Correcting for Selective Compliance in a Re-employment Bonus Experiment

Govert Bijwaard
Geert Ridder [1]

[1]Faculteit der Economische Wetenschappen en Econometrie, Vrije Universiteit Amsterdam and Department of Economics, Johns Hopkins University,Baltimore, MD 21218-2685, fax. 410-5167600, E-mail: gbijwaard@econ.vu.nl, gridder@jhu.edu

Abstract

We propose a two-stage instrumental variable estimator that is consistent if there is selective compliance in the treatment group of a randomized experiment and the outcome variable is a censored duration. The estimator assumes full compliance in the control group. We use the estimator to reanalyze data from the Illinois re-employment bonus experiment.

1 Introduction

In theory, data from a randomized experiment produce an unbiased estimate of the effect of an intervention or program on an outcome variable. The difference of the average outcomes of the treatment and control samples estimates the average treatment effect. In practice, a randomized experiment may suffer from the same problems that affect behavioral studies. In particular, the random assignment of the intervention is often compromised by non-compliance to the assigned intervention, *i.e.* members of the treatment sample may drop out from the program, and members of the control group may participate. Non-compliance complicates the analysis of data from a randomized experiment in the same way as non-response in (random) sample surveys and panel attrition in longitudinal studies. If the non-compliance is selective, *i.e.* is correlated with the outcome variable, then the difference of the average outcomes is a biased estimate of the average effect of the intervention.

Sample selectivity is a familiar problem for economists, and over the years a number of approaches have been suggested to reduce selectivity bias. Since Heckman (1979) the dominant approach has been to model the selection process. This is the natural approach if the selection process is of independent interest, and the econometrician understands the process well enough to propose a reasonably accurate model. The first generation of these models required an assumption on the joint distribution of the response variable and the (latent) variable that determined participation in the program. In the second generation (Dehejia and Wahba (1995), Heckman, Ichimura, Smith, and Todd (1995)) this assumption is replaced by an elaborate model of the selection process under the assumption that an unbiased estimate of the intervention effect is obtained by comparing units with an (approximately) equal probability of participation.

Often, and the application considered below is a good example, there is not enough information to specify a model of the selection process. Moreover, the available characteristics of the units, although significantly correlated with compliance, do not explain compliance well enough to enable a comparison between members of the treatment and control groups in a sub-sample with the same probability of compliance. Under these circumstances an approach that does not require a model of the selection process is preferable.

The method of instrumental variables (IV) gives an unbiased estimate of the intervention effect and does not require a model of participation. This method assumes that the treatment assignment results from a two-stage process, where in the first stage the sample is divided randomly in two (or more) groups and in the second stage units are free to decide whether to participate in the program or not. In the clinical literature this experimental design is called the intention-to-treat (ITT) design. As shown by Angrist, Imbens, and Rubin (1996), the IV method does not estimate the average treatment effect, if the effect of the intervention varies in the population, and this variation can not be captured by observed characteristics. In the sequel, we assume that the effect heterogeneity can be fully explained by observables, so that IV estimates the usual average treatment effect (see also Heckman (1997)).

Another method that does not require a model of compliance are the nonparametric bounds proposed by Horowitz and Manski (1997). These bounds are reasonably narrow, if the response variable is binary. The generalization to the type of intervention and outcome variable considered below is nontrivial, and is not attempted here.

An issue that is usually ignored in discussions of the identification of treatment effects is the nature of the outcome variable. Often it is implicitly assumed that the intervention has an additive effect on the outcome. The response variable in the re-employment bonus experiment is the duration of unemployment. A duration or waiting time is often considered to be the outcome of a complex random experiment that consists of a sequence of Bernoulli or 0-1 experiments with a 0 for the experiment at time t (the spell starts at time 0) meaning that the spell of unemployment continues until $t + dt$ and a 1 indicating that spell ends in that interval. The (conditional upon survival) probability of a 1 in any of these experiments is the hazard rate (times dt) at t. Economic models for durations, *e.g.* search models, often have direct implications for the hazard rate. Another reason to consider the effect of an intervention on the hazard rate is that duration data are usually censored. Censoring limits the observation period, but is not a feature of the program. Hence, the estimated effect should be independent of the censoring time. Because the hazard rate is invariant to censoring, it is natural to relate the intervention to this quantity. Also interventions may be time-varying. In the case of the bonus experiment only individuals who find a job within 11 weeks qualify for the bonus. In general, the intervention may even depend on information that accumulates during the unemployment spell. With a time-varying intervention, the effect of the intervention becomes dependent on the outcome. Relating the time-varying intervention to the hazard instead of *e.g.* the mean seems the natural solution.

If the specification of the hazard rate model is incomplete, then randomization at the start of the spell does not ensure an unbiased estimate of the treatment effect, if the intervention affects the hazard rate multiplicatively. To see this we consider a population that consists of two types: H types with a large hazard who leave unemployment quickly and L types with a small hazard who typically have long unemployment spells. A simple intervention that lasts until the end of the spell is assigned at random at time 0. At the start the fraction H types is the same in the treatment and the control group. Clearly H types leave at a faster rate than L types, so that over time L types are an increasing fraction of the survivors. Moreover, because of the multiplicative effect on the leaving rate, H types in the treatment group leave even faster than H types in the control group, and this induces a correlation between the type and the intervention indicator. It is not difficult to see that the resulting bias in the intervention effect is toward zero. Note that there is no bias if the intervention has an additive effect on the hazard rate.

Finally, duration data are often not observed from the start. For instance, the durations may be a sample of ongoing spells. Such spells are not directly comparable to unemployment spells that are observed from the beginning. Ham and LaLonde (1996) document the large biases that result if one ignores the observation plan. Again, these biases can be dealt with if the duration distribution is described by its hazard rate.

In this paper we propose an estimator that is consistent for the intervention effect, if there is selective compliance to the intervention, the outcome variable is a censored duration, and the intervention varies over the duration. The duration model is the popular mixed proportional hazard (MPH) model, although we need not impose the restriction on the disturbance distribution that is implicit in the MPH model. The estimator is a generalization of the linear rank estimator of Tsiatis (1990) and Robins and Tsiatis (1991). In particular, we allow for a nonconstant baseline hazard, which amounts to a transformation of the dependent duration in the regression representation of the MPH model. The estimator requires preliminary estimates of the baseline hazard. If there is compliance in the control group, which is often the case in social experiments where the number of participants is a small fraction of the population and controls are not informed of the existence of the experimental program or can easily be excluded from participation, then these preliminary estimates can be obtained from the control group sample. The preliminary estimates are substituted in the second stage estimating equation of the intervention effect. The two stage procedure that we call the two-stage linear rank (2SLR) estimator, is computationally attractive, because it avoids the choice of weighting functions for the estimation of the parameters of the baseline hazard. If the control group sample is also used in the second stage, the additional variability due to the preliminary estimates and the induced correlation between the preliminary estimates and the second stage estimating equation complicates the computation of the asymptotic variance. Using the counting process representation of the first stage score we can obtain an estimable expression of this rather complicated variance.

The estimator is used to reanalyze data from the Illinois unemployment bonus experiment. These data have been analyzed before with increasing sophistication by Woodbury and Spiegelman (1987) and Meyer (1996). 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 treatment 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 other treatment group, the employer bonus group, the bonus was paid to their employer. The members of the two treatment 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 not known, and it is hard to think of an economic model for this decision[1]. The refusal is not completely random, because it is significantly related to some characteristics of the participants, characteristics that are also important determinants of the re-employment hazard. Hence, we can not exclude the possibility that some unobserved variables affect both the compliance decision and the re-employment hazard. Indeed we find that the corrected estimates differ substantially from the uncorrected estimates. In other words, our estimates show that compliance in the Illinois bonus experiment was indeed selective. We also investigate whether evidence of effect heterogeneity by income before unemployment (Meyer (1996)) and by the probability of benefit exhaustion (O'Leary, Decker, and Wandner (1998)) is biased by selective compliance.

The plan of the paper is as follows. In section 2 we introduce the 2SLR estimator and we discuss some of its properties. Section 3 gives the results of a sampling experiments. The application of the 2SLR estimator to the re-employment bonus experiment is in section 4. We draw some conclusions in section 5.

