

EVALUATING MARGINAL POLICY CHANGES AND THE AVERAGE EFFECT OF TREATMENT FOR INDIVIDUALS AT THE MARGIN

BY PEDRO CARNEIRO, JAMES J. HECKMAN, AND EDWARD VYTLACIL¹

This paper develops methods for evaluating marginal policy changes. We characterize how the effects of marginal policy changes depend on the direction of the policy change, and show that marginal policy effects are fundamentally easier to identify and to estimate than conventional treatment parameters. We develop the connection between marginal policy effects and the average effect of treatment for persons on the margin of indifference between participation in treatment and nonparticipation, and use this connection to analyze both parameters. We apply our analysis to estimate the effect of marginal changes in tuition on the return to going to college.

KEYWORDS: Marginal treatment effect, effects of marginal policy changes, marginal policy relevant treatment effect, average marginal treatment effect.

1. INTRODUCTION

THE POLICY RELEVANT TREATMENT EFFECT (PRTE) is the mean effect of changing from a baseline policy to an alternative policy that provides different incentives to participate in treatment (Heckman and Vytlačil (2001b, 2005)). Identification and estimation of the PRTE are generally difficult tasks. Identification of the PRTE typically requires large support conditions. \sqrt{N} -consistent estimation of the PRTE parameter is generally not possible.

In many cases, proposed policy reforms are incremental in nature and a marginal version of the PRTE (MPRTE) is all that is required to answer questions of economic interest. This paper develops the MPRTE and establishes how the MPRTE depends on the direction of a proposed marginal policy change. We establish that the support conditions required for identifying the MPRTE are very weak. The essential requirement is availability of a continuous instrument. The MPRTE parameter can be represented as a weighted average derivative with weights determined by the marginal policy of interest. The parameter is \sqrt{N} -estimable under standard regularity conditions. Thus, the MPRTE parameter is fundamentally easier to identify and estimate than the PRTE parameter.

We connect the MPRTE to the average marginal treatment effect (AMTE): the mean benefit of treatment for people at the margin of indifference between

¹This research was supported by NIH R01-HD32058-3, NSF SES-0832845, NSF SES-024158, NSF SES-05-51089, ESRC RES-000-22-2542, the Geary Institute at University College Dublin, the American Bar Foundation, the Pritzker Consortium on Early Childhood Development, the Leverhulme Trust, and ESRC (RES-589-28-0001) through the funding of the Centre for Microdata Methods and Practice. The research was conducted in part while Edward Vytlačil was a Visiting Professor at Hitotsubashi University. We thank the editor, two anonymous referees, Hidehiko Ichimura, Richard Robb, Daniel Schmierer, and Azeem Shaikh for very helpful comments. We would like to thank Erica Blom and Sukjin Han for research assistance.

participation in treatment and nonparticipation. AMTE is compared to marginal cost in an aggregate benefit–cost analysis of a program. We establish a correspondence between MP RTE parameters and AMTE parameters, showing that the effect of a marginal policy change in a particular direction is the same as the average effect of treatment for those at the margin of indifference. We use this correspondence to produce new insights about the AMTE parameter.

The paper proceeds as follows. Section 2 presents the nonparametric selection model that underlies our analysis. Section 3 reviews the analysis of the PRTE by Heckman and Vytlacil (2005) and presents a new interpretation of the PRTE as a function from the space of all possible policies to the space of effects of policies. This interpretation is key to the analysis in Section 4, which introduces and analyzes the MP RTE. We discuss identification and estimation of the MP RTE in Sections 5 and 6, respectively. We define and analyze the AMTE in Section 7. Section 8 presents an empirical application of our analysis. Section 9 concludes.

2. NONPARAMETRIC SELECTION MODEL AND OUR ASSUMPTIONS

Assume that there are two potential outcomes (Y_0, Y_1) and a binary treatment choice indicator D . Outcome Y is written in switching regression form $Y = DY_1 + (1 - D)Y_0$, where Y_1 is the potential outcome that is observed if the agent chooses treatment 1 and Y_0 is the potential outcome that is observed if the agent chooses treatment 0. $Y_1 - Y_0$ is the individual level treatment effect. We keep implicit the conditioning on observed variables X that determine Y_1 or Y_0 , and maintain the assumption that D does not determine X (see Heckman and Vytlacil (2005)).

Program participation is voluntary. To link this framework to standard choice models, we characterize the decision rule for program participation by a latent index model:

$$(2.1) \quad D = \mathbf{1}[\mu(Z) - V \geq 0],$$

where $\mathbf{1}[\cdot]$ is the indicator function taking the value 1 if its argument is true and the value 0 otherwise. From the point of view of the econometrician, Z is observed and V is unobserved.

We maintain the following assumptions:

ASSUMPTION A-1: (Y_0, Y_1, V) is independent of Z .

ASSUMPTION A-2: The distribution of V is absolutely continuous with respect to Lebesgue measure.

ASSUMPTION A-3: The distribution of $\mu_D(Z)$ is absolutely continuous with respect to Lebesgue measure.

ASSUMPTION A-4: $\sup_v E(|Y_1||V = v) < \infty$ and $\sup_v E(|Y_0||V = v) < \infty$.

ASSUMPTION A-5: $0 < \Pr(D = 1) < 1$.

These conditions are discussed in Heckman and Vytlačil (2005, 2007). Under them, Vytlačil (2002) establishes the equivalence between the nonparametric latent index model (2.1) and the monotonicity assumption used by Imbens and Angrist (1994). A necessary condition for A-3 is that Z contains a continuous variable (i.e., that there is a continuous instrument for D). We use this assumption to analyze marginal policy changes and to identify marginal policy treatment effects. These conditions should be interpreted as conditional on X , which we have kept implicit. For example, A-1 should be interpreted as saying that (Y_0, Y_1, V) is independent of Z conditional on X , while there can be dependence between (Y_0, Y_1, V) and X .

Define $P(Z)$ as the probability of receiving treatment given Z : $P(Z) \equiv \Pr(D = 1|Z) = F_V(\mu(Z))$, where $F_V(\cdot)$ is the distribution function of V . We sometimes denote $P(Z)$ by P , suppressing the Z argument. We also use U , defined by $U = F_V(V)$ so that U is distributed unit uniform. In this notation, applying a monotonic transformation to both sides of the argument in equation (2.1) allows us to write that equation as $D = \mathbf{1}[P(Z) \geq U]$.

