

Removing the Veil of Ignorance in assessing the distributional impacts of social policies

Pedro Carneiro, Karsten T. Hansen
and James J. Heckman*

Summary

■ This paper summarizes our recent research on evaluating the distributional consequences of social programs. This research advances the economic policy evaluation literature beyond estimating assorted mean impacts to estimate distributions of outcomes generated by different policies and to determine how those policies shift persons across the distributions of potential outcomes produced by them. Our approach enables analysts to evaluate the distributional effects of social programs without invoking the “Veil of Ignorance” assumption often used in the literature in applied welfare economics. Our methods determine which persons are affected by a given policy, where they come from in the ex-ante outcome distribution and what their gains are. We apply our methods to analyze two proposed policy reforms in American education. These reforms benefit the middle class and not the poor.■

JEL Classification: D63, D33, I20.

Keywords: Veil of Ignorance, income inequality, counterfactuals, distributions of counterfactuals.

* Pedro Carneiro is a graduate student at the University of Chicago. Karsten Hansen is a research fellow at the University of Chicago. James Heckman is a professor at the University of Chicago and a senior fellow at the American Bar Foundation.

Removing the Veil of Ignorance in assessing the distributional impacts of social policies

Pedro Carneiro, Karsten T. Hansen
and James J. Heckman^{*}

Evaluating public policy is a central task of economics. Welfare economics presents different criteria. Research on program evaluation develops and applies a variety of different econometric estimators. Traditional empirical methods focus on mean impacts. Yet modern welfare economics emphasizes the importance of accounting for the impact of public policy on *distributions* of outcomes (Sen, 1997, 2000). A large body of empirical evidence indicates that people differ in their responses to the same policy and act on those differences, and that the representative agent paradigm is a poor approximation to reality because the marginal entrant into a social program is often different from the average participant. (Heckman, 2001a). This evidence highlights the importance of going beyond the representative agent framework when evaluating public policies.

This paper summarizes our recent research on evaluating the distributional consequences of public policy.¹ Our research advances the economic policy evaluation literature beyond estimating assorted mean impacts to estimate the distributions of outcomes generated by different policies and to determine how those policies shift persons across the distributions of potential outcomes produced by them. We distinguish the average participant in a program from the marginal

** This paper is a summary of some of the main ideas contained in our source papers "Educational Attainment and Labor Market Outcomes: Estimating Distributions of The Return to Educational Interventions" (2001) and "Estimating Distributions of Treatment Effects with an Application to the Return to Schooling" (2002). This research was supported by the following grants: NSF 97-09-873, NSF SES-0099195, and NICHD-40-4043-000-85-261. Heckman's work was also supported by the American Bar Foundation and the Donner Foundation. Carneiro's work was also supported by Fundação Ciência e Tecnologia and Fundação Calouste Gulbenkian. We thank Salvador Navarro-Lozano for assistance. The first version of this paper was presented at the Midwest Econometrics Group Conference, October, 2000.*

¹ Carneiro, Hansen and Heckman (2001, 2002).

entrant. Our research advances the existing literature on evaluating the distributional consequences of alternative policies beyond the “Veil of Ignorance” assumption used in modern welfare economics (See Atkinson, 1970; Sen 1997, 2000). Approaches based on that assumption compare two social states by assuming that the position of any particular individual in one distribution should be treated as irrelevant. In this approach the overall distribution of outcomes is all that matters. This is a consequence of the *anonymity postulate* that is fundamental to that literature. Anonymity is the property that only the distribution of outcomes matters and that reversing the positions of any two persons in the overall distribution does not affect the evaluation placed on the policy (or state of affairs) that produces the distribution.

There are normative arguments that support this criterion. (See Harsanyi, 1955; Vickery, 1960; and Roemer, 1996). As a positive description of actual social choice processes, the “Veil of Ignorance” seems implausible. Participants in the political process are likely to forecast their outcomes under alternative economic policies, and assess policies in this light (Heckman, 2001b). This paper extends current practice by developing and applying methods that forecast how people fare under different policies. We link the literature in modern welfare economics to the treatment effect literature.

This paper proceeds as follows. We briefly present the evaluation problem for an economy with two sectors (e.g. schooled and un-schooled) where agents select or are selected into “treatment” (one of the two sectors). We consider policies that affect choices of treatment (e.g. schooling) but not potential outcomes (the outcomes they experience under different treatments). We compare outcomes across two policy regimes that affect treatment choices. This task is much easier when individuals respond in the same way to treatment than when they differ in their response to treatment, and act on those differences in making treatment choice decisions. In the latter case, the marginal entrant into schooling is not the same as the average participant in treatment and the representative agent paradigm breaks down. In the Appendix, we show how to generate the counterfactual distributions of outcomes produced by alternative policies.

