^* to those who flow in after τ^* is hampered by the same problems as in the previous subsection. However, we can now also examine the effect at duration $\tau^* - \tau_0$ of a treatment that starts at duration S , as compared to the case where at duration $\tau^* - \tau_0$ no treatment is assigned yet. To this purpose, we may define average treatment effects measures by analogy to those in the previous subsection. For example,

$$\text{ATTS}(s', s, t|X) := \mathbb{E} \left[\theta_{T(s')} (t|X, V) - \theta_{T(s)} (t|X, V) \mid X, T(s') \geq t \right] \quad \text{with } s' \leq t, s.$$

$$\text{ATNTS}(s', s, t|X) := \mathbb{E} \left[\theta_{T(s')} (t|X, V) - \theta_{T(s)} (t|X, V) \mid X, T(s) \geq t \right] \quad \text{with } s' \leq t, s.$$

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

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

Proof: Note that $t_0 < t_1$. Let \Pr be a general symbol for a density as well as a probability. The density $\Pr(V|T \geq t_0, X, S = t_i)$ (with $i = 0, 1$) can be written as

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

In this expression, $\Pr(T \geq t_0|V, X, S = t_i)$ equals $\Pr(T(t_i) \geq t_0|V, X)$ due to the randomized assignment assumption (Assumption 1: $S \perp\!\!\!\perp \{T(s)\} \mid (X, V)$). Moreover, $\Pr(V|X, S = t_i)$ equals $\Pr(V|X)$ due to the second part of Assumption 1 ($S \perp\!\!\!\perp V \mid X$).¹⁰ This means that the density $\Pr(V|T \geq t_0, X, S = t_i)$ as a function of V is proportional to $\Pr(T(t_i) \geq t_0|V, X)\Pr(V|X)$ which is proportional to $\Pr(V|T(t_i) \geq t_0, X)$.

Next, we show that $\Pr(V|T(s) \geq t_0, X)$ is the same for every $s \geq t_0$ including $s = t_1$. By analogy to the previous paragraph, the second part of Assumption 1 implies that the density $\Pr(V|T(s) \geq t_0, X)$ as a function of V is proportional to $\Pr(T(s) \geq t_0|V, X)\Pr(V|X)$. The term $\Pr(T(s) \geq t_0|V, X)$ can be expressed as

¹⁰In the setting of this subsection, the assumptions entail that the policy or treatment status is randomized among the stock of subjects in the state of interest, given X . See e.g. Ridder (1984) for an extensive discussion of stock samples.

$\exp(-\Theta_{T(s)}(t_0|X, V))$. By Assumption 2, this equals $\exp(-\Theta_{T(t_0)}(t_0|X, V))$ since $s \geq t_0$. This implies that the density $\Pr(V|T(s) \geq t_0, X)$ as a function of V is proportional to $\Pr(T(t_0) \geq t_0|V, X)\Pr(V|X)$, where the latter is proportional to $\Pr(V|T(t_0) \geq t_0, X)$. Thus, $\Pr(V|T(s) \geq t_0, X)$ is the same for every $s \geq t_0$. \square

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

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

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

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

Proof:

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

The first equality follows from the application of equation (5) to each term of the left-hand side of the first line. By Proposition 1, the distributions over which the expectations are taken in the second line are the same for any $t_1 \geq t_0$ and are

equal to the distribution of $[V|T(s) \geq t_0, X]$. This explains the second equality. As a result,

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

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

The $\text{ATTS}(t_0, t_1, t_0|X)$ and $\text{ATNTS}(t_0, t_1, t_0|X)$ capture the instantaneous causal effect of exposure to the policy (i.e., the instantaneous causal effect of the treatment) at elapsed durations t_0 , compared to when the assigned moment of exposure takes place at a higher duration t_1 . It follows that these measures are identified without any functional-form restriction on the individual hazard rates and without the need to assume independence of 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.

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

Since $\text{ATTS}(t_0, t_1, t_0|X)$ and $\text{ATNTS}(t_0, t_1, t_0|X)$ are equal and do not depend on t_1 as long as $t_1 > t_0$, we may replace them by a short-hand measure $\text{ATS}(t_0|X)$ giving the average instantaneous effect of the policy reform on the survivors with elapsed duration t_0 at the moment of the reform. The effect is measured in deviation from the hazard rate among subpopulations who attained elapsed duration t_0 strictly before the reform. The latter individuals may have acted under the belief that a reform would never take place. Alternatively, they may have acted under the belief that a reform might take place at a later point in time without knowing in advance exactly when (if at all) it would take place. By analogy to the econometric dynamic evaluation literature we may say that this




Figure 3. *Identification based on two cohorts.*

alternative setting allows for an “ex ante” effect of the reform (see e.g. Abbring and Van den Berg, 2003, 2005). Ex ante effects should not be confused with anticipation effects that violate Assumption 2. Ex ante effects are not necessarily incompatible with Assumption 2. However, whether ex ante effects are present or not affects the interpretation of the estimated $ATS(t_0|X)$. If they are present then the estimated $ATS(t_0|X)$ captures the effect of an instantaneous reform exposure versus a reform exposure at some unknown later point in time.

If the reform involves the immediate compulsory participation in some scheme then $ATS(t_0|X)$ also measures the effect of participation on the hazard rate, in deviation from the hazard rate that applies in absence of the scheme. In case of ex ante effects, $ATS(t_0|X)$ measures the effect of participation now versus participation at some unknown later point in time. If the reform does not involve the immediate compulsory participation in some scheme but rather involves the universal arrival of new information then the interpretation is essentially analogous. If the arrival of new information concerns the future date of the individual treat-

ment then $\text{ATS}(t_0|X)$ can be said to measure anticipation effects of the future treatment. Again, the presence or absence of ex ante effects determines whether the anticipation is measured in deviation such an ex ante effect.

In the inference, 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, without further assumptions, it is not possible to combine the average treatment effects for different t_0 in order to estimate how the average effect on the hazard changes over time for a given (sub-)population. Dynamic matching estimators share this limitation (see Crépon et al., 2009).

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

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

with $s' \leq t, s$ and $a > 0$.

