Endogenous Attrition in Panels

Xavier D'HAULTFOEUILLE(*), Laurent DAVEZIES(*)
(*) CREST

1 Introduction

Panel data are very useful to distinguish between state dependence and unobserved heterogeneity (see, e.g., [11]), to analyze the dynamics of variables such as income (see, e.g., [9]) or spells in duration analysis (see, e.g., [15]). However, these advantages may be counterbalanced by attrition, which can be especially severe when units are observed over a long period of time. Besides, attrition is often considered more problematic than standard nonresponse, because the reasons of attrition are often related to the outcomes of interest, or variations in these outcomes. Several solutions have been considered in the literature to handle these issues. A first model is to suppose that attrition is exogenous, i.e. depends on lagged values that are observed by the econometrician (see, e.g., [16]). This, however, rules out a dependence between attrition and current outcomes, and is thus likely to fail in many cases. A second model takes the opposite point of view by assuming attrition to depend on contemporaneous values only (see [10]). Such an assumption fails to hold if attrition is related to transitions in outcomes. In the French labor force survey, for instance, households that move during the period are lost. But house moving is likely to be related to changes in the employment status. To handle more complex attrition patterns, [12] generalize the two previous models by allowing attrition to depend both on contemporaneous and lagged values. This generalization is made possible when a refreshment sample, i.e. a sample of new units surveyed at each period, is available.

In this paper, we consider still another approach, based on instruments. Contrary to [12], we do not impose any functional restrictions on the probability of attrition conditional on lagged and contemporaneous values. Refreshment sample are not needed either. On the other hand, we suppose to have in hand an instrument which is independent of attrition conditional on past and contemporaneous outcomes. A rank condition between the instrument and the contemporaneous outcome, which can be stated in terms of completeness, is also needed. Hence, the instrument is typically a lagged variable which affects the contemporaneous outcome but not directly attrition. We can use for instance past outcomes obtained from a retrospective questionnaire. We show indeed that under a nonlinear fixed effect model, such a variable is likely to meet the nonparametric rank condition, and satisfies also the conditional independence condition if attrition only depends on transitions on the outcome.

An advantage of our method is that even if no more instruments than outcomes are available, we can test for the conditional independence assumption. Another way of testing this assumption is to use refreshment samples, even though they are unnecessary in our setting. Indeed, the marginal distribution of the contemporaneous outcome is directly identified with such samples. Thus, we can compare this distribution with the one obtained under our identifying restriction.

We also conduct inference under such an attrition process. In the case of discrete outcomes and instruments, the model is parametric and a straightforward constrained maximum likelihood estimation procedure is proposed. In the continuous case, the model is semiparametric and estimation is more involved. We first provide a necessary and sufficient condition for root-n estimability of linear functionals, and compute under this condition the asymptotic efficiency bound. Our results are close to those obtained recently by [23] in the case of nonparametric instrumental regressions. Second, we propose two estimation methods to estimate such linear functionals. The first is efficient but relies on rather restrictive conditions, whereas the second is not efficient but consistent under mild assumptions.

Finally, we apply our results to study transitions on the French labor market, using the labor force survey of the French national institute of statistics (INSEE). This survey is one of the most important survey conducted by INSEE, but its reliability has been much questioned inside the institute by the end of 2006 and the beginning of 2007 (i.e., during the French presidential elections campaign), as the discrepancy between the INSEE unemployment rate estimate and the one coming from administrative data started to increase. We reinvestigate this issue by studying the nature of attrition in this survey. Using the refreshment sample, we test and accept on the subsample of women the conditional independence assumption with past employment status used as an instrument. Our estimates indicate that attrition is highly related to transitions in the labor market, in a way that violated the additive restriction considered by [12]. We show that this has important implications for the estimation of the probabilities of transition on the labor market.

The paper is organized as follows. In the second section, we study identification and testability under endogenous attrition, and compare our model with the existing literature. In the third section, we develop inference for both discrete or continuous outcomes. The fourth section is devoted to our application. Finally, the fifth section concludes. All proofs are gathered in the appendix.

2 Identification

2.1 The setting and main result

For simplicity, we consider a panel dataset with two dates $t = 1, 2$, and also suppose that there is no or ignorable nonresponse at date 1. We let $D = 1$ if the unit is observed at date 2, $D = 0$ otherwise. We let Y_t denote the outcome at t and $Y = (Y_1, Y_2)$. We also consider an instrument Z_1 whose role will be explained below, and let $Z = (Y_1, Z_1)$. We focus hereafter on the identification of either the joint distribution of (D, Y, Z) or on a parameter $\beta_0 = E(g(Y, Z))$. Our first assumption states the observational problem.

Assumption 2.1. *We observe* (D, Z) *and* Y_2 *when* $D = 1$.

To satisfy this requirement, Z_1 should be observed at the first period, or at the second period if some information on nonrespondents at the second period is available.[1] Of course, to achieve full identification of (D, Y, Z), restrictions are needed on the distribution of (D, Y, Z). If attrition directly depends on the outcome Y, the usual assumption of exogenous selection fails, and it may be difficult to find an instrument that affects the selection variable but not the outcome. On the other hand, a variable Z_1 related to Y but not directly to D may be available in this case. We thus assume the following :[2]

Assumption 2.2. *$D \perp\!\!\!\perp Z_1 | Y$.*

This assumption is identical to the one considered by [6] in the case of endogenous selection. This assumption was also considered by [5], [24] and [19] in a nonresponse framework. Intuitively, this assumption states that the attrition equation depends on Y_1 and Y_2 but not on Z_1. If Y_2 was endogenous (but always observed) in this equation, we could instrument it by Z_1 to identify the causal effect of Y_2 on D. Here the problem is actually slightly different : Y_2 is observed only when $D = 1$. The identification strategy will be similar, however, as we will use the instrument to recover the conditional distribution of attrition.

Let $P(Y) = P(D = 1|Y)$. Because identification is based on inverse probability weighted moment conditions, we assume the following :

Assumption 2.3. *(i) $P(Y) > 0$ almost surely.*

This assumption is similar to the common support condition in the treatment effects literature. It does not hold if D is a deterministic function of Y, as in simple truncation models where $D = \mathbb{1}\{g(Y) \geq y_0\}$, y_0 denoting a fixed threshold.