We apply our analysis to estimate the distributional consequences of two proposed policy reforms in American education. Even though the two policies barely affect the overall distribution of outcomes, and so would be judged to be equivalent to the pre-policy origin state un-

der the Veil of Ignorance criterion, they have substantial effects on a small group of people concentrated in the middle to the high end of the pre-policy wage distribution. Marginal entrants attracted into college get smaller gains than average college students suggesting diminishing returns to programs that encourage college enrollment. Marginal entrants into junior college are about the same as average entrants, suggesting constant returns for that schooling level. Since most of the people affected by the policies come from the middle to the high end of the original wage distribution, there is little impact of these policies on the poor.

1. The evaluation problem for means and distributions

In order to place our work in the context of the current literature on social program evaluation, and to link it to the economics of education, it is helpful to consider a simple generalized Roy (1951) economy with two sectors. Let $S=1$ denote college and $S=0$ be high school. Persons (or their agents, such as their parents) can choose to be in either sector. There are two potential outcomes for each person (Y_0, Y_1), only one of which is observed, since it is assumed that only one option can be chosen. For simplicity, we assume that the decision rule governing sectoral choices is

$$S = \begin{cases} 1 & \text{if } I = Y_1 - Y_0 - C \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

Here C is the cost of choosing $S=1$. In the context of a schooling model, C is tuition or monetized psychic cost, while $Y_1 - Y_0$ is the net gain from schooling expressed, say, in present value terms.

We decompose Y_1 and Y_0 in terms of their means μ_1 and μ_0 and mean zero idiosyncratic deviations (U_1, U_0) or residuals:

$$\begin{aligned} Y_1 &= \mu_1 + U_1 \\ Y_0 &= \mu_0 + U_0. \end{aligned}$$

We condition on X variables, but for notational simplicity we keep this dependence implicit. Decomposing C in a similar fashion, we may write:

$$C = \mu_C + U_C,$$

so that

$$I = \mu_1 - \mu_0 - \mu_C + (U_1 - U_0 - U_C).$$

It is fruitful to distinguish two kinds of policies: (a) those that affect potential outcomes (Y_0, Y_1) through price and quality effects and (b) those that affect sectoral choices (through C) but do not affect potential outcomes. Tuition and access policies that do not have general equilibrium effects fall into the second category of policy. Policies with general equilibrium effects and policies that directly affect rewards to potential outcomes and quality are examples of the first kind of policy. It is the second kind of policy that receives the most attention in empirical work on estimating economic returns to schooling (see e.g., the survey by Card, 1999) or in evaluating schooling policies (see e.g., Kane, 1994).

Consider two policy environments denoted A and B . These produce two social states for outcomes that we wish to compare. In the general case, we may distinguish an economy operating under policy A with associated cost and outcome vector (Y_0^A, Y_1^A, C^A) for each person, from an economy operating under policy B with associated cost and outcome vector (Y_0^B, Y_1^B, C^B) . Policy interventions with no effect on potential outcomes can be described as producing two choice sets (Y_0, Y_1, C^A) and (Y_0, Y_1, C^B) for each person. In this paper we focus on evaluating the second kind of policy that keeps invariant the distribution of potential outcomes across policy states but affects the cost of choosing sector 1 within each state.

Our framework differs in its emphasis from the standard model of modern welfare economics. Analysts writing in that tradition focus on the distribution of outcomes produced by each policy without inquiring how those outcomes are produced. All policies that produce the same aggregate outcome distributions are judged to be equally good. The details of how the observed distribution is produced are deemed irrelevant. The distinctions we make between policies that affect potential outcomes and policies that affect which potential outcomes are selected are also ignored in that literature. There is no explicit discus-

sion of sectoral choice within policy states. The literature starts, and stops, with an analysis of distributions of the observed outcomes for each person in each policy state (Y^A, Y^B) defined as

$$Y^A = S_A Y_1^A + (1 - S_A) Y_0^A, \text{ and } Y^B = S_B Y_1^B + (1 - S_B) Y_0^B,$$

where S_A and S_B are schooling choice indicators under policies A and B respectively, without inquiring more deeply into the sources of the differences in the distributions of outcomes.

The modern treatment effect literature focuses on these details and distinguishes choice of treatments from the treatment outcomes. However, it only inquires about certain mean treatment effects. The operating assumption in the literature is that policies do not affect potential outcomes (so $(Y_0^A, Y_1^A) = (Y_0^B, Y_1^B)$), but do affect choices of sectors.

