Instrumental Variable Estimation of Treatment Effects for Duration Outcomes

Govert E. Bijwaard*

Econometric Institute

Erasmus University Rotterdam

Econometric Institute Report EI 2007-20

*Erasmus University Econometric Institute, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands; Phone: (+31) 10 40 81424; Fax: (+31) 10 40 89162; E-mail: bijwaard@few.eur.nl. This research is financially supported by the Netherlands Organization for Scientific Research (NWO) no. 451-04-011. I am very grateful to Geert Ridder for his comments and suggestions. I also benefited from the comments of participants at the 2003 Extending the Tinbergen Heritage conference in Rotterdam, at the 2003 ESEM Meeting, at the 14^{th} EC^2 conference and, at the Econometric Evaluation of Public Policies : Methods and Applications in Paris 2005. I thank Steve Woodbury and the W. E. Upjohn Institute for making the data available to me.

Instrumental Variable Estimation of Treatment Effects for Duration Outcomes

Abstract

In this article we propose and implement an instrumental variable estimation procedure to obtain treatment effects on duration outcomes. The method can handle the typical complications that arise with duration data of time-varying treatment and censoring. The treatment effect we define is in terms of shifting the quantiles of the outcome distribution based on the Generalized Accelerated Failure Time (GAFT) model. The GAFT model encompasses two competing approaches to duration data; the (Mixed) Proportional Hazard (MPH) model and the Accelerated Failure Time (AFT) model. We discuss the large sample properties of the proposed Instrumental Variable Linear Rank (IVLR), and show how we can, with one additional step, improve upon its efficiency. We discuss the empirical implementation of the estimator and apply it to the Illinois re-employment bonus experiment.

JEL classification: C21, C41, J64.
Key words: Treatment effect; Duration model; Censoring; Instrumental Variable.

1 Introduction

In recent years, social experiments have gained popularity as a method for evaluating social and labor market programs (see e.g. Meyer (1995), Heckman et al. (1999) and Angrist and Krueger (1999)). In experiments the assignment of individuals to the treatment can be manipulated. If assignment is random, the average impact of the treatment can be estimated. However, a randomized assignment may be compromised, if the individuals can refuse to participate, either by dropping out, if they are to receive the treatment, or by obtaining the treatment, if they are in the control group. If this non–compliance to the assigned treatment is correlated with the outcome variable, i.e. is selective, then the difference of the average outcomes is a biased estimate of the average effect of the treatment.

Most of the evaluation literature has focused on static treatments, *i.e.* treatment that is administered at a particular point in time or in a particular time interval. If the outcome is a duration the treatment or its effect can be dynamic, *i.e.* it can be switched on and off over time. A further complication with duration data is that the durations are often censored. Examples of evaluation studies based on duration outcomes are Ham and LaLonde (1996), Meyer (1995), Van den Berg et al. (2004), Ashenfelter et al. (2005) and Bijwaard and Ridder (2005). These studies show that simple comparisons between average durations in the various treatment regimes are confounded by dynamic selection. We specify a causal model for duration outcomes that can handle censoring and time-varying treatments. The treatment effect in this model is defined in terms of shifting the quantiles of the survival rate.

Two competing approaches to the estimation of the effect of a, possible time–varying, covariate on survival has been the (Mixed) Proportional Hazard (MPH) model (for a recent survey see Van den Berg (2001)) and the Accelerated Failure Time (AFT) model (see a.o. Kalbfleisch and Prentice (2002), Brännäs (1992), and Klein and Moeschberger (1997)). In the Mixed Proportional Hazards (MPH) model the hazard is written as the product of the baseline hazard, a non-negative regression function, and a non-negative random variable that represents the covariates that are omitted from the regression function. Ridder (1990) introduced the Generalized Accelerated Failure Time (GAFT) model, a generalization of the AFT models that also includes the MPH model.

The GAFT model is based on transforming the duration and assuming some distribution for this transformed duration. The transformation is the integrated hazard of a PH model. The

AFT model is obtained by restricting the transformation but it does not restrict the distribution of the transformed duration, while the MPH model restricts the this distribution to a mixed of exponentials. The GAFT model is also related to the generalized regression model of Han (1987). The regression coefficients in a GAFT model can be interpreted in terms of the effect of regressing on the quantiles of the distribution of the transformed duration for the reference individual. In an AFT model the relation between the quantile of a individual with observed characteristics X and the quantile of the reference individual is the acceleration factor. In a GAFT this acceleration factor is multiplied by the ratio of the duration dependence at the two quantile durations. We use the GAFT model to construct a causal model for duration outcomes.

Our approach to obtain the treatment effect in a duration model is related to rank estimation procedures of Robins and Tsiatis (1991) and of Bijwaard and Ridder (2005). The procedure of Bijwaard and Ridder (2005) is based on a semi-parametric MPH and requires preliminary estimates of the baseline hazard. The procedure of Robins and Tsiatis (1991) is based on the strong version of the Accelerated Failure Time model. Their model imposes a strong non-interaction assumption. This implies that if two individuals have the identical observed durations and observed treatment histories then they would have had identical durations had treatment always been withheld.

We propose an instrumental variable procedure for duration models that does neither impose the non-interaction assumption nor requires full compliance in the control group. It enables us to estimate the treatment effect in terms of changes in the quantiles. This is accomplished by using a GAFT model that transforms the duration such that for the true, population, parameter the transformed duration is independent of the instrument. Thus, for a binary instrument, R, the proportion of the people with $R = 1$ remains the same over the survivors on the transformed duration time. This implies that the rank test statistic for the instrument on the relevant time scale, i.e. the transformed duration time scale, should go to zero. The inverse of this rank test,or a weighted version, provides the estimation procedure we call the Instrumental Variable Linear Rank (IVLR) estimation.

The existence of an endogenous treatment implies (possible) dependence between the transformed duration and the censoring time. This implies that the IVLR estimator, which exploits the independence between the transformed durations and the instruments, may give biased results. However, with a small adjustment we can still apply the IVLR approach, based upon the assumption that the (potential) censoring time is known at the start of the study. For example, in our empirical data, we only observed the unemployed while receiving UI benefits. In these data the potential censoring time for all individuals is at 26 weeks, the maximum duration of receiving UI benefits. With known (potential) censoring time we can modify the GAFT transformation by introducing additional censoring such that this modified transformation and the instruments become independent for the uncensored observations. Then, the IVLR estimator on this modified transformation leads to consistent estimators.

The IVLR estimation is based on a vector of mean restrictions on weight functions of the covariates, instrument and the transformed durations. Thus the IVLR is also related to GMM estimation. In GMM estimation it is feasible to get the most efficient, optimal, GMM estimator in just two steps. At the first step a GMM estimator is obtained using suboptimal, but feasible, weighting matrices. From this consistent estimator the optimal weighting matrices can be estimated consistently. With these estimated matrices we can obtain an efficient estimate of the parameters involved in just one additional step. A similar reasoning applies to the IVLR-estimator. At the first step we use feasible weighting functions to obtain consistent estimates of the parameters of the GAFT model. From these parameters we estimate the distribution of the transformed durations, which is needed to calculate the most efficient weighting functions. Then, in just one additional step the efficient IVLR estimator is obtained.

Our approach identifies the treatment effect in terms of the effect of the quantiles of the transformed duration. It is therefore not surprising that our approach is also related to quantile regression. Koenker and Geling (2001) and Koenker and Bilias (2001) use quantile survival regression to identify the treatment effect as the difference between the two marginal survival distributions of treated and untreated individuals. It is, however, unclear how these methods can handle time-varying treatment.

Various other papers consider methods for Instrumental Variable estimation for duration models that are related to the approach of this paper. We mentioned already the methods developed by Robins and Tsiatis (1991) and by Bijwaard and Ridder (2005). Recently the application of the timing-of-events approach, introduced by Abbring and van den Berg (2003) has emerged. This method identified the causal effect on the hazard rate in a MPH model. The important identifying assumption is that the time to the treatment randomly varies between the individuals. Thus, when the timing of the treatment follows a predescribed protocol, as is the case for the bonus data, the timing-of-events method cannot be used to obtain the treatment effect. Abbring and van den Berg (2005) show that for a binary time-invariant treatment the (fixed) treatment effect on the hazard rate in an MPH model is identified. The treatment effect can then consistently be estimated based on the behaviour of very short durations. The practical use of this estimator is limited because very short durations are often ill-recorded. Chesher (2002) considers an MPH-type model with an endogenous continuous treatment as well as exogenous variables, a continuous instrument, and a latent variable relating the treatment and the instrument.

For our empirical application we use data from the Illinois unemployment bonus experiment. These data have been analysed before with increasing sophistication by Woodbury and Spiegelman (1987), Meyer (1996) and Bijwaard and Ridder (2005). In this experiment a group of individuals who became unemployed during four months in 1984 were divided at random in three groups of about equal size: two bonus groups and a control group. The unemployed in the claimant bonus group qualified for a cash bonus if they found a job within 11 weeks and would hold this job for at least four months. In the employer bonus group, the bonus was paid to their employer. The members of the two bonus groups were asked whether they were prepared to participate in the experiment. About 15% of the claimant bonus and 35% of the employer bonus groups refused participation. The reason for this refusal is unknown, and it is very likely that the decision to participate is related to the unemployment duration. This makes the indicator of being eligible for a bonus an endogenous treatment variable.

The outline of the article is as follows. Section 2 provides the definition of treatment effect in the Generalized Accelerated Failure Time model. We also discuss identification. Section 3 provides the intuition behind the Instrumental Variable Linear Rank estimator that is introduced in Section 4. We prove the consistency and asymptotic normality of the estimator and discuss the efficiency and the practical implementation of the IVLR. In Section 5 we apply the IVLR estimator to estimate the treatment effect of a re-employment bonus based on the Illinois bonus experiment. We conclude with a summary of the proposed procedure and empirical results, and discuss avenues for further research.

2 Treatment effects in Duration Models

We consider the population of individuals flowing into unemployment, and the durations these individuals subsequently spend in unemployment. Each individual is either assigned to a prescribed treatment protocol or to a control group. Thus, we know for each individual when the treatment is turned on or off. We are interested in the causal effect of the treatment on the unemployment duration.

It can rarely be defended that a study on unemployment durations includes all the relevant characteristics of the individuals looking for a job. For example, consider a study on the effect of a re-employment bonus on finding a job within a certain period. Because such a bonus increases the reward of leaving unemployment it gives an incentive to search more intensively and therefore it is expected to increase the re–employment hazard. However, the search intensity of the unemployed individuals is usually not observed.

Suppose that the unemployed have to choose at the start of their unemployment spell whether they want to be eligible for a bonus. If they choose to be eligible they have to fill in some forms, notify their new employer and provide a proof that they held that job for at least 4 months. Thus, joining the bonus program implies some administrative duties and cooperation with their new employer for the unemployed. This might refrain some individuals from joining (as we see in our application). It is very likely that the unobserved motivation to return to work has an impact on both the decision to be eligible for a bonus and the search intensity. This makes the indicator of joining the bonus program an endogenous variable. Without adjusting for this (self)-selection standard duration analysis give biased results of the effect of the bonus on unemployment duration.

A way of adjusting for an endogenous treatment is the conventional instrumental variable assumptions of instrument–error independence and an exclusion restriction. A familiar example of an instrumental variable is a randomly assigned treatment for a policy in which the actual treatment still depends on a decision by the agents (or on decisions made by those who execute the program). For instance, long–term unemployed can be randomly assigned to a training program, but for many programs they can still decide not to join, or the training manager can decide to withhold some training from some people. Then, the assignment indicator is an instrument for the actual indicator of training received.