Before stating our main result, let us introduce some notations. For any random variable U and $p > 0$, let $L^p(U)$ (respectively $L^p(U|D = 1)$) denote the space of functions q satisfying $E(|q(U)|^p) < +\infty$ (respectively $E(|q(U)|^p|D = 1) < +\infty$). Note that $1/P \in L^1(Y|D = 1)$ because $E(1/P(Y)|D = 1) = 1/E(D)$. For any set $A \subset L^1(U|D = 1)$, let also

$$A^\perp = \{q \in L^1(U|D = 1) : \forall a \in A, E(|q(U)a(U)||D = 1) < \infty \; E(q(U)a(U)|D = 1) = 0\}.$$

The following operator will be important for identification issues :

$$\begin{array}{rlcl} T: & L^1(Y|D=1) & \to & L^1(Z|D=1) \\ & q & \mapsto & (z \mapsto E(q(Y)|D=1, Z=z)) \,. \end{array}$$

Because Y is observed when $D = 1$, T is identified. Besides, and as indicated previously, identification hinges upon dependence conditions between Y_2 and Z, which are actually related to the null space Ker (T) of T. Let $\mathcal{F} = \{q \in L^1(Y|D = 1) : q(Y) \geq 1 - 1/P(Y) \text{ a.s.}\}$ and for $f \in L^1(Y, Z)$,

$$\mathcal{F}_f = \left\{q \in L^1(Y|D = 1) : q(Y) \geq 1 - 1/P(Y) \text{ a.s. and } E(|q(Y)f(Y, Z)||D = 1) < \infty\right\}.$$

Finally, in the case where $g \in L^1(Y, Z)$ we denote $\beta(Y) = E[g(Y, Z)|Y]$. Our main result is the following.

1. Actually, our results would hold if Z_1 is observed only when $D = 1$, provided that the distribution of (Y_1, Z) is identified (through another dataset for instance).

2. Assumption 2.2 could be extended to include covariates, i.e. $D \perp\!\!\!\perp Z_1|Y, X$, provided that X is always observed. We do not include them for the sake of simplicity.

Théorème 2.1. *If assumptions 2.1-2.3 hold, then :*

1. *The distribution of (D, Y, Z) is identified if and only if $Ker(T) \cap \mathcal{F} = \{0\}$.*

Moreover, if $g \in L^1(Y, Z)$ we have :

2. *The set of identification of β_0 is $\{\beta_0 + E(D)E[\beta(Y)h(Y)|D = 1] : h \in Ker(T) \cap \mathcal{F}_g\}$.*
3. *β_0 is identified if and only if $\beta(.) \in (Ker(T) \cap \mathcal{F}_g)^{\perp}$.*

Let us provide the intuition for the easiest result, i.e. the "if" part of the first statement. We rely on the fact that under Assumptions 2.2 and 2.3(i), it is sufficient to identify $P(Y)$ to recover the whole distribution of (D, Y, Z). Besides, we show that this function satisfies

$$T\left(\frac{1}{P}\right) = w, \tag{2.1}$$

where $w(Z) = 1/P(D = 1|Z)$. Because T and w are identified, P is identified if there is a unique solution in $(0, 1]$ of this equation. This uniqueness can be established if $\text{Ker}(T) \cap \mathcal{F} = \{0\}$.

The identifying condition $\text{Ker}(T) \cap \mathcal{F} = \{0\}$ is related to various completeness conditions considered in the literature (see, e.g., [18], [22], [4], and [7]).[3] When Y and Z have a finite support (respectively by $(1, ..., I)$ and $(1, ..., J)$), this assumption amounts to $\text{rank}(M) = I$, where M is the matrix of typical element $P(Y = i|D = 1, Z = j)$ (see [18]). Hence, the support of Z must be at least as rich as the one of Y $(J \geq I)$ and the dependence between the two variables must be strong enough for I linearly independent conditional distributions to exist. Because the matrix M is identified, it is straightforward to test for this condition, using for instance the determinant of MM' (see Subsection 3.1 below). When Y and Z are continuous, it is far more difficult to characterize them. Conditions have been provided by [18], [6] and [7]. We consider below another example, related to our panel framework, where the restriction $\text{Ker}(T) \cap \mathcal{F} = \{0\}$ is satisfied.

The third statement of the theorem shows that when we specialize in one parameter rather than on the full distribution of (D, Y, Z), identification is achieved under weaker restrictions. Indeed, $\mathcal{F}_g \subset \mathcal{F}$ and then $(\text{Ker}(T) \cap \mathcal{F})^{\perp} \subset (\text{Ker}(T) \cap \mathcal{F}_g)^{\perp}$. Thus, $\text{Ker}(T) \cap \mathcal{F} = \{0\}$ implies that $\beta(.) \in (\text{Ker}(T) \cap \mathcal{F}_g)^{\perp}$ but the contrary needs not be true. This result is closely related to Lemma 2.1 of [23], who consider identification of linear functionals related to a nonparametric instrumental regression. Finally, the second statement of the theorem describes the identification set of β_0 in general.

As an illustration of Theorem 2.1 with continuous outcomes, suppose that we observe at the first date a past outcome Y_0, thanks to a retrospective questionnaire. This will be the case in the application considered in Section 4. Suppose also that the outcomes satisfy the following nonlinear fixed effect model :

$$\Lambda(Y_t) = U + \varepsilon_t, \tag{2.2}$$

where $\Lambda(.)$ is a strictly increasing real function and $(U, \varepsilon_0, \varepsilon_1, \varepsilon_2)$ are independent. Such a model generalizes standard linear fixed effect model $Y_t = U + \varepsilon_t$ and is close to the accelerated failure time model in duration analysis. Note that we do not introduce covariates here for simplicity, but our result can be extended to the more realistic model considered

3. Our condition is intermediate between the stronger "standard" completeness condition $\text{Ker}(T) = \{0\}$ and the bounded completeness condition $\text{Ker}(T) \cap \mathcal{B} = \{0\}$, where $\mathcal{B}$ is the set of bounded functions.

by [8], namely $\Lambda(Y_t, X_t) = \psi(U, X_t) + \varepsilon_t$ with Λ strictly increasing in Y_t, provided that the covariates X_t are always observed at each period.

We also suppose that attrition only depends on current outcomes and transitions :

$$D = g(Y_1, Y_2, \eta), \quad \eta \perp\!\!\!\perp (Y_0, Y_1, Y_2). \tag{2.3}$$

Finally, we impose the following technical restriction on U, ε_0 and ε_2. For any random variable V, we let Ψ_V denote its characteristic function.

