Essential Heterogeneity in the Impact of Community Driven Development*

Jean-Louis Arcand[†] Léandre Bassole[‡]

October 16, 2011

Work in progress. Comments welcome.

Abstract

We consider *essential heterogeneity* in the impact of a large-scale Community-Driven Development (CDD) program in Senegal. *Essential heterogeneity* arises when unobservables determine the idiosyncratic gains from participation in a program, thereby generating correlation between treatment effects and selection. Standard instrumental variables estimates are shown to provide an extremely poor estimate of the impact of the program on child nutrition. Application of the local instrumental variables (LIV) estimator using semi-parametric techniques reveals extensive heterogeneity in the impact of the program, with children living in villages with unobservables that make them more likely to receive a completed project benefitting much more than children in villages whose unobservables induce a low likelihood of receiving a project. These techniques provide a sensible and easily applicable manner of assessing whether the decentralized allocation of projects that is inherent to CDD programs is resulting in those who benefit most from treatment actually receiving it.

Keywords: Impact evaluation, Community Driven Development, Essential heterogeneity, Marginal treatment effect, Local instrumental variables.

JEL Classification numbers: O19, H43, I12, I38.

1 Introduction

Traditional "top-down" approaches, for years the standard implementation mechanisms for development and poverty reduction by governments, as well as bilateral and multilateral donors, have increasingly been called into question. Instead of viewing the poor as being *targets* of poverty reduction, the focus has shifted towards treating the poor and the civil society organizations within which they operate as both assets and partners in the development process. World Bank funded Community Driven Development (CDD) programs are based on this premise. Decisionmaking, and often the control of financial resources, is increasingly vested in local communities.

*We thank seminar participants at CERDI, Shanghai University, Fudan University, the University of Sussex, the Universidad Publica de Navarra, The Graduate Institute Geneva, and the University of Hannover for comments. The usual disclaimer applies.

[†]The Graduate Institute, Geneva. Email: Jean-Louis.Arcand@graduateinstitute.ch

[‡]African Development Bank.

As with many poverty reduction programs, there are many options in terms of the specific institutional arrangements through which CDD programs can be implemented. One of these is through the support of social funds.[1] One distinguishing feature of social fund projects is that, rather than implementing investment decisions that have been predetermined at the project appraisal level, they allow local stakeholders to make these decisions through "subproject proposals" that they submit themselves during the course of overall project implementation. In theory, the flexibility of financing instruments makes is possible for both governments and the World Bank to provide resources for community initiatives that will be closely tailored to local needs, in particular when it comes to the provision of basic infrastructure, such as primary health care and education, or potable water.

A major weakness of traditional impact evaluations of development programs is that they usually limit their attention to average causal effects. While this provides both the researcher and, ultimately, decisionmakers, with an estimate of the average impact of the program being evaluated on the response variable(s) of interest, it is not difficult to imagine this response being heterogeneous. Another way of putting this is that it is difficult to believe that all individuals, households or villages benefit by the same amount from a given social program.

Heterogeneity in treatment effects is intimately related to the selection issue that lies at the root of the evaluation problem. Self-selection arises when individuals, households or villages select into treatment based on characteristics that also influence outcomes. As it is put by Heckman and Vytlacil (2001), heterogeneity is an integral part of human existence, while self-selection is an integral part of human actions.

Accounting for heterogeneity in treatment effects has been the focus of an important corpus of work in the causal effects literature.[2] In the literature dealing with the micro-econometrics of program evaluation, the interaction of heterogeneity and self-selection has recently been dubbed *essential heterogeneity* by Heckman, Urzua, and Vytlacil (2006).[3] The existence of essential heterogeneity implies that the effect of the treatment may vary across individuals, households or villages even after conditioning on observables characteristics. This form of heterogeneity is different from that addressed by quantile regression techniques, which stems from the outcome variable.

The basic source of essential heterogeneity lies in unobservables that determine the idiosyncratic gains from participation in a program. In this context, the estimation of treatment effects is far from being straightforward because of the correlation between idiosyncratic unobservables and treatment status *per se*. Another way of putting this is that while traditional instrumental variables (IV) methods are geared towards dealing with correlation between unobservables and treatment status, they break down completely when the gains to treatment (i.e. the "parameter" one is seeking to estimate)

[1] Social funds are defined as "Agencies that finance small projects in several sectors targeted to benefit a country's poor and vulnerable groups based on a participatory manner of demand generated by local groups and screened against a set of eligibility criteria."

[2] In economics, heterogeneity in the effects of treatment was introduced in the literature on switching regression models by Quandt (1972), Heckman (1978), Heckman (1979), while in the statistics literature the causal model of potential outcomes of Rubin (1974) also allows for heterogeneity in treatment effects.

[3] Heckman, Smith, and Clements (1997) were the first to analyze essential heterogeneity using data from social experiments.

are correlated with unobservables. In this case, there is no single "parameter" to be estimated. Rather, a whole *distribution* of parameters must be estimated, with each one corresponding to a particular value of the unobservables that determine treatment status.[4]

2 Heterogeneity in the returns to treatment by CDD

2.1 The program

In early 2001, the Senegalese government, within the context of its overall development strategy, and in collaboration with World Bank and IFAD, initiated a major CDD program, the *Programme national d'infrastructures rurales* ("National Rural Infrastructures Program", henceforth, PNIR) geared towards improving infrastructure on a demand-driven basis, thereby enabling economic development amongst the rural poor. The PNIR operates at the level of the smallest sub-regional administrative unit in Senegal —the *Communauté rurale* ("rural community", henceforth, CR).