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

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

We end this subsection with a brief discussion of the identification of other interesting average treatment effects. Clearly, one cannot hope to identify a full

model, that is, the unknown functions $\theta_{T(s)}(t|X, V)$ for all s and the distribution of $V|X$. Now consider average treatment effects on the individual hazard rate $\text{ATTS}(s', s, t|X)$ and $\text{ATNTS}(s', s, t|X)$ if s' is strictly smaller than t and s . In such cases, inference is subject to the same problem as in Subsection 2.2: the dynamic selection between s' and t causes the sub-population with $S = s'$ among the survivors at t to be systematically different from the sub-population with $S = s$ among the survivors at t . This also implies that, without additional exogenous variation in the treatment duration, and without any functional form assumptions, we can not identify 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 the latter shortcoming is averted if the conditional survival probability is the object of interest, by considering effects at time intervals when accumulation effects or delayed responses may kick in.

3 Non-parametric estimation

3.1 Boundary kernel hazard estimation

From Subsection 2.3, the identification of average causal effects of the policy change on the individual hazard rates is based on the comparison of observable hazard rates from different entry cohorts into the state of interest. Each observable hazard rate is trivially identified from the corresponding cohort-specific duration data. It is therefore natural to non-parametrically estimate these hazard rates.

Specifically, we are interested in $\theta_T(t_0|X, S = t_0)$ and $\theta_T(t_0|X, S = t_1)$ for some $t_1 > t_0$. In Subsection 3.2 below we consider alternative estimators based on $\lim_{t \uparrow t_0} \theta_T(t|X, S = t_0)$ and $\theta_T(t_0|X, S \geq t_1)$ for some $t_1 > t_0$. In every case, the relevant estimate concerns the hazard at the boundary t_0 . There is no reason to assume a connection between the shape of the individual hazard rate before the policy change at t_0 and the shape after t_0 , so estimation of the hazard rate at one side of the boundary only uses outcomes from that particular side of the boundary. Standard non-parametric hazard estimators are heavily biased at the boundary point. We therefore apply methods that deal with this. Specifically, we use boundary kernel hazard estimators and local linear kernel smoothing estimators.¹¹ See Andersen et al. (1993) for an introduction to these estimators.

¹¹Most 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

In the remainder of this subsection we discuss the second-order boundary kernel hazard estimator of Müller and Wang (1994) in detail. We use this estimator in the empirical analysis in Section 4. For expositional convenience we restrict attention to hazard estimation at t_0 , and we transform the truncated duration distribution $T|T \geq t_0, X, S$ to the left such that our ultimate interest is in the hazard rate at the boundary 0 when evaluating it from above. Similarly, in the current subsection, we may suppress S in the notation. Since in the empirical analysis we subsume X into V , we do not consider observed explanatory variables X in the current subsection either.¹²

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

We assume that the true hazard rate is twice continuously differentiable in an interval A starting at 0. To explain the kernel estimator, consider first the case in which the bandwidth b is global. We distinguish between the boundary region $B = \{t : 0 \leq t < b\}$ and an interior region I which is adjacent to B (we need not discuss the right boundary of A here). In I , the kernel hazard estimator is the standard Ramlau-Hansen kernel hazard estimator,¹³

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

where K is taken to be the Epanechnikov kernel,

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

is often only evaluated at interior points.

¹²If X is exogenously time-varying on $(0, t_0)$ in a specific way across cohorts (e.g., as in a linear time trend), then this may cause the “treatment” and “control” hazard rates to have common determinants, but we do not pursue this here.

¹³This smoothes the increments of the Nelson-Aalen estimator $\Lambda_n(t)$ of the integrated hazard based on a random sample of n subjects,

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

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

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

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

where $q \in [0, 1]$ and $z \in [-1, q]$. The latter implies that the support of the boundary kernel does not extend beyond the left boundary. Müller and Wang (1994) plot $K(q, z)$ as a function of z for various values of q . As expected $K(1, z)$ is again the Epanechnikov kernel. As q decreases, the kernel becomes more and more skewed, and the weight assigned to values close to the boundary increases strongly. At the left boundary, q equals zero, so our estimator of $\theta(0)$ equals

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

with

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

There is a positive probability that the above $\tilde{\theta}(0)$ is negative, in which case it is replaced by zero.

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

1994). In Appendix 2 we present the algorithm, slightly modified in response to our experiences regarding its performance.

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

Instead of boundary kernel approaches, one may use local linear smoothing (or local linear fitting, or locally weighted least squares) as a non-parametric approach to deal with estimation at a boundary. Wang (2005) gives an intuitive overview of local linear hazard rate estimation, while Nielsen and Tanggaard (2001), Jiang and Doksum (2003) and Bagkavos and Patil (2008) provide details.¹⁴ The asymptotic properties of the estimator are qualitatively identical to those of the boundary kernel hazard estimator. Jiang and Doksum (2003) compare both methods with data-adaptive local bandwidths, in some Monte-Carlo simulation experiments. Both methods give similar results and both perform very well at the boundary, where their relative ranking depends on the shape of the true hazard rate.

The results of this subsection can be straightforwardly applied for inference on the difference of two independently estimated hazard rates. In Appendix 1.2 we discuss inference of the ratio of two independently estimated hazard rates.

3.2 Implementation issues

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

¹⁴Local linear estimation of hazard rates is related to fixed design non-parametric regression.

(i) **The “comparison” cohort(s).** In Subsection 3.1, we used a boundary-corrected estimator for the observed hazard $\theta_T(t_0|X, S = t_1)$ at t_0 in the cohort that is eventually exposed to the reform at some elapsed duration $t_1 > t_0$. Instead, one may use a standard kernel (or local linear or local constant) hazard estimator, if one is prepared to assume that this hazard is smooth in an interval around t_0 , since then the estimation concerns the interior of an interval on which the hazard is smooth. Whether this assumption makes sense depends on the setting at hand. At certain elapsed durations t_0 of interest, the eligibility to other policy measures may change, causing the individual hazard rates $\theta_{T(s)}(t|X, V)$ to be discontinuous at $t = t_0$ for all s . The application in Section 4 is a case in point. To rule out that this affects the estimated effects, one needs to resort to boundary correction methods.