2 The Two-stage Linear Rank estimator

Our exposition begins with a discussion of the type of evaluation design, in which the Two-stage Linear Rank (2SLR) estimator can be used. Next, we introduce the model for the outcome variable. Finally, we introduce the estimator. Because the emphasis is on the application, the derivation of the properties of the estimator is heuristic.

2.1 Randomized experiments with non-compliance to treatment

As in the re-employment bonus experiment, we assume that the evaluation design has random assignment of treatment. The randomization indicator is denoted by R with $R = 1$ for the treatment group and $R = 0$ for the control group. The outcome variable is a duration, in the case of the re-employment bonus experiment

[1]Following Moffitt (1983), Meyer (1995) suggests that partial take–up (which in addition to nonparticipation includes failure to collect the bonus) can be explained by stigma or transaction costs. However, these explanations do not provide identifying restrictions in the bonus experiment. It should be stressed that the estimated program effects in this study correct for bias due to nonparticipation. We do not try to estimate the effect under the (counterfactual) assumption that all eligible individuals indeed collect the bonus

an unemployment duration. Randomization occurs at time 0, which corresponds to the start of the spell. In the sequel time is duration time, not calendar time. The treatment indicator is denoted by $D(t)$, *i.e.* we allow for treatments that vary over time. If the unit is in the treatment regime at time t, then $D(t) = 1$, and $D(t) = 0$ if not. We assume that $D(t)$ only depends on events that occurred during the time interval $[0, t]$, *i.e.* the treatment indicator is assumed to be weakly exogenous or predictable in the terminology of the general theory of counting processes. This allows for dependence on observables and unobservables that are random variables with a realization dated at time 0. In the re-employment bonus experiment the unemployed who participate in the program, qualify for a bonus during the first 11 weeks of unemployment. Hence, the treatment path $\{D(t), t \geq 0\}$ is determined at time 0, and the duration of the intervention is known and exogenous. This particular restriction on the treatment path simplifies the computation of the asymptotic variance of the 2SLR estimator.

The 2SLR estimator can be used in experiments in which a randomly selected sub-population is excluded from participation in the program. This condition is met if there is full compliance in the control group, *i.e.* if for all units in the population

$$R = 0 \Longrightarrow D(t) = 0, t \geq 0 \tag{2.1}$$

If a randomly selected sub-population is excluded from the program, we can use a sample from this sub-population to estimate some nuisance parameters that enter the distribution of the outcome variable for both the treatment and control group. In the re-employment bonus experiment full compliance in the control group is likely, because members of the control group were not informed of the existence of the re-employment bonus[2].

In many evaluation studies, and the bonus experiment is no exception, some units that are assigned to the treatment regime, do not participate in the program, *i.e.* they do not comply to their assigned treatment. In section 4 we show that in the bonus experiment non-compliance depends on observed characteristics of the units. Because predictions of compliance conditional on these characteristics are not perfect, unobservables must also play a role. The observable characteristics that are correlated with compliance also have a significant impact on re-employment. The resulting potential bias in the estimate of the program effect can be eliminated by conditioning on these observed characteristics. However, if compliance and re-employment have common unobserved determinants, (non)compliance is selective and conditioning on observables does not remove the selectivity bias.

2.2 The Mixed Proportional Hazard model for duration data

The response variable is the waiting time to an event, in the bonus experiment the event of re-employment. If the response variable is a waiting time or duration, a comparison of means is usually not the appropriate method to estimate the program effect. The two main reasons are censoring and time-varying interventions. Censoring is a feature of the observation plan and not of the program. For that reason the estimate of the program effect should be independent of censoring. Although the 2SLR estimator can deal with general independent censoring, *i.e.* censoring times that are stochastically independent of the durations (conditionally on observed covariates), the exposition is simplified by the assumption of fixed censoring. Hence, we assume that there is a fixed censoring time C that is common to all units. As in the bonus experiment this censoring time usually corresponds to the end of the observation period. Fixed censoring is a special case of independent censoring.

[2] In medical experiments there may be non-compliance in the control group, in particular if the treatment is thought to be effective.

With censoring the observed variables are

$$\Delta = I(T \leq C) \tag{2.2}$$
$$\tilde{T} = \min(T, C) \tag{2.3}$$

where $I(.)$ denotes the indicator of the event between parentheses, and Δ is one if T is observed.

Censoring and time-varying interventions can both be handled if we focus on the effect of the intervention on the hazard rate of the duration distribution. It is well-known that there is a one-to-one correspondence between hazard rates and (censored) distributions. Although it is not essential for the 2SLR estimator we choose a semi-parametric model for the duration. The reason is that we want to allow for both a vector of observable and unobservable determinants of the hazard. Moreover, we do not want to impose restrictions on evolution of the hazard rate over the duration. A model that satisfies these requirements is the Mixed Proportional Hazard (MPH) model, a generalization of the Proportional Hazard model originally introduced by Cox (1972)[3]. The hazard of the MPH model is

$$\theta(t \mid x, D(t), V) = \lambda(t) e^{\beta' x + \gamma D(t)} V \tag{2.4}$$

In this equation, x is a vector of time-constant characteristics of the unit[4], λ is the baseline hazard, and V is a random variable that captures variables not in x. Both x and V are thought to be determined at time 0 (or before). If $D(t)$ is the observed treatment indicator, then with selective compliance $D(t)$ or its path $\{D(t), t \geq 0\}$ may depend on x or V or both. Note that if the (path of the) treatment indicator only depends on x, we can estimate the parameter of interest γ by a number of procedures: parametric maximum likelihood after specifying λ and the distribution of V, nonparametric maximum likelihood (Heckman and Singer (1984)) after specifying λ or the Gamma–frailty estimator of Nielsen et al. (Nielsen, Gill, Andersen, and Sørensen (1992)) that does not require a specification of λ.

In the sequel we allow for dependence on both x and V. The distribution of V is left unspecified, because the 2SLR estimator does not require assumptions on this distribution. This is a major advantage, because inference on the distribution of V is notoriously unreliable (see e.g. Baker and Melino (1997) and the simulation results in section 3) and this unreliability may bias the estimate of the treatment effect in samples as large as those in the bonus experiment (about 4000 observations in both the control and treatment groups).

The c.d.f. and p.d.f. of the distribution of the duration T can be expressed as functions of the hazard rate. These expressions can be used to obtain MLE's of the parameters of the model. To understand the 2SLR estimator 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 for repeated events. 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. Regression models with martingale difference errors are the basis for inference in time series models with dependent observations. Hence, it is not surprising that inference is much simplified by using a similar representation in duration models.

The hazard in equation (2.4) 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

[3]This generalization was introduced independently by Lancaster (1979) and Manton, Stallard, and Vaupel (1981).

[4]The restriction to time-constant covariates is not essential. It only simplifies the notation.

process has a jump +1 at the time of occurrence of the event[5]. A jump occurs if and only if $dN(t) = N(t) - N(t-) = 1$. In the bonus experiment, the event can only occur once, because the unemployed are only observed until re-employment. Therefore we introduce the observation indicator $Y(t) = I(\tilde{T} \geq t)$ that is 0 after re-employment. By multiplying the intensity by this observation indicator we effectively limit the number of occurrences of the event to 1. It is essential that the observation indicator only depends on events up to time t. We also define the history of the process up to time t by

$$
\begin{aligned}
H(t) &= \{Y(s), D(s); 0 \leq s \leq t\} && (2.5)\\
H^V(t) &= \{H(t), V\} && (2.6)
\end{aligned}
$$

The history $H(t)$ only contains observable events, while the history $H^V(t)$ also includes the unobservable V. Note that, without loss of generality, we have omitted x.