Assumption 2.4. *U admits a density with respect to the Lebesgue measure, whose support is the real line. Ψ_{ε_0} vanishes only on isolated points. The distribution of ε_2 admits a continuous density f_{ε_2} with respect to the Lebesgue measure. Moreover, $f_{\varepsilon_2}(0) > 0$ and there exists $\alpha > 2$ such that $t \mapsto t^\alpha f_{\varepsilon_2}(t)$ is bounded. Lastly, Ψ_{ε_2} does not vanish and is infinitely often differentiable in $\mathbb{R} \backslash A$ for some finite set A.*

The assumption imposed on the characteristic function of ε_0 is very mild and satisfied by all standard distributions. The conditions on ε_2 are more restrictive but hold for many distributions such as the normal, the student with degrees of freedom greater than one [4] and the stable distributions with characteristic exponent greater than one. The following proposition shows that under these conditions, the model is fully identified using Y_0 as the instrument.

Proposition 2.2. *Let $Z = (Y_0, Y_1)$, and suppose that Assumptions 2.3(i), 2.4, Equations (2.2) and (2.3) hold. Then Assumption 2.2 holds and $Ker(T) \cap \mathcal{F} = \{0\}$. Thus, the distribution of (D, Y, Z) is identified.*

2.2 Partial identification and testability

Apart from point identification under various completeness conditions, our attrition model displays two interesting features. First, Assumption 2.2 is refutable, contrary to the ignorable attrition assumption considered above. Second, we can obtain bounds on parameters of interest when the model is underidentified, i.e. when the above completeness condition fails to hold. Both are due to the fact that solutions to Equation (2.1) must lie in $[0, 1]$. These inequality constraints can be used both for testing and bound parameters of interest.

To see this, consider the case where (Y, Z) has a finite support. If Y and Z take respectively I and J distinct values, then (2.1) can be written as a linear system of J equations with I unknown parameters and the inequality constraints xx. Of course, the model is overidentified and thus testable when $I > J$, but we can also test for the inequality constraints when $I \leq J$. We derive a formal statistical test of this condition in Subsection 3.1 below. We can also partially identify parameters of interest in the underidentified case $I < J$. DETAILLER

Finally, a stronger test of the conditional independence assumption can be derived if a refreshment sample is available, as in [12]. In this case, the marginal distribution of Y_2 is identified. But then we can reject the conditional independence assumption if for all Q satisfying $T(1/Q) = w$, there exists t such that

$$E\left[\frac{D\mathbb{1}\{Y_2 \leq t\}}{Q(Y)}\right] \neq P(Y_2 \leq t).$$

4. See e.g. [17] for a proof that the conditions on the characteristic function of student distributions are indeed satisfied.

2.3 Comparison with the literature

We compare our approach with the most usual models of attrition.

2.3.1 Missing at random attrition

This model, which has been considered by, e.g., [21] and [1], posits that D only depends on Y_1 :

$$D \perp\!\!\!\perp Y_2 | Y_1. \tag{2.4}$$

Identification of the joint distribution of (Y_1, Y_2) follows directly from the fact that, letting f_{D,Y_1,Y_2} denote the density of (D, Y_1, Y_2) with respect to an appropriate measure,

$$f_{Y_1,Y_2}(y_1, y_2) = \frac{f_{D,Y_1,Y_2}(1, y_1, y_2)}{P(D = 1|Y_1 = y_1)}.$$

Condition (2.4) is the equivalent, in a panel setting, of the so-called missing at random assumption (see, e.g., [16]) or the unconfoundedness assumption in the treatment effect literature (see for instance [13]). Because it rules out any dependence between attrition and current outcomes, it is likely to fail in many cases. In a labor force survey, for instance, house moving is a common source of attrition, and is itself related to changes in employment and/or earnings.

2.3.2 Dependence on current values

Compared to the first, the logic of this model is the opposite, as attrition is related to current values only :

$$D \perp\!\!\!\perp Y_1 | Y_2. \tag{2.5}$$

This assumption has been considered by [10] in a parametric model. This assumption takes into account nonignorable attrition, but in a special way. Indeed, it rules out the possibility that transitions (i.e., functions of (Y_1, Y_2)) explain attrition. Abstracting from the parametric restrictions of [10], identification can be proved along the same lines as previously. It suffices indeed to solve in g the functional equation

$$E\left[g(Y_2)|D = 1, Y_1\right] = 1/P(D = 1|Y_1).$$

Under completeness conditions similar to the one above, this equation admits a unique solution in g, namely $1/P(D = 1|Y_2 = .)$.

2.3.3 Additive restriction on the probability of attrition

[12] propose a two period framework which generalize both previous examples in the sense that D may depend on both Y_1 and Y_2. This generalization is possible when a refreshment sample, which allows one to identify directly the distribution of Y_2, is available. [5] They also suppose that

$$1/P(D = 1|Y_1, Y_2) = g(\alpha + k_1(Y_1) + k_2(Y_2)), \tag{2.6}$$

5. Note that because, the distribution of Y_1 is also identified from the panel at date 1, the problem reduces to recover the copula of (Y_1, Y_2).

where g is a known function, but $k_1(.)$ and $k_2(.)$ are unknown. They show that $k_1(.)$ and $k_2(.)$ are identified. This allows them to recover the joint distribution of (Y_1, Y_2), since by the Bayes' rule,

$$f_{Y_1,Y_2}(y_1, y_2) = f_{Y_1,Y_2|D=1}(y_1, y_2)P(D = 1)g(\alpha + k_1(y_1) + k_2(y_2)),$$

Compared to our approach, [12] do not rely on any exclusion restriction. This comes at the cost of imposing the additive restriction on $P(D = 1|Y_1, Y_2)$, which may be restrictive (see below), and having a refreshment sample, which is not needed in our case.

Though the identification proof of [12] is much different from ours, the two frameworks are actually related. As shown by [2], identification in this additive model can be directly obtained from the functional equations

$$E\left[g(k_1(Y_1) + k_2(Y_2))|D = 1, Y_i\right] = 1/P(D = 1|Y_i).$$

Thus, identification is actually achieved along similar lines as above, the instrument being equal to (Y_1, Y_2). The difference here is that only the marginal distributions of the instrument is identified. This is the reason why they have to impose Model (2.6) to the attrition process. Note that such a restriction is not innocuous. If attrition depends on transitions, then their restriction is likely fails to hold. If, as in our application, attrition occurs for individuals who move, and that moving itself occurs with a large probability when employment status changes, then $P(D = 1|Y_1, Y_2)$ depends on $\mathbb{1}\{Y_1 = Y_2\}$. Model (2.6) cannot handle such an attrition process.

3 Estimation

We now turn to inference within our framework of endogenous attrition. As previously, we focus on the estimation of the distribution of (D, Y, Z), but also on the parameter $\beta_0 = E(g(Y, Z))$, which can be estimated under restrictions detailed before. We first posit an i.i.d. sample of n observations.

Assumption 3.1. *We observe an iid sample* $((D_1, D_1Y_{21}, Z_1), ..., (D_n, D_nY_{2n}, Z_n))$.