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

Note that one may widen the “control group” and increase the precision of the estimates of interest, by estimating $\theta_T(t_0|X, t_2 > S \geq t_1)$ with $t_0 < t_1 < t_2 \leq \infty$, instead of $\theta_T(t_0|X, S = t_1)$. This does come at a price, namely that Assumption 1, ruling out the absence of cohort effects, needs to be extended to multiple comparison cohorts flowing in at or before $\tau^* - t_1$. Recall that we require unobserved cohort effects to be absent, since otherwise $S \not\perp V|X$ so that Assumption 1 is violated. Observable 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 “control group”, one may use the availability of multiple potential comparison cohorts in order to select the most similar cohort (or set of cohorts) among the cohorts flowing in before $\tau^* - t_0$. We do not observe the distribution of $V|X$ in a cohort, but we observe outcomes that are informative on it, namely the duration distribution on the duration interval $[0, \tau^*)$ in the corresponding cohort. As a selection mechanism, one may match on the survival probability in the cohort at duration τ^* , or, even stronger, on the shape of the duration distribution in the cohort on the duration interval $[0, \tau^*)$. The more similar this shape, the more similar the composition of survivors at the duration τ^* .

If one comparison cohort is to be selected, then one may consider a cohort that flowed in only marginally earlier than the “treated” cohort, following the line of thought that unobserved changes of the entry composition of the cohorts are a smooth function of the moment of entry. However, such a choice of t_1

being almost equal to t_0 has a practical disadvantage. To see this, notice that $\theta_T(t|X, S = t_1)$ may display a discontinuity at t_1 , so the value $\theta_T(t_0|X, S = t_1)$ at the elapsed duration $t_0 < t_1$ can only be estimated from observed realized durations in an interval to the right of t_0 that does not stretch beyond t_1 . Spells in the comparison cohort with durations exceeding t_1 should be treated as right-censored at t_1 . Consequently, the measure of realized duration outcomes that is informative on $\theta_T(t_0|X, S = t_1)$ is very small if t_1 is only marginally larger than t_0 .

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

(iii) Discrete time. Now let us reconsider the continuous nature of the duration variable. Sometimes a continuous-time analysis may be unfeasible. For example, the data may be time-aggregated in the sense that events are recorded in time intervals (e.g. unemployment duration is collected in months even though individuals may 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 may 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. We examine this empirically in Section 4 below.¹⁵

(iv) Reduced-form model estimation. In econometric duration analysis it has been common to view identification results as a justification for the estimation of parameterized reduced-form model specifications, following the line of thought that identification that does not rely on functional-form assumptions entails that the estimates are also not fundamentally driven by functional-form assumptions (see Van den Berg, 2001). This way of reasoning can also be applied to the results in Section 2, by specifying parameterized models to estimate the objects of interest. An obvious choice is to estimate separate Proportional Hazard (PH) models by whether $\tau_0 + t \geq \tau^*$, using all available cohorts, including in each case calendar time as a time-varying covariate.

Indeed, one may go one step further and specify a single PH model for the full distribution of $T|X, S$ with S being reduced to a simple time-varying covariate $I(t \geq \tau^* - \tau_0)$,

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

where α is the parameter of interest. Such a model may be regarded as a simple representation of a distribution that is generated by an underlying model for individual hazard rates with unobserved heterogeneity. Clearly, the estimates may be affected by the PH model structure which may be too restrictive. If we abstract from that issue then the estimated coefficient of S in this model captures the policy effect of interest. So, the coefficient in a model without unobserved heterogeneity is estimated correctly even if in reality there is unobserved heterogeneity. It follows that estimates obtained under the assumption of no unobserved heterogeneity may also be valid without this assumption. Obviously, it is essential that the analysis uses data that include spells that are ongoing at the moment of the policy change.

Hall and Hartman (2010) provide an example of a study in which a PH model is estimated using spells interrupted by a policy change. Specifically, they estimate a PH model for the transition rate from unemployment into sickness absence as

¹⁵In a discrete-time setting it is possible to identify anticipation effects under certain assumptions if sufficiently rich instrumental variables are available; see Heckman and Navarro (2007).

a function of the sickness benefit policy regime, using unemployment spells that cover a date at which a policy regime change was implemented. They find that a reduced cap for sickness benefits lowers the transition rate to sickness absence by about 35% in the treated population. In their study, they also estimate MPH-type model extensions that allow for unobserved heterogeneity as proportional fixed effects in the individual hazard rates, exploiting the fact that the data contain multiple unemployment spells for many subjects. Interestingly, they find that the estimated policy effect is virtually identical to that in the main analysis, suggesting that, indeed, this coefficient is estimated correctly even when ignoring unobserved heterogeneity.¹⁶

(v) Dynamic treatment evaluation. The results of Section 2 can be applied to dynamic treatment evaluation settings. In such settings, the exposure to a treatment is not necessarily due to some institutional change at a fixed point in time. Rather, different individuals in the same cohort are exposed to a treatment at different elapsed durations, where the treatment may affect the individual hazard from the moment of exposure onwards. Typically, S is only observed if $S \leq T$. It is interesting to reassess Assumptions 1 and 2 in such settings. Abbring and Van den Berg (2003) demonstrate that Assumptions 1 and 2 are fundamental in the following sense: any causal model is observationally equivalent to a model in which these two assumptions are satisfied. If one of the assumptions is relaxed then point-identification requires additional structure.

Notice that Assumption 1 entails that the treatment is exogenous conditional on X . Assumption 2 rules out that survivors at $T = t$ who are not yet treated at t use information on their assigned future treatment date that is unobserved to the researcher. Recall from Subsection 2.3 that our approach allows for ex ante effects. If S varies at the individual level e.g. because of discretionary behavior of the case worker assigning the treatment then it is likely that ex ante effects exist. In that case, if Assumption 2 is satisfied, $ATS(t_0|X)$ involves a comparison of treatment at t_0 versus treatment after t_0 where the latter treatment date is unknown in advance but subjects may use a probability distribution of the future treatment date given X . The ex ante effect of the program then affects the hazard in the comparison group.

We repeat that the arrival of information about a future treatment date can be cast as the actual treatment of interest, provided that this information is also

¹⁶Such a finding is in marked contrast to the literature on single-spell duration analysis with time-invariant covariates, where ignoring unobserved heterogeneity typically leads to attenuation of covariate effects (see Van den Berg, 2001, for an overview).

observed by the researcher. This allows for inference on anticipatory effects of the future treatment (see e.g. Crépon et al., 2013).

4 Empirical application

4.1 The New Deal for Young People: policy regime and treatment

The New Deal has been the flagship welfare-to-work program in the UK since the late 1990s. Over the years 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 at least 6 months. Participation was compulsory upon reaching 6 months in the claimant count. Refusal to participate was punished by a temporary benefits withdrawal.