As with dynamic regressors in time-series models, the time-varying $D(t)$ may depend on the dependent variable up to time t but not after time t (conditionally on V), *i.e.* $D(t)$ only depends on $H^V(t - dt)$. In the counting process literature such a time-varying covariate is called predictable. This corresponds to the assumption of weak exogeneity in econometric time-series models. If the conditional distributions of $N(t)$ given either $H^V(t)$ or $H(t)$ are well-defined (see Andersen, Borgan, Gill, and Keiding (1993) for assumptions that ensure this, we have can express the probability of an event in $(t - dt, t]$ as

$$\Pr(dN(t) = 1 \mid H^V(t - dt)) = Y(t - dt)\theta(t - dt \mid x, D(t - dt), V)\, dt \qquad (2.7)$$

and using this representation of the intensity of the counting process we have by the Doob-Meier decomposition[6]

$$dN(t) = Y(t)\theta(t \mid x, D(t), V)\, dt + dM(t) \qquad (2.8)$$

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

$$\mathrm{E}(dM(t) \mid H^V(t)) = 0 \qquad (2.9)$$

$$\mathrm{Var}(dM(t) \mid H^V(t)) = Y(t)\theta(t \mid x, D(t), V)\, dt \qquad (2.10)$$

The (conditional) mean and variance of the counting process are equal, so that the disturbances in equation (2.8) are heteroscedastic. The probability in equation (2.7) is 0, if the unit is no longer under observation. A counting process can be considered as a sequence of Bernoulli experiments, because if dt is small equations (2.7) and (2.10) 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 (2.8), that 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.

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

[6]Because the sample paths of $\{Y(t), D(t), t \geq 0\}$ are assumed to be left-continuous (as is the baseline hazard), we may substitute t for $t - dt$ in (2.7).

The 2SLR estimator is defined on transformed durations. The intensity (2.7) varies between units, because of its dependence on the covariates x_i and over time, because of the variation on the baseline hazard. We consider the transformation

$$U = b \int_0^T \lambda(s) e^{\gamma D(s)}\, ds \equiv h(T) \tag{2.11}$$

with $b = \exp\{\beta' x\}$. This transformation is increasing and leaves the origin of time unchanged. Moreover, if we denote the population b, γ, λ by b_0, γ_0, λ_0, and we let functions that depend on these population parameters also have a superscript 0, then it is easily seen that

$$U_0 \stackrel{d}{=} \frac{A}{V} \tag{2.12}$$

where A has a standard exponential distribution, which has a constant hazard rate equal to 1. Hence, in the population the transformed duration has a constant (over time) hazard rate that varies only over the population by V.

Just as the distribution of T, that of U can be represented by a (transformed) counting process $\{N^U(u); u \geq 0\}$. The relation between the original and transformed counting process, observation indicator, and time-varying treatment is

$$N^U(u) = N(h^{-1}(u)) \tag{2.13}$$
$$Y^U(u) = Y(h^{-1}(u)) \tag{2.14}$$
$$D^U(u) = D(h^{-1}(u)) \tag{2.15}$$

The corresponding history is $H^U(u) = \{Y^U(s), D^U(s); 0 \leq s \leq u\}$ or $H^{UV}(u) = \{H^U(u), V\}$ if V is included. The intensity of the transformed counting process (with respect to history H^{UV} is (see Andersen, Borgan, Gill, and Keiding (1993), p. 87)[7]

$$\Pr(dN^U(u) = 1 \mid H^{UV}(u)) = Y^U(u) \frac{b_0}{b} \frac{\lambda_0(h^{-1}(u))}{\lambda(h^{-1}(u))} e^{(\gamma_0 - \gamma) D^U(u)} V\, du \tag{2.16}$$

which for the population parameters simplifies to V which is consistent with (2.12).

This simplification only occurs if we substitute the population parameters. The observed intensity of the transformed counting process for other values of the parameters is obtained by the innovation theorem (Andersen, Borgan, Gill, and Keiding (1993), p. 80)

$$\begin{aligned} &\Pr(dN^U(u) = 1 \mid H^U(u)) = \\ &\qquad Y^U(u) \frac{b_0}{b} \frac{\lambda_0(h^{-1}(u))}{\lambda(h^{-1}(u))} e^{(\gamma_0 - \gamma) D^U(u)} E(V \mid H^U(u))\, du \end{aligned} \tag{2.17}$$

i.e. we integrate with respect to the conditional distribution of V given $H^U(u)$. If we consider (transformed) durations from a randomized experiment, we must add the randomization indicator R that is determined at time 0 to the conditioning variables in (2.16). Another application of the innovation theorem to the treatment indicator

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

$$\lambda_U(u) = \lambda_T(h^{-1}(u)) \frac{1}{h'(h^{-1}(u))}$$

$D(t)$, gives the intensity of the transformed process if the history is restricted to $\overline{Y}^U(u) = \{Y^U(s); 0 \le s \le u\}$

$$\begin{aligned} &\Pr(dN^U(u) = 1 \mid \overline{Y}^U(u), R) = \\ &\quad Y^U(u)\frac{b_0}{b}\mathrm{E}\left(\frac{\lambda_0(h^{-1}(u))}{\lambda(h^{-1}(u))}e^{(\gamma_0-\gamma)D^U(u)}V\middle|\overline{Y}^U(u), R\right)du \end{aligned} \tag{2.18}$$

Note that because by randomization R and V are stochastically independent, this intensity 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 parameters, and in particular of the treatment effect γ_0. Independence of R and the hazard rate of U_0 implies (infinitely) many orthogonality restrictions or moment conditions. For instance, if we assume independence of the hazard of U_0 and R for $u \le C_U$, we can choose C_U so that it corresponds to a specific quantile of the distribution of U_0. In other words, the moment conditions can be seen as simultaneous quantile independence restrictions.

If the durations are censored at C, the observed transformed durations are

$$\tilde{U} = \min(h(T), h(C)) = h(\tilde{T}) \tag{2.19}$$

with the observation indicator Δ unchanged. One is tempted to assume that the same orthogonality conditions apply to the intensity of the censored transformed durations. However, with censoring some of these orthogonality conditions are no longer valid. To see why we consider the case of fixed censoring and we assume that b and λ are as in the population. Moreover, we consider a time-constant intervention D. Hence

$$U = b_0 e^{\gamma D}\Lambda_0(T) \tag{2.20}$$

with Λ the integral of λ. For all units, irrespective of treatment regime, censoring occurs at time C. Hence, if $D = 0$ the censoring in the transformed time is at $b_0\Lambda_0(C)$, but if $D = 1$ the censoring time is $b_0 e^{\gamma}\Lambda_0(C)$. Hence, if $\gamma > 0$, then all transformed durations in the interval $[b_0\Lambda_0(C), e^{\gamma}b_0\Lambda_0(C)]$ have $D = 1$, *i.e.* belong to the treatment group (for $\gamma < 0$ the boundaries are reversed). The intensity of U_0 on this interval is clearly not independent of D and hence of R. The independence of the intensity of U_0 and R only holds up to the lower bound of the interval. This implies that in the 2SLR estimator, which exploits this independence, the transformed durations that fall in the problematic interval have to be censored. Note that there is an asymmetry between the cases of a positive and a negative treatment effect. In the case of a negative treatment effect a lower bound on γ that should be below the population value is used to determine which transformed durations must be censored.

2.3 The Two-Stage Linear Rank estimator

The two-stage linear rank estimator (2SLR) exploits the independence of the hazard rate of U_0 and R. The data are i.i.d. observations of $(\Delta_i, \tilde{T}_i, \overline{D}(\tilde{T}_i, R_i, V_i))$ with $i = 1, \ldots, n+m$ and $\overline{D}(t, R, V) = \{D(s, R, V); 0 \le s \le t\}$, the treatment path up to t. The observations $1, \ldots, n$ correspond to the control group, and the observations $n+1, \ldots, n+m$ to the treatment group. In the first stage of the estimation the data from the control group are used to estimate the regression parameters β and the parameters of the baseline hazard. The estimates are obtained by maximum likelihood using a discrete mixture for the distribution of V (Heckman and Singer (1984)). The estimates of the regression parameters and the parameters of the baseline hazard are used to compute the transformed durations and the new uninformative censoring times $C_U(x_i)$ for the transformed durations (obtained by setting γ either equal to 0 or to some negative lower bound). These transformed durations are a function of γ, *i.e.* they are denoted by $\tilde{U}_i(\gamma), i = 1, \ldots, n+m$. Note that the observations from the control group are re-used in the second stage.