An average CR includes 42 villages (the number varies between 3 and 132 villages over the 320 CRs in Senegal), and has a population of 13,391 inhabitants (std. = 12,799). For the purposes of the PNIR, 90 CRs were chosen from amongst the poorest in the nine rural regions of Senegal for treatment; 78% of the poor in Senegal live in rural areas, where the average incidence of poverty is about 40%, as compared with 16% in urban areas. The rural population that was to benefit from the project was estimated at nearly two million people, more than half of whom were then poor.

Our purpose in this paper is to assess whether there is essential heterogeneity in the impact of the PNIR on the welfare of rural populations, or whether traditional impact evaluation methods do indeed paint a faithful picture in terms of the average causal effect of treatment. We concentrate on a single outcome variable Y_{ivt} given by the weight-to-age (WAZ) $z-$scores of children under 5 years of age, a standard measure of short-term nutritional status. In what follows, i will index children, v villages and t time.

CDD social funds appear to be a particularly topical area in which to apply these methods, in that selection into treatment (i.e. receiving or not receiving completed projects) is likely to be, at least in part, driven by the unobserved characteristics of the target populations and villages which, simultaneously, affect outcomes. To the best of our knowledge, this paper is the first attempt to apply these methods to CDD impact evaluation context.

2.2 The underlying Roy model

We begin by defining a pair of results equations, for treated and untreated children:

$$Y_{1ivt} = \alpha_1 + X_{ivt}\beta_1 + U_{1ivt}, \quad (1)$$
$$Y_{0ivt} = \alpha_0 + X_{ivt}\beta_0 + U_{0ivt}, \quad (2)$$

[4]These methods are fully described in Heckman and Vytlacil (1999), Florens, Heckman, Meghir, and Vytlacil (2004), and Heckman, Urzua, and Vytlacil (2006).

where X is a matrix of observed covariates that affect outcomes, and U represents a disturbance term. The latent index model that determines treatment is given by:

$$D^*_{vt} = Z_{vt}\gamma - V_{vt}, \tag{3}$$

where D^* is a latent variable, Z is a matrix of observed covariates that determines treatment status, and V is the corresponding disturbance term. The basic problem is that (U_{0ivt}, U_{1ivt}) and V_{vt} are likely to be correlated. Treatment obtains when:

$$D_{vt} = \begin{cases} 1 & \text{if } D^*_{vt} > 0, \\ 0 & \text{if } D^*_{vt} \leqslant 0, \end{cases}$$

which one can rewrite in more compact notation in terms of an indicator function as:

$$D_{vt} = \mathbf{1}\left[D^*_{vt} > 0\right]. \tag{4}$$

We assume that the error structure in the results equations is given by a factor representation:

$$\begin{aligned} U_{1ivt} &= \lambda_i + \sigma_1 V_{vt} + \varepsilon_{1ivt}, & (5) \\ U_{0ivt} &= \lambda_i + \sigma_0 V_{vt} + \varepsilon_{0ivt}, & (6) \end{aligned}$$

where the common factor that generates the correlation between treatment status and the disturbance of the structural equation is the time-varying village level unobservable V_{vt} that affects both the outcome and treatment status. When $\sigma_1 < 0$, for example, a village-level shock that increases the likelihood of the village receiving a completed project also improves the outcome of children when they are treated.[5]

Our identifying restrictions are that there exists a matrix of instrumental variables Z_{vt} such that:

$$(U_{1ivt}, U_{0ivt}, V_{vt}) \perp Z_{vt} | \, X_{ivt}.$$

Write the results equation in terms of observables

$$\begin{aligned} Y_{ivt} &= D_{vt}Y_{1ivt} + (1 - D_{vt})\,Y_{0ivt} \\ &= Y_{0ivt} + D_{vt}\left(Y_{1ivt} - Y_{0ivt}\right) \\ &= \underbrace{\alpha_0 + X_{ivt}\beta_0 + U_{0ivt}}_{=Y_{0ivt}} + D_{vt}[\underbrace{\alpha_1 + X_{ivt}\beta_1 + U_{1ivt}}_{=Y_{1ivt}} - (\underbrace{\alpha_0 + X_{ivt}\beta_0 + U_{0ivt}}_{=Y_{0ivt}})] \\ &= \alpha_0 + X_{ivt}\beta_0 + D_{vt}\left(\alpha_1 - \alpha_0\right) + D_{vt}X_{ivt}\left(\beta_1 - \beta_0\right) + U_{0ivt} + D_{vt}\left(U_{1ivt} - U_{0ivt}\right). \end{aligned}$$