Since in the UK the entitlement to UB is neither time-limited nor dependent on past working history, and since 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,¹⁷ the job search assistance treatment was split into three stages. It comprised a first period of up to 4 months of intensive job search assistance, with fortnightly 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: (i) subsidized employment, (ii) full-time education or training, (iii) working in an organization in the voluntary sector and (iv) working in an environment-focused organization. Participation in one of these four options was compulsory for individuals having completed 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. The latter was called the Follow Through. If perceived beneficial to the worker, repeated participation in the four alternative options could be arranged.¹⁸

¹⁷Throughout the section we use “enrollment” to denote actual mandatory participation in the job search assistance program and subsequent programs. As we shall see, actual participation may start strictly later than the moment at which the NDYP was introduced.

¹⁸More details on the program can be found in White and Knight (2002), Podivinsky and

The NDYP treated millions of people before being replaced by another program in 2009, the Flexible New Deal. To give an impression of the size of the NDYP, over 2006, 172,000 new participants entered the NDYP, and the average number of participants at any month during that year was 93,000. 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 (DWP (Department for Work and Pensions), 2006). However, a large proportion of this concerns UB that would be due independently of the program, for as long as individuals remain unemployed.

4.2 The introduction of the policy

The NDYP was released nation-wide on April 1, 1998. This corresponds to the reform date τ^* in the framework of Section 2. The existing stock of those who were unemployed for at least 6 months at τ^* was gradually moved into the program. At τ^* , only those whose elapsed unemployment duration was an integer multiple of 6 months were enrolled. (Enrollment occurred during the intensive job-focused interviews scheduled every 6 months within unemployment spells.) If the elapsed duration was not an integer multiple of 6 months, then the individual was enrolled at the moment that his or her elapsed duration attained an integer multiple of 6 months, provided that he or she was not yet 25 years old at that point in time. In the empirical analysis we do not exploit the age eligibility criterion except for robustness checks.

Figure 4 depicts the enrollment scheme in the years around τ^* . Obviously, this scheme is somewhat more complicated than the scheme in Subsection 2.3. However, it allows for the identification and non-parametric estimation of average causal effects of the arrival of information about the new policy regime, for all elapsed durations $t \geq 0$. For elapsed durations of 6 months as well as integer multiples of 6 months, this translates into the effects of enrolling in the NDYP. Similarly, for elapsed durations $t < 6$, the average causal effect translates into the effect of learning that program enrollment is scheduled to take place in $6-t$ months. For example, individuals who are unemployed for 1 month at τ^* learn that if they stay unemployed for another 5 months they will have to enroll, so they anticipate future participation and may react in advance. These are anticipation effects that are identified because we observe the moment of arrival of information.

Thus, we infer the average causal effects of receiving information about the moment of the future enrollment in the job search assistance program, evaluated

McVicar (2002), Blundell et al. (2004), Van Reenen (2004), or Dorsett (2006).




Figure 4. *Introduction of NDYP.* (Note: in the grey area, individuals are enrolled in the job search assistance program.)

at specific elapsed durations $t < 6$. Our methodology allows us to measure such anticipatory responses. In other words, we identify average anticipation effects of the job search assistance treatment on the individual hazard rate. To avoid misunderstandings, recall that this approach is not in conflict with Assumption 2, and we still require a no-anticipation assumption. Specifically, individuals are not allowed to anticipate the moment at which the information arrives.¹⁹

Notice that it is not possible to identify effects of the actual participation in the NDYP at 6 months among those who are unemployed for less than 6 months at τ^* . Since such individuals may act upon the information about the future treatment throughout the time interval between τ^* and the moment that $t = 6$, the dynamic selection from τ^* up to $t = 6$ will differ from that in any other

¹⁹Blundell et al. (2004) study anticipation of the reform (and hence of the moment of the corresponding information arrival) on April 1, 1998, by exploiting spatial and age discontinuities. No significant anticipatory effects are found.

cohort.

Two minor institutional features constitute deviations from the above description of the introduction of the policy. First, individuals with $t < 6$ at τ^* can try to apply for early enrollment, especially if they are disadvantaged (disabled, former convict or lacking basic skills). However, such applications seem unlikely to have been prevalent at the initial stages of the program, when information about NDYP details was still limited. Secondly, the NDYP was introduced in a few small pilot areas on January 1, 1998, i.e., three months before the national release. We use the data from these areas and shift calendar time with 3 months when combining these data with the data from the rest of the country. Since the pilot study did not receive massive attention before April 1, and the evaluation of the pilot was not completed on April 1, we feel that the risk of knowledge spillovers from the pilot areas to the rest of the country (and the ensuing violation of Assumption 2) is small.

4.3 Data

The data are 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. Therefore, in what follows, the outcome of interest is “all exits from the claimant count”, independently of destination.

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

4.4 The choice of treatment and comparison groups

Consider the $ATS(t_0)$ for t_0 equal to 6 months, or more precisely, 182 days. The estimation relies on comparing the survivors among the cohort attaining the elapsed duration t_0 at τ^* (which we call the treatment group, or the treated)

with a similar sample of survivors from earlier cohorts (the comparison groups). Following Subsection 3.2 we do not pursue estimation of ATS conditional on covariates X .

The continuous-time framework must be reconciled with the requirement of a positive sample size. In practice, we need samples of cohorts flowing into unemployment within time intervals rather than at two singular points in time. To proceed, instead of restricting attention to those individuals reaching 6 months of unemployment on a particular calendar day, we consider a full monthly cohort. For instance, the treated sample includes all spells starting in October 1997 (or July 1997 in pilot areas), lasting for at least 6 months.