We consider the estimating equation

$$S_{n+m}(\gamma) = \sum_{i=1}^{n+m} \tilde{\Delta}_i \left\{ R_i - \frac{\sum_{j=1}^{n+m} Y_j^U(\tilde{U}_i(\gamma)) R_j}{\sum_{j=1}^{n+m} Y_j^U(\tilde{U}_i(\gamma))} \right\} \tag{2.21}$$

with $\tilde{\Delta}_i$ the observation indicator for the adjusted censoring times. This can be expressed as an integral with respect to the counting process $\{N_i^{\tilde{U}}(t); t \geq 0\}$ (the counting measure can be seen as a discrete 'probability distribution' that assigns weight 1 to uncensored transformed durations and is zero elsewhere)

$$S_{n+m}(\gamma) = \sum_{i=1}^{n+m} \int_0^{C_U(x_i)} \left\{ R_i - \frac{\sum_{j=1}^{n+m} Y_j^U(u) R_j}{\sum_{j=1}^{n+m} Y_j^U(u)} \right\} dN_i^{\tilde{U}}(u) \tag{2.22}$$

Substituting the Doob-Meier decomposition for the counting process $N^{\tilde{U}}$ and history $(\overline{Y}^U(u) = \{Y^U(s); 0 \leq s \leq u\}, R)$ so that the intensity is given by (2.18), in this equation results in the sum of two integrals. One of the integrals is with respect to a martingale difference and has expectation 0. For the population parameters the second integral converges to the covariance of R and the intensity of U_0 on the interval $[0, C_U(x_i)]$. This covariance is 0 because of the randomized assignment. For that reason, the second stage estimate of the treatment effect is obtained by setting (2.22) equal to 0. Because $S_{n+m}(\gamma)$ is a discontinuous function of γ, the estimator is the solution to

$$\hat{\gamma} = \inf\{\gamma \mid S_{n+m}(\gamma-).S_{n+m}(\gamma+) \leq 0\} \tag{2.23}$$

The name of the estimator derives from $S_{n+m}(\gamma)$, which is the linear rank test statistic (Prentice (1978)). Tsiatis (1990) suggested the use of this statistic as an estimating equation for a censored linear regression. The $-$ and $+$ in the definition are the left- and right-hand limit at $\hat{\gamma}$: $\hat{\gamma}$ is the value at which the function changes sign. In case that the function is 0 on an interval, we choose the lower bound of that interval.

It is illuminating to express $S_{n+m}(\gamma)$ in the ordered (censored) transformed durations that we denote by

$$\tilde{U}_{(1)}(\gamma) \leq \tilde{U}_{(2)}(\gamma) \leq \ldots \leq \tilde{U}_{(n+m)}(\gamma) \tag{2.24}$$

to obtain

$$S_{n+m}(\gamma) = \sum_{i=1}^{n+m} \tilde{\Delta}_{(i)} \left\{ R_{(i)} - \frac{\sum_{j=i}^{n+m} R_{(j)}}{n+m-i+1} \right\} \tag{2.25}$$

The indicator $\tilde{\Delta}_{(i)}$ does not depend on γ. $S_{n+m}(\gamma)$ is constant on intervals that leave the order of the residuals unaltered. A discontinuity occurs if γ makes two residuals, say $\tilde{U}_{(k)}(\gamma)$ and $\tilde{U}_{(k+1)}(\gamma)$, equal. To concentrate on essentials, we take $D(t, R, V) = D$ where we suppress the dependence on R, V. Consider a value of γ with $\tilde{U}_{(k+1)}(\gamma) > \tilde{U}_{(k)}(\gamma)$. Hence

$$\frac{\tilde{U}_{(k+1)}(\gamma)}{\tilde{U}_{(k)}(\gamma)} = \frac{b_{(k+1)}\Lambda(\tilde{T}_{(k+1)})}{b_{(k)}\Lambda(\tilde{T}_{(k)})} \exp\{\gamma(D_{(k+1)} - D_{(k)})\} > 1 \tag{2.26}$$

If $D_{(k+1)} = D_{(k)}$, then this ratio does not depend on γ, and the residuals can not be made equal by increasing or decreasing γ. The order of such a pair of residuals does not depend on γ, and it does not result in a discontinuity in $S_n(\gamma)$. If $D_{(k+1)} > D_{(k)}$, then equality of residuals, and a discontinuity in $S_{n+m}(\gamma)$, occurs if γ is decreased to the value that makes the ratio of the residuals equal to 1. If $\Delta_{(k+1)} = \Delta_{(k)} = 1$, the jump is

$$\frac{R_{(k)} - R_{(k+1)}}{n+m-k} \tag{2.27}$$

If $D_{(k+1)} < D_{(k)}$, γ has to be increased, and the jump is again given by equation (2.27). If $\Delta_{(k+1)} = \Delta_{(k)} = 1$ and $D_{(k+1)} > D_{(k)}$ implies $R_{(k+1)} \geq R_{(k)}$ or $D_{(k+1)} < D_{(k)}$ implies $R_{(k+1)} \leq R_{(k)}$[8], then the function has a negative jump at the value of γ that makes the two (uncensored) residuals equal. Hence, if there is no censoring and no strict disagreement between censoring and randomization indicators, then $S_{n+m}(\gamma)$ is a decreasing step function of γ.

If these conditions do not hold, the monotonicity is lost. However, it is easily shown (Tsiatis (1990)), that the discontinuities are bounded by $2/(n+m-k)$ with $n+m-k$ the number of units under observation. Hence for k fixed, the discontinuities become small if the number of observations increases. This makes it possible to find the 2SLR estimator by search methods for the root of a continuous equation, which makes the second stage estimator computationally easier than other semi-parametric estimators, as *e.g.* the Maximum Score estimator (Manski (1975))[9] This result could have been anticipated from the observation, that equation (2.22) converges to a function that is continuous in γ.

In the derivation of the asymptotic distribution we use the linearization of $S_{n+m}(\gamma)$ suggested by Tsiatis (1990). Denote the parameters that are estimated in the first stage by θ, and let $d_{\gamma,i}(\gamma_0, \theta_0)$ and $d_{\theta,i}(\gamma_0, \theta_0)$ be the derivative of the hazard rate in (2.18) with respect to γ and θ, respectively, both evaluated for the population values of the parameters. Then by a linear approximation of the hazard rate, we obtain the following expression ($\overline{R}(u,\gamma)$ is the average of the R_i among the survivors at u)

$$\sqrt{n+m}(\hat{\gamma} - \gamma_0) = \frac{\frac{1}{\sqrt{n+m}} \sum_{i=1}^{n+m} \int_0^{C_U(x_i)} \left\{ R_i - \overline{R}(u,\gamma_0) \right\} \, dM_i^0(u) + a(\theta_0)\sqrt{n}(\hat{\theta} - \theta_0)}{\frac{1}{n+m} \sum_{i=1}^{n+m} \int_0^{C_U(x_i)} \left\{ R_i - \overline{R}(u,\gamma_0) \right\} d_{\gamma,i}(\gamma_0, \theta_0)\, du} \tag{2.28}$$

where

$$a(\theta_0) = \frac{\sqrt{n+m}}{\sqrt{n}} \frac{1}{n+m} \sum_{i=1}^{n+m} \int_0^{C_U(x_i)} \left\{ R_i - \overline{R}(u,\gamma_0) \right\} d_{\theta,i}(\gamma_0, \theta_0)\, du \tag{2.29}$$

Because the first stage score vector can be expressed as an integral with respect to the martingale difference $dM_i^0(u)$ an explicit expression for the covariance between the estimating equation and the first-stage MLE can be obtained. Here it suffices that closed form and estimable expressions for both the variance of the numerator and for the denominator can be obtained[10]. We conclude that the 2SLR estimator is consistent and asymptotically normal with an estimable variance[11].

[8]In other words, there is no strict disagreement between the treatment and randomization indicators

[9]With the 2SLR estimator there is no need for smoothing as proposed by Horowitz (1996) for the Maximum Score estimator.

[10]The expressions are rather lengthy and can be obtained from the authors upon request

[11]In the derivation of the variance we use certain simplifications that are specific to the bonus experiment. In a more general setting it may be preferable to use the bootstrap to approximate the sampling distribution of the estimator.