We consider two cases subsequently. The first one, in line with our application, assume that the support of (Y, Z) is finite. Under this case we derive a simple and efficient estimator and a test of the exclusion condition. In a second time, we relax the finite support assumption and we exhibit a necessary condition for existence of a root-N estimator. Under this condition we derive the semi-parametric efficiency bounds and finally we propose a root-N estimator.

3.1 The finite support case

We denote the support of Y_t and Z_1 by respectively $\{1, ..., I\}$ and $\{1, ..., J\}$, with $I \leq J$.[6] In this case, the data (D, DY_2, Z) are distributed according to a multinomial distribution. To get asymptotic efficient estimators, we consider constrained maximum likelihood estimation hereafter.

6. Of course, our setting readily extends to a case where discrete covariates X are available, and Assumption 2.2 holds conditionally (i.e., $D \perp\!\!\!\perp Z|Y, X$).

For a fixed y, let $p_{1ij} = P(D = 1, Y_2 = i, Z_1 = j|Y_1 = y)$ and $p_{0.j} = P(D = 0, Z_1 = j|Y_1 = y)$ denote the probabilities corresponding to the observations, and define $p_1 = (p_{111}, ..., p_{1IJ})$, $p_{0.} = (p_{0.1}, ..., p_{0.J})$ and $p = (p_1, p_{0.})$. Note that we let the dependence in y implicit hereafter. p is the natural parameter of the statistical model here, as it fully describes the distribution of (D, DY_2, Z_1) (conditional on Y_1).[7] However, it does not directly allow us to recover the whole distribution of (D, Y_2, Z_1). This is why we also introduce $p_{0ij} = P(D = 0, Y_2 = i, Z_1 = j|Y_1 = y)$, and define p_0 as p_1. Then any parameter θ_0 of the distribution of (D, Y_2, Z_1) is a function of (p_0, p_1), and we write $\theta_0 = g(p_0, p_1)$.[8]

Finally, we adopt the same notations for the constrained maximum likelihood estimator $\widehat{p}$ as for p, and we let $n_{1ij} = \sum_{k:Y_{1k}=y} D_k \mathbb{1}\{Y_{2k} = i\}\mathbb{1}\{Z_{1k} = j\}$ and define $n_{0.j}$ accordingly. The following proposition shows how to compute $\widehat{p}$ and an efficient estimator of θ_0 in our attrition model.

Proposition 3.1. *Suppose that Assumptions 2.1-2.3(i) hold. Then the maximum likelihood estimator $\widehat{p}$ satisfies*

$$\widehat{p} = \arg \max_{(q,b)\in[0,1]^{(I+1)J}\times\mathbb{R}^I} \sum_{j=1}^{J} \left[n_{0.j} \ln q_{0.j} + \sum_{i=1}^{I} n_{1ij} \ln q_{1ij} \right]$$

$$s.t. \left| \begin{array}{ll} \sum_{j=1}^{J} \left[q_{0.j} + \sum_{i=1}^{I} q_{1ij} \right] = 1, & \\ b_i \geq 0 & i = 1, ..., I, \\ \sum_{i=1}^{I} q_{1ij} b_i = q_{0.j} & j = 1, ..., J. \end{array} \right. \quad (C)$$

Suppose moreover that the matrix P_1 of typical element p_{1ij} has rank I and g is differentiable. Then θ_0 is identifiable and can be estimated efficiently by

$$\widehat{\theta} = g(\widehat{p}_0, \widehat{p}_1),$$

where $\widehat{p}_{0ij} = \widehat{b}_i \widehat{p}_{1ij}$, and $\widehat{b} = (\widehat{b}_1, ..., \widehat{b}_I)$ is a solution of constraints (C) taken at $q = \widehat{p}$.

Proposition 3.1 establishes that the maximum likelihood of p can be obtained by a constrained maximization with quite simple (although nonlinear) constraints. It also shows how to compute an efficient estimator of θ_0. The idea behind the introduction of the $(b_i)_i$ is that, by Bayes' rule and Assumption 2.2,

$$p_{0ij} = \frac{P(D = 0|Y_1 = y, Y_2 = i)}{P(D = 1|Y_1 = y, Y_2 = i)} p_{1ij},$$

and b_i represents the odds $P(D = 0|Y_1 = y, Y_2 = i)/P(D = 1|Y_1 = y, Y_2 = i)$. The inequality constraints $b_i \geq 0$ then ensure that $P(D = 1|Y = i)$ is indeed a probability, while the equality constraints are a rewriting of Equation (2.1) in the discrete context.

The identifying condition $\text{rank}(P_1) = I$ is the equivalent of $\text{Ker}\,(T) \cap \mathcal{F} = \{0\}$ here. It can be easily tested in the data because under the null hypothesis that $\text{rank}(P_1) < I$,

7. This parametrization is also convenient for the unconstrained model where Assumption 2.2 does not necessarily hold.

8. We thus consider here implicitly parameters that depend on the distribution of (D, Y_2, Z_1) conditional on Y_1. To estimate unconditional parameters, it suffices to integrate conditional parameters over the empirical distribution of Y_1. Because Assumption 2.2 does not impose any restriction on the distribution of Y_1, this results in asymptotically efficient estimators.

we have $\mu_0 \equiv \det(P_1 P_1') = 0$. Then, letting $\widehat{\mu} = \det(\widehat{P}_1 \widehat{P}_1')$, $\sqrt{n}\widehat{\mu}$ tends to a zero mean normal variable under the null by the delta method. We use this result to test for the rank condition in our application Section 4.

Finally, as noted before, we can test for Assumption 2.2 by two ways. The first and standard one is that the equality constraints in (C) may not hold when $J > I$, because there is no $(b_i)_{1\leq i\leq I}$ such that $\sum_{i=1}^I b_i p_{1ij} = p_{0.j}$. Basically, this arises when the different values of Z are not "compatible", as with the Sargan test in linear IV models. The second is that the $(b_i)_{1\leq i\leq I}$ which satisfy these equality constraints must be nonnegative. This may not hold in general, even when $I = J$. To test for both conditions simultaneously, we use the same Wald statistic as the one of [14]. In our framework, the unconstrained model where Assumption 2.2 does not necessarily hold is simply the multinomial model on (D, DY_2, Z) parameterized by p, and the maximum likelihood estimator $\widehat{p}_U$ simply corresponds to the sample proportions. The contraints (C) corresponding to Assumption 2.2 hold if and only if there exists $b \geq 0$ (understood componentwise) such that $P_1' b = p_{0.}$. This condition is equivalent to [9]

$$\left[P_1'(P_1P_1')^{-1}P_1 - I\right] p_{0.} = 0, \quad (P_1P_1')^{-1}P_1 p_{0.} \geq 0.$$