It is important to realize that the usage of such a fixed inflow time interval may not be innocuous. In particular, those who started a spell towards the end of October 1997 will have had some weeks to react to the new information becoming available on April 1st, 1998, before they actually enroll. By analogy to the discussion in Subsection 4.2 on the effect of receiving information at an elapsed duration of say 1 month about enrollment at 6 months, such information may have an immediate impact on job search behavior before t_0 in the treated sample. This may lead to biased inference. The bias should be negligible if anticipatory effects within intervals of at most a few weeks are much smaller than the effect of actual participation in the NDYP. We argue below that the distortion may lead to an under-estimation of $ATS(t_0)$ at 6 months.²⁰

We define comparison groups in an analogous way, selecting individuals reaching 182 days in unemployment over an entire calendar month prior to April 1998. As candidate groups we consider the cohorts flowing in during June 1997 (pilot areas) and September 1997 (non-pilot areas), or May 1997 and August 1997, or July 1996 and October 1996, or the combination of June and September 1997 with July and October 1996. For simplicity, we designate each cohort by the month of inflow in non-pilot areas as these represent a larger proportion of the population. However, we include data on both pilot and non-pilot regions in all that follows. In accordance to Subsection 3.2, different candidate groups are assessed based on two types of outcomes: the distribution of T on days 1 to 181, and the distribution of observed characteristics X among survivors at 182 days.

Figure 5 displays the survival functions for the treatment and comparison groups up to 181 days into unemployment, prior to the release of NDYP at

²⁰Clearly, it is preferable to apply an estimator in which the inflow time interval shrinks as the sample size increases, such that observations from cohorts close to the inflow date of interest are given more weight. Given our modest sample sizes we do not pursue such an approach, and we leave this as a topic for further research.

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




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

Table 1 compares the empirical distributions of observed covariates among the survivors in the treatment and comparison groups. The September 1997 cohort displays no discernible differences to the treatment group (column 1 in the table).

The combined cohort does not perform as well, with systematic differences in the history of unemployment up to three years prior to inflow (column 4 in the table).

We conclude that the empirical evidence favors the September 1997 cohort, and we henceforth discard the other candidate comparison groups. In any case, it turns out that the estimation results are robust to the choice of the comparison group (estimates available upon request). In total, the sample size of individuals completing 182 days in unemployment during March and April 1998 while aged 20 to 24 is 902. This is almost equally split between the treatment (April 1998) and comparison (March 1998) groups.

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

	Comparison cohort			
	September 97 (1)	August 97 (2)	October 96 (3)	Sep97 + Oct96 (4)
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: The treatment group is the October 1997 cohort. The variables in rows 5 to 7 describe the UB claiming history in the 3 years preceding inflow into current unemployment spells. Numbers in bold highlight statistically significant differences in the distribution of the covariate, at the 5% level.

Individuals in the September 1997 comparison cohort do not enroll before an elapsed duration of 12 months, as they are past the 6 months threshold at the time of the reform τ^* . Shortly after τ^* , at elapsed durations exceeding 6 months, the behavior of some members of this group may be affected by the information becoming available at τ^* , possibly confounding the estimate of $ATS(t_0)$. This source of bias can be simply eliminated by right-censoring spells in the comparison group at τ^* . 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 in the comparison group is a long distance away. We check the

sensitivity of the results, and it turns out that they are robust to whether right-censoring is adopted or not (estimates available upon request). The robustness to the choice of the comparison group also suggests that this source of bias may be irrelevant in the analysis.

As noted above, we also consider anticipatory effects of the actual arrival of information about the NDYP reform among those approaching enrollment. These are interesting per se. But they are also informative on the accuracy of estimates of the impact of program participation that ignore anticipation, by exposing the extent to which anticipation affects the composition of the treatment group prior to participation.

We may estimate the anticipatory effects of future enrollment evaluated at each duration t_0 shorter than 6 months (182 days). We take individuals who are aged 20 to 24 at the moment they reach 6 months of unemployment. One should be careful with the terminology, since now “treatment” means exposure to information about future enrollment while “comparison” means the absence of such exposure. Thus, the treatment and comparison groups are now defined in reference to whether they are exposed to the information arrival at the reform date τ^* . For a given $t_0 < 6$ months, the treatment group consists of individuals who reach t_0 during April 1998.

Like above, we need to select appropriate comparison groups. Clearly, each duration t_0 requires its own comparison group. Adopting the selection procedure used above would then involve a very large number of comparisons. We therefore follow a slightly simpler procedure. Recall that the procedure used above for the enrollment effects led to the choice of a comparison group that flowed in one month before the treatment group. In accordance to this, we choose comparison groups of individuals who flowed in one month before $\tau^* - t_0$ and who hence reach the duration $t_0 < 6$ during March 1998. It remains to be seen whether these comparison groups meet the checks for aligned dynamic selection up to t_0 .

Figure 6 displays the survival functions up to t_0 for the treatment and comparison groups for each of four values of t_0 (namely, 2, 3, 4 and 5 months). There are some signs of differential selection, apparently due to conditions in December/January. In later cohorts, which cross December/January earlier in their spells (panels B and C), the survival functions diverge throughout the duration interval $(0, t_0)$, especially at the end of the period, when approaching April (treatment) or March (comparison groups). Post December/January cohorts (panel A), unaffected by conditions in those months, exhibit very similar survival functions. Earlier cohorts (panel D) are also affected but return quickly to a common path. The latter finding echoes the observed patterns for the October and September

cohorts in Figure 5.