It is instructive to consider the asymptotic variance for a special case. Let us assume that the population value θ_0 is known, that the intervention has a fixed and exogenous duration (as in the bonus experiment), and that the distribution of U_0 is standard exponential. In that case the asymptotic variance is equal to

$$\frac{1}{\overline{R}(1-\overline{R})\Pr(D=1 \mid R=1)^2} \frac{\frac{1}{n+m}\sum_{i=1}^{m+n} e^{-C_U(x_i)}}{\left(\frac{1}{m}\sum_{i=n+1}^{m+n} e^{-P_U(x_i)}\right)^2} \tag{2.30}$$

The effect of the censoring and the finite duration of the intervention is in the second factor. The first factor is minimal if there is an equal number of treatments and controls and it decreases in the compliance rate. In the second factor $P_U(x_i)$ is defined as $C_U(x_i)$ if we replace the fixed censoring time by the fixed time at which the intervention ends. The variance decreases if the intervention lasts longer. It seems that the variance can be reduced by decreasing the censoring time. This is true as long as the censoring time is later than the end of the intervention. If censoring is before this time, $P_U(x_i)$ is equal to $C_U(x_i)$ and the variance increases if C decreases. The asymptotic variance shows that it is optimal to let the censoring time coincide with the end of the intervention.

3 Sampling Experiments

We performed two sampling experiments to study the small sample performance of the 2SLR estimator. Besides the usual concerns of small sample bias and the approximation of the sampling distribution, the computation of the estimator is of particular interest. As stressed in section 2 the 2SLR estimator is the root of a discontinuous function. The discontinuities decrease with the number of observations, and the function is approximately linear near this root. So standard root finding algorithms, as the quasi-Newton algorithm, may work well, even though they use numerical derivatives in the computation of the search step. Moreover, if there is no censoring, the function is decreasing, so that there is only one root[12]. With censoring the function is not monotone, and hence there may be multiple solutions. In the sampling experiments we check for multiple roots under conditions that are close to those in our empirical application.

The sampling experiments are designed to resemble the bonus experiment in section 4. The response model is an MPH model and we consider censored and uncensored data. The compliance depends on observed and unobserved characteristics of the units. In the first experiment the compliance only depends on observed characteristics (exogenous intervention), while in the second experiment the unobserved characteristics also play a role (endogenous intervention). We compare the 2SLR estimator with two alternative estimators: the Intention–To–Treat (ITT) estimator and the ML estimator. The ML estimator is the estimator that ignores selective compliance, *i.e.* we estimate the parameters of an MPH model with an endogenous covariate. The ITT estimator is obtained by replacing the endogenous treatment indicator by the instrumental variable R. Although R is exogenous, the ITT estimator suffers from an errors-in-variables bias. The first stage estimates of the 2SLR procedure are ML estimates on the duration data from the sub-sample of the control group.

In both experiments we have 8000 individuals randomly divided into two groups of equal size, the treatment group and the control group. The individuals in the control group are excluded from participation in the program. For a comparison with the empirical application, we take the unit of time as one week. The durations

[12] With probability 0 the function is 0 on an interval.

are right–censored at a fixed censoring time of 26 weeks and are generated according to an MPH model with hazard rate

$$\theta(t \mid x, D, V) = V\lambda_0(t)\exp(\beta_1 x + \gamma D), \tag{3.1}$$

where V follows a discrete distribution with three points of support 0.25, 2.5, and 5.5 with probabilities 0.8, 0.1, and 0.1. The mean of this distribution is equal to 1. The baseline hazard is constant on intervals: the intervals are 0-4, 4-11, 11-24, 24-, and the hazard values are 0.09072, 0.06721, 0.06721, 0.1003. The intervention indicator is equal to $D(x, V)I(t \leq 11)R$. For the individuals in the treatment group compliance with their assigned treatment is determined by

$$D(x_i, V_i) = I(\xi_1 x_i + \xi_2 V_i > \xi_3) \tag{3.2}$$

where the ξ's are chosen to obtain a level of compliance close to that observed in our application (65%, the compliance rate in the employer bonus experiment), and to obtain nonzero correlations between (1) the treatment indicator and x, and (2) the treatment indicator and V. These correlations are 0.8 and 0 and 0.7 and -0.4, in the first and second experiments, respectively. The exogenous x is normal with mean 0 and variance 8. The regression coefficient of x is 0.2 and the treatment effect is equal to 0.25. The mean and standard deviation of the durations are 16.6 and 10.5 in both experiments. The fraction of durations that is censored is 47%. If compliance is exogenous (experiment 1), the MLE is the efficient estimator, and we can estimate the relative efficiency of the 2SLR estimator. The number of replications is 100.

Table 3.1 reports the results for the parameter of interest, the treatment effect, and for the first-stage MLE's of the other parameters[13]. For the two experiments we give the average bias and its standard error, the standard deviation of the estimate, the average asymptotic standard error, and the root mean squared error (RMSE) of the estimate. The 2SLR estimate performs well in both experiments. The relative bias is small (less than 7%) and not significantly different from 0. The asymptotic standard errors are a reasonable approximation to the true standard errors. The latter result also holds for the ML and ITT estimators. The biases in the ML and ITT estimators of the program effect are large. The bias in the ITT estimator is always significantly negative, as one would expect. In the second experiment, with endogenous compliance, the relative bias in the ML estimator is large and highly significant. Its size and sign depend on the magnitude and sign of the correlation between D and V. With a positive correlation, individuals with a large V have a higher probability of complying with their assigned treatment, and the MLE over-estimates the treatment effect, because individuals with characteristics favorable to re-employment are more likely to participate. A negative correlation between compliance and the unobserved component induces a negative bias in the MLE.

In the first experiment (exogenous compliance) the MLE is consistent and asymptotically efficient. A comparison of the sampling variances of the MLE and 2SLR shows that the relative efficiency of the 2SLR estimator is about .44.

The sampling distribution of the 2SLR estimator in the two experiments is depicted in figure 3.1. We also plot the asymptotic normal approximation to this sampling distribution (the mean is the true value and the variance is the average asymptotic variance).

[13] We use a reparameterization: the α_k for $k = 1, 2, 3, 4$ are the log of the hazard values on the duration intervals, a log transformation is also applied to the points of support of the discrete distribution of V, and $\delta_j = -\ln(-ln(p_j))$ for $j = 1, 2, 3$ are the parameters associated with the probabilities of this distribution. In both experiments we could not compute the variance of the latter parameters in every simulation, 7 respectively 5 times. Therefore, for those parameters we only averaged over the simulations in which we could compute the variance.

The MLE for the regression coefficient of the exogenous regressor and the parameters of the baseline hazard are unbiased with exogenous compliance, but significantly biased with endogenous compliance. The parameters of the distribution of V are biased even if the specification is correct (exogenous compliance). Even in large samples inference on this distribution is inaccurate. In experiments with smaller samples we could not find evidence of unobserved heterogeneity, *i.e.* the distribution of V became degenerate, in a significant fraction of the replications. The bias in the MLE is independent of the sample size. Hence, an empirical strategy in which failure to find evidence of unobserved heterogeneity leads to the conclusion that compliance is exogenous can result in a large bias in the estimated treatment effect.

We used two methods to compute the 2SLR estimator: the Brent algorithm, an improved bisection method, that does not use numerical derivatives, and the quasi-Newton algorithm that does. In all instances the quasi-Newton and Brent algorithm converged to the same solution. We conclude that, as argued before, the discontinuities in the 2SLR equation are indeed inessential. We also did not encounter multiple solutions.

4 Application to re-employment bonus experiment

4.1 The re-employment bonus experiment

Between mid–1984 and mid–1985, the Illinois Department of Employment Security conducted a controlled social experiment[14]. The experiment provides the opportunity to explore, within a controlled experimental setting, whether bonuses paid to Unemployment Insurance (UI) beneficiaries (treatment 1) or their employers (treatment 2) reduce the unemployment of beneficiaries relative to a randomly selected control group. Both treatments consisted of a $500 bonus payment, which was about four times the average weekly unemployment insurance benefit.

In the experiment, newly unemployed claimants were randomly divided into three groups[15]:

1. The Claimant Bonus Group. The members of this group were instructed that they 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. 4186 individuals were selected for this group. Of those 3527 (84%) agreed to participate.

2. The Employer Bonus Group. The members of this group were told that their next employer would qualify for a cash bonus of $500 if they, the claimants, found a job (of at least 30 hours) within 11 weeks and, if they held that job for at least four months. 3963 were selected for this group and 2586 (65%) agreed to participate.

3. The Control Group, *i.e.* all claimants not assigned to one of the other groups. This group consisted of 3952 individuals.