This literature distinguishes three cases. Case I arises when everyone (with the same X) gets the same effect from treatment ($Y_1 - Y_0$ is the same for everyone). Case II occurs when $Y_1 - Y_0$ differs among people of the same X but decisions to enroll in the program are not affected by these differences:

$$\Pr(S = 1 | Y_1 - Y_0) = \Pr(S = 1).$$

Case III occurs when $Y_1 - Y_0$ differs among people and people act on these differences. In cases I and II, the marginal entrant into a program is the same as the average entrant. In case III, this is not so. People select in part on gains. If they select solely on gains, then the marginal entrant gets a lower return than those participants (in $S=1$) who are inframarginal; that is, the marginal treatment effect (MTE)

$$E(Y_1 - Y_0 | I = 0) < E(Y_1 - Y_0 | S = 1).$$

See Heckman (2001a) for more detailed discussion of the various treatment effects.²

² Björklund and Moffitt (1987) introduced the marginal treatment effect into the evaluation literature. See Heckman (2001a) for a summary of extensions of this literature.

2. Comparing two policy states

Consider two policies, A and B , that affect sectoral choices without affecting the distributions of potential outcomes. For concreteness, we can think of these as policies that affect C (e.g. tuition or access to school) by shifting its mean, changing its variance or changing the covariance between C and (Y_0, Y_1) . Each policy produces a distribution of outcomes. For concreteness, think of the outcome as wages associated with different schooling levels.

In the literature on evaluating inequality, comparisons of policies are made in terms of comparisons of distributions. If policy B produces an aggregate distribution of wages that second order stochastically dominates that produced from policy A , B is preferred.³ The details of who benefits or loses from the policy are considered to be irrelevant as a consequence of the anonymity postulate.

The literature on evaluating inequality in modern welfare economics compares two aggregate outcome distributions. If policy A has been implemented, but policy B has not, evaluation of B entails construction of a counterfactual aggregate outcome distribution. Under the assumptions used in the treatment effect literature, all that is required is determination of how policy B sorts persons into sectors “0” and “1”, and how such sorting affects *observed* outcome distributions in sectors “0” and “1”. In our example, what is required is a schooling choice equation and a selection model to identify the invariant potential outcome distributions. The selection model enables analysts to go from observed (selected) distributions of Y_0 and Y_1 to the population potential distributions. With sufficient individual variation in C within an economy governed by policy A , it is possible to accurately forecast the effect of policy B on the overall distribution without previously observing it, as we demonstrate in this paper.

Our approach to the evaluation of public policy is more ambitious in some respects than the recent literature in welfare economics and is more in line with the objectives of modern political economy (Persson and Tabellini, 2000). We relax the anonymity postulate and determine how individuals at different positions within the initial overall distribution respond to policies in terms of their treatment choices and gains. We estimate the number of people directly affected by the

³ See Sen (1997).

policy, where they start, and where they end up in the overall distribution.

In the context of the treatment effect framework, this task is broken down into two sub-tasks. The first sub-task is to determine who shifts treatment state in response to the policy and where they are located in the initial overall distribution. The second sub-task is to determine where they end up in the overall distribution after taking the treatment, and how much they gain. Since this approach assumes that potential outcome distributions are not affected by the policies, it is less ambitious, in this respect, than the approach advocated in modern welfare economics which entertains that possibility.

Under case I, this task is greatly simplified. Everyone who shifts from “0” to “1” gets the same gain Δ . The only problem is to find where the switchers are located in the initial overall distribution. Under case II, Δ varies among observationally identical people. The gain is not necessarily the same for persons with different initial Y_0 values. However, *on average*, across all movers, the gain is the same as the mean difference between the two potential outcome distributions *within* policy regime \mathcal{A} . Hence the marginal entrant has the same mean as the average person and the average participant:

$$E(Y_1 - Y_0 | I = 0) = E(Y_1 - Y_0) = E(Y_1 - Y_0 | S = 1).$$

Case III differs from case II in that in general the gains to the average switcher are not the same as the gains to the previous participants. If $(Y_1 - Y_0)$ is positively correlated with $I = Y_1 - Y_0 - C$, the marginal entrant receives lower gains on average than does the average participant. The details of constructing the transition densities for the switchers are presented in our companion paper.