Complications arise if the outcome variable of interest is a duration variable, like the unemployment duration. Models for duration data are usually non-linear. In general the value of the treatment variable may depend on information that accumulates during the evolution of the duration. The common approach to accommodate such time-varying variables is to relate them to the hazard rate or survival rate. Another reason not to consider the effect on the mean is that duration data are usually (right)-censored, due to a limited observation window. The hazard rate and survival rate are invariant to censoring. We shall define the treatment effect in terms of a change in the quantiles of the distribution based on a Generalized Accelerated Failure Time Model.

2.1 The Generalized Accelerated Failure Time Model

The key variables in duration analysis are the duration till the next event, T, and the indicator of censoring, δ. The observed durations may be right–censored, i.e. we observe $\tilde{T} = \min(T, C)$ with C the censoring time. The possible time–varying treatment indicators are given by the vector $D_i(t)$ where i refers to a member of the population. The path of the treatment is predetermined. Thus $\overline{D}(t) = \{D(s); 0 \leq s \leq t\}$ does not depend on future, after t events.

Two competing approaches for the analysis of duration data has been the (Mixed) Proportional Hazard (MPH) model and the Accelerated Failure Time (AFT) model. The MPH model assumes that the hazard, the instantaneous probability of an event at duration t, given that no event occurred before t can be written as

$$\lambda(t|\overline{D}(t), V) = v\lambda_0(s;\alpha)e^{\gamma D(t)}$$

where $\lambda_0(t)$ represents the baseline hazard, that is, the duration dependence of the hazard common to all individuals. The treatment affect, $e^{\gamma' D(s)}$, and the unobserved heterogeneity ,V, enter the hazard proportionally. The integrated hazard, $\Lambda_v(t) = \int_0^t \lambda(s|\overline{D}(s), V)\, ds = v\int_0^t \lambda_0(s;\alpha)e^{\gamma D(s)}\, ds$ is unit exponentially distributed, $\mathcal{E}(1)$ and therefore the MPH model is also equivalent to

$$\int_0^t \lambda_0(s;\alpha)e^{\gamma D(s)}\, \mathrm{d}s \sim \frac{\mathcal{E}(1)}{V} \tag{1}$$

Thus the distribution of $\Lambda(t) = \int_0^t \lambda_0(s;\alpha)e^{\gamma' D(s)}\, \mathrm{d}s$ is a mixture of exponential distributions.

The AFT model assumes that the survival function of an individual with covariate path $\overline{D}(t)$ is related to the survival function of the untreated individual, the individual with $\overline{D}(t) = 0$, by

$$S\big(t|\overline{D}(t)\big) = S_0\Big(\int_0^t e^{\gamma' D(s)}\, \mathrm{d}s\Big)$$

with $S_0(\cdot)$ is the survival function of the untreated individual. Let T_0 be the duration of an untreated individual. If the treatment is time-invariant then an AFT model implies that the distribution of T_D, the duration of an individual with treatment D, and the distribution of $e^{-\gamma' D}T_0$ are the same. Thus, the treatment D accelerates, $\gamma < 0$, or decelerates, $\gamma > 0$ the duration. This is equivalent with a linear regression model for the log-duration

$$\ln(T) = -\gamma' D + \epsilon$$

where ϵ is a random variable with same distribution as $\ln(T_0)$.

A class of duration models that generalizes the AFT models in such a way that it also includes the MPH models is the Generalized Accelerated Failure Time (GAFT) model, introduced by Ridder (1990). The GAFT model is not specified by the distribution of the log-duration. Instead, we transform the duration, and assume that this transformed duration has some distribution. The transformation of the duration is related to the integrated hazard in a PH-model. The GAFT model is related to the generalized regression model proposed by Han (1987).

The GAFT model assumes that the relation between the duration T and the treatment is specified as

$$\int_0^T \lambda(s;\alpha)e^{\psi(s,D(s),\gamma)}\,\mathrm{d}s = U = h\big(T,\overline{D}(T),\theta\big) \tag{2}$$

where $\lambda(t;\alpha)$ is a non–negative function on $[0,\infty)$ and $\psi(t,D(t),\gamma)$ is the treatment function with treatment coefficient (vector) γ and $\theta=(\alpha',\gamma')'$. For the true population parameter-vector θ_0 we have $U_0 = h\big(T,\overline{D}(T),\theta_0\big)$. We assume that the treatment is a binary variable that only changes at predescribed durations and that the treatment effect may change over the duration. Then a flexible functional form for ψ is

$$\psi(t,D(t),\gamma) = \sum_{j=0}^{J} \gamma_j \cdot D_j \cdot I_j(t) \tag{3}$$

where $I_j(t) = I(t_j < t \le t_{j+1})$ with $t_0 = 0$ and $t_{J+1} = \infty$.

A flexible functional form for λ is the piecewise constant function. The GAFT model is characterized by these functions and by the distribution of the non–negative random variable U_0. We denote the survivor function of U_0 by $\overline{G}_0(u)$, which does not depend on the treatment, and its hazard function by $\kappa_0(u)$. We assume that the distribution of U_0 is absolutely continuous. The semi–parametric estimators considered in this article avoid assumptions on the distribution of U_0.

As mentioned, the GAFT model contains as special cases the AFT, the PH and the MPH models. The AFT model restricts the transformation to $\lambda(t;\alpha)\equiv 1$, but leaves the distribution of U_0 unrestricted (with the exception of that U_0 should be non–negative, see e.g. Cox and Oakes (1984)). The (M)PH model restricts the distribution of U_0, but leaves the λ unrestricted (non–negative). The distribution of U_0 is an unit exponential distribution (PH) or a mixture of exponential distributions (MPH). A convenient assumption is that the unobserved heterogeneity has a gamma distribution with variance σ^2. Then, U_0 has a Burr distribution with density $g_0(u) = (1+u\sigma^2)^{-1-1/\sigma^2}$, mean $1/(1-\sigma^2)$ and hazard function $(1+u\sigma^2)^{-1}$.

2.2 Counterfactual analysis

We model the causal effect of the treatment on the unemployment duration using the potential-outcome framework, as formulated by Holland (1986) for static interventions, and extended by Robins (1986) to time-varying interventions.

Equation (2) implicitly defines the outcome, *i.e.* duration, distributions for the treatment and control groups. We indicate these regimes by the subscripts (1) and (0), respectively. The outcome distributions for the two regimes for $T^*_{(1)}, T^*_{(0)}$ are given by

$$U^*_{(0)} = \int_0^{T^*_{(0)}} \lambda(s;\alpha_0)\, ds \tag{4}$$

$$U^*_{(1)} = \int_0^{T^*_{(1)}} \lambda(s;\alpha_0) e^{\psi(s,1;\gamma_0)} ds \tag{5}$$

The fundamental assumption that assures identification is that $U^*_{(0)}$ and $U^*_{(1)}$ have the same distribution. They may be dependent, but the joint distribution is only identified in special cases and not of particular interest. A special case is that $U^*_{(1)} \equiv U^*_{(0)}$. In that case there is a deterministic relation between the outcomes in the two regimes and,

$$\int_0^{T^*_{(0)}} \lambda(s;\alpha_0) ds = \int_0^{T^*_{(1)}} \lambda(s;\alpha_0) e^{\psi(s,1;\gamma_0)} ds \tag{6}$$

This model corresponds, but is not identical to the Rank Preserving Structural Failure Time (RPSFT) model of Robins and Tsiatis (1991). In the RPSFT model $T^*_{(0)i}$ is the potential duration of individual i, if she/he was never to receive the treatment. Thus, Robins and Tsiatis think of $T^*_{(0)i}$ is the baseline duration for the never treated, as a pre–treatment characteristic of the individual. An example of an RPSFT model is to assume that

$$T^*_{(0)} = \int_0^{T_i} e^{\psi(s,D_i(s);\gamma_0)} ds \tag{7}$$

The outcome model defined in (4) and (5) is also rank preserving. If we consider two individuals with potential outcomes $T^*_{(0)i}, T^*_{(1)i},\ i = 1, 2$, then $T^*_{(0)1} > T^*_{(0)2}$ implies $T^*_{(1)1} > T^*_{(1)2}$ and vice versa. However,

$$T^*_{(0)} \neq \int_0^{T^*_{(1)}} \lambda(s;\alpha_0) e^{\psi(s,D(s);\gamma_0)} ds \Big|_{\psi(t,D(t);\gamma_0)\equiv 0} \tag{8}$$

and $T^*_{(0)}$ cannot be interpreted as the latent baseline duration. An unattractive feature of the RPSFT model is that it implies that individuals with the same outcome in the treatment regime have the same outcome in the no treatment regime, and the other way round. This equality of counterfactuals is a strong assumption.

Under the weaker assumption that $U^*_{(0)}$ and $U^*_{(1)}$ have the same distribution, the relation in (6) still holds but now for the quantiles of the outcome distribution. Thus we can interpret the treatment effect in terms of regressing on the $U^*_{(0)}$ quantiles. Let $t_q(0)$ and $t_q(1)$ be the quantiles of $U^*_{(0)}$ and $U^*_{(1)}$ then,

$$\overline{G}_0\Big(h\big(t_q(1), 1; \theta_0\big)\Big) = 1 - q = \overline{G}_0\big(h\big(t_q(0), 0; \theta_0\big)\big) \tag{9}$$

Hence we obtain a relation between $t_q(1)$ and $t_q(0)$ that is defined implicitly by

$$\int_0^{t_q(1)} \lambda(s; \alpha_0) e^{\psi(s,1;\gamma_0)}\, \mathrm{d}s = \int_0^{t_q(0)} \lambda(s; \alpha_0)\, ds \tag{10}$$

This implies that

$$\frac{\mathrm{d}t_q(1)}{\mathrm{d}t_q(0)} = e^{-\psi\left(t_q(1), D(t_q(1)); \gamma_0\right)} \times \frac{\lambda(t_q(0); \alpha_0)}{\lambda(t_q(1); \alpha_0)} \tag{11}$$

In the MPH model we can interpret $\lambda(t)$ as the baseline hazard, i.e. the factor in the proportional hazard that captures the (duration) time variation in the hazard function. Thus, in the MPH model the ratio in (11) can be interpreted as the ratio of baseline hazards.

Note that if the treatment is time-invariant and has a fixed coefficient, $\psi(t, D, \gamma_0) = \gamma_0 D$ and $\gamma_0 > 0$ then an AFT model would overestimate the treatment effect if $\lambda\big(t_q(0); \alpha\big) < \lambda\big(t_q(1); \alpha\big)$ and underestimate the treatment effect if $\lambda\big(t_q(0); \alpha\big) > \lambda\big(t_q(1); \alpha\big)$. Hence failure to correct for variation in λ may bias the treatment coefficient.

It may be instructive to give an example. The hazard or re–employment rate of unemployment durations often exhibits a spike just before the time that unemployment benefits are exhausted. This corresponds to a large increase in λ at that spike. Now let us assume that all unemployed who find a job receive a bonus data and that the bonus has a small impact on the job finding rate, so that $t_q(0) > t_q(1)$. If $t_q(0)$ is in the spike while $t_q(1)$ is not, then the left–hand side of (11) is greater than one and the AFT treatment effect at $t_q(1)$ is negative. We conclude that if there is (substantial) variation in λ, the uncorrected treatment effects that can be directly estimated from the marginal outcome distributions are misleading.

2.3 Identification of the GAFT model

In the GAFT model with an exogenous treatment and a log–linear regression function, the model is characterized by the non–negative function $\lambda(t; \alpha)$ defined on $[0, \infty)$, the distribution of U_0 and the regression parameter γ. Ridder (1990) has shown that if the covariates are time constant, all observationally equivalent GAFT models, i.e. models that give the same conditional distribution of T given D, have regression parameters $c_1\gamma$, integrated transformation $c_2\Big(\int_0^t \lambda_0(s;\alpha)\,ds\Big)^{c_1}$ and U_0 distribution $G_0\Big(\big(\frac{u}{c_2}\big)^{1/c_1}\Big)$ for some constants $c_1, c_2 > 0$. The equivalent class follows from the fact that a GAFT model with time constant covariates can be expressed as a transformation model

$$\ln\Big(\int_0^T \lambda_0(s;\alpha)\,\mathrm{d}s\Big) \stackrel{d}{=} -\gamma D + \ln U,$$

and the constants c_1, c_2 correspond to addition of e^{c_2} to and division by c_1 of the left– and right–hand sides.