Substituting for U_{1ivt} and U_{0ivt} from (5) and (6) yields:

$$\begin{aligned} Y_{ivt} = \; & \alpha_0 + X_{ivt}\beta_0 + D_{vt}\left(\alpha_1 - \alpha_0\right) + D_{vt}X_{ivt}\left(\beta_1 - \beta_0\right) \\ & + \underbrace{\lambda_i + \sigma_0 V_{vt} + \varepsilon_{0ivt}}_{=U_{0ivt}} + D_{vt}[\underbrace{\lambda_i + \sigma_1 V_{vt} + \varepsilon_{1ivt}}_{=U_{1ivt}} - (\underbrace{\lambda_i + \sigma_0 V_{vt} + \varepsilon_{0ivt}}_{=U_{0ivt}})], \end{aligned}$$

[5]If $V_{vt} \downarrow$, $D^*_{vt} \uparrow$ so the likelihood of treatment increases, but under the assumption that $\sigma_1 < 0$, this also implies that $U_{1ivt} \uparrow (= \lambda_i + \sigma_1 V_{vt} + \varepsilon_{1ivt})$, so the outcome improves under treatment.

and thus:

$$\begin{aligned} Y_{ivt} &= \alpha_0 + X_{ivt}\beta_0 + D_{vt}\left(\alpha_1 - \alpha_0\right) + D_{vt}X_{ivt}\left(\beta_1 - \beta_0\right) \\ &\quad + \lambda_i + \sigma_0 V_{vt} + \varepsilon_{0ivt} + D_{vt}\left[\left(\sigma_1 - \sigma_0\right)V_{vt} + \varepsilon_{1ivt} - \varepsilon_{0ivt}\right], \end{aligned} \tag{7}$$

In the "classical" case in which the disturbance terms in the results equations are identical ($U_{1ivt} = U_{0ivt} = \lambda_i + \sigma V_{vt} + \varepsilon_{ivt}$) and $\beta_1 = \beta_0$, this boils down to:

$$Y_{ivt} = \alpha_0 + X_{ivt}\beta_0 + D_{vt}\left(\alpha_1 - \alpha_0\right) + \lambda_i + \sigma V_{vt} + \varepsilon_{ivt}, \tag{8}$$

which is simply a panel regression where the correlation between treatment status D_{vt} and the disturbance in the structural equation ($\lambda_i + \sigma V_{vt} + \varepsilon_{ivt}$) is generated by the time-varying village-specific shocks V_{vt}, when $\sigma \neq 0$. Consistent parameter estimates can then be obtained by running a 2SLS panel regression with child-specific effects (in order to control for λ_i), using Z_{vt} as the excluded IVs. The "treatment effect" is then given by the estimate of the parameter $\alpha_1 - \alpha_0$. As a benchmark, and for comparison purposes, we will present such an estimate for the case of the PNIR in the results section below.

Define the propensity score as:

$$P\left(Z_{vt}\right) = \Pr\left(Z_{vt}\gamma - V_{vt} > 0\right). \tag{9}$$

By the *Probability Integral Transform*, we can, without loss of generality (Heckman and Vytlacil 2005), arbitrarily normalize the disturbance term V_{vt} so that it is distributed according to the uniform density over the unit interval, implying that treatment occurs when:

$$Z_{vt}\gamma > U_{Dvt},$$

where U_{Dvt} is distributed according to the uniform density. It follows that the probability of treatment is given by:

$$\Pr(D_{vt} = 1) = P\left(Z_{vt}\right) = \int_0^{P(Z_{vt})} dU_{Dvt}.$$

Taking expectations of (7), with respect to unobservables, conditional on $(X_{ivt}, P\left(Z_{vt}\right))$, yields:

$$E\left[Y_{ivt}\left|X_{ivt} = x_{ivt}, P\left(Z_{vt}\right) = p_{vt}\right.\right] = \alpha_0 + x_{ivt}\beta_0 + p_{vt}x_{ivt}\left(\beta_1 - \beta_0\right) + K\left(p_{vt}\right), \tag{10}$$

where we have grouped all of the non-linear terms under the $K\left(p_{vt}\right)$ heading, and where:

$$\begin{aligned} K\left(p_{vt}\right) &= p_{vt}\left(\alpha_1 - \alpha_0\right) + E\left[\lambda_i + \sigma_0 V_{vt} + \varepsilon_{0ivt}\left|P\left(Z_{vt}\right) = p_{vt}\right.\right] \\ &\quad + E\left[\left(\sigma_1 - \sigma_0\right)V_{vt} + \varepsilon_{1ivt} - \varepsilon_{0ivt}\right]p_{vt}. \end{aligned}$$

The marginal treatment effect or local instrumental variables estimator (Heckman and Vytlacil 1999) is then defined as:

$$\Delta^{MTE}\left(x_{ivt}, u_{Dvt}\right) = \left.\frac{\partial E\left[Y_{ivt}\left|X_{ivt} = x_{ivt}, P\left(Z_{vt}\right) = p_{vt}\right.\right]}{\partial p_{vt}}\right|_{p_{vt} = u_{Dvt}} = \Delta^{LIV}\left(x_{ivt}, u_{Dvt}\right). \tag{11}$$