Let us rewrite these constraints as $h_1(p) = 0$ and $h_2(p) \geq 0$, and let $h(p) = (h_1(p), h_2(p))$. Let also $\mathcal{H}_0 = 0^J \times \mathbb{R}^{+I}$ denote the set of $h = (h_1, h_2)$ satisfying these constraints. Denote by Σ_{ii} (resp. Σ_{12}) the asymptotic variance of $\widehat{h}_i \equiv h_i(\widehat{p}_U)$ (resp. covariance of $h_1(\widehat{p}_U)$ and $h_2(\widehat{p}_U)$), and by Σ the asymptotic variance of $\widehat{h} \equiv h(\widehat{p}_U)$. Finally, let $\widehat{\Sigma}$ denote a consistent estimator of Σ. The test statistic W_n is then defined as

$$W_n = n \min_{h \in \mathcal{H}_0} \left(h - \widehat{h}\right)' \widehat{\Sigma}^- \left(h - \widehat{h}\right),$$

where $\widehat{\Sigma}^-$ denotes the Moore-Penrose inverse of $\widehat{\Sigma}$. [10] Computing W_n is straightforward as it corresponds to a quadratic programming problem.

To derive the asymptotic distribution of W_n, we cannot apply directly the results of [14] [11] and must introduce additional notations. Let $(h_{10}, h_{20}) = h_0 = h(p_1, p_{0.})$ denote the true parameter. The asymptotic distribution of W_n depends on whether the components $(h_{20i})_{1\leq i\leq I}$ are equal to zero or not. Let $\mathcal{R}_j$ be equal to $\mathbb{R}^+$ if $h_{20j} = 0$, and to $\mathbb{R}$ otherwise. Then let

$$\mathcal{H}(h_0) = 0^J \times \mathcal{R}_1 \times ... \times \mathcal{R}_I.$$

We show in the proof of Proposition 3.2 below that

$$\lim_{n\to\infty} \Pr(W_n \geq w) = P\left(\min_{h\in\mathcal{H}(h_0)} (h - U)' \Sigma^- (h - U) \geq w\right) \tag{3.1}$$

where $U \sim \mathcal{N}(0, \Sigma)$. To compute the level of the test based on this asymptotic distribution, we thus need to estimate $\mathcal{H}(h_0)$. Following [20], we consider a sequence $(c_n)_{n\in\mathbb{N}}$ such that

9. Indeed, the existence of $b \in \mathbb{R}^{+I}$ such that $P_1'b = p_{0.}$ is equivalent to the fact that the least square solution $(P_1P_1')^{-1}P_1p_{0.}$ to $\min_{b\in\mathbb{R}^I} \left\|P_1 b - p_{0.}\right\|$ satisfies the equation exactly and belongs to $\mathbb{R}^{+I}$.

10. $\widehat{\Sigma}$ is not full rank in general, because the rank of Σ_{11} is $J - I$, while $h_1(p) \in \mathbb{R}^J$. This is logical, since we only have $J - I$ overidentifying equality constraints here.

11. Technically, [14] compute $\sup_{h\in\mathcal{H}_{0\Sigma}} \lim_{n\to\infty}$, where $\mathcal{H}_{0\Sigma}$ is the subset of $\mathcal{H}_0$ such that the asymptotic variance of $\widehat{h}$ is equal to Σ. For their results to apply, $\mathcal{H}_{0\Sigma}$ should be a convex cone. Unfortunately, Σ depends on the true parameter $h_0 \in \mathcal{H}_0$ here, making this latter condition fail even though $\mathcal{H}_0$ is a convex cone.

$c_n \to \infty$ and $c_n/\sqrt{n} \to 0$. We let $\widehat{\mathcal{R}}_j$ be equal to $\mathbb{R}^+$ if $\widehat{h}_{2j} \leq c_n/\sqrt{n}$, and to $\mathbb{R}$ otherwise, and

$$\widehat{\mathcal{H}(h_0)} = 0^J \times \widehat{\mathcal{R}}_1 \times ... \times \widehat{\mathcal{R}}_I.$$

Finally, let $\widehat{c}_\alpha$ satisfy

$$P\left(\min_{h\in\widehat{\mathcal{H}(h_0)}} \left(h - \widehat{U}\right)' \widehat{\Sigma}^- \left(h - \widehat{U}\right) \geq \widehat{c}_\alpha\right) = \alpha,$$

where $\widehat{U} \sim \mathcal{N}(0, \widehat{\Sigma})$. $\widehat{c}_\alpha$ (or, similarly, the p-value of the test) can be obtained by simulations.

Proposition 3.2. *The test defined by the critical region $\{W_n > \widehat{c}_\alpha\}$ has asymptotic level α and is consistent.*

Following the analysis of [14], it is also possible to express this asymptotic distribution as a mixture of chi-square. The corresponding weights, however, do not have a closed form in general, so that it is actually easier to approximate the asymptotic distribution using (3.1) rather than their expression. [12] We use such simulations to compute our p-values in the application below.

3.2 The continuous case

The situation is more involved when (Y, Z) is continuous, because we are not in a parametric setting anymore. We first obtain a necessary and sufficient condition for root-n estimability and provide the efficiency in this case. These results are closely related to those of [23]. We then develop a root-n consistent estimator of β_0.

3.2.1 Semi-parametric efficiency bounds

The first issue is to determine if β_0 can be estimated at a root-n rate, and if so, to compute the asymptotic efficiency bound. In the estimation of linear functionals depending on a nonparametric regression with an endogenous regressor (see [23]), root-n estimability is related to conditions on two dual operators. Let T^* denote the operator defined by :

$$\begin{array}{llll} T^*: & L^2(Z|D=1) & \to & L^2(Y|D=1) \\ & q & \mapsto & (y \mapsto E(q(Z)|D=1, Y=y)) \,. \end{array}$$

This notation stems from the fact that T^* is the adjoint of T (defined on $L^2(Y|D=1)$) if one uses the scalar products associated with $L^2(Z|D=1)$ and $L^2(Y|D=1)$. Actually, we only need considering the restriction $T^*_{\mathcal{Y}_0}(q)$ of $T^*(q)$ on $\mathcal{Y}_0 = \text{Supp}(Y|D=0)$. By Assumptions 2.2 and 2.3, $E(q(Z)|D=1, Y) = E(q(Z)|D=0, Y)$ $P^{Y|D=0}$ almost surely. This allows us to extend $T^*_{\mathcal{Y}_0}(q)$ on $L^2(Z|D=0)$. By a slight abuse of notation, this extension is also denoted $T^*_{\mathcal{Y}_0}$. Finally, let $\beta_{\mathcal{Y}_0}$ denote the restriction of $\beta(.)$ on $\mathcal{Y}_0$. The condition for root-n estimability is the following.