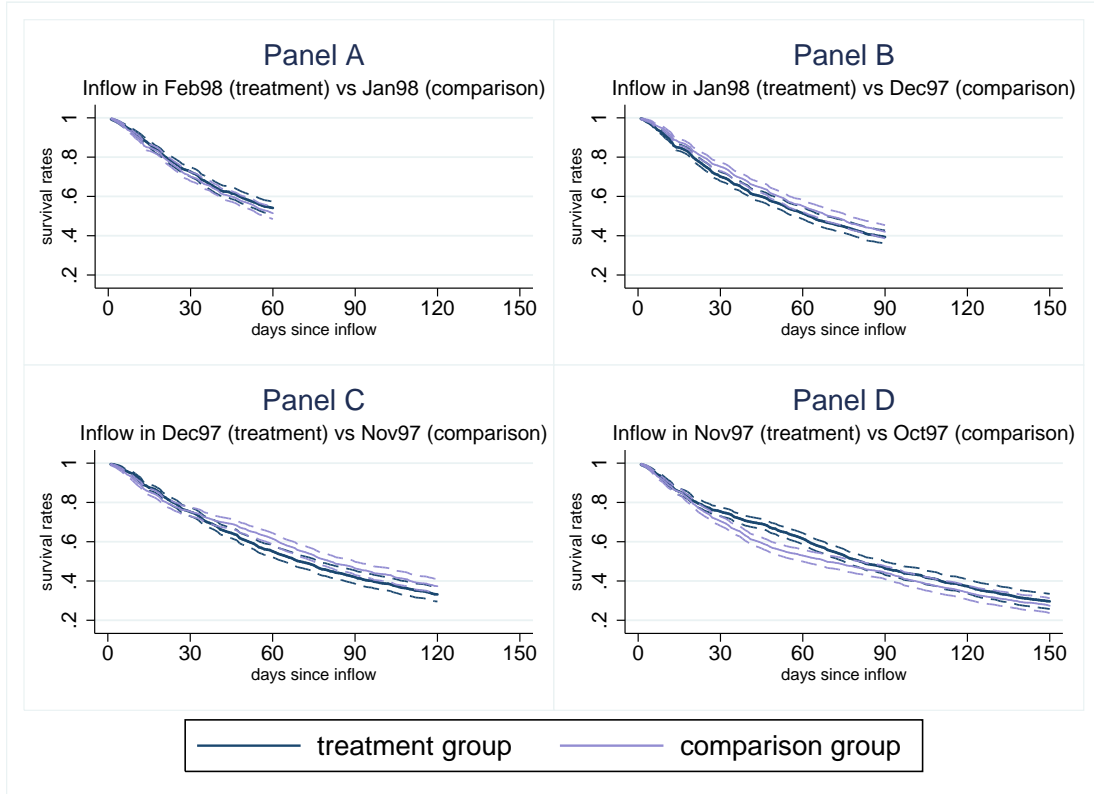


Figure 6. Empirical survival functions among cohorts reaching durations of 2 (panel A) to 5 (panel D) months in April 1998 (treatment group) or March 1998 (comparison group).

Note: Dashed lines represent 95% confidence intervals.

Table 2 compares the empirical distributions of observed covariates among the survivors in the treatment and comparison groups. Column 2 shows that the December 1997 and January 1998 cohorts are compositionally different upon having reached 3 months in unemployment. For earlier cohorts, the absence of statistically significant differences further supports their comparability (columns 3 and 4). In the light of these findings, our analysis of anticipatory effects focuses on durations from 4 up to and including 5 months. Note that one may expect anticipatory effects to be larger at durations close to enrollment at 6 months than at lower durations.

In contrast to our earlier discussion on the estimation of enrollment effects, right-censoring shortly at calendar time τ^* is expected to be important for the estimation results of anticipatory effects. First, the comparison group will itself

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

		Month of inflow			
(1)	treatment group	Feb 98	Jan 98	Dec 97	Nov 97
(2)	comparison group	Jan 98	Dec 97	Nov 97	Oct 97
(3)	elapsed duration t_0	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: Row 1 (2) details the inflow date of the treatment (comparison) group for the evaluation of the effect at the elapsed duration in row 3. The variables in rows 8 to 10 describe the UB claiming history in the 3 years preceding inflow into current unemployment spells. Numbers in bold highlight statistically significant differences in the distribution of the covariate, at the 5% level.

be subject to the information arrival on April 1st, 1998 (i.e., shortly after the moment at which the treatment group receives the information), and they may react to this in-between $t_0 + 1$ and 6 months. This may be particularly relevant if t_0 itself is close to 6 months. Secondly, the treatment group will enroll into job search assistance upon 6 months of unemployment, with potential causal effects on their hazard rate from that moment onwards. We examine these two issues in Subsection 4.6.

In all cases, we estimate effects in discrete as well as in continuous time, by varying the length of the time unit. Estimates in discrete time capture effects on aggregate monthly conditional transition probabilities while estimates in continuous time do the same for daily hazard rates. Both sets of estimates are based on the same treated and comparison samples.

4.5 Results for the average causal effect of NDYP program enrollment

Table 3 presents the main estimates in the discrete-time setting. In particular, the estimate in column (1) is the estimated $ATS(t_0)$ capturing the average causal additive effect of enrolling into the NDYP program at an elapsed unemployment duration of $t_0 = 6$ months on the conditional probability of leaving unemployment within one month. We find that the latter probability increases by 4.5 percentage points in the first month after enrollment, which is quite substantial. This estimate is significantly positive at the 5% level. The corresponding relative increase in the conditional probability is about 35%. The fact that the enrollment in NDYP has a positive effect on employment is in line with the results in Blundell et al. (2004) who use conditional difference-in-differences and before-after observations to evaluate the NDYP.

The entry in column (2) of Table 3 is the $ATS(t_0)$ estimate obtained *as if* the policy reform had taken place on April 1, 1997 instead of April 1, 1998, using treatment and comparison groups based on cohorts flowing into unemployment one year before those used for the actual $ATS(t_0)$. Similarly, the entry in column (3) is the $ATS(t_0)$ estimate obtained *as if* the NDYP reform on April 1, 1998 was designed for 25-29 year olds, using treatment and comparison groups based on the corresponding contemporaneous age cohorts. These two exercises can be interpreted as placebo analyses. If the methodology is appropriate then the estimates in columns (2) and (3) should be insignificantly different from zero. Alternatively, in the presence of seasonal effects in the inflow composition, the entry in column (2) reflects this. Similarly, in the presence of macro-economic labor market changes around April 1998, the entry in column (3) reflects this. The same applies in case of substitution or crowding out effects across age categories in response to the NDYP. Note that this involves the use of the age discontinuity in eligibility. As it turns out, neither of the two estimates is statistically significant. This confirms the validity of our approach and facilitates the interpretation of our main estimates.

Figure 7 displays the continuous-time counterparts of the estimated treatment effect $ATS(t_0)$ of Table 3, using the Müller and Wang estimator with optimal local bandwidths. We display both average additive and average multiplicative effects, together with 95% confidence intervals based on the analytic asymptotic variance without bias correction.²¹ Although $t_0 = 182$ days is the minimum elapsed un-

²¹With bootstrapping we obtain virtually the same intervals. The estimated optimal local bandwidth for the additive effect at the boundary of 182 days is 80 days with a standard error

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

	Treatment effect	Placebo effects	
	20-24 years	20-24 years	25-29 years
age			
τ^*	April 1, 1998	April 1, 1997	April 1, 1998
	(1)	(2)	(3)
estimate	.045	.014	-.009
standard error	(.023)	(.022)	.021
# individuals	911	1118	1365

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

employment duration for enrollment into job search assistance, it is conceivable that program participation requires a positive amount 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. For this reason, Figure 7 shows estimates at elapsed durations from 182 to 212 days. A zero effect in the early stage of the enrollment period implies that the dynamic selection process does not differ between the treatment and comparison groups in this stage, so that the $ATS(t_0)$ is also identified at the value of t_0 at the end of this early stage.