For the case at hand, differentiating (10) with respect to p_{ivt} yields:

$$\Delta^{MTE}\left(x_{ivt}, u_{Dvt}\right) = x_{ivt}\left(\beta_1 - \beta_0\right) + \left.\frac{\partial K\left(p_{vt}\right)}{\partial p_{vt}}\right|_{p_{vt}=u_{Dvt}}. \tag{12}$$

The usual treatment parameters are then defined as weighted averages of (12). For the case of the impact of "treatment on the treated" (Δ^{TT} or TT in the standard Heckman notation), for example, we have:

$$\begin{aligned}
\Delta^{TT} &= \int \underbrace{\left(\int_0^1 \Delta^{MTE}\left(x_{ivt}, u_{Dvt}\right)\omega_{TT}(x_{ivt}, u_{Dvt})du_{Dvt}\right)}_{\Delta^{TT}(x_{ivt})} dF\left(x_{ivt}\right| D_{vt}=1) \\
&= \int \left(\int_0^1 \left(x_{ivt}\left(\beta_1 - \beta_0\right) + \left.\frac{\partial K\left(p_{vt}\right)}{\partial p_{vt}}\right|_{p_{vt}=u_{Dvt}}\right)\omega_{TT}(x_{ivt}, u_{Dvt})du_{Dvt}\right) \\
&\quad \times dF\left(x_{ivt}\right| D_{vt}=1),
\end{aligned} \tag{13}$$

where the appropriate weights $\omega_{TT}(x_{ivt}, u_{Dvt})$ are defined in Heckman (2001). Similar expressions can be computed for the average treatment effect (ATE) and for treatment on the untreated (TUT).

3 Identification strategy and results

3.1 Identification strategy

Our identification strategy is based on variations in the stock of political capital enjoyed by different villages within the *Conseil rural* of the CRs in which they are located. The intuition is as follows: the *Conseil rural* in a CR is one of the main institutional actors that determines whether PNIR projects proposed by various villages within the CR obtain PNIR funding, and it is the *Conseil rural* that must arbitrate between the competing claims of several villages. The *Conseil rural* is constituted by individuals who originate from different villages within the CR.[6] Our hypothesis is that the probability of a village obtaining a project, when it is eligible (i.e. when the village is within a CR that is treated by the PNIR), is an increasing function of its "stock of political capital" within the *Conseil rural.*

Given the vibrant nature of party politics in Senegal at the sub-regional level, we construct two instruments on the basis of this intuition. First, we consider whether at least one of the village's councilors has a party affiliation that corresponds to the party which controls the *greatest number* of seats on the *Conseil rural* (the party in question may therefore not possess an absolute majority). On the one hand, it is possible that belonging to the political party that controls the largest block of votes within the *Conseil rural* may increase the political leverage of the village's councilors. On the other hand, standard political economy arguments suggest that one might uncover a "dictatorship of the minority" effect, in which belonging to a minority group increases one's power through one's ability to *block* proposals. Note that there are a total of 16 different political parties represented in the *Conseil ruraux* in our dataset, though 2 —the ruling

[6] See Senegal (1998) for the institutional details.

Liberal party of President Wade, and the Socialist party of former President Diouf— are by far the most important. Second, we consider the same variable, but defined in terms of the *absolute majority* on the *Conseil rural*, when such a majority exists.

A third political variable that we use in order to achieve identification is constituted by the number of female villagers on the *Conseil rural*: given the relative rarity of female *Conseil rural* members, it is likely that their presence on the *Conseil rural* will have an impact on the likelihood of projects being attributed to a given village. In particular, since the philosophy that underlies CDD programs such as the PNIR is to provide "voice" to disenfranchised groups, such as women, one would expect the number of female villagers sent to the *Conseil rural* to significantly increase the likelihood of the village receiving a completed PNIR project.

For all of these politically-motivated instrumental variables, the exclusion restrictions are that they do not have any effect on child health within the village, apart from the indirect effect that obtains through the probability that the village receives a PNIR project. Given that we control for household- or child-specific effects, these exclusion restrictions are robust to time-invariant unobservables that would affect both the political instruments and household income or child health. For our exclusion restrictions on these IVs to be invalid, one would therefore need time-varying shocks to affect both the probability of obtaining a project at the village level and the political makeup of the *Conseil rural* at the CR level. Though possible, such a configuration strikes us as being highly unlikely, especially given that the main features of the makeup of the various *Conseil ruraux* in our sample were determined in the 2001 local elections, two years prior to our baseline survey. Variations across time in the makeup of the *Conseil ruraux* are thus, to a large extent, the result of random deaths and retirements. Summary statistics on the excluded IVs are presented in Table 1.

Note also that village political power at the *Conseil rural* level could be used to attract other programmes which might then confound our identification of the effect attributable to treatment by the PNIR. This last point is entirely possible and indeed likely, but given that the other interventions were time-invariant at the village level in our dataset, the child-specific effects (village-specific effects would suffice and are subsumed under the fixed effects at lower levels of aggregation) already purge this potential confounding factor. Given the progressive phasing-in of PNIR eligibility and village projects, and the variation over time in the political capital of the villages, this collinearity does not apply to the PNIR projects and allows one to identify their effects.