Assumption 3.2. *(i) $g \in L^2(Y, Z)$*
(ii) There exists $q \in L^2(Z|D=0)$ such that $T^*_{\mathcal{Y}_0}(q) = \beta_{\mathcal{Y}_0}(.)$ and

$$E\left[\frac{1-P(Y)}{P(Y)}\left(q(Z) - g(Y,Z)\right)\left(q(Z) - g(Y,Z)\right)'\right] < \infty.$$

12. It is also possible to derive lower and upper bounds on the critical values of this test, see, e.g., [14].

The first condition is standard to derive the asymptotic of $\frac{1}{n}\sum_{i=1}^n g(Y_i, Z_i)$ even in case of perfect observability of Y. The second condition is similar to the one considered by [23]. Indeed, if the standard completeness condition holds, then $\text{Ker}\,(T) = \{0\}$, so that $\text{Ker}\,(T) = \{0\}$ and $\overline{\mathcal{R}\,(T^*)} = L^2(Y|D = 1)$, where $\mathcal{R}\,(T^*)$ denotes the range of T^*. As a consequence if $g \in L^2(Y, Z)$, then $\beta(.)$ lies in $\overline{\mathcal{R}\,(T^*)}$ and $\beta_{\mathcal{Y}_0}(.)$ lies in $\overline{\mathcal{R}\,(T^*)}_{\mathcal{Y}_0}$. However, when Y is continuous, $\mathcal{R}\,(T^*)$ is not closed in general, so that even if the standard completeness holds, it may happen that $\beta_{\mathcal{Y}_0}(.) \notin \mathcal{R}\,(T^*)_{\mathcal{Y}_0}$. When $\beta_{\mathcal{Y}_0}(.) \in \mathcal{R}\,(T^*)_{\mathcal{Y}_0}$ on the other hand, Theorem 3.3 shows that β_0 is root-n estimable, and provide the corresponding asymptotic efficiency bound.

Théorème 3.3. *Suppose that Assumptions 2.1-2.3(i) hold, and $\beta(.) \in (Ker\,(T) \cap \mathcal{F})^{\perp}$. Then a regular root-n estimator of β_0 exists only if Assumption 3.2 holds, and in this case the semiparametric efficiency bound V^* satisfies*

$$V^* = V\,(g(Y, Z)) + \min_{q(.)\in L^2(Z|D=0)\cap T_{\mathcal{Y}_0}^{*-1}(\{\beta_{\mathcal{Y}_0}(.)\})} E\left[\frac{1-P(Y)}{P(Y)}\,(q(Z)-g(Y,Z))\,(q(Z)-g(Y,Z))'\right].$$

The second part of the theorem shows that asymptotic efficiency bound comprises two terms. The first corresponds to the standard estimation of β_0 without any attrition, i.e. when $D = 1$. The second accounts for attrition, and is indeed, loosely speaking, increasing with $P(Y)$. It is also related to the approximation of $g(Y, Z)$ by a function of Z (among elements of $T_{\mathcal{Y}_0}^{*-1}(\{\beta_{\mathcal{Y}_0}(.)\})$) . Intuitively, if $g(Y, Z) \equiv g(Z)$, then $q(.) = g(.) \in T_{\mathcal{Y}_0}^{*-1}(\{\beta_{\mathcal{Y}_0}(.)\})$, so that the second term is zero. In this case, the sample average $\sum_{i=1}^n g(Z_i)/n$ is asymptotically efficient.

4 Application

4.1 Introduction

In this section, we apply the previous results to estimate transitions on employment status in the French labor market. Beyond the unemployment rate, measuring such transitions is important to assess, for instance, the importance of short and long-term unemployment. We use for that purpose the Labor Force Survey (LFS) conducted by the French national institute of statistics (INSEE). This survey is probably the best tool to measure such transitions in France. Indeed, and compared to administrative data or other surveys, it properly measures unemployment with respect to the standard ILO definition, has a comprehensive coverage of the population and its sample size is large. Since 2003, it is a rotating panel where approximately 5,900 new households are surveyed each quarter. These new households are then questioned the five following quarters. On the first and sixth wave, interrogations are face to face, while on the others they are conducted by telephone. It has been argued that the use of phone may introduce specific measurement errors (see, e.g., [3]), so we focus on the first and last interrogations hereafter. We also restrict ourselves to people between 15 and 65 and pool together all labor force surveys on the period 2003-2005.

Table 1 – Summary statistics on the French LFS (first waves during 2003-2005).

Statistics	All	Men	Women
Main sample :			
Number of individuals	107,031	52,245	54,786
Attrition rate on last waves	21.78%	22.26%	21.31%
Participation rate on first waves	68.17%	73.91%	62.69%
(Uncorrected) participation rate on last waves	67.38%	72.75%	62.32%
Unemployment rate on first waves	9.68%	9.05%	10.39%
(Uncorrected) unemployment rate on last waves	8.02%	7.22%	8.90%
Refreshment sample for last waves :			
Number of observations	109,404	53,337	56,067
Participation rate on the refreshment sample	67.92%	73.31%	62.78%
Unemployment rate on the refreshment sample	9.97%	9.43%	10.57%

Table 1 provides some summary statistics on our dataset, which emphasize that attrition may be problematic in the LFS survey. This is especially striking when we compare the (uncorrected) participation and unemployment rate on last waves and the one on the refreshment sample (i.e., entrants at the same time). We observe differences around 1.5 percent points on participation rates, and around 2 percent points on unemployment rates.To understand these differences, recall that in the French LFS, moving households are not followed by interviewers, who stick instead on housings which were selected in the first waves. This is likely to affect activity rates and transition estimates on the labor market, because transitions are very different for moving and non-moving households.

As suggested in Section 2, we propose to correct for potentially endogenous attrition by using past employment status, measured by a retrospective question asked on the first waves. The underlying assumption is that attrition depends on the current transition on this outcome, but not on previous ones. This assumption is plausible if most of the endogeneity in attrition stems from the moving of households. The instrument Z we use is employment status 6 months before the first wave. We choose to divide this variable in three categories (unemployed, employed, and out of labour force) as our outcome which is contemporary employment status. To assess the strength of our instrument we implement rank test between Y_2 and Z conditionally to Y. We also test our instrument following the approach developed in Proposition 3.2.We also assess the plausibility of our instrument by comparing the corrected participation and unemployment rates we obtain on last waves with the one on the refreshment sample. Finally, as a matter of comparison, we also correct for attrition under ignorability or using the method of [12] with a logistic cumulative distribution for the link function.