With time–varying covariates, the set of observationally equivalent GAFT models is generally smaller. In particular, the power transformation that gives an observationally equivalent model if the covariates are time constant, in general does not result in a GAFT model. As an example consider the GAFT model with time–varying regressors that differ between two groups. In group I

$$D(t) = \begin{cases} 1 & \text{if } 0 \le t \le 1, \\ 0 & \text{if } t > 1. \end{cases}$$

and in group II, $D(t) = 0;\ t \ge 0$. Moreover $\lambda_0(t;\alpha) = \alpha t^{\alpha-1}$. With time constant regressors the parameter α is not identified. It can be shown that the observationally equivalent GAFT models have transformation $c_2 t^{\alpha}$ and U_0–distribution with c.d.f. $G_u\big(\frac{u}{c_2}\big)$. Hence, with time–varying covariates α is identified (and so is γ).

We conclude that identification depends on whether the covariates are time constant or time–varying. If the covariates are time constant we can identify the transformation $h(T, \overline{D}(T);\theta)$ up to a power and γ up to scale (with the power and the scale being equal). Moreover, if we fix the power we can identify $h(T, \overline{D}(T);\theta)^{c_1}$ up to scale and the distribution of U up to the same scale parameter.

If the covariates are time–varying, we can, except in special cases, identify $h(T, \overline{D}(T);\theta)$ and the distribution of U_0 up to a common scale parameter. Because we leave the distribution of U_0 unspecified in our estimation method, we can not use restrictions on this distribution to find the scale parameter. For that reason we normalize $h(T, \overline{D}(T);\theta)$ by setting $h(T, 0;\theta) = 1$ for some $t_0 > 0$. With time constant regressors we need the same normalisation, but in addition we need to set one regression coefficient equal to one. Of course, we could choose a class of transformations that is not closed under the power transformation. This amounts to identification by functional form.

Finally, we need a condition on the sample paths of D in the population. If we rewrite (2) as

$$\int_0^T e^{\ln \lambda(s;\alpha)+\psi(s,D(s),\gamma)} \mathrm{d}s = U \tag{12}$$

we require that

$$\Pr\Big(\ln \lambda(s;\alpha) + \ln \psi(s, D(s), \gamma) = 0\Big) = 0 \tag{13}$$

where the probability is computed over the distribution of D as a random function of t and 0 is the zero function. In other words, $\ln \lambda$ is not collinear with D.

The identification issues in the GAFT model with endogenous treatment are the same as in the GAFT model with only exogenous variables, except that we need additional assumptions on the instrument. First, the instrument should only affect the duration through the treatment and not directly. Second, the value of the instrument should influence the value of the treatment in a non-trivial way, for example if both the instrument and the treatment are binary $\Pr(D = 1|R = 1) > 0$ and $\Pr(D = 0|R = 0) > 0$.

3 Instrumental Variable estimation for Duration data

If treatment choice is exogenous to the duration, standard techniques for the analysis of survival time data can be used to estimate the treatment coefficient. For example, we can use a Mixed Proportional Hazards model and estimate the treatment coefficient using (semi–parametric) Maximum Likelihood procedures, depending on the assumptions we make about the distribution of the unobserved heterogeneity, V, and the baseline hazard. If the model is correctly specified the MLE yields a consistent estimate. However, standard MLE will give biased estimation results if the treatment choice is endogenous.

3.1 Common solutions to endogeneity in duration models

Consider, for example, a randomized trail with selective compliance to the assigned time-constant treatment. Let $R = 0, 1$ be the random assignment and $D = 0, 1$ the actual treatment. Since physical randomization implies that at time zero all attributes of the two groups are (in expectation) identical, a commonly used solution to the endogeneity problem, is to ignore the post–randomization compliance and rely on the analysis of the treatment assignment groups. This intention–to–treat analysis replaces the actual treatment D, by the treatment assignment indicator R, in the estimation procedure. Further, if the model is correctly specified the estimated treatment coefficients will correspond to the overall effect that would be realized in the whole population, under the assumption that the compliance rate and the factors influencing compliance in the sample are identical to those that would occur in the whole population.

The major drawback of the intention–to–treat analysis is that the estimated effect is a mixture of the population effect and the effect on the compliance. Hence, if the treatment effectively raises the re–employment hazard, the intention–to–treat measure of this effect will diminish as non–compliance increases. Another disadvantage is that compliance is very likely to depend on the perceived effects of the treatment. If, for example, the unemployed know that being eligible for a re–employment bonus does not stigmatize them, they will be more prone to participate. Thus, when the pattern of compliance is a function of the perceived efficacy of the treatment the estimated intention–to–treat will not represent the overall effect of the treatment had it been adopted in the whole population.

Recently an alternative approach introduced by Abbring and van den Berg (2003) to the analysis of treatment effects in duration models has emerged. This timing-of-events method identified the causal effect on the hazard rate in a MPH model. The important identifying assumption is that the time to the treatment randomly varies between the individuals. Thus, when the timing of the treatment follows a predescribed protocol, as is the case for the bonus data, the timing-of-events method cannot be used to obtain the treatment effect.

Abbring and van den Berg (2005) show that for a binary time-invariant treatment the (fixed) treatment effect on the hazard rate in an MPH model is identified. In the MPH model the hazard rate is $\lambda(t|D,V) = v\lambda_0(s;\alpha)e^{\gamma' D}$. The treatment effect γ can then consistently be estimated based on the behaviour of very short durations. The practical use of this estimator is limited because very short durations are often ill-recorded.

When a continuous instrument is available more general models of the treatment effect are identified. Chesher (2002) considers an MPH-type model with an endogenous continuous treatment D as well as exogenous variables X, a continuous instrument R, and a latent variable relating D and R. He demonstrates local identification of ratios of the derivatives of the individual hazard rate with respect to D and R. Abbring and van den Berg (2005) prove that for an MPH model with perfect compliance in the control group and a binary treatment a time-varying treatment effect that possibly depends on X is identified.

Our approach to obtain the treatment effect in a duration model is related to rank estimation procedures of Robins and Tsiatis (1991) and of Bijwaard and Ridder (2005). The procedure of Bijwaard and Ridder (2005) is based on a semi-parametric MPH and requires preliminary estimates of the baseline hazard. If there is full compliance in the control group these preliminary estimates can be obtained from the control sample. These estimates are substituted in the second stage estimation equation. Their 2 Stage Linear Rank Estimator (2SLR) is a device to reduce the computational burden by dividing the computation into two steps. This is appealing because in the 2SLR is the solution to a discontinuous estimating function. However, the first stage of the 2SLR estimator is based on a MPH model for the control group and can only be estimated if the individuals in that group do not have any possibility to get a treatment.

As discussed in Section 2.2, the procedure of Robins and Tsiatis (1991) is based on the strong version of the Accelerated Failure Time model. Their model imposes a strong non-interaction assumption. This implies that if two individuals have the identical observed durations and observed treatment histories then they would have had identical durations had treatment always been withheld.

We propose an Instrumental Variable (IV) procedure for duration models that does neither impose the non-interaction assumption nor requires full compliance in the control group. It enables us to estimate the treatment effect in terms of change in the quantiles, as defined in (11). This is accomplished by using a GAFT model that transforms the duration such that for the true, population, parameter the transformed duration is independent of the instrument.

The intuition behind the idea of transforming can be clarified by considering a simple example. Suppose we have data from an experiment on a treatment that may reduce unemployment duration. The assignment to treatment is random, but the compliance to the assigned treatment is selective. If the treatment has no impact on the re-employment hazard the probability of observing an individual with assigned to treatment, $R = 1$, among the survivors at some unemployment duration t is equal to the treatment assignment probability at the start, $\Pr(R = 1)$. However, if the treatment has an effect, say positive, on the hazard this does not hold, because in that case the treated individuals will find a job faster. But the distributions of the counterfactual transformed durations $U^*_{(0)}$ and $U^*_{(1)}$ are both independent of the treatment assignment. Thus, on the transformed duration time-scale, using the true, population, parameters, $U_0 = h\big(T, \overline{D}(T), \theta_0\big)$, in the GAFT model in (2) the proportion of the individuals in assignment group remains the same

$$\Pr\big(R = 1 \mid U_0 \geq u\big) = \Pr\big(R = 1 \mid T \geq 0\big), \tag{14}$$

In other words: the basic assumption underlying the IV estimator is that for the right GAFT model the distribution of the duration on the transformed time scale is independent of the instrument. This implies that the hazard of the *population* transformed duration is independent of the instrument. This independence only holds for the population parameters and therefore we can build an estimation procedure that exploits this conditional independence[1]. Independence of R and the hazard rate of U_0 implies (infinitely) many orthogonality restrictions. For instance, if we assume independence of the hazard of U_0 and R for $u \leq C_u$, we can choose C_u so that it corresponds to a specific quantile of the distribution of U_0. In section 4 we explore this moment conditions that lead to the Instrumental Variable Linear Rank (IVLR) estimator.

3.2 Censoring and endogenous treatment in GAFT

A common feature of duration data is that some of the observations are censored. Assume the censoring time, C, is (potentially) known. For example, in the analysis of unemployment duration based on administrative data the duration is often only observed while the individual receives unemployment benefits. Usually, the maximum duration of receiving benefits is based on the labor market history of the individual and is recorded in the data. Then, the potential censoring time is known and the observed durations are $\tilde{T} = \min(T, C)$ and $\Delta = I(T \leq C)$, where Δ is one if T is observed.

One is tempted to define the censored transformed durations by the minimum of the transformed time till (potential) censoring and the transformed time till the event occurs, $\tilde{U}(\theta) = \min\big(h(T;\theta), h(C;\theta)\big) = h(\tilde{T};\theta)$. However, with an endogenous treatment censoring makes some of the orthogonality conditions implied by the independence of the instrument and the transformed durations fail to hold. This can be illustrated by a simple example: Consider a fixed censoring time, all individuals have the same maximum duration of receiving benefits. Then for all individuals, irrespective of their treatment status, censoring occurs at time C. Suppose that treatment is time-invariant and has a constant treatment function, $\psi(\cdot) = \gamma D$. Finally, we assume that except for the treatment parameter, γ, all the parameters of λ are known. Hence

$$U_0 = e^{\gamma_0 D} \Lambda_0(T) \tag{15}$$

with $\Lambda_0(t) = \int_0^t \lambda(s, \alpha_0)\, \mathrm{d}s$. Hence, if $D = 0$ censoring in the transformed time occurs at $\Lambda_0(C)$,

[1]Here we only concentrate on a static binary instrument and a discrete, but possible time-varying according to a prescribed protocol, treatment. It is not difficult to extend the analysis to more, discrete, levels of both the instrument and the treatment and to have a sequential instruments.

but if $D = 1$ censoring occurs at $e^{\gamma_0}\Lambda_0(C)$. Thus, if $\gamma_0 > 0$, then all transformed durations in the interval $[\Lambda_0(C), e^{\gamma_0}\Lambda_0(C)]$ have $D = 1$, i.e belong to the treatment group (for $\gamma_0 < 0$ the boundaries are reversed). The hazard of U_0 on this interval clearly depends on D and hence on R. The independence of the hazard of U_0 and R only holds up to the lower bound of the interval. This implies that in the IVLR, which exploits this independence, the transformed durations that fall in the problematic interval have to be censored.