3.2 The marginal treatment effect of CDD

In Part I of Table 2, we present within-child least squares and within-child IV estimates of the impact of completed PNIR projects on WAZ. The IV point estimate, using the excluded instruments described in the preceding section, is equal to 1.086 ($se = 0.61$) indicating that treatment by the PNIR erases all of the nutritional shortfall of children in the sample, as can be seen from the summary statistics presented in Table 3. These estimates are extremely robust to changes in the exclusion restrictions, and readily pass the standard test of the overidentifying restrictions as well as the Hahn and Hausman (2002) joint test of instrument strength and instrument orthogonality. The question posed in this paper is whether this treatment effect is constant, or whether it varies across villages with *unobservables* that render them more or less likely to receive a completed project.

In order to test for the presence of essential heterogeneity, we estimate a simple discrete choice model (probit) that corresponds to the indicator function representation given in equation (4). Results are presented in Table 4, with all three excluded IVs being highly significant determinants, at usual levels of confidence, of the presence of a completed PNIR project in the village. Histograms of the estimated propensity scores, for treated and untreated villages, are presented in Figure 1. Overlap of the two histograms is satisfactory, implying that we are able to estimate the marginal treatment effect (and the associated standard measures of treatment) over almost the entire unit interval.

We then estimate the same basic specification as in Part I of Table 2, while replacing the treatment dummy by the estimated propensity score stemming from Table 4. The first test for the presence of essential heterogeneity involves including the estimated propensity score in squared and cubed form, and testing for the joint significance of these two terms. Results are presented in Part II of Table 2, and indicate rejection of the homogeneous treatment effect null. We obtain roughly the same result when adding interactions between the estimated propensity score and all of the X_{ivt} (which corresponds to adding the third term on the right-hand-side of equation (10)), as we do using a test based on a semiparametric specification (which corresponds exactly to the specification given in equation (10)). The results of these tests imply that one should be extremely wary of standard IV point estimates such as the one presented in Part I of Table 2, given that the underlying maintained hypothesis of a constant marginal treatment effect is soundly rejected by the data.

Having established that essential heterogeneity exists in the effect of treatment by the PNIR on child health, we adopt the semiparametric approach using local linear regression methods first introduced in Heckman, Ichimura, Smith, and Todd (1998) and Heckman and Vytlacil (1999) and developed further in Heckman, Urzua, and Vytlacil (2006).[7] The first step involves estimating equation (10) in semiparametric form (Ruppert, Wand, and Carroll 2003) using the procedure described in Heckman, Urzua, and Vytlacil (2006). Results are presented in Figure 2, revealing a good deal of heterogeneity in the slope of $K(p_{vt})$. This is a first indication that the impact of treatment by the PNIR may not be the same for all individuals in the sample.

The second step involves a semiparametric estimation of $\frac{\partial K(p_{vt})}{\partial p_{vt}}$, which corresponds to the second element on the right-hand-side of equation (12), the result of which is presented in Figure 3.[8] Figure 3 provides concrete evidence, to the extent that our identification strategy is valid, that there are large variations in the marginal treatment effect as u_{Dvt} varies. Finally, the estimated "standard" treatment effects, as defined for the example of treatment on the treated in equation (13) are presented in Part III of Table 2. Given the great heterogeneity revealed in Figure 3, it is not surprizing that the standard errors associated with the three standard treatment effects (ATE, TT and TUT) are extremely large.

[7] See Carneiro, Heckman, and Vytlacil (2003) for an application of these methods to estimating the returns to college, as well as Heckman and Navarro-Lozano (2004) for a general discussion of the relative strengths and weaknesses of IV, matching and control function approaches.

[8] On derivative estimation, see Newell and Einbeck (2007).

The interpretation of Figure 3, and its correspondence with equation (12), is as follows: given that the only statistically significant $x_{ivt}(\beta_1 - \beta_0)$ term involved in the underlying semiparametric estimation is that associated with child age ($\beta_1 - \beta_0$ =0.159, $se = 0.01$), and given that mean child age in the sample is equal to 18.44 months, the corresponding *linear* portion (see equation (12)) of the marginal treatment effect is estimated at $0.159 \times 18.44 = 2.93$. As such, shifting up the $\frac{\partial K(p_{vt})}{\partial p_{vt}}$ curve by 2.93 reveals that the *total* marginal treatment effect of the program ($\Delta^{MTE}(x_{ivt}, u_{Dvt})$) is positive and often relatively large until a value of $u_D \approx 0.35$ is reached. For $u_D \in [0.35, 0.75]$ the marginal impact of the program is roughly zero (given the magnitude of the ± 2 standard error confidence interval), whereas for $u_D \in [0.75, 1]$ the impact of the program is negative. In words, these results imply that: (i) when unobservables are such that it is highly likely that a village will receive a completed PNIR project, the gains to the program are positive and quantitatively important, and this is confirmed by the estimated value of treatment on the treated (TT= 1.712, see Part III of Table 2), which places more weight on observations close to $u_D = 0$; (ii) when unobservables are such that it is unlikely that the village will receive a project ($u_D = 1$), the gains to the program are either statistically insignificant or negative; again, this is reflected by the value of treatment on the untreated (TUT= -1.944). In all cases, the estimated standard treatment effects are in no way similar to the linear IV point estimate, highlighting the basic point that such estimates can often be misleading.