The patterns in the estimation results are almost identical whether these concern average additive or average multiplicative effects.²² We therefore discuss the former only. The main interest is in the result at the first duration beyond 182 days at which the additive effect is significant. Any features after that may be due to duration dependence of the treatment effect or to differential dynamic selection, or both. We find significant effects of enrollment only after about a week into the program. At that moment, the estimated effect as a function of the elapsed duration jumps rather abruptly to a positive level of about 0.006 per day. This amounts to more than doubling the hazard rate in the absence of NDYP, as can be seen from the figure for the multiplicative effect. The estimated effect then drops to a lower positive level that just misses the 95% significance level. However, at that stage we can no longer separate causal and confounding

of 30 days.

²²Recall that inference on the multiplicative effect warrants an additional assumption (Assumption 3 in Appendix 1.1) whereas inference on the additive effect does not.




Figure 7. *Non-parametric continuous-time estimation of the average causal effect $ATS(t_0)$ of enrolling into the NDYP program at the elapsed unemployment duration of $t_0 = 6$ months on the hazard rate of leaving unemployment.*

Note: Dashed lines represent 95% confidence intervals.

compositional effects. We conclude that among those who enter the new policy regime at 6 months of unemployment duration, the program has a significant and sizeable positive average causal effect on the hazard rate at 6 months.

4.6 Results for the average causal effect of receiving information about future enrollment

By analogy to the previous subsection, Table 4 presents the discrete-time effects $ATS(t_0)$ of receiving information at durations $t_0 < 6$ months about enrollment at 6 months. In particular, the estimates in column (1) capture the average causal additive effects at 4 months and 5 months. Both are negative, but none is significantly different from zero. This suggests that individuals do not react in advance to the prospect of future job search assistance in the NDYP. In other words, the information does not induce them to modify their behavior, and in that sense they do not anticipate the future enrollment. This is a useful finding. It sheds

insight into the behavior of the unemployed individuals. In addition, it means that estimates of the impact of job search assistance that use methodologies that ignore anticipation are not subject to bias due to anticipation effects.

The estimates in columns (2) and (3) are obtained using the same approach as for the corresponding columns of Table 3. The estimates are also statistically zero, except for the 20–24 year olds 5 months after inflow in the case in which no actual reform had taken place. The latter estimate is based on a comparison of those who flowed in in November 1996 to those who flowed in in October 1996. This might reflect a seasonal effect in the composition of the inflow of the 20–24 year olds, in the following specific sense: the inflow in November contains more individuals who return to work in April of the subsequent year than the inflow in October contains individuals who return to work in March. Recall from Figure 6 that we did not find a differential survival probability between the October and November cohorts upon reaching 5 months of unemployment, so any effect of the month of inflow can only become visible during the fifth month. Under this scenario, the estimate of $ATS(t_0)$ at $t_0 = 5$ months is the net result of a positive seasonal effect and a negative anticipatory effect. The latter suggests that individuals in the month before enrollment in job search assistance prefer to hold out until enrollment. However, it is rather odd that a seasonal effect only occurs when going from October to November and not when going from September to October or from November to December. Moreover, we do not find any evidence of seasonal effects for the 25–29 year olds.

As explained in Subsection 4.4, the estimates in Table 4 may be biased due to the fact that the one-month time interval used to estimate the conditional outflow probability for the comparison group crosses April 1, 1998. In effect, part of the comparison group is exposed to the 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 enrollment. Therefore it makes sense to artificially right-censor spells in the comparison group once they cross April 1, 1998. In addition, we right-censor spells in the treatment group once they enter enrollment, in order to prevent that the estimate is affected by the causal effect of enrollment. As we have seen in Subsection 4.4, the latter causal effect kicks in at an elapsed duration of 189 days, so we use this as the right-censoring value. Recall that we aim to estimate effects conditional on durations of 4 or 5 months, i.e., of 123 to 181 days. For each elapsed duration we require a sufficiently large number of informative spells, and for this reason we restrict the continuous-time analysis to durations below 172 days.

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

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

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

Figure 8 shows the continuous-time estimates of the causal $ATS(t_0)$ effects of the information arrival. For convenience we only display the average additive effects. Notice that the computational burden to produce Figure 8 is much higher than for Figure 7. Figure 7 is based on the estimation of two non-parametric hazard rates. In contrast, in Figure 8, every day in the interval $(123, 172)$ of elapsed durations on the horizontal axis requires a separate estimation of two non-parametric hazard rates. Thus, Figure 8 does not visualize the difference of two hazard rates but rather the values of the differences of two hazard rates each evaluated at the boundary t_0 for a *range* of values of t_0 . In accordance to previous subsections, each value t_0 gives rise to a monthly cohort.

Clearly, the results provide evidence of anticipatory behavior. This behavior leads to a drop of the hazard rate after the beginning of the 5th month and gains importance within 15 days before enrollment into job search assistance. Despite the wide 95% confidence intervals towards the end of the period (due to the bias corrections discussed above), the anticipatory effect is statistically significant at high durations. Such a finding may not be (and, indeed, was not) detected in a discrete-time analysis with a monthly time unit.

This evidence of anticipatory behavior is new. Previous NDYP studies did




Figure 8. *Non-parametric continuous-time estimation of the average causal effect $ATS(t_0)$ of receiving information at elapsed durations t_0 between 123 and 172 days about enrollment at 6 months, on the hazard rate of leaving unemployment.*

Note: Dashed lines represent 95% confidence intervals.

not consider changes in behavior closely before the moment of enrollment.²³

The result has implications for other types of evaluation approaches. It is likely that those who postpone job search until after the enrollment into intensified job search assistance at 6 months are on average more work-prone than those who remain unemployed for 6 months in the absence of the NDYP regime. In that case,

²³De Giorgi (2005) estimates $\Pr(T < 6|X, \text{inflow after Apr98}) - \Pr(T < 6|X, \text{inflow before Oct97})$. This method is only applicable to unconditional survival probabilities. The study reports no significant anticipatory effects. The pivotal study of Black et al. (2003) provides evidence that many unemployed workers in the U.S. dislike being an unemployment insurance claimant if it involves mandatory participation in programs of job search counselling. This dislike causes increased exits out of unemployment in the initial weeks of the spells. In our analysis, those who are unemployed for 5.5 months prefer to wait until the enrollment into the program at 6 months. This suggests that the response to the prospect of future job search assistance differs between short-term unemployed in the U.S. and long-term unemployed youth in the UK. Indeed, it is conceivable that this is the result of a dynamic selection where those who find a job relatively fast anyway leave unemployment very fast and hence are underrepresented in the sub-populations over which our estimated ATS are defined.

a before-after comparison of spells with elapsed durations of 6 months (censoring any spells crossing the reform date) would lead to an upward bias of the effect of the job search assistance.