The individuals assigned to the control group were excluded from participation in the experiment. In fact, they did not know that the experiment took place. Table 4.1

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

[15] The eligible population for either the Claimant Experiment or the Employer Experiment consisted of those who filed an initial claim for UI between July 29, 1984 and November 17, 1984 and who registered with one of the 22 Job Service offices in northern and central Illinois. Individuals had to be eligible for 26 weeks of UI benefits, had to be between ages 20 and 55 and, had to have no (non)monetary eligible claims.

reports some descriptive statistics for the three groups. The table confirms that the randomization resulted in three similar groups.

4.2 Results of previous analyses

Woodbury and Spiegelman (1987) concluded from a direct comparison of the control group and the two treatment 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 0. These results are confirmed in table 4.2. 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. Benefit spells are censored (at time of benefit exhaustion) unemployment spells. In table 4.2 no allowance is made for censoring. In the table we distinguish between compliers 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.

Meyer (1988 and 1996) analyzed the same data with a PH model with a flexible specification of the baseline hazard. He used the randomization indicator as an explanatory variable, *i.e.* he used the ITT estimator. He found a significantly positive effect of the claimant bonus. In his PH model with flexible baseline hazard, he did not find evidence of unobserved heterogeneity. The sampling experiments in section 3 indicate that the ITT estimator has a downward bias. Moreover, inference on unobserved heterogeneity is inaccurate, and tests for unobserved heterogeneity are likely to be biased, so that failure to find evidence of unobserved heterogeneity does not imply that compliance is exogenous.

The reported unemployment durations are affected by measurement error. In Illinois UI recipients must confirm their unemployed status every two weeks by sending in a certification form. This induces rounding to even weeks, so that the hazard is higher in such weeks. To deal with this rounding, it is important to allow for a flexible baseline hazard, *e.g.* a piecewise constant baseline hazard with weekly duration intervals (this is also Meyer's choice).

In table 4.3 we report the results of a probit analysis of the compliance decision. The estimates show that compliance is related to observed characteristics of the unemployed. The model does not predict compliance well[16]. Other (unobserved) variables seem to be important. Whether unobserved characteristics that also affect the re–employment rate play a role in the compliance decision can only be investigated with an estimator that is consistent even with selective compliance.

4.3 Estimation

We use the 2SLR estimator to estimate the effect of the claimant and employer bonus on the unemployment duration, and we compare these estimates with the ML and ITT estimates. In the first stage of the 2SLR procedure we use the control group data to estimate the regression parameters and the parameters of the baseline hazard. We include the following explanatory variables: age (AGE), the logarithm of the pre–unemployment earnings (LNBPE), gender (MALE= 1), ethnicity (BLACK= 1), and the logarithm of the weekly amount of UI benefits plus dependence allowance (LNBEN).

We employ two specifications for the baseline hazard. In the flexible specification we choose a piecewise constant baseline hazard on small duration intervals. In the data, spells are observed in weeks and all spells are censored at 26 weeks. The

[16] McFadden's R^2 is .007 for the claimant and .014 for the employer bonus sample

most flexible baseline hazard has 26 parameters α_k that are the logs of the (time-average) baseline hazard values in the first 26 weeks. We consider two specifications for the effect of the bonus: (1) the bonus affects the re-employment rate during the whole unemployment period and (2) the bonus has an effect on the re-employment hazard only during the first 11 weeks of unemployment. The second specification is consistent with the design of the program.

The results with a flexible specification of the duration dependence are reported in tables 4.4 and 4.5. The parameters of the baseline hazard are reported in table 4.9, and the implied baseline hazard is depicted in figure 4.1. Again, the MLE for the baseline hazard and the regression parameters for the 2SLR procedure are the first-stage MLE in the control group. These estimates are not directly comparable to the other columns.

A comparison of the results in table 4.4 and table 4.5 (the last two columns) show that both the ML and ITT estimators underestimate the program effect. In particular, the bias in the effect of the employer bonus is sizable. The 2SLR estimator of the effect is significantly different from 0, while the ITT and ML estimates are not. Although the effect of the claimant bonus is still larger, the relative difference between the 2SLR estimates of the effects is much smaller than that of the biased MLE's. In the MPH model with flexible duration dependence, we cannot reject the hypothesis of no unobserved heterogeneity. As argued in section 3, this preliminary test may be biased and says little about the bias in the ML estimate of the program effect. Note that the estimated effects are larger, if we allow for the finite duration of the bonus offer.

To check the sensitivity of the 2SLR estimate to the specification of the baseline hazard, we also estimate the program effects using a simpler MPH model in the first stage: an exponential model without unobserved heterogeneity. The results for this specification are reported in the first two columns of table 4.5. The top of the table shows the first-stage MLE. We find that, with the Illinois data, the program estimates are not very sensitive to the specification of the first stage MPH model.

4.4 Effect heterogeneity: program effects in subgroups

Until now we have assumed that the program effects are identical for all individuals. We now allow for effect heterogeneity by observed characteristics of the individuals. In particular, we stratify the sample on observed characteristics, and we compute separate 2SLR estimates on the subgroups. In our choice of observed characteristics, we shall follow previous research. We are particularly interested in the question, whether some of the conclusions are affected by selectivity bias. This could be the result of differences in the selectivity of compliance in subgroups.

Previous studies have considered effect heterogeneity by previous earnings (Meyer (1996)) and by the probability of benefit exhaustion. Meyer (1996) shows that both labor supply theory and search theory predict a larger effect of the bonus on lower income groups. O'Leary, Decker, and Wandner (1998) study the targeting of the bonus to groups for which the bonus has a relatively large effect, in order to increase the cost effectiveness of the bonus program. They distinguish the unemployed by their probability of benefit exhaustion, and they search for the subgroup with the largest program effect.

Meyer could not find a compelling relationship between previous earnings and the effect of the (claimant) bonus. However, previous earnings are most likely correlated with the unobserved characteristics that influence both the compliance decision and the unemployment duration. Therefore, the predicted log wage, computed from an OLS regression of the logarithm of pre–unemployment earnings on characteristics, may be a better choice to study the relation between the program effect and previous earnings.

The OLS estimates are shown in table 4.6. We divide the individuals into 3 groups based on quantiles of the distribution of the predicted log wage. We estimate the program effects by 2SLR and by ML on the three subsamples. The results in table 4.8 show that the effect of the claimant bonus does not vary with predicted previous earnings. The smaller effect for the top 25% in the MLE is spurious. The effect of the employer bonus decreases with predicted previous earnings.

Profiling is now used in all states as part of the Worker Profiling and Reemployment Services system. It involves predicting an individual's probability of exhausting UI benefits based on a logit or probit model on historical data for the state. O'Leary, Decker, and Wandner (1998) simulate the use of a profiling mechanism for the re–employment bonus experiment. In our estimates, the subgroups are selected with the predicted benefit exhaustion probabilities computed with the estimates in table 4.7. We use quantiles of the distribution of predicted probabilities to divide the sample. We distinguish the bottom 50%, the next 50% to 75% and the top 25% of this distribution. The results of the ML and 2SLR estimates on these three subsets are shown in the bottom half of table 4.8. For both the claimant and employer bonus the effect is largest for individuals with a benefit exhaustion probability between the median and the third quartile. Note that the relation between the exhaustion probability and the bonus effect is rather different in the MLE. Finally, note that the effect of the claimant and employer bonus is almost identical for the individuals who have a benefit exhaustion probability above the median. The difference between the estimated program effects is concentrated in the lower half of the distribution.

5 Conclusion

In this paper we have proposed and implemented an instrumental variable estimator that generalizes an estimator proposed by Robins and Tsiatis (1991) to MPH models. Sampling experiments indicate that the estimator performs well for sample sizes that are encountered in applications. A re-analysis of data from the Illinois reemployment bonus experiment shows that the ML and ITT estimates are downward biased, most likely due to selective non-compliance. We also find that conclusions on effect heterogeneity are different for the 2SLR and ML estimators.

Some issues have to be addressed in future research. First, the efficiency of the 2SLR procedure can can be increased by replacing the instrument, the randomization indicator R, by an appropriately chosen function $g(u, R)$. The choice $g(u, R) = R$ gives an efficient estimator if the transformed duration has a standard exponential distribution, *e.g.* if there is no unobserved heterogeneity. Second, the 2SLR estimator is a device to reduce the computational burden by dividing the computation in two steps. This is appealing because in the 2SLR is the solution to an equation that is discontinuous in the parameters. It is possible to estimate the program effect, the regression parameters, and the parameters of the base–line hazard simultaneously. However, the choice of instruments requires knowledge of the distribution of the transformed duration.