4 Concluding remarks

Application of Heckam and Vytlacil's local instrumental variables method to estimate the marginal treatment effect of a CDD program in Senegal has revealed that the marginal treatment effect is a decreasing function of unobservables inversely associated with the probability of a given village receiving a completed project. In terms of the efficiency of the decentralized allocation of projects by this particular CDD program, it is apparent that villages that benefit the most from projects are those whose unobservables also make them more likely to receive one.

This finding provides extremely strong support for the decentralized allocation mechanism set up by the program initiators. In particular, had we found an upward-sloping marginal treatment effect, the opposite conclusion would have to be drawn, namely that villages who would benefit the most from the program are those whose unobservables are such that they are the least likely to receive a completed project. It is noteworthy that our result does not stem from any centralized form of targeting, but rather from decentralized decision-making allowing the local *Conseil rural* discretion in terms of arbitrage among various competing project proposals. Our paper demonstrates that application of local instrumental variables is feasible in a developing country context, and provides a simple diagnostic, crystallized in the form of the slope of the marginal treatment effect, that allows one to assess whether project allocation is or is not efficient.

References

Carneiro, P., J. J. Heckman, and E. Vytlacil (2003): "Understanding What Instrumental Variables Estimate: Estimating Marginal and Average Returns to Education," processed, University of Chicago, The American Bar Foundation and Stanford University, July 19.

Florens, J.-P., J. J. Heckman, C. Meghir, and E. Vytlacil (2004): "Instrumental Variables, Local Instrumental Variables and Control Functions," Centre for Microdata Methods and Practice, Institute for Fiscal Studies, CWP15/02.

Hahn, J., and J. A. Hausman (2002): "A New Specification Test for the Validity of Instrumental Variables," *Econometrica*, 70(1), 163–189.

Heckman, J. J. (1978): "Dummy Endogenous Variables in a Simultaneous Equation System," *Econometrica*, 46(4), 931–959.

——— (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47(1), 153–161.

——— (2001): "Micro Data, Heterogeneity and the Evaluation of Public Policy: Nobel Lecture," *Journal of Political Economy*, 109(4), 673–748.

Heckman, J. J., H. Ichimura, J. Smith, and P. Todd (1998): "Characterizing Selection Bias Using Experimental Data," *Econometrica*, 66(5), 1017–1098.

Heckman, J. J., and S. Navarro-Lozano (2004): "Using Matching, Instrumental Variables and Control Functions to Estimate Economic Choice Models," *Review of Economics and Statistics*, 86(1), 30–57.

Heckman, J. J., J. Smith, and N. Clements (1997): "Making the Most of Programme Evaluations and Social Experiments: Accounting for Heterogeneity in Programme Impacts," *Review of Economic Studies*, 64(4, Special Issue), 487–535.

Heckman, J. J., S. Urzua, and E. Vytlacil (2006): "Understanding Instrumental Variables in Models with Essential Heterogeneity," *Review of Economics and Statistics*, 88(3), 389–432.

Heckman, J. J., and E. J. Vytlacil (1999): "Local Instrumental Variables and Latent Variable Models for Identifying and Bounding Treatment Effects," *Proceedings of the National Academy of Sciences*, 96(8), 4730–4734.

——— (2001): "Policy-Relevant Treatment Effects," *American Economic Review, Papers and Proceedings*, 91(2), 107–111.

——— (2005): "Structural Equations, Treatment Effects and Econometric Policy Evaluation," *Econometrica*, 73(3), 669–738.

Newell, J., and J. Einbeck (2007): "A Comparative Study of Nonparametric Derivative Estimators," Proceedings of the 22th IWSM, Barcelona.

Quandt, R. (1972): "A New Approach for Estimating Switching Regressions," *Journal of the American Statistical Association*, 67(338), 306–310.

Rubin, D. B. (1974): "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Education Psychology*, 66, 688–701.

Ruppert, D., M. P. Wand, and R. J. Carroll (2003): *Semiparametric Regression*. Cambridge University Press, New York, NY.

Senegal (1998): "Le Guide Du Conseiller Rural," Dakar, Sénégal, Electronic edition: Friedrich Ebert Stiftung, Bonn, Germany. Available online at www.fes.de.

	Full sample	22 CRs that are treated by $t = 4$
	Mean (std.)	Mean (std.)
Village is PNIR-eligible	0.565	0.775
	(0.49)	(0.41)
Village political capital:		
Villager is a member of the biggest		0.502
party on the Conseil rural		(0.50)
Villager is a member of the majority		0.429
party on the Conseil rural		(0.49)
Number of female villagers		0.457
on the Conseil rural		(0.82)
Observations	341	249
Time periods	5	5
Villages	71	52

Table 1: Summary statistics on village political capital IVs.