The marginal treatment effect (MTE) plays a fundamental role in our analysis. MTE is defined as $\text{MTE}(u) \equiv E(Y_1 - Y_0|U = u)$, that is, the expected treatment effect conditional on the unobservables that determine participation. For values of u close to zero, $\text{MTE}(u)$ is the expected effect of treatment on individuals who have unobservables that make them most likely to participate in treatment and who would participate even if the mean scale utility $\mu(Z)$ were small. See Heckman and Vytlačil (2005, 2007) for more discussion and interpretation of the MTE.

In some of our examples, we consider the following special case of our general model:

ASSUMPTION B-1: *Suppose $\mu(Z) = Z\gamma$ and suppose that F_V is strictly increasing. Suppose that the k th component of Z , $Z^{[k]}$, has a strictly positive coefficient, $\gamma^{[k]} > 0$. Let $\tilde{Z}\tilde{\gamma} = Z\gamma - Z^{[k]}\gamma^{[k]}$. Suppose that the distribution of $Z^{[k]}\gamma^{[k]}$ conditional on $\tilde{Z}\tilde{\gamma}$ has a density with respect to Lebesgue measure.*

3. THE POLICY RELEVANT TREATMENT EFFECT

We first present the PRTE of Heckman and Vytlačil (2001b, 2005) to motivate our marginal version of it. We then reformulate the PRTE as a function of the proposed policy change. This enables us to define a sequence of PRTEs corresponding to a sequence of proposed policy changes.

3.1. Review of PRTE

Following Heckman and Vytlacil (2001b, 2005), consider a class of policies that affect P —the probability of participation in a program—but that do not affect potential outcomes or unobservables related to the selection process, (Y_1, Y_0, U) .² An example from the literature on the economic returns to schooling would be policies that change tuition or distance to school but that do not directly affect potential wages (Card (2001)). We ignore general equilibrium effects.

Let D^* be the treatment choice that would be made after the policy change. Let P^* be the corresponding probability that $D^* = 1$ after the policy change. D^* is defined by $D^* = \mathbf{1}[P^* \geq U]$. Let $Y^* = D^*Y_1 + (1 - D^*)Y_0$ be the outcome under the alternative policy. Following Heckman and Vytlacil (2005), the mean effect of going from a baseline policy to an alternative policy per net person shifted is the PRTE, defined when $E(D) \neq E(D^*)$ as

$$(3.1) \quad \frac{E(Y|\text{alternative policy}) - E(Y|\text{baseline policy})}{E(D|\text{alternative policy}) - E(D|\text{baseline policy})} \\ = \frac{E(Y^*) - E(Y)}{E(D^*) - E(D)} = \int_0^1 \text{MTE}(u) \omega_{\text{PRTE}}(u) du,$$

where

$$(3.2) \quad \omega_{\text{PRTE}}(u) = \frac{F_P(u) - F_{P^*}(u)}{E_{F_{P^*}}(P) - E_{F_P}(P)}.$$

The condition $E(D) \neq E(D^*)$ is consistent with the program having a non-monotonic effect on participation as long as the fraction switching into treatment is not exactly offset by the fraction switching out of treatment. The PRTE parameter gives the normalized effect of a change from a baseline policy to an alternative policy and depends on the alternative being considered.³ Heckman and Vytlacil (2005) show that the PRTE can be identified under strong support conditions. In Section 5, we establish that the marginal version of the PRTE parameter can be identified under much weaker conditions than are required to identify the PRTE. In Section 6, we establish that, unlike the PRTE parameter, the MPRTTE parameter is generally estimable at a \sqrt{N} rate.

We define PRTE as the average effect per net person shifted into treatment. With this definition, no normalization is required when taking limits to define

²This restriction can be relaxed to a weaker policy invariance for the distribution of (Y_1, Y_0, U) ; see Heckman and Vytlacil (2005, 2007).

³The PRTE can be interpreted as an economically more explicit version of Stock's (1989) nonparametric policy analysis parameter for a class of policy interventions with explicit agent preferences where the policies evaluated operate solely on agent choice sets.

the MP RTE, and we obtain an equivalence between the MP RTE and the average effect for individuals at the margin of indifference (AMTE). Under this definition, the MP RTE corresponds to a weighted average derivative and can be estimated by a weighted average derivative estimator.

Alternatively, we could follow Heckman and Vytlacil (2001b) in defining PRTE without normalizing it by the net change in treatment status, that is, we could define PRTE as $E(Y^*) - E(Y)$. However, some normalization is required when taking limits. If we were to define PRTE in that manner and then normalize by $E(D^*) - E(D)$ when taking limits, we would obtain the same limits as would be obtained from analyzing the marginal version of PRTE defined by equation (3.1). Using the unnormalized version of PRTE, we can analyze alternative normalizations when taking limits. In footnote 5 below, we discuss one such normalization which results in an alternative version of MP RTE that is a rescaled version of the marginal version of PRTE defined by equation (3.1).

3.2. PRTE as a Function of Proposed Policy Changes

The PRTE depends on a policy change only through the distribution of P^* after the policy change. Given our assumptions, F_{P^*} is sufficient to summarize everything about the proposed policy change that is relevant for calculating the average effect of the policy change. We can thus define the PRTE function as a function mapping the proposed policy change (corresponding to a distribution of P^*) to the resulting per-person change in outcomes. Expressing the PRTE this way is important in the next step of our analysis that uses sequences of PRTEs to define a marginal version of the PRTE.

Let \mathcal{G} denote the space of all cumulative distribution functions for random variables that lie in the unit interval such that $\int t dG(t) \neq \int t dF_P(t)$, that is, all right-continuous, nondecreasing functions that satisfy $G(t) = 1$ for $t \geq 1$ and $G(t) = 0$ for $t < 0$, and such that the first moment of G does not equal the first moment of F_P . Any $G \in \mathcal{G}$ corresponds to a potential distribution of P^* and thus corresponds to a potential alternative policy, with $\text{PRTE}(G)$ denoting the corresponding policy effect. We define the PRTE function, $\text{PRTE} : \mathcal{G} \mapsto \mathbb{R}$, by

$$(3.3) \quad \text{PRTE}(G) = \int_0^1 \text{MTE}(u) \omega_{\text{PRTE}}(u; G) du,$$

where

$$(3.4) \quad \omega_{\text{PRTE}}(u, G) = \frac{F_P(u) - G(u)}{E_G(P) - E_{F_P}(P)}.$$