4.2 The results

We first check the rank condition between Z_1 and Y_2 conditional on gender and Y_1, relying on the determinant test proposed in Subsection 3.1. Results are displayed in Table 2. The p-value of the rank test associated to any state Y_1 are always smaller than 10% for both men and women. We also implement the test developed in the Proposition 3.2. Though some inequality constraints are binding in our estimates, we do not reject the independence assumption $Z \perp\!\!\!\perp D|Y_1, Y_2$ here (see Table 3).

TABLE 2 – Rank test between Z and Y_2 conditional on gender and Y_1.

	P-value (Men)	P-value (Women)
Y_1 = Empl.	0.0047	0.0031
Y_1 = Unempl.	0.0695	0.0539
Y_1 = Out L.F.	0.0498	0.0867

TABLE 3 – Test of $Z \perp\!\!\!\perp D|Y_1, Y_2$ by gender.

	P-value (Men)	P-value (Women)
Y_1 = Empl.	0.7430	0.8818
Y_1 = Unempl.	0.5724	0.5418
Y_1 = Out L.F.	0.7186	0.7088

Second, we estimate the probabilities of attrition (or non-attrition) conditional on (Y_1, Y_2). Our results, displayed in Table 4, confirm that (under the validity of our instrument), attrition is related to transitions on employment status. People who remain stable on the labor market have always a significant larger probability to respond in the second wave than people who change. In particular, we observe a large attrition for those who move from employment to unemployment or inactivity whereas attrition seems negligible for those who remain unemployed at both periods. As suggested above, such transitions are likely to be related to house movings. For instance, transitions from inactivity to employment or unemployment mostly correspond to students who enter the labor market and move at the same time. Such features cannot be captured under the missing at random (MAR) scheme $D \perp\!\!\!\perp Y_2|Y_1$, or the additive model of [12]. In particular, they tend to underestimate the probability of attrition for people whose status change on the labor market, and to overestimate them for stable trajectories (see Table **??** in appendix for tests on the difference between our IV models and the two others).

TABLE 4 – $\widehat{P}(D=1|Y_1,Y_2)$ for women under various assumption

	Men			Women		
	Y_2 = Empl.	Y_2 = Unempl.	Y_2 = Out L.F.	Y_2 = Empl.	Y_2 = Unempl.	Y_2 = Out L.F.
IV						
Y_1 = Empl.	84.33 (0.85)	34.33 (3.92)	46.31 (7.04)	85.72 (0.92)	46.45 (6.42)	44.57 (4.46)
Y_1 = Unempl.	55.56 (4.33)	100.00 (3.91)	51.01 (9.27)	52.46 (4.25)	100.00 (2.06)	76.38 (13.43)
Y_1 = Out L.F.	54.83 (11.35)	55.85 (11.35)	85.72 (1.99)	56.43 (11.16)	67.10 (16.77)	84.34 (1.10)
MAR						
Y_1 = Empl.	78.22 (0.22)	78.22 (0.22)	78.22 (0.22)	79.00 (0.24)	79.00 (0.24)	79.00 (0.24)
Y_1 = Unempl.	65.90 (0.79)	65.90 (0.79)	65.90 (0.79)	69.77 (0.77)	69.77 (0.77)	69.77 (0.77)
Y_1 = Out L.F.	79.52 (0.34)	79.52 (0.34)	79.52 (0.34)	79.77 (0.28)	79.77 (0.28)	79.77 (0.28)
HIRR						
Y_1 = Empl.	79.01 (0.28)	59.98 (2.07)	76.44 (2.25)	79.57 (0.32)	66.59 (2.21)	77.57 (2.02)
Y_1 = Unempl.	75.84 (1.43)	55.55 (1.24)	73.01 (2.04)	76.04 (1.46)	61.89 (1.40)	73.81 (1.77)
Y_1 = Out L.F.	82.41 (1.69)	65.09 (2.33)	80.15 (0.47)	82.01 (1.67)	69.99 (2.07)	80.18 (0.39)

Note : Standard error in brackets computed with 1000 bootstrap samples.

Before presenting our results on transitions, we estimate the distribution of Y_2 with our IV method and compare it with the one of the refreshment sample. We also estimate this distribution supposing that data are missing at random (MAR), i.e. $D \perp\!\!\!\perp Y_2|Y_1$. Table 5 shows that on the five statistics related to the distribution of Y_2, our estimator is close, and not statistically significant at usual levels, to the one based on the refreshment sample. Those based on the MAR assumptions, on the other hand, do differ significantly for several features of Y_2. In other words, we can reject, using the refreshment sample, the hypothesis that attrition only depends on past outcomes, while our independence condition is not rejected in the data. Note that we cannot use the refreshment sample to properly compare our method with the one of [12] because by construction, their estimator exactly matches the distribution of Y_2 on the refreshment sample.

Table 5 – Comparison of the methods with the refreshment sample

	Men			Women		
	REF.	MAR	IV	REF.	MAR	IV
$P(Y_2 = \text{Empl.})$	66.40	67.47 (<0.0001)	64.59 (0.0545)	56.15	56.81 (0.0055)	55.07 (0.2718)
$P(Y_2 = \text{Unempl.})$	6.92	5.62 (<0.0001)	7.53 (0.2851)	6.63	5.78 (<0.0001)	6.51 (0.8529)
$P(Y_2 = \text{Out L.F.})$	26.69	26.92 (0.2825)	27.88 (0.2033)	37.22	37.40 (0.4300)	38.42 (0.1373)
Participation rate	73.31	73.08 (0.2825)	72.12 (0.2033)	62.78	62.60 (0.4300)	61.58 (0.1373)
Unemployment rate	9.43	7.68 (<0.0001)	10.44 (0.1884)	10.57	9.24 (<0.0001)	10.58 (0.9913)

Note : we indicate the p-values of the difference with the refreshment sample under parentheses. Computation based on 1000 bootstrap samples.

Finally, we compute transitions on the labor market using our IV method, the MAR assumption and the additive method of [12] (see Table 6). Not surprisingly given the discrepancies on the probabilities of attrition, our results differ significantly from those obtained by the other methods (see Table **??** for the tests of differences). Other methods tend in particular to overestimate stability on the labor market. If it is impossible to discriminate between our IV method and the one of [12] without extra information on these transitions, some patterns on unemployment seem to support our model over theirs in this application. In particular, the estimated probability of staying unemployed after 15 months are respectively equal to 25% with our IV method and around 44% with the one of [12] (41% for women and 47% for the men). These latter figures seem particularly high, compared to the rate we observe between the first wave and 11 months before, namely 44,9% for women and 49.6% for men. It is also notably at odds with the ? ?% rate of long term (i.e., one year or more) unemployment directly observed on the LFS.