This can be generalized to a model with a time–varying treatment parameter and/or a time–varying treatment. Assume the piecewise constant structure for the treatment function in (3). This implies that for a duration in the interval $(t_k, t_{k+1}]$, the treatment parameter is e^{γ_k}. We define the transformed censoring time $C^U(\theta)$ such that: (a) $T \geq C$ implies $h(T;\theta) \geq C^U(\theta)$ and (b) U_0 and R are independent on the interval bounded above by $C^U(\theta)$.

Note that we either observe $T \leq C$ and $\Delta = 1$, or $T > C$ and $\Delta = 0$. The transformed censoring times (conditional on $T, C > t_k$) that take all these considerations into account are the sum of the transformed duration up to t_k, $h(t_k;\theta)$ and the censoring adjustment, *i.e*

$$C^U(\theta) = \begin{cases} \int_0^C \lambda(s;\alpha)P(s;\gamma)\,\mathrm{d}s & \text{if } T > C, \\ \int_0^T \lambda(s;\alpha)P(s;\gamma)\,\mathrm{d}s + \int_T^C \lambda(s;\alpha)\,ds & \text{if } T \leq C. \end{cases} \tag{16}$$

where $P(s;\gamma) = I(s \geq t_k)\prod_{j=0}^k \min(e^{\gamma_j}, 1)$. From the last term on the right–hand side of (16) we see why we need to know C even for the uncensored observations. Otherwise we can not compute $C^U(\theta)$ for these observations. We can estimate the treatment coefficient γ and the parameters of λ from the following observed data

$$\tilde{U}(\theta) = \min\bigl(U(\theta), C^U(\theta)\bigr), \qquad \Delta^U(\theta) = I\bigl(U(\theta) < C^U(\theta)\bigr)$$

Now $\tilde{U}(\theta_0)$ is independent of R for $\Delta^U(\theta_0) = 1$. Note that if, at least, one of the γ's is different from zero, we introduce extra censoring on the transformed durations, because then some units with $\Delta = 1$ have $\Delta^U(\theta) = 0$.

4 Instrumental Variable Linear Rank Estimation

In this section we introduce the Instrumental Variable Linear Rank (IVLR) estimator of the GAFT model. The estimator is based on independence of the transformed durations $\{\tilde{U}(\theta_0), \Delta^U(\theta_0)\}$ and R. The IVLR is motivated by (14) and is defined on the possibly censored durations $\tilde{U}(\theta)$.

4.1 The IVLR estimator

For the population parameter vector θ_0 the hazard of U_0, $\kappa_0(u)$, is independent of the instrument history up to $h_0^{-1}(u)$. This independence can be used to construct test statistics close to the linear rank test (see Prentice (1978)).[2] The IVLR also exploits this independence and is the estimation procedure derived from these rank tests.

The data is a random sample $\tilde{T}_i, \Delta_i, \overline{D}_i(T_i), R_i, i = 1, \ldots, N$. For some θ this random sample can be transformed to $\tilde{U}_i(\theta), \Delta_i^U(\theta), \overline{D}_i^U\big(\tilde{U}_i(\theta)\big), i = 1, \ldots, N$ and this is the sample that is used in the statistic

$$S_n(\theta; W) = \sum_{i=1}^{n} \Delta_i^U(\theta)\Big\{W\big(\tilde{U}_i(\theta), R_i\big) - \overline{W}\big(\tilde{U}_i(\theta)\big)\Big\} \tag{17}$$

where

$$\overline{W}\big((\tilde{U}_i(\theta)\big) = \frac{\sum_{j=1}^{n} I\big(\tilde{U}_j(\theta) \geq \tilde{U}_i(\theta)\big) W\big(\tilde{U}_i(\theta), R_j\big)}{\sum_{j=1}^{n} I\big(\tilde{U}_j(\theta) \geq \tilde{U}_i(\theta)\big)},$$

the average of the weight function evaluated at $\tilde{U}_i(\theta)$ among the individuals still under observation at that (transformed) duration. Note that we used $\Delta_i^U(\theta)$ instead of Δ_i to assure independence of the instruments and the transformed durations for all uncensored observations. The estimating equation that defines the IVLR estimator contains a left–continuous weight function W. The dimension of W is greater than or equal to the dimension of θ_0. The weight function may depend on $\tilde{U}_i(\theta)$ and R. The variance of the IVLR estimator depends on W and in section 4.2 we discuss the optimal choice of this function.

The interpretation of $S_n(\theta; W)$ is that it compares the weight function for a transformed duration that ends at $\tilde{U}_i(\theta)$ to the average of the weight functions at that transformed duration for those individuals that are still under observation. The independence suggests that the difference of the weight function for individual i and the average weight function for the individuals still under observation is zero at the population parameter value θ_0. In large samples this is correct if we choose for instance $W\big(\tilde{U}_i(\theta), R_i\big) = R_i$, because for $\theta = \theta_0$ the transformed duration U_0 is independent of R_i. For $\theta = \theta_0$ the transformed durations U_{0i} are identically distributed and this implies that the rank statistic is zero in large samples for this choice of W.

The expression for the rank statistic simplifies if we order the observations by increasing transformed duration

$$\tilde{U}_{(1)}(\theta) \leq \tilde{U}_{(2)}(\theta) \leq \ldots \leq \tilde{U}_{(N)}(\theta)$$

[2]Hausman and Woutersen (2005) derive an instrumental variable estimator based on Kendall's rank-correlation for the MPH model.

In the ordered transformed durations we obtain

$$S_N(\theta, W) = \sum_{i=1}^{N} \Delta^U_{(i)}(\theta)\left\{W\big(\tilde{U}_{(i)}(\theta), R_{(i)}\big) - \frac{\sum_{j=i}^{N} W\big(\tilde{U}_{(i)}(\theta), R_{(j)}\big)}{N+i-1}\right\} \tag{18}$$

The function S_N is not continuous in θ. The points of discontinuity are values of θ that make e.g. $\tilde{U}_{(k)}(\theta)$ and $\tilde{U}_{(k+1)}(\theta)$ equal. If $\Delta_{(k)}(\theta) = \Delta_{(k+1)}(\theta) = 1$, the discontinuity is

$$\frac{W\big(\tilde{U}_{(k)}(\theta), R_{(k+1)}\big) - W\big(\tilde{U}_{(k)}(\theta), R_{(k)}\big)}{N-k} \tag{19}$$

and this goes to 0 if N increases. Lemma 1 in the appendix shows that S_N is asymptotically equivalent to a linear (and hence continuous) function in θ.

Thus, the statistic $S_n(\theta; W)$ has mean zero at the population parameters and, therefore, we base our estimator on the roots of $S_n(\theta; W) = 0$. However, the estimating functions are discontinuous, piecewise constant, functions of θ and a solution may not exist. For that reason we define the Instrumental Linear Rank estimator (IVLR) $\hat{\theta}_n(W)$ as the minimizer of the quadratic form, *i.e.*

$$\hat{\theta}_n(W) = \inf\{\theta \mid S_n(\theta; W)' S_n(\theta; W)\} \tag{20}$$

If $S_n(\theta; W)$ were differentiable with respect to θ, then asymptotic normality can be proved using Taylor series expansion in a neighborhood of θ_0. Tsiatis (1990) showed that, if $S_n(\theta; W)$ is not differentiable, as in the current problem, we can still use a linear approximation of $n^{-1/2} S_n(\theta; W)$. Using this approximation and the asymptotic normality of $S_n(\theta_0; W)$, we can show that $\sqrt{n}(\hat{\theta}_n(W) - \theta_0)$ is asymptotically normal. Let $a(u; \theta_0)$ be the probability limit of the average weight function (see assumption C6 in the appendix) evaluated at u, C_0 the transformed censoring time for $\theta = \theta_0$. The asymptotic properties of the IVLR estimator are summarized in the following two theorems that use the following terms: Let $d_{i0}(u)$ the derivative of the hazard of $U(\theta)$ w.r.t. θ evaluated at $\theta = \theta_0$ and $V(u, \theta)$ is the probability limit of the covariance between $d_{i0}(u)$ and $W(u, R_i)$ at u. The assumptions are given in the appendix.

Theorem 1 (Consistency)**.**

If assumptions C1 to C7 hold $\hat{\theta}_n(W)$ converges in probability to θ_0.

Proof: See the appendix.

Theorem 2 (Asymptotic Normality)**.**

If assumptions C1 to C9 hold and $Q(W)$ has full rank, then

$$\sqrt{n}(\hat{\theta}_n(W) - \theta_0) \xrightarrow{d} N\big(0, Q^{-1}(W)\Omega(W)Q'^{-1}(W)\big) \tag{21}$$

where

$$\Omega(W) = \int_0^{C_0} a(u;\theta_0)\kappa_0(u)\, du \tag{22}$$

is the asymptotic variance of $n^{-1/2}S_n(\theta_0;W)$ and,

$$Q(W) = \int_0^{C_0} V(u,\theta_0)\, du \tag{23}$$

the limiting covariance matrix of the processes $W(u,R_i)$ and $d_{i0}(u)/\kappa_0(u)$.

Proof: See the appendix.

4.2 Efficiency of the IVLR estimator

Many different choices of the weight functions lead to consistent estimates of the parameters. By properly choosing the weight function the asymptotic variance of the IVLR can be minimized. Tsiatis (1990) has shown that for the AFT model with exogenous covariates weight functions proportional to $u\kappa_0'(u)/\kappa_0(u)X$ minimize the asymptotic variance of the estimated regression parameters. In general the distribution of U_0 is unknown. This distribution can, however, consistently be estimated from the implied U from any IVLR with a weight function not involving these functionals.

The IVLR estimation is based on a vector of mean restrictions on weight functions of the covariates, instrument and the transformed durations. GMM estimation is also based on moment conditions and in GMM estimation it is feasible to get the most efficient GMM estimator in just two steps. A similar reasoning applies to the IVLR-estimator. This justifies an adaptive construction of an efficient estimator. In the next section we address the practical implementation of an adaptive estimation procedure. First, we introduce the optimal weight function.

Theorem 3 (Optimal weight function in IVLR)**.**

The weight–function that gives the smallest asymptotic variance for $\hat{\theta}_n(W)$ is

$$W_{\text{opt}}(u,R)) \propto \frac{d_{i0}(u)}{\kappa_0(u)} \tag{24}$$

The asymptotic covariance matrix of the optimal IVLR estimator reduces to

$$\Omega^{-1}(W_{\text{opt}}) = Q^{-1}(W_{\text{opt}}). \tag{25}$$

Proof of theorem 3. From

$$\frac{1}{\sqrt{n}}\begin{pmatrix} S_n(\vartheta_0;W) \\ S_n(\vartheta_0;W_{\text{opt}}) \end{pmatrix} \xrightarrow{d} N\left(0, \begin{pmatrix} \Omega(W) & Q(W)' \\ Q(W) & \Omega(W_{\text{opt}}) \end{pmatrix}\right)$$

follows that the matrix

$$Z = \begin{pmatrix} \Omega(W) & Q(W)' \\ Q(W) & \Omega(W_{\text{opt}}) \end{pmatrix}$$

is non–negative definite, the same is true for its inverse. In particular, the submatrices on the main diagonal of the inverse are non-negative definite. Hence the matrix

$$Q^{-1}(W)\Omega(W)Q'^{-1}(W) - \Omega^{-1}(W_{\text{opt}})$$

is a non-negative definite matrix. □

Consider, for example, a GAFT model with a piecewise constant λ function,

$$\lambda(t,\alpha) = \sum_{j=0}^{J} e^{\alpha_j} I(t_j < t \leq t_{j+1}) \tag{26}$$

with $t_0 = 0$ and $t_{L+1} = \infty$ and the hazard on the last interval is normalized to 1, $\alpha_L = 0$. Assume that the model has a constant treatment coefficient then by (24) the optimal weight functions are