References

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., G. W. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables (with discussion). *Journal of the American Statistical Association 91*, 444–472.

Baker, M. and A. Melino (1997). Duration dependence and nonparametric heterogeneity: A Monte Carlo study. Working paper, Department of Economics, University of Toronto.

Cox, D. R. (1972). Regression models and life–tables (with discussion). *Journal of the Royal Statistical Society: Series B 34*, 187–220.

Dehejia, R. H. and S. Wahba (1995). Causal effect in non–experimental studies. Working paper, Department of Economics, Harvard University.

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.

Heckman, J., Ichimura, J. A. Smith, and Todd (1995). Non–parametric characterization of selection bias using experimental data: A case study of adult males in JTPA. Working paper, University of Chicago.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica 47*, 153–161.

Heckman, J. J. (1997). Instrumental variables: A study of implicit behavioral assumptions used in making program evaluations. *Journal of Human Resources 32*, 441–462.

Heckman, J. J. and B. Singer (1984). A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica 52*, 271–320.

Horowitz, J. (1996). Semiparametric estimation of a regression model with an unknown transformation of the dependent variable. *Econometrica 64*, 103–137.

Horowitz, J. L. and C. F. Manski (1997). Nonparametric analysis of randomized experiments with missing covariate and outcome data. Working paper 97–12, Department of Economics, College of Business Administration, University of Iowa.

Lancaster, T. (1979). Econometric methods for the duration of unemployment. *Econometrica 47*, 939–956.

Manski, C. F. (1975). Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics 3*, 205–228.

Manton, K. G., E. Stallard, and J. W. Vaupel (1981). Methods for the mortality experience of heterogeous populations. *Demography 18*, 389–410.

Meyer, B. D. (1988). Implications of the Illinois reemployment bonus experiments for theories of unemployment and policy design. Working paper, no. 2783, NBER.

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.

Moffitt, R. (1983). An economic model of welfare stigma. *American Economic Review 73*, 1023–1035.

Nielsen, G. G., R. D. Gill, P. K. Andersen, and T. A. I. Sørensen (1992). A counting process approach to maximum likelihood estimation in frailty models. *Scandinavian Journal of Statistics 19*, 25–43.

O'Leary, C. J., P. Decker, and S. A. Wandner (1998). Reemployment bonuses and profiling. Staff working paper 98–51, W. E. Upjohn Institute, Kalamazoo, Michigan.

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

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

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.

Table 3.1: Sampling experiments: 2SLR, ITT and ML estimator of program effect; 8000 observations and 100 replications; true values between parentheses.

	ave. bias	std. error bias	std. error est.	asymp. std. error	RMSE
			Exogenous Compliance		
2SLR (.25)	0.0136	0.0089	0.0894	0.1018	0.0905
MLE (.25)	0.0023	0.0060	0.0597	0.0580	0.0597
ITT (.25)	-0.0594	0.0056	0.0559	0.0541	0.0816
MLE					
β (.2)	0.0008	0.0011	0.0110	0.0103	0.0110
α_1 (-2.4)	0.0327	0.0131	0.1311	0.0895	0.1351
α_2 (-2.7)	0.0421	0.0150	0.1499	0.1096	0.1557
α_3 (-2.7)	0.0418	0.0154	0.1535	0.1228	0.1591
α_4 (-2.3)	0.0390	0.0170	0.1704	0.1448	0.1748
$\ln(v_1)$ (-1.39)	-0.1320	0.0364	0.3473	2.7190	0.3715
$\ln(v_2)$ (.92)	-0.2973	0.0711	0.6779	1.1070	0.7402
$\ln(v_3)$ (1.70)	0.2152	0.0979	0.9339	0.7910	0.9584
δ_1 (1.50)	-0.1709	0.0476	0.4539	0.7027	0.4850
δ_2 (-.83)	0.1694	0.0423	0.4035	0.6521	0.4377
			Endogenous Compliance		
2SLR (.25)	-0.0102	0.0088	0.0882	0.0942	0.0887
MLE (.25)	-0.6833	0.0061	0.0613	0.0626	0.6861
ITT (.25)	-0.0878	0.0057	0.0565	0.0532	0.1044
MLE					
β (.2)	0.0341	0.0013	0.0134	0.0115	0.0366
α_1(-2.4)	0.1699	0.0095	0.0946	0.0762	0.1944
α_2 (-2.7)	0.1733	0.0103	0.1034	0.1012	0.2018
α_3 (-2.7)	0.0121	0.0122	0.1217	0.1163	0.1223
α_4 (-2.3)	0.0566	0.0151	0.1510	0.1431	0.1612
$\ln(v_1)$ (-1.39)	-0.1928	0.0471	0.4595	0.3444	0.4983
$\ln(v_2)$ (.92)	-0.5927	0.0817	0.7960	0.5228	0.9924
$\ln(v_3)$ (1.70)	-0.0690	0.0927	0.9040	0.4823	0.9066
δ_1 (1.50)	-0.5994	0.0464	0.4522	0.4209	0.7509
δ_2 (-.83)	0.4314	0.0346	0.3373	0.3809	0.5476

Table 4.1: Descriptive statistics for Control, Claimant Bonus and Employer Bonus group (standard error of average).

	Control Group		Claimant Bonus		Employer Bonus	
	N	fraction	N	fraction	N	fraction
White	2497	0.632	2723	0.651	2565	0.647
Black	1072	0.271	1050	0.251	1014	0.256
Other	383	0.097	413	0.099	384	0.097
Male	2162	0.547	2357	0.563	2131	0.538
Age 20–29	1680	0.425	1827	0.436	1679	0.424
Age 30–39	1315	0.333	1357	0.324	1292	0.326
Age 40–49	708	0.179	776	0.185	740	0.187
Age 50–54	248	0.063	226	0.054	252	0.064
Weekly benefit						
-$51	347	0.088	355	0.085	333	0.084
$52–$90	794	0.201	887	0.212	861	0.217
$91–$120	666	0.169	738	0.176	711	0.179
$121–$160	749	0.190	822	0.196	716	0.181
$161–	1396	0.353	1384	0.331	1342	0.339
Dependence allowance	1834	0.323	1955	0.345	1883	0.332
Average pre–claim earnings		3188 (35.89)		3222 (36.91)		3215 (37.83)
Average age		33.0 (0.14)		32.9 (0.20)		33.1 (0.21)
Average weekly benefit		119.9 (0.65)		118.8 (0.63)		118.5 (0.64)

Table 4.2: Average unemployment durations: control group and (non-)compliers (standard error of average).

	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

Table 4.3: Probit analysis of the compliance decision.

	Claimant Bonus	Employer Bonus
Constant	0.9973	0.3108
	(0.0384)	(0.0329)
AGE	-0.0029	-0.0025
	(0.0027)	(0.0024)
LNBPE	0.0763	-0.1328
	(0.0544)	(0.0559)
BLACK	-0.1866	0.0565
	(0.0527)	(0.0484)
MALE	0.1073	0.1360
	(0.0478)	(0.0421)
LNBEN	-0.1973	-0.1535
	(0.0958)	(0.0939)
Log likelihood	-1810.21	-2524.71
LR test (5 d.f.)	24.60	69.69
N	4186	3963
No. of compliers	3527	2586

Table 4.4: Proportional hazard model with flexible duration dependence: regression parameters and program effects for ITT and ML (parameters baseline hazard in Table 4.9). (1) Model with time-constant treatment effect (2) Model with treatment effect only during first 11 weeks.