In many cases, the class of policy alternatives under consideration can be indexed by a scalar variable. Let $P_0 = P$ denote the baseline probability for $D = 1$. Let \mathbf{M} denote a subset of \mathbb{R} with $0 \in \mathbf{M}$. Let $\{P_\alpha : \alpha \in \mathbf{M}\}$ denote a class of alternative probabilities corresponding to alternative policy regimes with

associated cumulative distribution functions F_α . For example, one alternative policy could increase the probability of attending college by an amount α , so that $P_\alpha = P_0 + \alpha$ and $F_\alpha(t) = F_0(t - \alpha)$. An alternative policy could change each person's probability of attending college by the proportion $(1 + \alpha)$, so that $P_\alpha = (1 + \alpha)P_0$ and $F_\alpha(t) = F_0(\frac{t}{1+\alpha})$. The policy intervention might have an effect similar to a shift in one of the components of Z , say $Z^{[k]}$. In particular, suppose $Z_\alpha^{[k]} = Z^{[k]} + \alpha$ and $Z_\alpha^{[j]} = Z^{[j]}$ for $j \neq k$. For example, the k th element of Z might be college tuition, and the policy under consideration subsidizes college tuition by the fixed amount α . Suppose that the linear latent index Assumption B-1 holds. Then $P_\alpha(Z) = F_V(Z\gamma + \alpha\gamma^{[k]})$ and $F_\alpha(t) = F_{Z\gamma}(F_V^{-1}(t) - \alpha\gamma^{[k]})$. Notice that the first two examples have the form $P_\alpha = q_\alpha(P_0)$ for some function q_α , while the last example has the form $P_\alpha = F_V(Z^{[-k]}\gamma^{[-k]} + q_\alpha(Z^{[k]})\gamma^{[k]})$, where $Z^{[-k]}$ denotes the elements of Z not in $Z^{[k]}$ and $\gamma^{[-k]}$ is the corresponding coefficient vector. We will explore more general examples in the next section.

4. MARGINAL POLICY CHANGES

The PRTE is defined for a discrete change from a baseline policy to a fixed alternative. We now consider a marginal version of the PRTE parameter that corresponds to a marginal change from a baseline policy.⁴ It is expositionally convenient to think of the treatment as college attendance and the policy as a change in tuition. The marginal version of the PRTE depends on the nature of the perturbation that defines the marginal change. For example, a policy change that subsidizes tuition by a fixed amount has different effects than a policy change that subsidizes tuition by a fixed proportion. The limits of these two policies for infinitesimally small subsidies are different.

To define the marginal version of the PRTE, we could consider the limit of $\Delta^{\text{PRTE}}(G)$ as G gets close to F_p in some metric. We could define the marginal PRTE as a directional derivative. For ease of exposition, we do not work with this more general formulation, but instead work with a one-dimensional version of it. Thus we do not analyze general perturbations within the function space \mathcal{G} , but only one-dimensional curves within \mathcal{G} .

Consider a class of alternative policies indexed by α , $\{F_\alpha : \alpha \in \mathbf{M}\}$, where 0 is a limit point of \mathbf{M} and 0 represents the baseline policy so that $F_0 = F_p$. Consider the effect of a marginal change in α in a neighborhood of the current base policy of $\alpha = 0$. When the limit exists, we can define the MP RTE as

$$\text{MP RTE}(\{F_\alpha\}) = \lim_{\tau \rightarrow 0} \text{PRTE}(F_\tau),$$

where the MP RTE parameter depends on the class of alternative policies $\{F_\alpha : \alpha \in \mathbf{M}\}$, that is, on the choice of a particular curve within \mathcal{G} . Thus, the

⁴Ichimura and Taber (2002) present a discussion of local policy analysis in a model without the MTE structure using a framework developed by Hurwicz (1962).

MPRTE can be seen as a path derivative along the path $\{F_\alpha : \alpha \in \mathbf{M}\}$. We impose the following sufficient conditions for the limit to exist:

ASSUMPTION A-6: For α in a neighborhood of 0, F_α has a density f_α with respect to Lebesgue measure, f_α is differentiable in α , $\sup_{t \in (0,1)} |\frac{\partial}{\partial \alpha} f_\alpha(t)| < \infty$, and $\int_0^1 t \frac{\partial}{\partial \alpha} f_\alpha(t) dt \neq 0$.

Under Assumptions A-1–A-6, $\text{MPRTE}(\{F_\alpha\}) = \lim_{\tau \rightarrow 0} \text{PRTE}(F_\tau)$ exists and is given by

$$(4.1) \quad \text{MPRTE}(\{F_\alpha\}) = \int_0^1 \text{MTE}(u) \omega_{\text{MPRTE}}(u; \{F_\alpha\}) du,$$

where $\omega_{\text{MPRTE}}(u; \{F_\alpha\})$ is

$$(4.2) \quad \omega_{\text{MPRTE}}(u; \{F_\alpha\}) = -\frac{\frac{\partial}{\partial \alpha} F_0(u)}{\frac{\partial}{\partial \alpha} E_{F_0}(P)} = -\frac{\int_0^u \left(\frac{\partial}{\partial \alpha} f_0(p)\right) dp}{\int_0^1 p \left(\frac{\partial}{\partial \alpha} f_0(p)\right) dp},$$

where we write $\frac{\partial}{\partial \alpha} F_0(p)$ and $\frac{\partial}{\partial \alpha} f_0(p)$ as shorthand expressions for $\frac{\partial}{\partial \alpha} F_\alpha(p)|_{\alpha=0}$ and $\frac{\partial}{\partial \alpha} f_\alpha(p)|_{\alpha=0}$, respectively. An alternative way to express the form of the weights uses the property that $0 \leq P_\alpha \leq 1$ so that $E(P_\alpha) = \int_0^1 (1 - F_\alpha(t)) dt$ to obtain

$$(4.3) \quad \omega_{\text{MPRTE}}(u; \{F_\alpha\}) = \frac{\frac{\partial}{\partial \alpha} F_0(u)}{\int_0^1 \left(\frac{\partial}{\partial \alpha} F_0(t)\right) dt},$$

which makes it clear that the weights always integrate to unity.⁵

Just as the PRTE parameter depends on which particular policy counterfactual is being considered, the marginal PRTE parameter depends on which particular class of policy perturbations is being considered. Just as the effect of a fixed dollar amount tuition subsidy will result in a different PRTE parameter than a proportional tuition subsidy, the limit parameter for a marginal change in an additive tuition subsidy will be different from the limit parameter for a marginal change in a proportional tuition subsidy. Of particular interest to us are the following two cases for which we will define classes of functions. Let

⁵An alternative way to define MPRTE that does not require the condition $E(D) \neq E(D^*)$ normalizes by α instead of $E(D) - E(D^*)$: $\lim_{\alpha \rightarrow 0} \int_0^1 \text{MTE}(u) ([F_\alpha(u) - F_0(u)]/\alpha) du$. With this definition the corresponding weights on MTE are given by $\omega_{\text{MPRTE}}(u; \{F_\alpha\}) = \frac{\partial}{\partial \alpha} F_0(u)$. This change in the normalization only affects the constant of integration for the weights and results in an alternative MPRTE that is a rescaled version of the MPRTE analyzed in this paper.