3. Identifying counterfactual distributions under treatment effect assumptions

Identifying the joint distribution of potential outcomes under treatment effect assumptions is more difficult than identifying the various mean treatment effects.⁴ The fundamental problem is that we never

⁴ A large econometric literature identifies the mean impacts under a variety of assumptions. See Heckman, LaLonde and Smith (1999) for one survey. Heckman and

observe both components of (Y_0, Y_1) for anyone.⁵ Thus we cannot directly form the joint distribution of potential outcomes (Y_0, Y_1) .

In the Appendix, and in our source papers, we review various approaches to estimating, or bounding counterfactual distributions that have appeared in the literature. In our source papers, we develop a new method for identifying these distributions. It is based on an idea common in factor analysis but applied to model counterfactual distributions. If potential outcomes are generated by a low dimensional set of factors, then it is possible to estimate the distributions of factors and generate distributions of the counterfactuals. Here, low dimensional refers to the number of factors relative to the number of measured outcomes. See the Appendix for the intuitive idea that motivates the analysis in our source papers. We next turn to an application of our analysis to American data.

4. Some evidence from America on two educational reforms

Our companion paper uses data on wages, schooling choices and covariates for white males from the National Longitudinal Study of Youth (NLSY) to estimate a three factor version of the model described in the Appendix using a Bayesian semiparametric mixture of normals econometric framework. We consider four schooling levels: dropout, high school graduate, junior college, and four year college. We use local labor market variables, tuition, and family background information to identify the model. The estimated model fits the data well. Observed wage distributions are closely approximated. There is no need for more than three factors to fit our data which includes panel data measurements on wages as well as indicators of ability and motivation.⁶

Our paper estimates models for a variety of schooling groups. Here, for the sake of brevity, we focus only on some key empirical results. We report the wage returns to college and high school, and selection on levels and gains into those schooling categories. We ana-

Vytlacil (2000, 2001) consider identification of marginal treatment effects and unify the treatment effect literature.

⁵ Panel data estimators sometimes enable analysts to observe both components. See Heckman, LaLonde and Smith (1999).

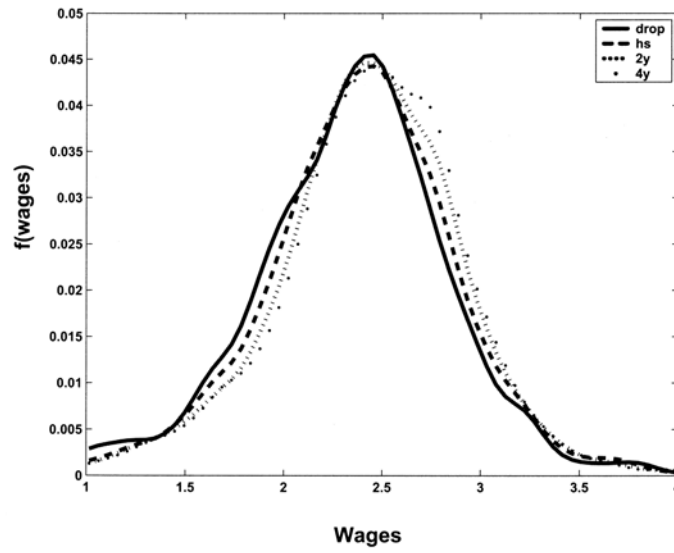
⁶ The factor model is strongly overidentified so that it would have been possible to estimate many more than three factors.

lyze two policies: (a) a full tuition subsidy for junior colleges and (b) a policy promoting access to four year colleges which places an institution in the immediate vicinity (the county of residence) of each American. We consider only partial equilibrium treatment effects and do not consider the full cost of financing the reforms.


Our evidence shows considerable dispersion in terms of levels and returns (gains) to various schooling categories. Indeed, *ex post* returns are negative for a substantial fraction of people. There is little evidence of selection either on levels or gains for high school graduates. There is a lot of evidence of selection on levels and gains for college graduates. The marginal entrants into four year colleges induced by the access policy we consider have wage outcomes below the average college participant both in terms of levels and gains. This is not true for the junior college tuition subsidy policy we also analyze. For that case, there is little impact on the overall quality of junior college graduates.

Figure 1 shows the potential high school wages for all four schooling groups (what people who actually attend various schooling levels would have earned had they gone to high school). The four densities are nearly the same suggesting that there is little evidence of selection on levels into high school. Three of these four densities are counterfactual. The density for high school graduates is factual. For college (Figure 2), there is strong evidence of selection on levels. Persons who attend college do better in college than dropouts, high school graduates, or junior college graduates would do. This result contrasts sharply with the corresponding result for the factual and counterfactual wage densities for high school graduates.