Table 6 – Estimated probability of transitions for women under various assumption

	Men			Women		
	Y_2 = Empl.	Y_2 = Unempl.	Y_2 = Out L.F.	Y_2 = Empl.	Y_2 = Unempl.	Y_2 = Out L.F.
IV						
Y_1 = Empl.	85.86 (0.85)	6.12 (0.70)	8.02 (1.13)	83.46 (0.89)	4.73 (0.63)	11.81 (1.15)
Y_1 = Unempl.	49.00 (3.71)	25.85 (1.46)	25.15 (3.74)	52.61 (3.92)	25.30 (0.95)	22.08 (3.91)
Y_1 = Out L.F.	13.81 (2.63)	6.47 (1.16)	79.72 (1.84)	12.74 (2.12)	5.92 (1.41)	81.34 (1.05)
MAR						
Y_1 = Empl.	92.56 (0.16)	2.69 (0.10)	4.75 (0.13)	90.56 (0.19)	2.78 (0.11)	6.66 (0.16)
Y_1 = Unempl.	41.31 (1.06)	39.23 (1.01)	19.46 (0.82)	39.56 (0.98)	36.27 (0.95)	24.18 (0.84)
Y_1 = Out L.F.	9.52 (0.27)	4.55 (0.20)	85.93 (0.33)	9.02 (0.23)	4.98 (0.17)	86.00 (0.28)
HIRR						
Y_1 = Empl.	91.64 (0.20)	3.50 (0.12)	4.86 (0.15)	89.92 (0.23)	3.30 (0.12)	6.79 (0.18)
Y_1 = Unempl.	35.90 (0.96)	46.54 (0.91)	17.57 (0.75)	36.29 (0.97)	40.87 (0.90)	22.85 (0.83)
Y_1 = Out L.F.	9.19 (0.27)	5.56 (0.24)	85.26 (0.36)	8.77 (0.23)	5.68 (0.17)	85.55 (0.28)

Note : Standard error in brackets computed with 1000 bootstrap samples.

5 Conclusion

In this paper, we develop an alternative method to correct for endogenous attrition in panel. We allow for both dependence on current and past outcomes and, thanks to the availability of an instrument, do not need to impose functional restrictions on the probability of attrition, contrary to [12]. The application suggests that our method may do a good job for handling attrition processes which mostly depend on transitions.

The paper raises two challenging issues, related to our main conditional independence assumption. The first is whether the refreshment sample could be used to weaken this assumption, rather than to test for it. This may be useful in settings where this condition is considered too stringent. The second is whether one can build bounds on parameters of interest if the conditional independence assumption is replaced by weaker conditions such as monotonicity ones. Although not considered here, these questions are clearly at the top of our research agenda.

Références

[1] Abowd, J. M., Crépon, B., and Kramarz, F. Moment estimation with attrition : An application to economic models. *Journal of the American Statistical Association 96* (1999), 1223–1231.

[2] Bhattacharya, D. Inference in panel data models under attrition caused by unobservables. *Journal of Econometrics 144* (2008), 430–446.

[3] Biemer, P. P. Nonresponse bias and measurement bias in a comparison of face to face and telephone interviewing. *Journal of Official Statistics 17* (2001), 295–320.

[4] Blundell, R., Chen, X., and Kristensen, D. Nonparametric iv estimation of shape-invariant engel curves. *Econometrica 75* (2007), 1613–1669.

[5] Chen, C. Parametric models for response-biased sampling. *Journal of the Royal Statistical Society, Series B 63* (2001), 775–789.

[6] D'Haultfœuille, X. A new instrumental method for dealing with endogenous selection. *Journal of Econometrics 154* (2010), 1–15.

[7] D'Haultfœuille, X. On the completeness condition in nonparametric instrumental regression. *Econometric Theory, forthcoming* (2011).

[8] Evdokimov, K. Nonparametric identification of a nonlinear panel model with application to duration analysis with multiple spells. Working paper, 2011.

[9] Hall, R., and Mishkin, F. S. The sensitivity of consumption to transitory income : estimates from panel data on households. *Econometrica 50* (1982), 461–481.

[10] Hausman, J. A., and Wise, D. A. Attrition bias in experimental and panel data : the Gary income maintenance experiment. *Econometrica 47* (1979), 455–473.

[11] Heckman, J. Micro data, heterogeneity, and the evaluation of public policy : Nobel lecture. *Journal of Political Economy 109* (2001), 673–748.

[12] Hirano, K., Imbens, G. W., Ridder, G., and Rubin, D. B. Combining panel data sets with attrition and refreshment samples. *Econometrica 69* (2001), 1645–1659.

[13] Imbens, G. Nonparametric estimation of average treatment effects under exogeneity : a review. *The Review of Economics and Statistics 86* (2004), 4–29.

[14] Kodde, D. A., and Palm, F. C. Wald criteria for jointly testing equality and inequality restrictions. *Econometrica 54* (1986), 1243–1248.

[15] Lancaster, T. *The Econometric Analysis of Transition Data*. Cambridge University Press, 1990.

[16] Little, R., and Rubin, D. B. *Statistical analysis with Missing Data*. John Wiley & Sons, New York, 1987.

[17] Mattner, L. Completeness of location families, translated moments, and uniqueness of charges. *Probability Theory and Related Fields 92* (1992), 137–149.

[18] Newey, W., and Powell, J. Instrumental variable estimation of nonparametric models. *Econometrica 71* (2003), 1565–1578.

[19] Ramalho, E. A., and Smith, R. J. Discrete choice nonresponse. *Review of Economic Studies, forthcoming* (2011).

[20] Rosen, A. Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities. *Journal of Econometrics 146* (2008), 107–117.

[21] Rubin, D. B. Inference and missing data. *Biometrika 63* (1976), 581–592.

[22] Severini, T., and Tripathi, G. Some identification issues in nonparametric linear models with endogenous regressors. *Econometric Theory 22* (2006), 258–278.

[23] Severini, T., and Tripathi, G. Efficiency bounds for estimating linear functionals of nonparametric regression models with endogenous regressors. *Journal of Econometrics, forthcoming* (2011).

[24] Tang, G., Little, R. J. A., and Raghunathan, T. E. Analysis of multivariate missing data with nonignorable nonresponse. *Biometrika 90* (2003), 747–764.