\mathcal{Q} denote the set of sequences of functions $\{q_\alpha(\cdot) : \alpha \in \mathbf{M}\}$ such that $q_\alpha(t)$ is strictly increasing in t for any $\alpha \in \mathbf{M}$; such that $q_0(\cdot)$ is the identity function; such that $q_\alpha(t)$ is differentiable in α with $\frac{\partial}{\partial \alpha} q_\alpha(t)$ bounded; such that r_α is a bounded function where r_α denotes the inverse of q_α ; and such that r_α is differentiable in α with $\frac{\partial}{\partial \alpha} r_\alpha(t)$ bounded. The following examples can be trivially modified to allow q_α to be monotonically decreasing.

EXAMPLE 1: Suppose that the alternative sequence of policies has the form $P_\alpha = q_\alpha(P_0)$ for some $\{q_\alpha\} \in \mathcal{Q}$. Then $F_\alpha(t) = F_P(r_\alpha(t))$ and

$$(4.4) \quad \omega_{\text{MPRTE}}(u; \{F_\alpha\}) = -\frac{f_P(u) \frac{\partial}{\partial \alpha} r_0(u)}{E_P\left(\frac{\partial}{\partial \alpha} q_0(P)\right)} = \frac{f_P(u) \frac{\partial}{\partial \alpha} r_0(u)}{\int_0^1 f_P(t) \frac{\partial}{\partial \alpha} r_0(t) dt}$$

EXAMPLE 2: Suppose that the alternative sequence of policies shifts the k th component of Z , $Z^{[k]}$, to $q_\alpha(Z^{[k]})$ for some $\{q_\alpha\} \in \mathcal{Q}$. Suppose that Assumption B-1 holds. Then

$$F_\alpha(t) = \int F_{Z^{[k]|\tilde{Z}\tilde{y}}} \left[r_\alpha \left(\frac{F_V^{-1}(t) - s}{\gamma^{[k]}} \right) \right] f_{\tilde{Z}\tilde{y}}(s) ds$$

and

$$(4.5) \quad \omega_{\text{MPRTE}}(u; \{F_\alpha\}) = \frac{E_{\tilde{Z}\tilde{y}} \left[f_{Z^{[k]|\tilde{Z}\tilde{y}}} \left[\frac{F_V^{-1}(u) - \tilde{Z}\tilde{y}}{\gamma^{[k]}} \right] \frac{\partial}{\partial \alpha} r_0 \left(\frac{F_V^{-1}(u) - \tilde{Z}\tilde{y}}{\gamma^{[k]}} \right) \right]}{\int_0^1 E_{\tilde{Z}\tilde{y}} \left[f_{Z^{[k]|\tilde{Z}\tilde{y}}} \left[\frac{F_V^{-1}(t) - \tilde{Z}\tilde{y}}{\gamma^{[k]}} \right] \frac{\partial}{\partial \alpha} r_0 \left(\frac{F_V^{-1}(t) - \tilde{Z}\tilde{y}}{\gamma^{[k]}} \right) \right] dt}$$

The expressions in equations (4.4) and (4.5) look different, but both can be represented as weighted densities

$$(4.6) \quad \omega_{\text{MPRTE}}(u; \{F_\alpha\}) = f_P(u)h(u)/E(h(P))$$

for some function $h(\cdot)$. For Example 1, this form is immediate, substituting $h(u) = \frac{\partial}{\partial \alpha} r_0(u)$. For Example 2, note that, for $\omega_{\text{MPRTE}}(u; \{F_\alpha\})$ given by equation (4.5),

$$\begin{aligned} \omega_{\text{MPRTE}}(u; \{F_\alpha\}) \neq 0 &\Rightarrow \int f_{Z^{[k]|\tilde{Z}\tilde{y}}} \left[\frac{F_V^{-1}(u) - s}{\gamma^{[k]}} \right] f_{\tilde{Z}\tilde{y}}(s) ds \neq 0 \\ &\Leftrightarrow f_{Z\gamma}(F_V^{-1}(u)) \neq 0 \\ &\Leftrightarrow f_P(u) f_V(F_V^{-1}(u)) \neq 0 \\ &\Rightarrow f_P(u) \neq 0, \end{aligned}$$

where we use the fact that $\Pr[P(Z) \leq u] = \Pr[Z\gamma \leq F_V^{-1}(u)]$ and the chain rule to obtain $f_P(u) = (f_{Z\gamma}(F_V^{-1}(u)))/(f_V(F_V^{-1}(u)))$. Thus, the weights for Example 2 are of the form $\omega_{\text{MPRTE}}(u; \{F_\alpha\}) = f_P(u)h(u)/C$ for some function h and some constant C . The weights integrate to 1, because $C = \int f_P(u)h(u) du = E(h(P))$. Thus the weights for both Example 1 and Example 2 are of the form (4.6). This expression plays an important role when we consider identification and estimation of the MPRTE.