**Figure 1. Distributions of wages, high school graduates
(white males, age 29 from NLSY)**




**Figure 2. Distributions of wages, college
(white males, age 29 from NLSY)**




There is also little evidence of selection on gains ($Y_1 - Y_0$) to high school (high school vs. dropout). See Figure 3 which plots the counterfactual returns to high school for all four schooling groups. The returns (high school vs. four year college) are greater for persons who become college graduates than for the other schooling groups, although there is a lot of overlap in the distributions (see Figure 4). *Ex post* many persons who actually stop their schooling at the high school level would make fine college graduates. Many college graduates experience negative returns. The marginal treatment effect comparing high school to college (Figure 5) suggests that as the unobservables that lead to a higher likelihood of attending college increase, (so $P(S=1)$ increases), the return to college increases. People most likely to attend college have the highest marginal returns. The corresponding figure for the return to high school is flat, suggesting that the marginal participant has the same return as the average participant.


**Figure 3. Distributions of returns to high school
(white males, age 29 from NLSY)**



**Figure 4. Distributions of returns to college vs. high school
(white males, age 29 from NLSY)**



**Figure 5. Marginal treatment effect, high school—college
(NLSY, white males)**



Notes: ^a Average return to persons at margin of attending school given characteristics u_s . ^b Variables related to schooling (higher u_s leads to a higher probability of attending college).

Figure 6. People affected by full subsidy to community college tuition by decile of initial overall wage distribution





Figure 7. People affected by making distance to 4y college = 0 by decile of wage distribution



Using the estimated model, we compare two policies: a full subsidy to community college tuition and a policy that places a four year college in each county in America. Table 1 shows the average log wages of participants before the policy change and their average return. It compares these levels and returns with what the marginal participant

attracted into the indicated schooling level by the policy would earn. Marginal and average log wages and returns are about the same for the community college policy. There is little decline in quality among the entrants. For the access policy, there is a sharp difference. Average participants in four year colleges earn more and have higher returns than marginal entrants. There is a sharp decline in the average quality of college graduates.

Table 1. Average logwages and returns for average and marginal person

	Full tuition subsidy (junior college)	Distance to 4y college = 0 (4y college)
Avg. log wage	2.4769	2.7258
std. error	0.0335	0.0245
95 % interval	(2.4107, 2.5447)	(2.6757, 2.7768)
Avg. log wage at margin	2.3898	2.4575
std. error	0.0253	0.0360
95% interval	(2.3400, 2.4419)	(2.3881, 2.5367)
	(junior college —high school)	(4y college— junior college)
Avg. return	0.0871	0.2722
std. error	0.0399	0.0386
95% interval	(0.0083, 0.1648)	(0.1934, 0.3505)
Avg. return at margin	0.0890	0.2471
std. error	0.0383	0.0500
95 interval	(0.0088, 0.1808)	(0.1587, 0.3632)

Despite the substantial sizes of the policy changes we consider, the effects on participation are small. The four year access policy only raises four year graduation rates by 1.3 percent. The junior college subsidy raises attendance at those institutions by 3.8 percent.

The policies operate unevenly over the deciles of the initial outcome distribution. The impacts are greatest at the center of the distribution for the community college policy. See Table 2 that presents transition probabilities by deciles from the pre-policy wage distribution to the post-policy wage distribution and the companion Figure 6. Mobility is from the top of the initial wage distribution for the four year college policy. See Table 3 which is in a format parallel to Table 2 and Figure 7. Neither policy benefits the poor.

Table 2. Mobility of people affected by full tuition subsidy
Fraction of total population affected by policy: 0.038

Deciles of origin	Fraction by decile of origin	Transition probabilities between deciles of pre-policy wage distribution and post-policy wage distribution (p_{ij}^a)									
		1	2	3	4	5	6	7	8	9	10
1	.088	.518	.209	.105	.058	.031	.028	.023	.010	.014	.003
2	.102	.142	.249	.210	.147	.090	.066	.043	.027	.014	.012
3	.117	.049	.137	.206	.190	.140	.124	.069	.040	.025	.020
4	.119	.024	.063	.135	.180	.168	.164	.126	.073	.042	.026
5	.108	.018	.037	.071	.123	.158	.192	.162	.132	.070	.036
6	.120	.009	.022	.047	.080	.121	.181	.193	.176	.114	.056
7	.106	.008	.013	.029	.048	.074	.138	.186	.214	.189	.101
8	.097	.007	.011	.020	.028	.050	.085	.148	.212	.261	.178
9	.079	.002	.009	.009	.017	.030	.060	.096	.163	.285	.326
10	.064	.002	.011	.014	.021	.024	.044	.062	.098	.178	.547

Note: ^a p_{ij} is the probability of being in decile j of the post-policy wage distribution conditional on being in decile i of the pre-policy wage distribution.