$$\begin{aligned}
W_{\text{opt},\alpha_j} &= \left(1 + u\frac{\kappa_0'(u)}{\kappa_0(u)}\right)\cdot\Big(RI_j^1(u) + (1-R)I_j^0(u)\Big) + \qquad (27)\\
&\quad + R\left[\big(1+u\kappa_0(u)\big)\frac{f_0(u|1,R)-f_0(u)}{f_0(u)} + u\frac{f_0'(u|1,R)-f_0(u)}{f_0(u)}\right]I_j^1(u) + \\
&\quad + (1-R)\left[\big(1+u\kappa_0(u)\big)\frac{f_0(u|0,R)-f_0(u)}{f_0(u)} + u\frac{f_0'(u|0,R)-f_0(u)}{f_0(u)}\right]I_j^0(u) \\
W_{\text{opt},\gamma} &= R\left[1 + u\frac{\kappa_0'(u)}{\kappa_0(u)}\right] + \qquad (28)\\
&\quad + R\left[\big(1+u\kappa_0(u)\big)\frac{f_0(u|1,R)-f_0(u)}{f_0(u)} + u\frac{f_0'(u|1,R)-f_0(u)}{f_0(u)}\right]
\end{aligned}$$

where $f_0(u|D,R)$ is the density of U_0 given D and R, $f_0'(\cdot)$ is the derivative of the density and $I_j^D(u) = I\big(m_j(D) < u \leq m_{j+1}(D)\big)$ for

$$m_j(D) = \int_0^{t_j} \lambda(s,\alpha)e^{\gamma D}\,\mathrm{d}s$$

4.3 Estimation in practice

The statistic $S_n(\theta; W)$ is a multi–dimensional step–function. Therefore, the standard Newton–Raphson algorithm cannot be used to solve (20). One of the alternative methods for finding a zero of a non–differentiable function is the Powell-method. This method is a multidimensional version of the Brent algorithm, see Press et al. (1986, §10.5) and Powell (1964).

Related to the computation of optimal weight function is the estimation of the variance matrix for an arbitrary weight function.[3] The difficulty in estimating the covariance matrix lies in the calculation of the matrix $Q(W)$ and not in the calculation of the variance matrix of the estimating equation. The latter can be consistently estimated by

$$\hat{\Omega} = \frac{1}{n} \sum_{i=1}^{n} \Delta_i^U(\hat{\theta}) \Big[W\big(u, R_i\big) - \overline{W}(u, \hat{\theta}) \Big] \Big[W\big(u, R_i\big) - \overline{W}(u, \hat{\theta}) \Big]' \tag{29}$$

The optimal weight functions, the covariance matrix and the most efficient estimators are estimated in two steps. The first step consists of obtaining a consistent estimate of θ_0 using a weight function that does not depend on the distribution of U_0. For the example of a GAFT model with a piecewise constant λ and a time-invariant treatment coefficient, a choice for the weight functions are $I_j(u)$ and R. The transformed durations for these parameter values are estimates of the unobserved population transformed durations. The second step concerns the estimation of the unknown distribution of U_0. Many different methods are available to get a reasonable estimate of an unknown distribution. We shall *not* apply the commonly used kernel based method. Although kernel–smoothed hazard rate estimators have been developed and adjusted to deal with the boundary problems inherent to hazard rates these methods can be difficult to implement due to the choice of the bandwidth. It is also unclear how the boundary corrections can be incorporated in the kernel estimates of the derivative of the hazard. We therefore choose to use a series approximation of the distribution.

Suppose the distribution of U_0 can be approximated arbitrary well using orthonormal polynomials. We base our approximation on Hermite polynomials using the exponential distribution as a weighting function:

$$g_0(u) = \frac{a e^{-au}}{\sum_{l=0}^{L} b_l^2} \Big[\sum_{l=0}^{L} b_l L_l(u) \Big]^2 \tag{30}$$

[3]Robins and Tsiatis (1991) suggested to use a numerical derivative of $n^{-1} S_n(\theta; W)$ that does not need an estimate of the optimal W–function to get $\hat{Q}(W)$. This numerical derivative is sensitive to the choice of the difference in θ. We found it hard to get stable results.

where

$$L_l(u) = \sum_{k=0}^{l} \binom{l}{k} \frac{(-au)^k}{k!} \tag{31}$$

are the Laguerre polynomials. The unknown parameters of this approximation are a and $b_0, \ldots, b_L$. If $b_l \equiv 0$ for all $l > 0$ the distribution of U_0 is exponential. Even for L as small as three (30) allows for many different shapes of $\kappa_0(u)$ and its derivative. Both can be derived analytically given the estimates of the parameters. The parameter estimators can be obtained from standard maximum likelihood procedures on the observed $(\tilde{U}_i(\hat{\theta}_n(W)), \Delta_i)$.

If a consistent but inefficient estimator $\hat{\theta}_n(W)$ of θ_0 is available and we have estimated the parameters of the polynomial approximation of the distribution of U_0 we can obtain an efficient estimator $\hat{\theta}_{\text{opt}}$ in just one additional step. From the linearization of the estimating equations, given in (43) in the appendix, we obtain an efficient estimator from

$$\hat{\theta}_{\text{opt}} = \hat{\theta}_n(W) - \hat{Q}(W)^{-1} S_n(\hat{\theta}_n(W); W_{\text{opt}})/n \tag{32}$$

This procedure is related to obtaining an efficient GMM estimator in two steps from a consistent, but possible, inefficient GMM estimator. It also possible to obtain the efficient estimator directly from minimizing the quadratic form. However, this involves again the minimization of a multi–dimensional step function.

5 Application to the Illinois Re-employment Bonus Experiment

Between mid–1984 and mid–1985, the Illinois Department of Employment Security conducted a controlled social experiment.[4] This experiment provides the opportunity to explore, within a controlled experimental setting, whether bonuses paid to Unemployment Insurance (UI) beneficiaries or their employers reduce the time spend in unemployment relative to a randomly selected control group. In the experiment, newly unemployed claimants were randomly divided into three groups: a *Claimant Bonus Group*, a *Employer Bonus Group* and, a *control group*. The members of both bonus groups were instructed that they (Claimant group) or their employer (Employer group) would qualify for a cash bonus of \$500 if they found a job (of at least 30 hours) within 11 weeks and, if they held that job for at least four months. Each newly

[4] A complete description of the experiment and a summary of its results can be found in Woodbury and Spiegelman (1987).

unemployed individual who was randomly assigned to one of the two bonus groups had the possibility to refuse participation in the experiment.

Woodbury and Spiegelman (1987) concluded from a direct comparison of the control group and the two bonus groups that the claimant bonus group had a significantly smaller average unemployment duration. The average unemployment duration was also smaller for the employer bonus group, but the difference was not significantly different from zero. These results are confirmed in table 1. Note that the response variable is insured weeks of unemployment. Because UI benefits end after 26 weeks, all unemployment durations are censored at 26 weeks. In table 1 no allowance is made for censoring. In the table we distinguish between compliers, those who agreed to be eligible for a bonus if assigned to a bonus group, and non-compliers. We see that the claimant bonus only affects the compliers and that the average unemployment duration of the non-compliers and the control group are almost equal.

Table 1: Average unemployment durations:control group and (non-)compliers.

	Control Group	Claimant Bonus			Employer Bonus		
		All	Compl.	Non-compl.	All	Compl.	Non-compl.
Benefit weeks	18.33	16.96	16.74	18.18	17.65	17.62	17.72
	(0.20)	(0.20)	(0.22)	(0.50)	(0.21)	(0.26)	(0.35)
N	3952	4186	3527	659	3963	2586	1377

standard error of average in brackets.

About 15% of Claimant group and 35% of the employer group declined participation. The reason for this refusal is unknown. Bijwaard and Ridder (2005) showed that the participation rate is significantly related to some observed characteristics of the individuals that also influence that re–employment hazard. Hence, we cannot exclude the possibility of unmeasured variables that affect both the compliance decision and the re–employment hazard. Meyer (1996) analyzed the same data with a PH model with a piecewise constant baseline hazard. He used the randomization indicator instead of the actual bonus-group agreement indicator as an explanatory variable. Thus he used the ITT estimator. He found a significantly positive effect of the claimant bonus. However, as shown by Bijwaard and Ridder (2005), the ITT has a downward bias.

We calculate the IVLR estimate of the effect of the claimant and employer bonus on the unemployment duration in a GAFT model and compare these estimates with the IVLR estimates of an AFT model, with ITT estimates in an MPH model and the ML estimates of an

MPH model that ignores the endogeneity of the decision to participate in the bonus group. We consider the two interventions separately: thus Claimant Bonus group versus Control group and Employer Bonus group versus Control.

We consider two alternative specifications for the treatment function of the bonus on unemployment duration: (i) constant treatment function and, (ii) a change in the treatment function after 10 weeks, in line with the end of the eligibility period of the bonuses. Thus, the implied transformed durations are

$$U(\theta) = \int_0^T \lambda(s;\alpha)e^{(\gamma_1 I_1(s)+\gamma_2 I_2(s))D}\mathrm{d}s \tag{33}$$

with $I_1(t) = I(0 \leq t < 11)$ and $I_2(t)$ is its complement. We employ two different specifications for $\lambda(t;\alpha_0)$: (i) AFT model, *i.e.* $\lambda(t;\alpha_0) \equiv 1$; and (ii) GAFT model with a piecewise constant λ on six intervals 0–2, 2–4, 4–6, 6–10, 10–25 and 25 and beyond.

For identification we need to set one of the parameters of the piecewise constant λ equal to one (or the log equal to zero). We let the base interval, the interval on which $\lambda = 1$, start on the last week before the end of the observation period, at 25 weeks. This allows us to capture the spike in the observed unemployment duration just before the UI eligibility period ends. The end of the UI eligibility period, at 26 weeks, is for all individuals the same and thus provides the potential censoring time.

For both the AFT and the GAFT specifications we estimate a first stage IVLR using the Powell-method and the one step optimal IVLR. The first stage IVLR uses the values of the bonus group assignment indicator (constant treatment function) or the bonus group assignment indicator times the interval indicators on the transformed duration, $R \cdot I_1(u)$ and $R \cdot I_2(u)$ (time-varying treatment function) and, (only for the GAFT-model) the interval indicators on the transformed duration, $I_j(u)$ as the weight functions. From these first stage IVLR's the implied transformed duration are obtained. Then, we estimate the parameters of the polynomial approximation of the distribution of U conditional on R and D as mentioned in section 4.3. From these estimated parameters we calculate the hazard and its derivative of the transformed duration. These functions are then used as inputs to derive the optimal weight functions (see Theorem 3), which in turn are necessary to calculate the covariance matrix. We also calculate the 1-step efficient estimates with these optimal weight functions. In the case of a constant treatment function, the optimal weight function are given in (27) and (28). For a time-varying treatment function the optimal weight function in (28) is more complicated and therefore not spelled out here.

Table 2: Instrumental Variable Linear Rank estimates for the coefficient of the Bonus

Claimant group

	AFT	GAFT[a]	MLE	ITT
Constant effect				
First stage	0.1446	0.1024	-	-
	(0.0493)	(0.0523)	-	-
1–step optimal	0.1596	0.0932	0.1039	0.1117
	(0.0460)	(0.0380)	(0.0285)	(0.0303)
Time varying effect				
First stage				
0-10	0.2955	0.1433	-	-
	(0.0523)	(0.0907)	-	-
10+	-0.0720	0.0063	-	-
	(0.0608)	(0.0886)	-	-
1–step optimal				
0-10	0.3865	0.1439	0.1601	0.1516
	(0.0486)	(0.0578)	(0.0361)	(0.0378)
10+	-0.0437	-0.0411	-	-
	(0.0572)	(0.0850)	-	-

Employer group

	AFT	GAFT[a]	MLE	ITT
Constant effect				
First stage	0.1011	0.0721	-	-
	(0.0646)	(0.0470)	-	-
1–step optimal	0.1332	0.0696	0.0387	0.0516
	(0.0612)	(0.0425)	(0.0318)	(0.0307)
Time varying effect				
First stage				
0-10	0.2304	0.1103	-	-
	(0.0710)	(0.0736)	-	-
10+	-0.0783	-0.0048	-	-
	(0.0836)	(0.1253)	-	-
1–step optimal				
0-10	0.6334	0.1279	0.0881	0.0800
	(0.0674)	(0.0521)	(0.0402)	(0.0384)
10+	0.0330	-0.0747	-	-
	(0.0745)	(0.0882)	-	-

[a] GAFT piecewise constant intervals: 0–2, 2–4, 4–6, 6–10, 10–25, 25 →; *Notes:* Standard error in brackets.