Special cases include $q_\alpha(t) = t + \alpha$ and $q_\alpha(t) = (1 + \alpha)t$ for policy changes that act like constant shifts or proportional shifts either in P or in a component of Z . For Example 1, if $q_\alpha(t) = t + \alpha$, then $\omega_{\text{MPRTE}}(u; \{F_\alpha\}) = f_P(u)$, so that the MPRTE weights MTE according to the density of $P(Z)$. In contrast, if $q_\alpha(t) = (1 + \alpha)t$, then $\omega_{\text{MPRTE}}(u; \{F_\alpha\}) = uf_P(u)/E(P)$. Thus, for example, the limit form associated with increasing the probability of college attendance by a fixed amount and the limit form associated with increasing the probability of college attendance by a proportional amount produce different weights on MTE. The limit of the latter puts higher weight on MTE for higher u evaluation points, that is, puts higher weight on MTE for individuals whose unobservables make them less likely to go to college. For Example 2, setting $q_\alpha(t) = t + \alpha$ results in

$$\omega_{\text{MPRTE}}(u; \{F_\alpha\}) = \frac{f_{Z\gamma}(F_V^{-1}(u))}{E(f_V(Z\gamma))} = \frac{f_P(u)f_V(F_V^{-1}(u))}{E(f_V(Z\gamma))},$$

which again will weight the MTE and thus weight people with different unobserved preferences for treatment differently from the way MTE is weighted in the other examples.

5. IDENTIFICATION

Equations (3.3)–(3.4) and (4.1)–(4.2) show that both the PRTE and the MPRTE treatment parameters can be expressed in the form

$$(5.1) \quad \text{treatment parameter}(j) = \int_0^1 \text{MTE}(u)\omega_j(u) du,$$

where $\omega_j(u)$ is the weighting function corresponding to treatment parameter j and where ω_j will depend on which policy change/marginal policy change is being considered. Heckman and Vytlačil (1999, 2005) show that, under our assumptions, the standard treatment parameters can also be expressed in this form and that the form can be used for identification. In particular, the weights are often easy to identify, in which case equation (5.1) implies identification of the parameter if we can identify the MTE at any evaluation point for which the weights are nonzero. Under our assumptions, Heckman and Vytlačil (1999,

2005) show that the MTE can be identified by the method of local instrumental variables (LIV)⁶ at any p in the support of $P(Z)$:

$$(5.2) \quad \frac{\partial}{\partial p} E(Y|P(Z) = p) = \text{MTE}(p).$$

Any parameter that can be represented as a weighted average of MTE can be identified by first estimating MTE over the appropriate support and then integrating the identified MTE function using the appropriate weights. To identify parameter j , not only does $P(Z)$ have to be a continuous random variable, but, in addition, the support of $P(Z)$ must contain all values of u such that $\omega_j(u) \neq 0$. For the standard treatment parameters, identification requires strong conditions on the support of the distribution of $P(Z)$. For example, for the average treatment effect, $E(Y_1 - Y_0)$, the weights are given by $\omega_{\text{ATE}}(u) = 1$ for $u \in [0, 1]$ so that the required support condition is that the support of $P(Z)$ is the full unit interval. It is possible to point-identify the parameters under only slightly weaker conditions than those required by this strategy, as shown in Heckman and Vytlacil (2001a, 2007).

Consider identification of PRTE. Suppose that we can identify F_{P^*} and thus the weights. To identify the parameter using the strategy just discussed, the support of $P(Z)$ must contain all values of u such that $F_P(u) - F_{P^*}(u) \neq 0$. Thus, if the support of P^* is not contained in the support of P , it is not possible to identify the PRTE.

For example, suppose that the largest estimated probability of attending college is strictly less than 1. For analyzing a tuition subsidy policy, it is possible that the largest probability of attending college under a tuition subsidy will be greater than the largest probability of attending college without a tuition subsidy, so the support condition for identifying the corresponding PRTE parameter is violated. More generally, in Examples 1 and 2 in Section 4, the PRTE parameters will not be identified unless the support of P is the full unit interval. Thus, the very strong support conditions required for identification of the standard treatment parameters are also required for identification of the PRTE parameter.

In contrast, the MPRTE parameter can generally be identified under weaker assumptions. The MPRTE weights are nonzero only if the density of $P(Z)$ is nonzero. Consider the classes of MPRTE parameters produced from Examples 1 and 2. For policy counterfactuals that act like transformations of $P(Z)$ as in Example 1, or act like transformations of Z as in Example 2, we have from

⁶LIV can be interpreted as the limit form of the Imbens and Angrist (1994) local average treatment effect (LATE) parameter (see Heckman and Vytlacil (1999)). The ideas of the marginal treatment effect and the limit form of LATE were first introduced in the context of a parametric normal generalized Roy model by Björklund and Moffitt (1987), and were analyzed more generally in Heckman (1997). Angrist, Graddy, and Imbens (2000) also defined and developed a limit form of LATE.

equation (4.6) that $\omega_{\text{MPRTE}}(u; \{F_\alpha\}) \neq 0 \Rightarrow f_P(u) \neq 0$, so that the MPRTE parameters of Examples 1 and 2 can be identified without any additional support conditions even though the corresponding PRTE parameters cannot be identified for any α . In these examples, the marginal PRTE parameters place positive weight on $\text{MTE}(u)$ for values of u where the density of P is positive, that is, they only place positive weight on $\text{MTE}(u)$ for values of u where we can use local instrumental variables to point-identify $\text{MTE}(u)$. Thus, even though in each example the PRTE is not identified for any value of $\alpha \neq 0$ without large support assumptions, the marginal PRTE is identified using the assumption that P is a continuous random variable. The MPRTE parameter is thus fundamentally easier to identify than either the conventional treatment parameters or the PRTE parameters.⁷

6. ISSUES IN ESTIMATION

In addition to the support requirement, an additional difficulty in estimating the standard treatment parameters is that under our assumptions they are not \sqrt{N} -estimable. Suppose that P is known. Using equations (5.1) and (5.2), we obtain

$$(6.1) \quad \begin{aligned} \text{treatment parameter}(j) &= \int \frac{\partial}{\partial p} E(Y|P(Z) = p) \omega_j(p) dp \\ &= E(g'(P)q_j(P)), \end{aligned}$$

where $g'(p) = \frac{\partial}{\partial p} E(Y|P(Z) = p)$ and $q_j(p) = \omega_j(p)/f_P(p)$, and we assume that $w_j(p) \neq 0 \Rightarrow f_P(p) \neq 0$ as is required for identification. A critical requirement for weighted average derivative estimators to be \sqrt{N} -consistent is that $q_j(p)f_P(p)$ vanish on the boundary of the support of P (see Newey and Stoker (1993)). Requiring that $q_j(p)f_P(p) = 0$ on the boundaries of the support of $P(Z)$ is equivalent to requiring that $\omega_j(p) = 0$ at the boundaries of the support of P . This requirement is not satisfied by conventional treatment parameters. For example, the average treatment effect imposes the requirement that $\omega_{\text{ATE}}(p) = 1 \forall p \in [0, 1]$, so that $\omega_{\text{ATE}}(p) \neq 0$ at the boundaries of the support and $E(g'(P)q_{\text{ATE}}(P))$ is not \sqrt{N} -estimable under our assumptions. A parallel analysis holds for the PRTE parameters. From equation (3.4), $\omega_{\text{PRTE}}(p; G) = 0$ at the boundaries of the support of P only if $F_P(t) - G(t) = 0$ at the boundaries of the support of P . This requires equality of the supremum of the supports of the two distributions, so that a necessary condition for \sqrt{N} -consistent estimation of the PRTE is that the policy counterfactual does not increase or decrease the largest value of the probability of participation.