5 Conclusions

In this paper we have merged regression discontinuity analysis and duration analysis. We have shown that, to study causal policy effects on hazard rates, one may usefully exploit spells crossing the moment of the introduction of the policy, even if the individual hazard rates depend on unobserved covariates. The approach does not need any functional form assumption on the hazard rate or its determinants. This stands in marked contrast to standard duration analysis which has been plagued by proportionality assumptions on the hazard rate, functional form assumptions on the duration dependence and the unobserved heterogeneity distribution, and a “random effects” assumption for observed and unobserved covariates. An additional advantage of the new approach is that it enables policy evaluation shortly after introduction of a new policy. If the outcome of interest is a duration variable then a comparison of pre and post reform outcomes requires an observation window stretching far beyond the reform date.

Our analysis shows that implementation details of the introduction of a new policy for subjects in a certain state have important consequences for the quality and timing of evaluation exercises regarding the hazard rate out of the state. A policy that applies to all subjects in the state at the time of the reform alleviates the need for strong identifying assumptions and supports the early production of evaluation results on the hazard rate. Conversely, a policy reform that applies only to new entrants into the state will have to deal with differential dynamic selection and possibly with differential selection at inflow once the new regime is announced, and wait for at least t periods before the impact at duration t can be evaluated. An example could be the introduction of a new active labor market policy for the unemployed with the objective to increase the reemployment rate. In the interest of evidence-based policy design, it is recommended to include the current unemployed in the target population exposed to the new policy.

Policy reforms in which subjects may choose between staying in the old regime and switching to the new regime cannot be evaluated with our approach. One may envisage an extension that incorporates this type of setting without having to compromise on the above-mentioned advantages of our approach. It is an obvious topic for further research to pursue this.

Our approach is also suitable to study the causal effect of the arrival of in-

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

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

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

well as on the propensity score.

References

- Abbring, J.H. and G.J. van den Berg (2003), “The non-parametric identification of treatment effects in duration models”, *Econometrica*, 71, 1491–1517.
- Abbring, J.H. and G.J. van den Berg (2005), “Social experiments and instrumental variables with duration outcomes”, Working paper, 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.
- Black, D.A., J.A. Smith, M.C. Berger and B.J. Noel (2003), “Is the threat of reemployment services more effective than the services themselves? Evidence from random assignment in the UI system”, *American Economic Review*, 93, 1313–1327.
- Blundell, R., M. Costa Dias, C. Meghir, and J. Van Reenen (2004), “Evaluating the employment impact of a mandatory job search program”, *Journal of the European Economic Association*, 2, 569–606.
- 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.
- Crépon, B., M. Ferracci, G. Jolivet, and G.J. van den Berg (2013), “Dynamic treatment evaluation using data on information shocks”, Working paper, University of Bristol.
- 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.
- Hall, C. and L. Hartman (2010), “Moral hazard among the sick and unemployed: evidence from a Swedish social insurance reform”, *Empirical Economics*, 39, 27–50.
- Härdle, W. (1994), *Applied Nonparametric Regression*, Cambridge University Press.
- Härdle, W., M. Müller, S. Sperlich and A. Werwatz (2004), *Nonparametric and Semiparametric Models*, Springer, Berlin.
- 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.
- Podivinsky, J.M. and D. McVicar (2002), “Unemployment duration before and after New Deal”, Working paper, Northern Ireland Economic Research Centre.
- Porter, J. (2003), “Estimation in the regression discontinuity model”, Working paper, Harvard University.
- Ridder, G. (1984), “The distribution of single-spell duration data”, in: G.R. Neumann and N. Westergård-Nielsen, eds., *Studies in labor market analysis*, Springer Verlag, Berlin.
- Rubin, D. (2000), “Discussion of causal inference without counterfactuals”, *Journal of the American Statistical Association*, 95, 435–438.
- 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.
- Vikström, J. (2014), “IPW estimation and related estimators for evaluation of active labor market policies in a dynamic setting”, Working paper, IFAU, Uppsala.
- Wang, J.L. (2005), “Smoothing hazard rates”, in: P. Armitage and T. Colton, eds., *Encyclopedia of Biostatistics*, Wiley, Chichester.
- White, M. and G. Knight (2002), “Benchmarking the effectiveness of NDYP”, Working paper, National Audit Office, London.

Appendix 1. Average multiplicative effects on individual hazard rates

A1.1 Identification

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

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

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

Assumption 3 (Multiplicative unobserved heterogeneity).

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

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

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

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

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

A1.2 Inference

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

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

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

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

to ratio estimators. First, we may perform bootstrapping simultaneously on both samples. Secondly, we may apply the delta method. If, following Tu (2007), we again ignore the asymptotic biases, then we obtain that the estimator $\tilde{r}(0)$ has an asymptotically normal distribution with mean $r(0)$ and variance

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

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

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

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

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

Step 1 Choose initial value of bandwidth and construct grids

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

3. If computation time is important and, as a consequence, \tilde{T} is sparse, construct a second, finer, equidistant grid for duration variable T in the domain $A = [0, R]$ to estimate the hazard functions. Call it $\tilde{\tilde{T}} = \{\tilde{\tilde{t}}_1, \dots, \tilde{\tilde{t}}_P\}$, where $P > M$.
4. Construct an equidistant grid for bandwidth b in $[b, \bar{b}]$, call it $\tilde{B} = \{\tilde{b}_1, \dots, \tilde{b}_L\}$. Müller and Wang (1994) propose using $\underline{b} = 2b_0/3$ and $\bar{b} = 4b_0$. In the empirical analysis in Section 4, this interval is too tight as the optimal choice often coincides with its boundaries. Therefore we use $[b, \bar{b}] = [b_0/6, 6b_0]$.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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