Part I: Least Squares and IV (Fuller) estimates	
Least squares estimate	0.089 (0.13)
IV estimate	1.086 (0.61)
Test of the OID restrictrictions [$p-$value]	0.677 [0.712]
Hahn-Hausman m_2 test statistic [$p-$value]	-0.143 [0.885]

Part II: Testing for the presence of essential heterogeneity	
joint significance of p_{ivt}^2 and p_{ivt}^3	$p-$value= 0.0007
joint significance of p_{ivt}^2 and p_{ivt}^3 and interactions between p_{ivt} and the X_{ivt}	$p-$value= 0.0000
significance of smooth terms in p_{ivt} in a semiparametric specification	$p-$value= 0.0000

Part III: Standard treatment effects	
ATE	0.207 (2.33)
TT	1.712 (2.59)
TUT	-1.944 (2.80)
Observations (treated)	1,315 (1,039)
Time periods	5
CRs	22
Villages	52
Households	343
Children	663

Table 2: Part I: Least squares and instrumental variables (Fuller) estimates of the impact of completed PNIR projects on child anthropometrics ($z-$scores). (standard errors clustered at the village level in parentheses). Child-specific effects included. Village-, household- and child-specific covariates included. Part II: Testing for essential heterogeneity. Part III: Standard treatment effects.

	Mean	Min	Max	SD
Child characteristics				
Height-for-age $z-$score	-1.25	-4.99	3.00	1.54
Weight-for-age $z-$score	-1.00	-4.89	4.54	1.34
Weight-for-height $z-$score	-0.24	-3.96	4.79	1.31
Age (months)	18.44	0.10	36.99	10.07
Female	0.491	0	1	0.500
Household characteristics				
Expenditures per capita	$13,708$	142	$152,500$	$13,175$
Age of head	53	17	92	14.1
Household size	10.7	1	34	4.9
Head literate	0.360	0	1	0.480
Female head	0.130	0	1	0.336
Ethnic group of head:				
Wolof	0.479	0	1	0.499
Pular	0.285	0	1	0.451
Serer	0.161	0	1	0.367
Diola	0.022	0	1	0.149
Other	0.017	0	1	0.132
Village characteristics				
Population of village	$1,331$	135	$10,046$	$1,538$
Electricity in village	0.252	0	1	0.434
Literacy program in village	0.527	0	1	0.499

Table 3: Summary statistics on the full sample: 5 time periods, 36 CRs, 71 villages. For child characteristics, the sample is constituted by 1,000 children that can be followed over at least two periods, living in 498 households yielding a total of 2,069 observations. For household characteristics, the sample is constituted by 756 households that can be followed over at least 2 periods (the sample is larger because a number of households have no children aged between 0 and 36 months) yielding a total of 3,458 observations.

Village political capital excluded IVs	
Villager is a member of the biggest party on the Conseil rural	-1.027 (0.19)
Villager is a member of the majority party on the Conseil rural	2.451 (0.20)
Number of female villagers on the Conseil rural	0.677 (0.07)
Akaike information criterion	961.23

Table 4: Determinants of completed PNIR projects: probit model. Village-, household- and child-specific covariates included (standard errors in parentheses).

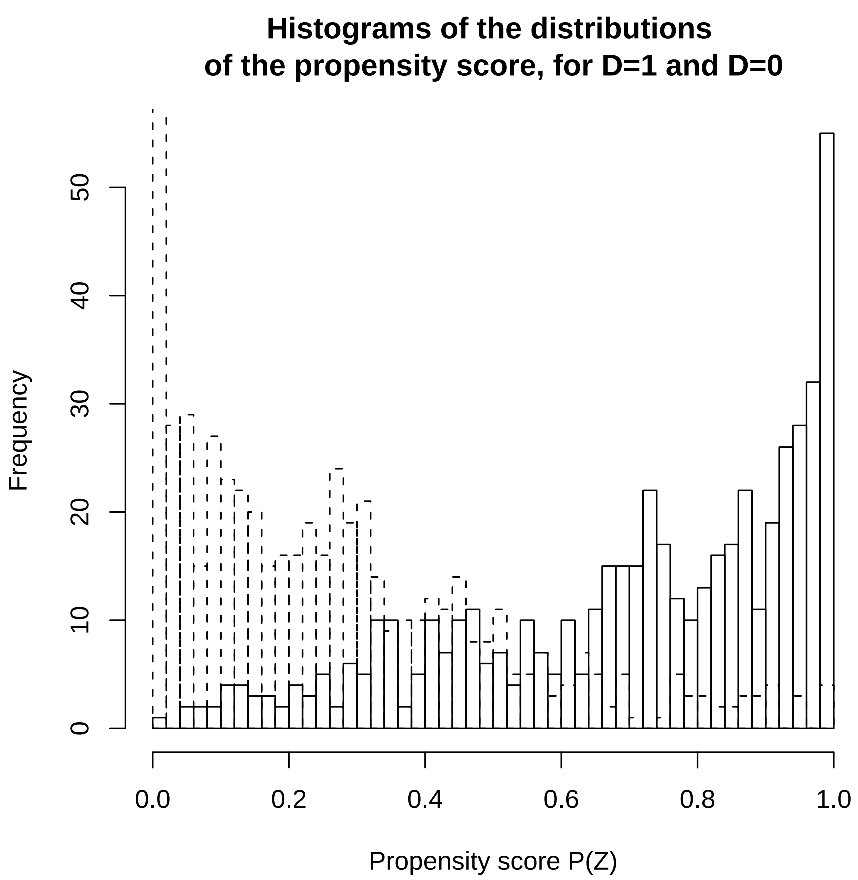

Figure 1: Histograms of the estimated propensity scores.