⁷As is clear from our analysis, Assumption A-4 can be relaxed to only require that $E(|Y_1||V = v)$ and $E(|Y_0||V = v)$ are bounded for v contained in any arbitrarily small enlargement of the support of $F_V^{-1}(P(Z))$.

In contrast, the MP RTE parameters are \sqrt{N} -estimable under standard regularity conditions. Consider examples of the form of Example 1—policy changes that act like transformations of the P —or examples of the form of Example 2—policy changes that act like transformations of Z . From equation (4.6), $\omega_{\text{MP RTE}}(\cdot; \{F_\alpha\}) \neq 0$ only if $f_P(\cdot) \neq 0$, so we again have that the weights will equal zero on the boundaries of the support of $P(Z)$ if f_P goes to zero on the boundary of the support. In each of these cases, the marginal PRTE parameter is given by a weighted average derivative $E(g'(P)q(P))$ with weights such that $q_j(p)f_P(p)$ equals zero on the boundary of the support of P if $f_P(p) = 0$ on the boundary of the support of P . Thus, each of these parameters is \sqrt{N} -estimable under appropriate regularity conditions (see Newey and Stoker (1993)).

Our discussion thus far ignores three issues. First, in general, $P(Z)$ is not known but must be estimated. Second, the weights sometimes need to be estimated. Third, we have not analyzed how to deal with observed regressors X that enter the outcome equations for Y_0 and Y_1 . If X contains only discrete elements, then the estimation theory just presented is still valid conditional on X . If X contains continuous elements, then we may instead work with marginal PRTE parameters averaged over the distribution of X to obtain \sqrt{N} estimates. One example of an estimator of the weighted average derivative that applies to our problem is the sieve minimum distance estimator of Ai and Chen (2007). We next consider identification and estimation of average marginal treatment effects.

7. THE AVERAGE MARGINAL TREATMENT EFFECT

We now relate the MP RTE parameter to the AMTE, the average effect of treatment for the marginal person who is indifferent between participation and nonparticipation. More precisely, for a given choice of how to measure the distance between P and U , and thus for a given choice of how to measure how close an individual is to being indifferent between treatment or not, we define the AMTE as the average effect of treatment for those who are arbitrarily close to being indifferent between treatment or not. While the MP RTE depends on the direction of the marginal policy change, the AMTE parameter depends on the metric by which one measures how close individuals are to being indifferent. We show an equivalence between the MP RTE and AMTE parameters. Choosing a particular distance measure for the AMTE is equivalent to examining a particular policy direction for the MP RTE. The effect of a marginal policy change in a particular direction is equal to the average effect of treatment for those at the margin of indifference in the precise sense that a marginal policy change in that direction would change their treatment participation decision.

Consider the average effect of treatment for those who are close to being indifferent between treatment or not. For any metric $m(\cdot, \cdot)$, we have

$$E(Y_1 - Y_0 | m(P, U) \leq \varepsilon) = \int_0^1 \text{MTE}(u) \frac{\Pr[m(P, u) \leq \varepsilon]}{\Pr[m(P, U) \leq \varepsilon]} du.$$

Suppose the metric $m(\cdot, \cdot)$ is such that, for some strictly monotonic and differentiable function q , $m(P, U) = |q(P) - q(U)|$. Let r denote the inverse of q . Then, under the regularity conditions that allow us to interchange limits and integration, the average effect for those arbitrarily close to indifference between treatment or not is

$$\lim_{\varepsilon \rightarrow 0} E(Y_1 - Y_0 | m(P, U) \leq \varepsilon) = \int_0^1 \text{MTE}(u) \omega_{\text{AMTE}}(u) du,$$

where

$$(7.1) \quad \omega_{\text{AMTE}}(u) = \frac{f_P(u)r'(u)}{\int f_P(u)r'(u) du} = \frac{f_P(u)r'(u)}{E(r'(P))}.$$

The form of the weights depends on the choice of a metric. Choosing different ways to measure the distance between P and U , and thus different ways to measure how close individuals are to being indifferent between treatment or not, produces different weighting functions. For any choice of q , the weights will be positive, will integrate to unity, and will only be nonzero where the density of P is nonzero, but otherwise the weights can be shifted around arbitrarily by picking alternative metrics. One might think that there is a natural way to pick a particular metric. Taking $m(P, U) = |P - U|$ seems more natural than taking $m(P, U) = |q(P) - q(U)|$ for general $q(\cdot)$ including $F_V^{-1}(\cdot)$. However, under Assumption B-1, it is not any more natural to pick $m(P, U) = |P - U|$ than it is to pick $m(P, U) = |F_V^{-1}(P) - F_V^{-1}(U)| = |Z\gamma - V|$. Yet, the two choices of metrics give different limit results.

Ambiguity over which metric to use for the AMTE can be resolved by connecting the AMTE to the marginal PRTE. The connection is natural. Marginal policy changes only affect people who are indifferent between treatment or not. Thus there is a close connection between the two types of parameters. An arbitrary choice of a metric to define AMTE lacks an economic motivation. The choice of which marginal policy to study to define the AMTE is well motivated. Comparing equations (4.6) and (7.1), we see a duality between MP RTE parameters and alternative definitions of the AMTE. For example, note that AMTE taking q to be the identity function is exactly the MP RTE expression for policy alternatives of the form $P_\alpha = P + \alpha$. If B-1 holds, then AMTE taking $q(\cdot) = F_V^{-1}(\cdot)$ is exactly the MP RTE expression for policy alternatives of the form of changing one component of Z by α . Different policy alternatives define different AMTE parameters.