REMOVING THE VEIL OF IGNORANCE, Pedro Carneiro, Karsten T. Hansen and James J. Heckman

Deciles of origin	Average pre-policy log wage	Average gain of moving	Average gain of moving, by initial decile									
			1	2	3	4	5	6	7	8	9	10
1	1.385	.309	-.011	.362	.613	.750	.914	1.107	1.220	1.358	1.599	1.943
2	1.892	.223	-.341	.009	.211	.355	.491	.610	.738	.894	1.077	1.407
3	2.110	.176	-.604	-.195	.005	.157	.290	.405	.536	.691	.869	1.233
4	2.270	.147	-.824	-.347	-.149	.006	.132	.253	.386	.521	.697	1.064
5	2.401	.122	-.911	-.492	-.278	-.125	.001	.126	.255	.392	.571	.955
6	2.528	.093	-1.058	-.631	-.402	-.247	-.120	.005	.134	.271	.438	.816
7	2.659	.070	-1.108	-.766	-.536	-.369	-.244	-.121	.006	.145	.311	.656
8	2.802	.027	-1.284	-.891	-.699	-.510	-.384	-.270	-.134	.006	.183	.489
9	2.981	-.017	-1.563	-1.101	-.870	-.713	-.570	-.436	-.310	-.162	.014	.346
10	3.366	-.236	-1.876	-1.498	-1.371	-1.089	-1.016	-.869	-.762	-.510	-.315	.074

Table 3. Mobility of people affected by changing distance to 4 year college to 0
Fraction of total population affected by policy: 0.013

Deciles of origin	Fraction by decile of origin	Transition probabilities between deciles of pre-policy wage distribution and post-policy wage distribution (p_{ij}^a)									
		1	2	3	4	5	6	7	8	9	10
1	.063	.366	.231	.133	.095	.059	.053	.023	.020	.010	.011
2	.081	.066	.145	.218	.180	.150	.105	.065	.036	.020	.015
3	.092	.029	.066	.128	.155	.170	.177	.117	.086	.048	.024
4	.103	.013	.032	.061	.116	.134	.193	.192	.121	.102	.036
5	.098	.013	.022	.030	.073	.116	.179	.199	.185	.125	.059
6	.117	.009	.011	.022	.046	.082	.131	.187	.212	.200	.101
7	.114	.005	.009	.012	.030	.049	.089	.145	.239	.275	.147
8	.119	.002	.006	.014	.021	.024	.054	.107	.214	.308	.250
9	.113	.001	.005	.003	.012	.018	.040	.063	.133	.305	.419
10	.100	.003	.003	.004	.004	.018	.022	.047	.081	.184	.634

Note: ^a p_{ij} is the probability of being in decile j of the post-policy wage distribution conditional on being in decile i of the pre-policy wage distribution.

REMOVING THE VEIL OF IGNORANCE, Pedro Carneiro, Karsten T. Hansen and James J. Heckman

Deciles of origin	Average pre- policy log wage	Average gain of moving	Average gain of moving, by initial decile									
			1	2	3	4	5	6	7	8	9	10
1	1.405	.494	.153	.424	.581	.762	.888	1.032	1.124	1.247	1.326	1.799
2	1.889	.359	-.391	.027	.219	.379	.491	.633	.752	.901	1.077	1.464
3	2.111	.310	-.614	-.185	-.003	.170	.287	.412	.543	.683	.847	1.254
4	2.269	.289	-1.093	-.367	-.160	.008	.145	.261	.387	.524	.697	1.259
5	2.403	.232	-.993	-.502	-.267	-.126	.003	.127	.263	.395	.571	.918
6	2.530	.210	-1.115	-.633	-.414	-.250	-.123	.009	.139	.277	.445	.825
7	2.661	.166	-1.441	-.751	-.517	-.381	-.248	-.119	.002	.143	.326	.661
8	2.803	.121	-1.193	-.888	-.700	-.520	-.378	-.270	-.133	.009	.190	.505
9	2.985	.061	-1.695	-1.092	-.871	-.678	-.584	-.446	-.303	-.161	.014	.347
10	3.350	-.127	-2.267	-1.351	-1.108	-1.152	-1.023	-.733	-.676	-.513	-.316	.093

Our approach to the evaluation of social policy is much richer, and more informative, than an analysis of aggregate outcomes of the sort contemplated in modern welfare theory. The overall Gini coefficient does not change (to two decimal places) when we implement the two policies. By the standards of the modern welfare economics pre- and post-policy distributions are the same. A focus on the aggregate outcome distribution masks important details which our approach reveals. Only a small group of persons are directly affected by the policy. The vast majority of persons would be unaffected by these policies, and presumably, would be indifferent to the policy.⁷ Our approach to policy evaluation lifts the Veil of Ignorance and provides a more complete interpretation of who benefits from the policy and where beneficiaries come from in the overall distribution of prepolicy income.