The estimation results for the bonus effects are reported in Table 2. The results for the piecewise constant λ can be found in Table 3. A comparison of the results that AFT overesti-

Table 3: Estimated λ in GAFT model for the Bonus data

Claimant				
	Constant Bonus effect		*Time varying Bonus effect*	
interval	first	opt.	first	opt.
0–2	0.8098	0.7500	0.8625	0.9328
	(0.4638)	(0.2052)	(0.5262)	(0.2409)
2–4	0.3146	0.2348	0.3542	0.2309
	(0.3691)	(0.1462)	(0.4048)	(0.1799)
4–6	-0.0782	-0.0415	-0.0390	0.0318
	(0.2646)	(0.1220)	(0.3015)	(0.1552)
6–10	-0.2743	-0.1859	-0.2341	-0.2085
	(0.2392)	(0.1133)	(0.2807)	(0.1369)
10–25	-0.6868	-0.6655	-0.6077	-0.6345
	(0.1626)	(0.1006)	(0.1758)	(0.1261)
Employer				
	Constant Bonus effect		*Time varying Bonus effect*	
interval	first	opt.	first	opt.
0–2	0.7095	0.8929	0.7088	0.5647
	(0.3063)	(0.1450)	(0.4375)	(0.1716)
2–4	0.2540	0.4451	0.2542	0.1464
	(0.2134)	(0.0939)	(0.3344)	(0.1227)
4–6	-0.1217	-0.1178	-0.1195	0.0875
	(0.2008)	(0.0925)	(0.2330)	(0.1050)
6–10	-0.4552	-0.2707	-0.4526	-0.4098
	(0.1516)	(0.0751)	(0.2255)	(0.0975)
10–25	-0.7492	-0.6826	-0.7180	-0.6057
	(0.0971)	(0.0372)	(0.1015)	(0.0491)

Notes: Standard error in brackets.

mates the effect and that both ML and ITT estimators underestimate the effect of the employer bonus. The results clearly indicate that the bonuses only influence the chances to find a job in the first ten weeks. This is in line with the bonus eligibility period: those who find a job after that period would not get the bonus. The effect of the Claimant Bonus increases from about 10% higher probability to find a job at every unemployment duration to about 15% higher probability to find a job in the first ten weeks (and no effect thereafter). The bonus for the Employer group raises the job finding probability with about 7% at every unemployment duration or with about 12% in the first ten weeks of unemployment. From Table 3 we can derive that the shape of the estimated λ's indicate a U–shaped λ.

An indication that the AFT is not the right model is the difference between the first stage and one–step optimal estimators for the AFT model. For a correctly specified model both estimators are consistent and, therefore, do not differ much. In the GAFT model the first stage and one–step estimator are of the same magnitude. The estimated standard errors of the latter are, as expected, substantially lower in most situations.

In the GAFT (and AFT) model the treatment effect of the bonus is defined in terms of the change in the quantiles $\mathrm{d}t_q(1)/\mathrm{d}t_q(0)$, see (11). In an AFT model with a time-constant treatment coefficient for the bonus the treatment effect is constant. In a GAFT model the λ function influences the treatment effect. Using the distribution of U_0, already calculated to estimate the optimal IVLR and the variance-covariance matrix, we can derive the treatment effect of the bonus in the GAFT depending on the quantile of the distribution. In Table 4 we present the effect for the 80%, 60% and 40% survival, together with the AFT treatment effect (first stage). Figure 1 and Figure 2 depict the change over the survival quantile of the treatment effect of the bonus in the GAFT model.

Table 4: Effect of the Bonus on the length of unemployment duration

		Claimant		Employer	
		Constant	*Time-varying*	*Constant*	*Time-varying*
AFT	0-10	0.865	0.744	0.904	0.794
	10+	0.865	1.075	0.904	1.081
GAFT		*reference individual*			
80%	$t_q(0)$	3.9	3.7	2.8	4.3
	$t_q(1)$	3.5	2.9	2.5	3.7
	effect	**0.911**	**0.866**	**0.933**	**0.823**
60%	$t_q(0)$	12.8	12.6	8.9	12.7
	$t_q(1)$	10.4	9.4	7.8	10.0
	effect	**0.911**	**0.571**	**0.933**	**1.078**
40%	$t_q(0)$	25.7	25.7	20.7	24.3
	$t_q(1)$	22.8	23.1	18.3	22.5
	effect	**1.772**	**1.973**	**0.933**	**1.078**

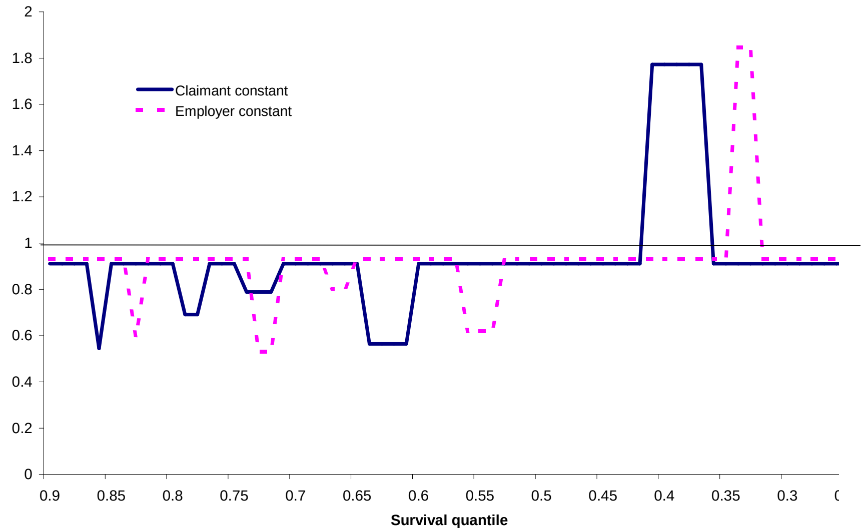

Figure 1: Treatment effect of Bonus on quantiles of unemployment duration, constant treatment coefficient

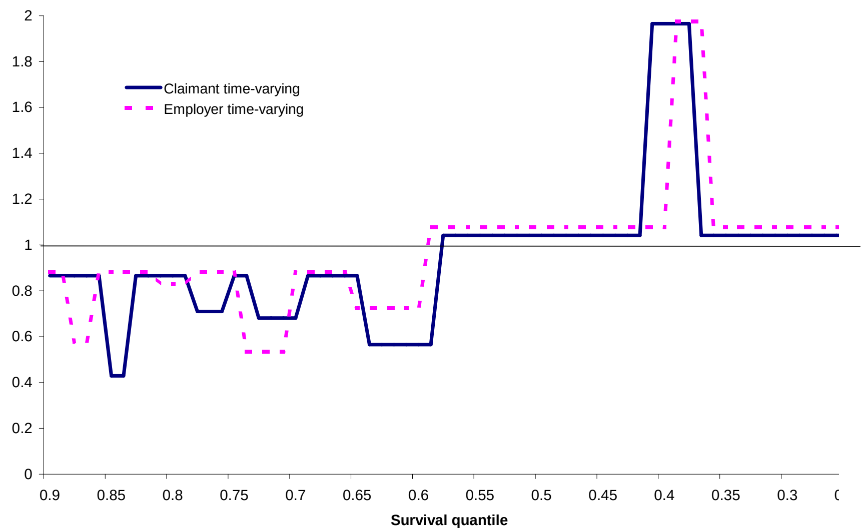

Figure 2: Treatment effect of Bonus on quantiles of unemployment duration, treatment coefficient that changes after 10 weeks

Note that an effect smaller than one indicates that the bonus decreases the duration till re-employment and an effect bigger than one increases the duration. We see from the table (and more pronounced in Figure 1) that even for a time-constant γ the effect of the bonus on the unemployment duration in the GAFT changes with the duration. The huge spike in the effect at the survival quantile of 40% for the claimant group is because re–employment rate of unemployment durations exhibits a spike just before the time that unemployment benefits are exhausted at 25 weeks. For the individuals in the control group the 40% survival time is just before 26 weeks, while in the claimant bonus group it is at 23 weeks. Thus the control group individuals are in the re-employment spike while the claimant bonus group are not. The interval boundaries of the other intervals of λ also cause, although not as pronounced, spikes. These spikes are downward because the λ is jumping to a lower level at these boundaries. The spikes are also visible in the effect of a time-varying coefficient of the bonus, see Figure 2. Here, the change in γ at a duration of 10 weeks, after which the coefficient is negative, is reflected is a upward shift of the effect curve.

6 Conclusion

In this article we proposed and implemented an instrumental variable estimation procedure to estimate treatment effects for duration outcomes based on a Generalized Accelerated Failure Time (GAFT) model. The GAFT model is based on a transformation of the durations that encompasses both the Accelerated Failure Time (AFT) and the Mixed Proportional Hazards (MPH) model. The interpretation of the treatment effect in the GAFT model is in terms of shifting the quantiles of the distribution.

The proposed Instrumental Variable Linear Rank (IVLR) estimation procedure is based on the inverse of an extended rank-test. It exploits that for the population parameters any (weight)function of the instrument is independent of the distribution of the transformed durations. It implies that the expected difference between the value of the weight function to the average of the weight functions for those individuals that are still under observation on the transformed duration scale is zero. The estimation procedure is related to the Rank Preserving Structural Failure Time Model estimator of Robins and Tsiatis (1991). The main difference with the RPSFT-model is that it assumes an AFT model while the IVLR allows for the more general GAFT-model, that includes duration dependence. The procedure is also related to the 2 stage Linear Rank (2SLR) estimation procedure of Bijwaard and Ridder (2005). The 2SLR assumes an MPH model and is a 2-steps procedure, while the IVLR is a one-step procedure.

The IVLR is based on a vector of mean restrictions and, therefore, it is related to the well-known GMM estimation procedure. Similar to the GMM estimation choosing the right weight functions can improve the efficiency. However, again similar to the GMM, these optimal weight functions are not directly observable. Fortunately, an adaptive (or even 2 step) procedure can provide the efficient IVLR. We show that a counting process framework simplifies the derivation and interpretation of the IVLR. The counting process framework also enables us to derive the large sample properties of the IVLR.

The empirical application shows that the ML and ITT estimates give downward biased treatment effects if the there is selective treatment choice. We also find that incorrectly assuming an AFT model can give misleading conclusions about the treatment effects of a bonus on the unemployment duration. In the Illinois bonus re-employment experiment many unemployed found a job just before their UI-benefits expires. This induces a spike in the re-employment hazard. In the GAFT model, even with a constant treatment coefficient, such a spike leads to an effect that changes over the quantiles. This has important implications for the evaluation of the treatment effect.

There are several issues that need further research. First, the current approach to adjust for endogenous censoring implies loss of information and depends on the (unknown) parameters of the model. An important improvement would be to find a method to adjust for endogenous censoring that is parameter independent and minimizes the loss of information. Further research on more general censoring patterns also deserve attention. A second issue for further research is the extension of the IVLR to recurrent duration data, like repeated unemployment spells.

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. Tinbergen Institute, discussion paper, TI 2005–47.