	ML		ITT	
	(1)	(2)	(1)	(2)
AGE	-0.4299	-0.4298	-0.4316	-0.4313
	(0.0472)	(0.0473)	(0.0472)	(0.0472)
LNBPE	0.2608	0.2607	0.2592	0.2591
	(0.0334)	(0.0335)	(0.0334)	(0.0334)
BLACK	-0.5064	-0.5062	-0.5069	-0.5067
	(0.0296)	(0.0296)	(0.0296)	(0.0296)
MALE	0.0657	0.0653	0.0675	0.0675
	(0.0242)	(0.0242)	(0.0242)	(0.0242)
LNBEN	-0.4718	-0.4710	-0.4710	-0.4712
	(0.0549)	(0.0549)	(0.0549)	(0.0549)
Claim. bonus	0.1029	0.1666	0.1104	0.1672
	(0.0274)	(0.0340)	(0.0290)	(0.0362)
Empl. bonus	0.0382	0.0961	0.0555	0.0981
	(0.0307)	(0.0384)	(0.0296)	(0.0371)
Log L	-29872	-29866	-29871	-29868

Table 4.5: Sensitivity of 2SLR to the specification of the MPH model; first-stage MLE of regression parameters and 2SLR of program effect (parameters of baseline hazard in Table 4.9). (1) Model time-constant treatment effect (2) Model with treatment effect only on first 11 weeks.

		Exponential		Duration dependence	
α		-3.2828	(0.0322)	.	.
AGE		-0.4476	(0.0790)	-0.4247	(0.0835)
LNBPE		0.3153	(0.0566)	0.2969	(0.0600)
BLACK		-0.5309	(0.0501)	-0.5071	(0.0521)
MALE		0.1062	(0.0409)	0.1027	(0.0432)
LNBEN		-0.5951	(0.0928)	-0.5648	(0.0982)
Log L		-9817.60		-9594.37	
Claim. bonus	(1)	0.1446	(0.0342)	0.1312	(0.0346)
	(2)	0.2205	(0.0529)	0.2140	(0.0483)
Empl. bonus	(1)	0.1011	(0.0451)	0.0926	(0.0452)
	(2)	0.1582	(0.0732)	0.1517	(0.0658)

Table 4.6: OLS regression of ln(BPE).

Constant	5.5312
	(0.0932)
AGE	0.1128
	(0.0055)
AGE^2	-0.0013
	(0.0001)
BLACK	-0.2206
	(0.0141)
MALE	0.2286
	(0.0124)
R^2	0.132
N	12074

Table 4.7: Probit analysis of benefit exhaustion probability.

Constant	-0.2640
	(0.0331)
AGE	0.3896
	(0.0809)
LNBPE	-0.2907
	(0.0576)
BLACK	0.4547
	(0.0463)
MALE	-0.1140
	(0.0417)
LNBEN	0.5185
	(0.0951)
Log L	-2596.98
LR test (5 d.f.)	180.11
N	3944
No. benefit exhaustion	1669

Table 4.8: The effect of bonus on subgroups. Groups distinguished by quantiles of (1) Predicted ln(BPE) and (2) Estimated probability of benefit exhaustion.

	2SLR		MLE	
	Claimant	Employer	Claimant	Employer
Sample used	Predicted ln(BPE)			
Bottom 25%	0.2140	0.2140	0.1730	0.1119
	(0.0990)	(0.1197)	(0.0678)	(0.0740)
25%-75%	0.2140	0.1517	0.2163	0.1268
	(0.0664)	(0.0938)	(0.0483)	(0.0547)
Top 25%	0.1975	0.0649	0.0649	-0.0001
	(0.0985)	(0.1456)	(0.0686)	(0.0796)
	Probability of benefit exhaustion			
Bottom 50%	0.1287	0.0077	0.1814	0.1375
	(0.0648)	(0.0971)	(0.0444)	(0.0504)
50% - 75%	0.3192	0.3122	0.1527	0.0651
	(0.0913)	(0.1145)	(0.0718)	(0.0791)
Top 25%	0.2655	0.2733	0.1568	0.0212
	(0.1174)	(0.1440)	(0.0797)	(0.0895)

Table 4.9: Parameters baseline hazard: Control group (first stage of 2SLR), MLE and ITT. (1) Model with time-constant treatment effect (2) Model with treatment effect during first 11 weeks. Log baseline hazard values in weeks 1-26.

	Control		ML				ITT			
			(1)		(2)		(1)		(2)	
α_1	-2.4584	(0.0629)	-2.4073	(0.0369)	-2.4400	(0.0381)	-2.4264	(0.0392)	-2.4614	(0.0414)
α_2	-2.8754	(0.0777)	-2.7897	(0.0440)	-2.8221	(0.0450)	-2.8088	(0.0460)	-2.8437	(0.0480)
α_3	-3.5240	(0.1063)	-3.3484	(0.0570)	-3.3806	(0.0580)	-3.3675	(0.0583)	-3.4023	(0.0600)
α_4	-2.8171	(0.0787)	-2.8267	(0.0464)	-2.8588	(0.0473)	-2.8459	(0.0481)	-2.8806	(0.0499)
α_5	-3.6049	(0.1152)	-3.4278	(0.0616)	-3.4598	(0.0625)	-3.4470	(0.0629)	-3.4816	(0.0645)
α_6	-3.0442	(0.0899)	-3.0142	(0.0520)	-3.0461	(0.0528)	-3.0334	(0.0536)	-3.0680	(0.0552)
α_7	-3.6219	(0.1198)	-3.5161	(0.0666)	-3.5479	(0.0673)	-3.5354	(0.0677)	-3.5700	(0.0691)
α_8	-3.1592	(0.0977)	-3.0753	(0.0554)	-3.1069	(0.0562)	-3.0945	(0.0570)	-3.1290	(0.0585)
α_9	-3.7841	(0.1340)	-3.6864	(0.0745)	-3.7180	(0.0751)	-3.7057	(0.0756)	-3.7401	(0.0768)
α_{10}	-3.1876	(0.1026)	-3.1399	(0.0589)	-3.1714	(0.0596)	-3.1591	(0.0604)	-3.1934	(0.0618)
α_{11}	-3.7723	(0.1378)	-3.5995	(0.0741)	-3.6310	(0.0747)	-3.6187	(0.0752)	-3.6530	(0.0764)
α_{12}	-3.4071	(0.1169)	-3.3951	(0.0682)	-3.3572	(0.0672)	-3.4142	(0.0694)	-3.3583	(0.0672)
α_{13}	-3.8524	(0.1466)	-3.8694	(0.0862)	-3.8314	(0.0854)	-3.8885	(0.0872)	-3.8326	(0.0854)
α_{14}	-3.2628	(0.1125)	-3.1994	(0.0642)	-3.1615	(0.0631)	-3.2185	(0.0656)	-3.1627	(0.0631)
α_{15}	-3.9029	(0.1546)	-3.8571	(0.0883)	-3.8192	(0.0876)	-3.8761	(0.0893)	-3.8204	(0.0876)
α_{16}	-3.5458	(0.1318)	-3.3996	(0.0723)	-3.3618	(0.0713)	-3.4187	(0.0735)	-3.3630	(0.0713)
α_{17}	-3.9537	(0.1624)	-3.9387	(0.0943)	-3.9010	(0.0936)	-3.9577	(0.0951)	-3.9021	(0.0936)
α_{18}	-3.4659	(0.1302)	-3.4386	(0.0755)	-3.4009	(0.0745)	-3.4576	(0.0765)	-3.4020	(0.0745)
α_{19}	-3.9843	(0.1684)	-3.9393	(0.0968)	-3.9016	(0.0961)	-3.9582	(0.0977)	-3.9027	(0.0961)
α_{20}	-3.3508	(0.1264)	-3.3857	(0.0755)	-3.3481	(0.0745)	-3.4047	(0.0766)	-3.3492	(0.0745)
α_{21}	-3.7736	(0.1566)	-3.8413	(0.0948)	-3.8036	(0.0940)	-3.8603	(0.0956)	-3.8047	(0.0940)
α_{22}	-3.0442	(0.1126)	-3.3075	(0.0747)	-3.2697	(0.0738)	-3.3266	(0.0757)	-3.2709	(0.0738)
α_{23}	-3.9743	(0.1785)	-3.9616	(0.1033)	-3.9238	(0.1026)	-3.9808	(0.1041)	-3.9249	(0.1025)
α_{24}	-3.1769	(0.1233)	-3.1990	(0.0729)	-3.1611	(0.0719)	-3.2181	(0.0740)	-3.1622	(0.0719)
α_{25}	-3.5932	(0.1532)	-3.6702	(0.0926)	-3.6323	(0.0918)	-3.6892	(0.0934)	-3.6334	(0.0918)
α_{26}	-2.7812	(0.1063)	-2.8768	(0.0650)	-2.8389	(0.0640)	-2.8958	(0.0662)	-2.8400	(0.0640)