5. Summary and conclusions

This paper summarizes our recent research on evaluating the distributional consequences of social programs. We move beyond the mean treatment effects that dominate discussion in the recent applied evaluation literature to analyze the impacts of policy on distributions of outcomes. We develop and apply methods for determining which persons are affected by the policy, where they come from in the initial distribution, and what their gains are.

We contrast the outcomes of participants in schooling before the policy change with the outcomes of marginal entrants induced into the treatment state by the policy. We compare our approach to the approach advocated in modern welfare economics. That approach focuses attention solely on the aggregate distribution and does not identify gainers and losers from a policy. Our approach identifies where gainers and losers are located in the overall distribution. The output produced from our approach generates the information required in modern positive political economy.

Our analysis has been conducted for a partial equilibrium treatment effect model that assumes that policies do not affect the distribution of potential outcomes, only the choice probabilities of particular treatments. It would be desirable to extend our framework to analyze the effects of more general policies that affect both outcome dis-

⁷ Counting their tax burden, they might even be hostile to these policies.

tributions and choices using the general equilibrium framework described in Heckman (2001b). We leave that task for another occasion.

References

- Aakvik, A., Heckman, J. and Vytlačil, E. (1999), Semiparametric program evaluation: Lessons from an evaluation of a Norwegian training program, Mimeo, University of Chicago, October 15.
- Atkinson, A. (1970), On the measurement of inequality, *Journal of Economic Theory* 2, 244-263.
- Björklund, A. and Moffitt, R. (1987), The estimation of wage gains and welfare gains in self-selection, *Review of Economics and Statistics* 69, 42-9.
- Card, D. (1999), Estimating the causal effect of education, in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, North-Holland, Amsterdam.
- Carneiro, P., Hansen, K.T. and Heckman, J.J. (2001), Educational attainment and labor market outcomes: estimating distributions of the returns to educational interventions, Unpublished paper presented as Klein Lecture, University of Pennsylvania, September 2001; First Draft presented at Midwest Econometrics Group (MEG), October, 2000.
- Carneiro, P., Hansen, K.T. and Heckman, J.J. (2002), Estimating distributions of treatment effects with an application to the returns to schooling, *International Economic Review*, forthcoming.
- Harsanyi, J.J. (1955), Cardinal welfare, individualistic ethics and interpersonal comparisons of utility, *Journal of Political Economy* 63, 309-21.
- Heckman, J.J. (1990), Varieties of selection bias, *American Economic Review* 80, 313-318.
- Heckman, J.J. (1992), Randomization and social policy evaluation, in C. Mawski and I. Garfinkel (eds.), *Evaluating Welfare and Training Programs*, Harvard University Press, Cambridge.
- Heckman, J.J. (2001a), Micro data, heterogeneity, and the evaluation of public policy: Nobel Lecture, *Journal of Political Economy* 109, 673-748.
- Heckman, J.J. (2001b), Accounting for heterogeneity, diversity and general equilibrium in evaluating social programs, *Economic Journal* VIII, 654-699.
- Heckman, J.J. and Honoré, B. (1990), The empirical content of the Roy Model, *Econometrica*, 58, 1121-1149.
- Heckman, J.J. and Smith, J. (1993), Assessing the case for randomized evaluations of social programs, in K. Jensen and P. Madsen (eds.), *Measuring Labor Market Outcomes*, Ministry of Labor, Copenhagen.

REMOVING THE VEIL OF IGNORANCE,
Pedro Carneiro, Karsten T. Hansen and James J. Heckman