Andersen, P. K., O. Borgan, R. D. Gill, and N. Keiding (1993). *Statistical Models Based on Counting Processes*. New York: Springer–Verlag.

Angrist, J. D. and A. B. Krueger (1999). Empirical strategies in labor economics. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics, Volume 3A*, Chapter 23, pp. 1277–1366. Amsterdam: North–Holland.

Ashenfelter, O., D. Ashmore, and O. Deschêne (2005). Do unemployment insurance recipients actively seek work? Evidence from randomized trials in four U.S. states. *Journal of Econometrics 125*, 53–75.

Bijwaard, G. E. and G. Ridder (2005). Correcting for selective compliance in a re–employment bonus experiment. *Journal of Econometrics 125*, 77–111.

Brännäs, K. (1992). *Econometrics of the Accelerated Duration Model*. Umeå: Solfjädern Offset AB.

Chesher, A. (2002). Semi–parametric identification in duration models. C*e*MMAP, working paper, CWP20/02.

Cox, D. R. and D. Oakes (1984). *Analysis of Survival Data*. London: Chapman and Hall.

Ham, J. C. and R. J. LaLonde (1996). The effect of sample selection and initial conditions in duration models: Evidence from experimental data on training. *Econometrica 64*, 175–205.

Han, A. K. (1987). Non–parametric analysis of a generalized regression model: The maximum rank correlation estimator. *Journal of Econometrics 35*, 303–316.

Hausman, J. A. and T. M. Woutersen (2005). Estimating a semi–parametric duration model without specifying heterogeneity. C*e*MMAP, working paper, CWP11/05.

Heckman, J. J., R. J. LaLonde, and J. A. Smith (1999). The economics and econometrics of active labor market programs. In O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics, Volume 3A*, Chapter 31, pp. 1866–2097. Amsterdam: North–Holland.

Holland, P. W. (1986). Statistics and causal inference (with discussion). *Journal of the American Statistical Association 81*, 945–970.

Kalbfleisch, J. D. and R. L. Prentice (2002). *The Statistical Analysis of Failure Time Data (second edition)*. John Wiley and Sons.

Klein, J. P. and M. L. Moeschberger (1997). *Survival Analysis: Techniques for Censored and Truncated Data*. New York: Springer–Verlag.

Koenker, R. and Y. Bilias (2001). Quantile regression for duration data: A reappraisal of the Pennsylvania reemployment bonus experiments. *Empirical Economics 26*, 199–220.

Koenker, R. and O. Geling (2001). Reappraising medfly longevity: A quantile regression survival analysis. *Journal of the American Statistical Association 96*, 458–468.

Meyer, B. D. (1995). Lessons from the U.S. unemeployment insurance experiments. *Journal of Economic Literature 33*, 91–131.

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

Powell, M. J. D. (1964). An efficient method for finding the minimum of a function of several variables without calculating derivatives. *The Computer Journal 7*, 155–162.

Prentice, R. L. (1978). Linear rank tests with right censored data. *Biometrika 65*, 167–179.

Press, W. H., B. P. Flannert, S. A. Teukolsky, and W. T. Vetterling (1986). *Numerical Recipes: The Art of Scientific Computing*. Cambridge: Cambridge UP.

Ridder, G. (1990). The non–parametric identification of generalized accelerated failure–time models. *Review of Economic Studies 57*, 167–182.

Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period: Application to control of the healthy worker survivor effect. *Mathematical Modelling 7*, 1393–1512.

Robins, J. M. and A. A. Tsiatis (1991). Correcting for non–compliance in randomized trials using rank-preserving structural failure time models. *Communications in Statistics Part A: Theory and Methods 20*, 2609–2631.

Tsiatis, A. A. (1990). Estimating regression parameters using linear rank tests for censored data. *Annals of Statistics 18*, 354–372.

Van den Berg, G. J. (2001). Duration models: Specification, identification, and multiple duration. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics, Volume V*, Chapter 55, pp. 3381–3460. Amsterdam: North–Holland.

Van den Berg, G. J., J. C. van Ours, and B. van der Klaauw (2004). Punitive sanctions and the transition rate from welfare to work. *Journal of Labor Economics 22*, 211–241.

Woodbury, S. A. and R. G. Spiegelman (1987). Bonuses to workers and employers to reduce unemployment: Randomized trials in Illinois. *American Economic Review 77*, 513–530.

Ying, Z. (1993). A large sample study of rank estimation for censored regression data. *Annals of Statistics 21*, 76–99.

A Counting process interpretation

The density and the survival function of a duration T can be expressed as functions of the hazard rate. These expressions can be used to obtain a likelihood function. We use a different (but of course equivalent) representation of the relation between the hazard rate and the random duration. In particular, we use the framework of counting processes (see e.g. Andersen et al. (1993) and Klein and Moeschberger (1997)). The main advantage of this framework is that it allows us to express the duration distribution as a regression model with an error term that is a martingale difference. This simplifies the analysis of estimators. The conditions for non selective observation can be precisely stated in this framework. The same is true for conditions on time–varying covariates.

The starting point is that the hazard of T is the intensity of the counting process $\{N(t); t \geq 0\}$ that counts the number of times that the event occurs during $[0, t]$. The counting process has a jump +1 at the time of occurrence of the event[5]. A jump occurs if and only if $\mathrm{d}N(t) = N(t) - N(t-) = 1$. For duration data, the event can only occur once. In many unemployment studies the unit are only observed until re-employment or censoring. So, at most one jump is observed for any unit. To account for this we introduce the observation indicator $Y(t) = I(\tilde{T} \geq t)$ that is zero after re–employment. By specifying the intensity as the product of this observation indicator and the hazard rate we effectively limit the number of occurrences of the event to one. We assume that the observation indicator only depends on events up to time t. The observation process is assumed to have left–continuous sample paths. We define the history of the process up to time t by $H(t) = \{\overline{Y}(t), \overline{D}(t)\}$, where $\overline{Y}(t) = \{Y(s), 0 \leq s \leq t\}$. The history $H(t)$ only contains observable events.

Let V be some unobserved variables that both influence the treatment choice and the duration. An example is the, usually, unobserved search intensity of unemployed looking for a job. Denote $H^V(t) = \{H(t), V\}$, the history that also includes the unobservables. As with dynamic regressors in time-series models, the time-varying treatment $D(t)$ may depend on the dependent variable up to time t but not after time t (conditionally on V). Thus D only depends on $H^V(t)$. In the counting process literature such a time-varying covariate is called predictable. We will use the econometric term predetermined.

If the conditional distributions of $N(t)$ given $H^V(t)$ or $H(t)$ are well-defined (see Andersen

[5] The sample paths are assumed to be right-continuous.

et al. (1993) for assumptions that ensure this) we can express the probability of an event in $(t - dt, t]$ as[6]

$$\Pr(\mathrm{d}N(t) = 1 \mid H^V(t)) = Y(t)\kappa\big(t \mid \overline{D}(t), V\big)\mathrm{d}t \tag{34}$$

with $\kappa(t \mid \cdot)$ is the hazard of T at t given $\overline{D}(t)$ and V. By the Doob-Meier decomposition

$$\mathrm{d}N(t) = Y(t)\kappa\big(t \mid \overline{D}(t), V\big)\mathrm{d}t + \mathrm{d}M(t) \tag{35}$$

with $\{M(t); t \geq 0\}$ a (local square integrable) martingale. The conditional mean and variance of this martingale are

$$\mathrm{E}\big(\mathrm{d}M(t) \mid H^V(t)\big) = 0 \tag{36}$$

$$\mathrm{Var}\big(\mathrm{d}M(t) \mid H^V(t)\big) = Y(t)\kappa\big(t \mid \overline{D}(t), V\big)\mathrm{d}t \tag{37}$$

The (conditional on $H(t)$) mean and variance of the counting process are equal, so that the disturbances in equation (35) are heteroscedastic. The probability in equation (34) is zero, if the individual is not at risk.

A counting process can be considered as a sequence of Bernoulli experiments, because if $\mathrm{d}t$ is small equations (34) and (37) give the mean and variance of a Bernoulli random variable. The relation between the counting process and the sequence of Bernoulli experiments is given in equation (35), which can be considered as a regression model with an additive error that is a martingale difference. This equation resembles a time-series regression model. The Doob-Meier decomposition is the key to the derivation of the distribution of the estimators, because the asymptotic behavior of partial sums of martingales is well-known.

A.1 Transformed Counting process

The GAFT model transforms the observed duration T to a baseline duration U_0. The transformation involved a parameter vector $\theta_0 = (\gamma_0', \alpha_0')'$. We denote the transformation for parameter vectors $\theta \neq \theta_0$ by $U(\theta)$ with $U_0 = U(\theta_0)$. The distribution of $U(\theta)$ can also be represented by a (transformed) counting process $\{N^U(u); u \geq 0\}$. The relation between the original and transformed counting process, the observation indicator, and the time-varying exogenous covariates is

$$N^U(u;\theta) = N\big(h^{-1}(u;\theta)\big) \qquad Y^U(u;\theta) = Y\big(h^{-1}(u;\theta)\big)$$

$$D^U(u;\theta) = D\big(h^{-1}(u;\theta)\big) \qquad I_k^U(u;\theta) = I_k\big(h^{-1}(u;\theta)\big)$$

[6]Because the sample paths of $\{Y(t), X(t), t \geq 0\}$ are assumed to be left-continuous (as is the baseline hazard), we can substitute t for $t - dt$ in (34).

with $h(T;\theta) = h\big(T,\overline{D}(T);\theta\big)$. For $\theta = \theta_0$ we denote $h_0(T) = h(T;\theta_0)$. The corresponding history is $H^U(u;\theta) = \{\overline{Y}^U(u;\theta), \overline{D}^U(u;\theta), \overline{I}_k^U(u;\theta)\}$. In the sequel we suppress θ and write $Y^U(u), N^U(u), \overline{D}^U(u), \overline{I}_k^U(u)$ and $H^U(u)$ for $\theta \neq \theta_0$ and $Y_0(u), N_0(u), \overline{D}_0(u), \overline{I}_{k0}(u)$ and $H_0(u)$ for $\theta = \theta_0$. The intensity of the transformed counting process with respect to history $H^U(u)$ is (see Andersen et al. (1993), p. 87)[7]

$$\Pr\big(\mathrm{d}N^U(u) = 1 \mid H^U(u)\big) = Y^U(u)\mathrm{E}\Bigg[\frac{\lambda\big(h^{-1}(u;\theta);\alpha_0\big)}{\lambda\big(h^{-1}(u;\theta);\alpha\big)} \times \exp\Big(\sum_{k=1}^{K}(\gamma_{k0}-\gamma_k)I_k^U(u)D^U(u)\Big)\kappa_0\Big(h_0\big(h^{-1}(u;\theta)\big)\Big)\Bigg|H^U(u)\Bigg]\mathrm{d}u \quad (38)$$

We implicitly integrate with respect to the distribution of the unobserved V conditional on $H^U(u)$. Note that these unobserved covariates are only introduced to ascertain the predictability of the endogenous treatment process. Although the distribution of those variables determines the distribution of U_0, the consistency of the IVLR is independent of that distribution. Unfortunately, even for the population parameters θ_0 the hazard of U_0, $\kappa_0(u)$, still depends on the treatment path (through the correlation with V). If we condition on the instrument instead of the actual treatment we do get the desired independence.