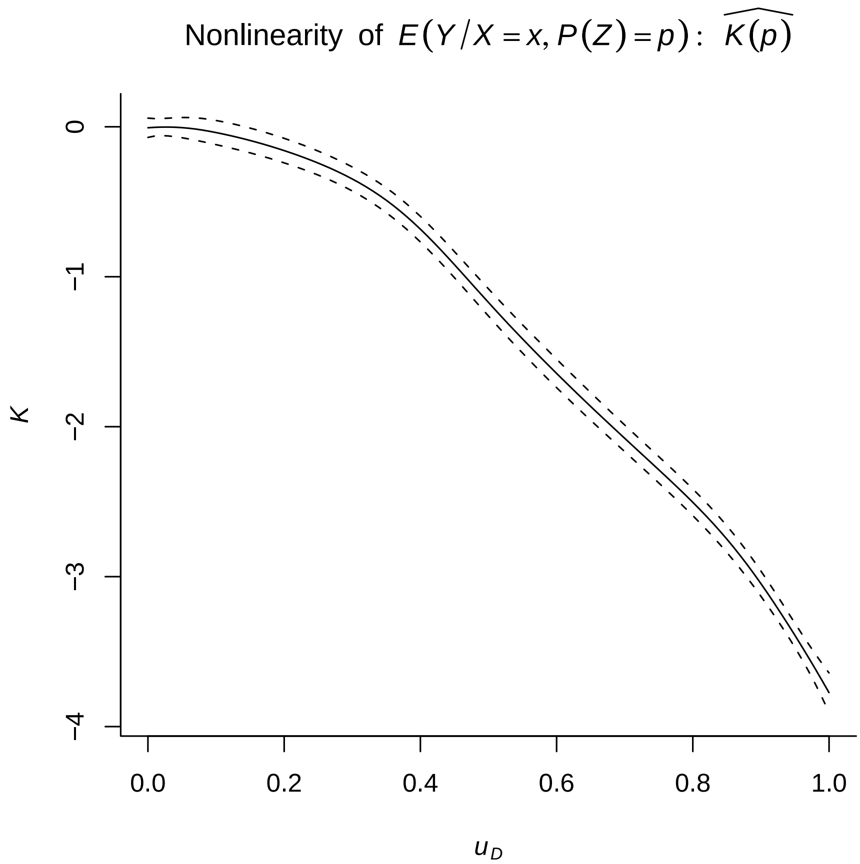

Figure 2: Semiparametric estimation of $K(p_{vt})$, with 95% confidence bands.

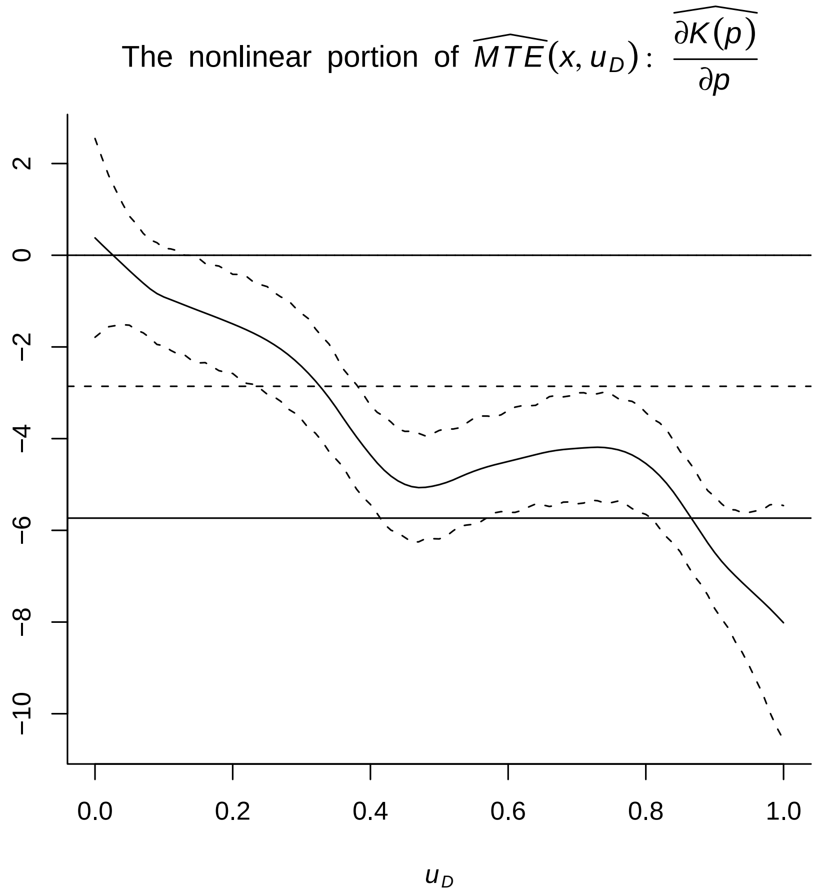

Figure 3: Semiparametric estimate of $\frac{\partial K(p_{vt})}{\partial p_{vt}}$, with 95% confidence bands.