Consider an analysis of college attendance. Each marginal PRTE parameter defines a different set of “marginal people” who are indifferent between college participation or not and who would change their college participation in response to a marginal change in the policy. Notice that any inframarginal individual with $\mu(Z) > V$ will not be affected by any marginal policy change. Only individuals with $\mu(Z) = V$ will be affected by a marginal change in the policy. Each marginal PRTE parameter will correspond to a different valid definition of the average marginal treatment effect. The average effect of college attendance on those who are on the margin defined by an infinitesimal level shift in college tuition is different from the average effect of college attendance on those who are on the margin defined by an infinitesimal proportional shift in college tuition.

Since each average marginal treatment effect can be equated with a marginal PRTE parameter, we can use the analysis of Section 4 to define the AMTE so that it can be identified without large support assumptions. In addition, the parameter will be \sqrt{N} -estimable under appropriate regularity conditions. Thus, under our assumptions, the AMTE parameter is fundamentally easier to identify and to estimate than are the conventional treatment parameters.

8. APPLYING THE ANALYSIS TO DATA

Following Carneiro, Heckman, and Vytlacil (2009), we estimate the MTE for a sample of white males from the National Longitudinal Survey of Youth (NLSY).⁸ Using this sample, we also estimate the weights needed to construct the MP RTE parameters that correspond to policy perturbations in alternative directions. We group individuals in two groups: persons with a high school education or below who do not go to college ($D = 0$) and persons with some college or above ($D = 1$). The outcome variable is the log of the average of nonmissing values of the hourly wage between 1989 and 1993, which we interpret as an estimate of the log hourly wage in 1991. We estimate the selection probability $P(Z)$ using a logit. We assume that $Y_1 = X\beta_1 + U_1$, $Y_0 = X\beta_0 + U_0$, and that U_1 and U_0 are independent of X , so that the marginal treatment effect is given by $MTE(x, u) = x(\beta_1 - \beta_0) + E(U_1 - U_0|U = u)$.⁹

⁸See the Supplemental Material Appendix.

⁹The components of Z are the Armed Forces Qualifying Test (AFQT) and its square, mother’s years of schooling and its square, permanent local earnings (average earnings between 1973 and 2000) in the county of residence at age 17 and its square, permanent local unemployment (average unemployment between 1973 and 2000) in state of residence at age 17 and its square, the presence of a four year college in the county of residence at age 17, average tuition in public four year colleges in the county of residence at age 17, log average wage in the county of residence at age 17, and unemployment rate in the state of residence at age 17. The latter four variables are excluded from the set of variables in the wage equations (X) for $D = 1$ and $D = 0$, which are years of experience and its square, log average wage in the county of residence in 1991, unemployment rate in the state of residence in 1991, AFQT and its square, mother’s years of schooling and its

To compute the MTE, we estimate $E(Y|X, P) = X\beta_0 + PX(\beta_1 - \beta_0) + K(P)$ and we take its derivative with respect to P : $\text{MTE}(x, p) = \frac{\partial E(Y|X=x, P=p)}{\partial p}$. We estimate β_1 and β_0 by applying the partially linear regression method of [Robinson \(1988\)](#). $K(\cdot)$ is estimated using locally quadratic regression.¹⁰ We trim 2% of the observations from the extremes of the distribution of P , which means that we are only able to identify the MTE at evaluation points between 0.1 and 0.93. Thus, we can only identify parameters over this support. We present annualized returns, obtained by dividing the MTE (and the parameters it generates) by 4 (corresponding to 4 years of college).

Figure 1 plots three alternative MP RTE weights. We take two cases (which we label A and C) based on Example 1, where $q_\alpha^A(P) = P + \alpha$ and $q_\alpha^C(P) = (1 + \alpha)P$. A third case (which we label B) is taken from Example 2, where $q_\alpha^B(Z^{[k]}) = Z^{[k]} + \alpha$.¹¹ $\omega_{\text{MP RTE}}^C$ puts more weight than $\omega_{\text{MP RTE}}^A$ on higher values of U , and both sets of weights are zero in the region of U for which we cannot identify the MTE. The intuitive reason for this pattern is that a proportional change in P results in larger absolute changes for high levels of P than for low levels of P , while an additive change in P results in uniform changes in P across the whole distribution. The MP RTE for a marginal proportional change in P is estimated to be 0.1296, while the MP RTE for a marginal additive change in P is estimated to be 0.0880, since the former puts more weight on high values of U where the MTE is estimated to be smaller. $\omega_{\text{MP RTE}}^B$ puts more weight than either $\omega_{\text{MP RTE}}^A$ or $\omega_{\text{MP RTE}}^C$ in the center of the support of U , and the MP RTE is estimated to be 0.1274. Figure 1 also shows that corresponding to each MP RTE weight there is an AMTE weight, defined by different choices of metric for measuring the distance between P and U . For case A , $m(P, U) = |P - U|$. For case B , $m(P, U) = |F_V^{-1}(P) - F_V^{-1}(U)| = |\mathcal{Z}\gamma - V|$. For case C , $m(P, U) = |\frac{P}{U} - 1|$.

square, permanent local earnings (average earnings between 1973 and 2000) in the county of residence at age 17 and its square, and permanent local unemployment (average unemployment between 1973 and 2000) in state of residence at age 17 and its square.

¹⁰We first run kernel regressions of each X on \hat{P} using a bandwidth of 0.05 and trimming 2% of the observations, and we compute the resulting residuals. We then estimate a linear regression of the outcome variable on these residuals and obtain estimates $\hat{\beta}_1$ and $\hat{\beta}_0$. To estimate $K(\cdot)$ and its first derivative, we run a locally quadratic regression of $Y - X\hat{\beta}_0 - \hat{P}X(\hat{\beta}_1 - \hat{\beta}_0)$ on \hat{P} , using a bandwidth of 0.322 (determined by cross validation) and trimming 2% of the observations. $\hat{K}'(\cdot)$ is constructed by taking the coefficient on the linear term of the locally quadratic regression.

¹¹In the empirical work, we impose an index sufficiency restriction when estimating $f_{P|X}: f_{P|X} = f_{P|X\delta}$. To estimate the scalar index $X\delta$, we use the fact that $f_{P|X} = f_{P|X\delta}$ implies that $E(P|X) = E(P|X\delta)$ to estimate δ using semiparametric least squares ([Ichimura \(1993\)](#)). We estimate $f_{P|X\delta}$ by running a locally linear regression of $(\frac{1}{h}) * K(\frac{\hat{P}-P}{h})$ on $X\hat{\delta}$, where $K(\cdot)$ is a standard normal density and $h = 1.06 * [\widehat{\text{Var}}(P)]^{1/2} * n^{(-1/5)}$. We use a bandwidth equal to $1.06 * [\widehat{\text{Var}}(X\delta)]^{1/2} * n^{(-1/5)}$. The figure fixes $X\delta$ at $\bar{X}\delta$. The weights are rescaled so that both the MTE and the weights fit in the figure.