We must add the instrument R to the conditioning variables in (38) if we consider instrumenting the endogenous treatment. Let the UR–history, $H^{UR}(u) = \{Y^U(s), R; 0 \leq s \leq u\}$, be the history on the transformed durations in which the treatment indicator D is replaced by the instrument, the treatment assignment indicator R. Then, another application of the innovation theorem gives the intensity of the transformed process on the UR–history

$$\Pr\big(\mathrm{d}N^U(u) = 1 \mid H^{UR}(u)\big) = Y^U(u)\mathrm{E}\Bigg[\frac{\lambda\big(h^{-1}(u;\theta);\alpha_0\big)}{\lambda\big(h^{-1}(u;\theta);\alpha\big)} \times \exp\Big(\sum_{k=1}^{K}(\gamma_{k0}-\gamma_k)I_k^U(u)D^U(u)\Big)\kappa_0\Big(h_0\big(h^{-1}(u;\theta)\big)\Big)\Bigg|H^{UR}(u)\Bigg]\mathrm{d}u \quad (39)$$

which for the population parameters simplifies to $Y_0^U(u)\kappa_0(u)\mathrm{d}u$ with $H_0^{UR}(u) = H^{UR}(u;\theta_0)$. Note that (38) and (39) only differ in the history the intensities are conditioned on.

The intensity in (39) is independent of R if we substitute the population parameter values, but not for other values of the parameters. This result is the basis for identification of the

[7] If $U = h(T)$ and κ_T is the hazard rate of the distribution of T, then the hazard rate of the distribution of U is

$$\kappa_U(u) = \kappa_T\big(h^{-1}(u)\big)\frac{1}{h'\big(h^{-1}(u)\big)}$$

parameters. Independence of R and the hazard rate of U_0 implies that the quantiles of the distribution of U_0 do not depend on R. By choosing C_0^U such that $\Pr(U_0 \leq C_0^U) = q$ we restrict the independence to the quantiles up to the q–th. For further reference we denote the intensity in (39) by $\kappa_i^U(u;\theta)$ such that

$$\Pr\big(\mathrm{d}N^U(u) = 1 \mid H^{UR}(u)\big) = Y^U(u)\kappa_i^U(u;\theta)\mathrm{d}u$$

which reduces to $\kappa_0(u)$ for the population parameters.

The counting process interpretation allows for an alternative formulation of the estimating equations in (17). The relevant counting measure, $N_i^U(u)$, can be seen as a discrete 'probability distribution' that assigns weight unity to uncensored transformed durations and is zero elsewhere. Then the estimating equations can be expressed as an integral with respect to that counting process

$$S_n(\theta;W) = \sum_{i=1}^{n} \int_0^{C_i^U} \Big\{W\big(u,R_i\big) - \overline{W}(u;\theta)\Big\}\,\mathrm{d}N_i^{\tilde{U}}(u) \tag{40}$$

where C_i^U is the transformed censoring time defined in (16).

B Asymptotic properties of the IVLR

In this section we discuss the asymptotic behavior of the Instrumental Variable Linear Rank estimator. The IVLR estimator is an extension of the linear rank estimator for the parameters in a censored regression of Tsiatis (1990). Robins and Tsiatis (1991) used similar estimation equations to deal with non-compliance in a randomized experiment. They consider a Accelerated Failure Time (AFT) model. This corresponds to the assumption of $\lambda = 1$ in the GAFT model. We assume a piecewise constant λ and the additional parameters α are estimated by adding extra estimation equations with weight functions based on $I_k(u)$. We also allow for a time-varying treatment coefficient.

We make the following assumptions:

C1: Random assignment of treatment at time 0. The randomization indicator is denoted by R with $R = 1$ for the treatment group and $R = 0$ for the control group. The path of the treatment $\{D(t); t \geq 0\}$ is predetermined at time 0 and the duration of the treatment is known and exogenous.

C2: The observation process $Y(t)$ is cadlag and $Y(t)$ is predetermined. Moreover,

$$\Pr\big(\mathrm{d}N(t) = 1 \mid Y(t) = 1, H(t)\big) = \Pr\big(\mathrm{d}N(t) = 0 \mid Y(t) = 0, H(t)\big)$$

C3: The population distribution of T given $\overline{X}$ and $\overline{D}$ satisfies

$$\int_0^T \lambda(s;\alpha_0)e^{\psi(s,D;\gamma_0)}\,\mathrm{d}s = U_0$$

The absolutely continuous distribution of U_0 does not depend on R. The p.d.f. of U_0 is bounded. $0 < \lim_{t\downarrow 0}\lambda(s;\alpha_0) < \infty$

C4: The transformed observation process $Y^U(u) = I\big(\tilde{U}(\theta) \geq u\big)$ is cadlag and predetermined, with $\tilde{U}(\theta) = \min\big(U(\theta), C^U\big)$ and C^U defined in (16).

C5: The weight function W is bounded and left–continuous.

C6: The intensity of $U(\theta)$, $\kappa_i^U(u)$ given history $H^{UR}(u)$ in (39) can be linearized in a neighborhood of θ_0 as a function of θ, *i.e.* there exist $\mu(u)$ and $\epsilon > 0$ such that for $\|\theta - \theta_0\| < \epsilon$

$$\big|\kappa_i^U(u;\theta) - \kappa_0(u) - (\theta-\theta_0)'d_{i0}(u)\big| \leq \|\theta-\theta_0\|^2\,\mu(u)$$

for $u \leq C_0 = C^U(\theta_0)$ with

$$d_{i0}(u) = \frac{\partial\kappa_i^U(u;\theta)}{\partial\theta}\bigg|_{\theta=\theta_0}$$

C7: There exists a continuous function $a(u;\theta)$ of θ in a neighborhood B of θ_0 such that

$$\sup_{u\leq C_0}\sup_{\theta\in B}\big\|\overline{W}(u;\theta) - a(u;\theta)\big\| \xrightarrow{p} 0$$

where

$$\overline{W}(u;\theta) = \frac{\sum_{j=1}^n Y_j^U(u)W\big(u,\overline{X}_j^U(u),R_j\big)}{\sum_{j=1}^n Y_j^U(u)}$$

C8: There exists a continuous matrix function $A(u;\theta)$ of θ in a neighborhood B of θ_0 such that

$$\sup_{u\leq C_0}\sup_{\theta\in B}\bigg\|\frac{1}{n}\sum_{i=1}^n\Big[W\big(u,R_i\big) - \overline{W}(u;\theta)\Big]\times\Big[W\big(u,R_i\big) - \overline{W}(u;\theta)\Big]'Y_i^U(u) - A(u;\theta)\bigg\| \xrightarrow{p} 0$$

C9: There exists a continuous matrix-function $V(u;\theta)$ of θ in a neighborhood B of θ_0 such that

$$\sup_{u\leq C_0}\sup_{\theta\in B}\bigg\|\frac{1}{n}\sum_{i=1}^n\Big[W\big(u,\overline{X}_i^U(u),R_i\big) - \overline{W}(u;\theta)\Big]\times d_{i0}(u)'Y_i^U(u) - V(u;\theta)\bigg\| \xrightarrow{p} 0$$

The assumptions are similar to those of Tsiatis (1990) and Robins and Tsiatis (1991). The starting point is (40), which can, for θ in a small neighborhood of θ_0, be rewritten as

$$S_n(\theta; W) = \sum_{i=1}^{n} \int_0^{C_{i0}} \Big\{ W\big(u, R_i\big) - \overline{W}(u; \theta) \Big\} \, \mathrm{d}N_i^U(u) + \sum_{i=1}^{n} \int_{C_i^U}^{C_{i0}} \Big\{ W\big(u, R_i\big) - \overline{W}(u; \theta) \Big\} \, \mathrm{d}N_i^U(u) \quad (41)$$

Substitution of the Doob–Meier composition in the first term on the right for N_i^U gives

$$S_n(\theta; W) = \sum_{i=1}^{n} \int_0^{C_{i0}} \Big\{ W\big(u, R_i\big) - \overline{W}(u; \theta) \Big\} \, dM_i^U(u) + \sum_{i=1}^{n} \int_0^{C_{i0}} \Big\{ W\big(u, R_i\big) - \overline{W}(u; \theta) \Big\} \kappa_i^U(u) Y_i^U(u) \, \mathrm{d}u \quad (42)$$

We consider both terms separately. The first term is, for θ close to θ_0, close to $S_n(\theta_0; W)$ and for the second term we have

$$(\theta - \theta_0) \cdot \sum_{i=1}^{n} \int_0^{C_{i0}} \Big\{ W\big(u, R_i\big) - \overline{W}(u; \theta) \Big\} \times \frac{\partial \kappa_i^U(u)'}{\partial \theta} Y_i^U(u) \, \mathrm{d}u + O_p\big(\|\theta - \theta_0\|^2\big)$$

Returning to (41) we note that the second term in this equation equals

$$\sum_{i=1}^{n} \left\{ \Big[W\Big(C_{i0}, \overline{X}_{i0}(C_{i0}), R_i\Big) - \overline{W}\big(C_{i0}; \theta_0\big) \Big] \times \theta_0(C_{i0}) Y_i(C_{i0}) \right\} + O_p\big(\|\theta - \theta_0\|^2\big)$$

The term between brackets is the covariance between $\theta_0(C_{i0})$ and $W\big(C_{i0}, \overline{X}_{i0}(C_{i0}), R_i\big)$ which is zero, because R is randomly assigned. Thus this whole term is zero for θ close to θ_0 and we have

$$\tilde{S}_n(\theta; W) = S_n(\theta_0; W) + n \int_0^{C_0} Z(u; \theta_0) \, du \cdot (\theta - \theta_0) \quad (43)$$

The proof of the consistency and asymptotic normality are both based upon this asymptotic linearity of $S_n(\theta; W)$ in the neighborhood of the true value θ_0. We follow the reasoning of Tsiatis (1990). Instead of a mean and variance condition, we have a mean and three covariance conditions. Let $\tilde{S}_n(\theta; W)$ be the right-hand side of (43). The following lemma shows that the linearization in (43) is uniformly close to the original estimating function

Lemma 1. *In neighbourhoods of $O(n^{-1/2})$ of θ_0*

$$n^{-1/2} \left\| \tilde{S}_n(\theta; W) - S_n(\theta; W) \right\|$$

converges uniformly to zero.

This lemma implies that $n^{-1/2}\tilde{S}_n(\theta; W)$ and $n^{-1/2}S_n(\theta; W)$ are asymptotically equivalent in a neighbourhood close to θ_0.

Proof: This can be proved in lines of Tsiatis (1990) Lemma (3.1) and (3.2) and theorem (3.2) and this is, because of the analogy, not repeated here.

Hence, approximately for the IVLR estimator $\hat{\theta}_n(W)$

$$\sqrt{n}(\hat{\theta}_n(W) - \theta_0) = \left[\int_0^{C_0} Z(u; \theta_0)\, du\right]^{-1} \frac{1}{\sqrt{n}} S_n(\theta_0; W) \tag{44}$$

Proof of theorem 1 and theorem 2. According to lemma 1 are $n^{-1/2}S_n(\theta; W)$ in a neighbourhood close to θ_0 asymptotically equivalent to $n^{-1/2}\tilde{S}_n(\theta; W)$. Then the estimates θ^* and $\hat{\theta}$, with $\tilde{S}_n(\theta^*; W) = 0$, will also be asymptotically equivalent. Clearly, θ^* converges in probability to θ_0. Hence, if we show that $\sqrt{n}(\hat{\theta}-\theta^*) \xrightarrow{p} 0$ then this would imply that $\hat{\theta}$ also converges in probability to θ_0. Tsiatis (1990) argues that lemma 1 suffices to proof this. This proves theorem 1.

According to the Mann–Wald theorem convergence in probability implies convergence in distribution. We note that $\sqrt{n}(\theta^* - \theta_0) = n^{-1/2}Q^{-1}(W)S_n(\theta_0; W)$ clearly converges to a normal distribution with mean zero and variance matrix $Q^{-1}(W)\Omega(W)Q'^{-1}(W)$. This completes the proof of theorem 2. $\square$

Remark. To establish detailed conditions on when $\tilde{S}_n(\theta; W)$ has a unique root is rather tedious; however Ying (1993) gave an excellent general treatment on rank estimation, which can also be used for the estimating equations in this article.