FIGURE 1.—Alternative definitions of the MPRT and the AMTE. This figure plots the marginal treatment effect ($MTE = E(Y_1 - Y_0 | X = x, U = u)$) and the marginal policy relevant treatment effect (MPRT) weights for three types of policy shifts: $P_\alpha^A = P + \alpha$, $P_\alpha^B(Z) = F_V(Z\gamma + \alpha\gamma^{(k)})$, and $P_\alpha^C = (1 + \alpha)P$, where P is the probability of receiving treatment conditional on the observed covariates (Z). The equivalent definitions of the AMTE correspond, respectively, to the following three metrics for measuring the distance between P and U ($m(P, U)$): $|P - U|$, $|Z\gamma - V|$, and $|\frac{P}{P_0} - 1|$. The MTE, the MPRT, and the AMTE weights are evaluated at values of X such that $E(Y_1 - Y_0 | X = x) = 0.13$. The average marginal policy effects are 0.1296 for the additive shift in P (A), 0.1274 for the additive shift in Z (B), and 0.0880 for the proportional shift in P (C). Estimated from NLSY data (see Carneiro, Heckman, and Vytlačil (2009), for details on estimating MTE).

9. SUMMARY AND CONCLUSIONS

This paper extends the analysis of Heckman and Vytlačil (1999, 2005, 2007) by using the MTE to identify the effect of a marginal policy change, and to identify the average effect of treatment on individuals who are indifferent between treatment or not. Conventional treatment parameters and the PRTE require large support conditions for identification and often are not \sqrt{N} -estimable. Under our assumptions, the marginal policy effect parameters and the average marginal treatment effects are generally identified without large support conditions and are \sqrt{N} -estimable. An analysis of the effect of

marginal changes in tuition policies on college attendance illustrates the empirical relevance of this analysis.

REFERENCES

- AI, C., AND X. CHEN (2007): "Estimation of Possibly Misspecified Semiparametric Conditional Moment Restriction Models With Different Conditioning Variables," *Journal of Econometrics*, 141, 5–43. [388]
- ANGRIST, J., K. GRADY, AND G. IMBENS (2000): "The Interpretation of Instrumental Variables Estimators in Simultaneous Equations Models With an Application to the Demand for Fish," *Review of Economic Studies*, 67, 499–527. [386]
- BJÖRKLUND, A., AND R. MOFFITT (1987): "The Estimation of Wage Gains and Welfare Gains in Self-Selection," *Review of Economics and Statistics*, 69, 42–49. [386]
- CARD, D. (2001): "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems," *Econometrica*, 69, 1127–1160. [380]
- CARNEIRO, P., J. J. HECKMAN, AND E. J. VYTLACIL (2009): "Estimating Marginal Returns to Education," Unpublished Manuscript, University College London, Department of Economics. [390,392]
- HECKMAN, J. J. (1997): "Instrumental Variables: A Study of Implicit Behavioral Assumptions Used in Making Program Evaluations," *Journal of Human Resources*, 32, 441–462; addendum published (1998), *Journal of Human Resources*, 33. [386]
- 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, 4730–4734. [385,386,392]
- (2001a): "Instrumental Variables, Selection Models, and Tight Bounds on the Average Treatment Effect," in *Econometric Evaluation of Labour Market Policies*, ed. by M. Lechner and F. Pfeiffer. New York: Center for European Economic Research, 1–15. [386]
- (2001b): "Policy-Relevant Treatment Effects," *American Economic Review*, 91, 107–111. [377,379–381]
- (2005): "Structural Equations, Treatment Effects and Econometric Policy Evaluation," *Econometrica*, 73, 669–738. [377–380,385,386,392]
- (2007): "Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Economic Estimators to Evaluate Social Programs and to Forecast Their Effects in New Environments," in *Handbook of Econometrics*, Vol. 6B, ed. by J. Heckman and E. Leamer. Amsterdam: Elsevier, 4875–5144. [379,380,386,392]
- HURWICZ, L. (1962): "On the Structural Form of Interdependent Systems," in *Logic, Methodology and Philosophy of Sciences*, ed. by E. Nagel, P. Suppes, and A. Tarski. Stanford, CA: Stanford University Press, 232–239. [382]
- ICHIMURA, H. (1993): "Semiparametric Least Squares (SLS) and Weighted SLS Estimation of Single-Index Models," *Journal of Econometrics*, 58, 71–120. [391]
- ICHIMURA, H., AND C. TABER (2002): "Direct Estimation of Policy Impacts," Unpublished Working Paper, University College London, Department of Economics. [382]
- IMBENS, G. W., AND J. D. ANGRIST (1994): "Identification and Estimation of Local Average Treatment Effects," *Econometrica*, 62, 467–475. [379,386]
- NEWBY, W., AND T. M. STOKER (1993): "Efficiency of Weighted Average Derivative Estimators and Index Models," *Econometrica*, 61, 1199–1223. [387,388]
- ROBINSON, P. M. (1988): "Root- N -Consistent Semiparametric Regression," *Econometrica*, 56, 931–954. [391]
- STOCK, J. H. (1989): "Nonparametric Policy Analysis," *Journal of the American Statistical Association*, 84, 565–575. [380]

VYTLACIL, E. J. (2002): "Independence, Monotonicity, and Latent Index Models: An Equivalence Result," *Econometrica*, 70, 331–341. [379]

Dept. of Economics, University College London, Gower Street, London WC1E 6BT, U.K. and IFS and CeMMAP; p.carneiro@ucl.ac.uk,

Dept. of Economics, University of Chicago, 1126 E. 59th Street, Chicago, IL 60637, U.S.A. and Geary Institute, University College Dublin; Cowles Foundation, Yale University; American Bar Foundation; jjh@uchicago.edu,

and

Dept. of Economics, Yale University, P.O. Box 208281, New Haven, CT 06520-8281, U.S.A.; edward.vytlacil@yale.edu.

Manuscript received April, 2007; final revision received July, 2009.