- Heckman, J.J. and Vytlacil, E. (2000), Local instrumental variables, in C. Hsiao, K. Morimune, and J. Powells (eds.), *Nonlinear Statistical Modeling: Proceedings of the Thirteenth International Symposium in Economic Theory and Econometrics: Essays in Honor of Takeshi Amemiya*, Cambridge University Press, Cambridge.
- Heckman, J.J. and Vytlacil, E. (2001), Structural equations, treatment effects and econometric policy evaluation, Fisher Schultz Lecture, World Meetings of the Econometric Society, Seattle, forthcoming in *Econometrica*.
- Heckman, J.J., Clements, N. and Smith, J. (1997), Making the most out of programme evaluations and social experiments: accounting for heterogeneity in program impacts, *Review of Economic Studies* 64, 487-535.
- Heckman, J.J., LaLonde, R. and Smith, J. (1999), The economics and econometrics of active labor market programs, in O. Ashenfelter and D. Cards (eds.), *Handbook of Labor Economics*, North-Holland, Amsterdam.
- Jöreskog, K. and Sörbom, D. (1979), *Advances in Factor Analysis and Structural Equations Models*, Abt Associates, Cambridge.
- Kane, T. (1994), College entry by blacks since 1970: the role of college costs, family background, and the returns to education, *Journal of Political Economy* 102, 878-911.
- Persson, T. and Tabellini, G. (2000), *Political Economics—Explaining Economic Policy*, MIT Press, Cambridge, Massachusetts.
- Rawls, J. (1971), *A Theory of Justice*, Harvard University Press, Cambridge, Massachusetts.
- Roemer, J. (1996), *Theories of Distributive Justice*, Harvard University Press, Cambridge.
- Roy, A. (1951), Some thoughts on the distribution of earnings, *Oxford Economic Papers* 3, 135-146.
- Sen, A.K. (1997), *On Economic Inequality*, 2nd Edition, Clarendon Press, Oxford.
- Sen, A.K. (2000), Social justice and the distribution of income, in A.B. Atkinson and F. Bourguignon (eds.), *Handbook of Income Distribution*, North Holland, Amsterdam.
- Vickery, W. (1960), Utility, strategy and social decision rules, *Quarterly Journal of Economics* 74, 507-35.

Appendix

A.1. Identifying counterfactual distributions under treatment effect assumptions

Heckman and Honoré (1990) show that in the context of the original Roy (1951) model under normality or exclusion restrictions, it is possible to identify the joint density of potential outcomes. The original Roy model sets $C = 0$. Sectoral choices are then determined solely by potential outcomes. This extra information identifies the full model and lets analysts identify the joint distributions of outcomes across policy states. If there is variation in C across persons, this method breaks down and it is only possible to identify $g(Y_0|S=0)$ and $g(Y_1|S=1)$, the conditional densities of the potential outcomes, as well as $\Pr(S=1)$, but not the joint density, $g(Y_0, Y_1)$ (Heckman, 1990). Another special case that is discussed in Heckman (1992), is case I where $Y_1 \equiv Y_0 + \Delta$, and Δ is a constant. Then from the marginal distribution of Y_0 or Y_1 it is possible to construct the joint distribution (Y_0, Y_1) which is degenerate. Heckman and Smith (1993) and Heckman, Smith and Clements (1997) generalize this case to assume that the persons at the q^{th} percentile in the density of Y_0 are at the q^{th} percentile of Y_1 . Even without imposing this information, it is also possible to bound the joint densities from the marginals using classical results in probability theory. In practice these bounds turn out to be rather wide (Heckman and Smith, 1993; Heckman, Smith and Clements, 1997).

In our source papers (Carneiro, Hansen and Heckman, 2001, 2002), we generate the distributions of potential outcomes using a panel data factor structure model. The paper builds on and generalizes an idea presented in a cross section setting in Aakvik, Heckman and Vytlačil (1999). For the details of our method we refer the reader to our source papers. Here we present the intuitive idea that underlies our method and in the text we report its application. We discuss the most elementary case, leaving a complete discussion of the more general case for our companion papers.

Suppose that the mean of C depends on shifter variables Z that do not affect (are independent of) potential outcomes (Y_0, Y_1) . These are instruments. Suppose that for some values of Z within available samples we observe

$$\Pr(S = 1|Z) \doteq 1 \quad Z \in Z_1$$

while for other values of Z

$$\Pr(S = 0|Z) \doteq 1 \quad Z \in Z_0$$

Thus if Z is tuition, people who face a low tuition cost (possibly even a large subsidy) are almost surely likely to go to college while those who face a very high tuition cost are almost certainly likely not to go to school.⁸ We assume that the distribution of potential outcomes is the same in these subsets as they are in the overall distribution. Thus we can identify the marginal distribution of Y_1 from the first sample and the marginal distribution of Y_0 from the second sample.

Within these samples, we observe post-schooling outcomes

$$Y_{0t}, \quad t = 1, \dots, T, \quad \text{for } Z \in Z_0,$$

$$Y_{1t}, \quad t = 1, \dots, T, \quad \text{for } Z \in Z_1.$$

From these data we can form the joint densities of each outcome over time on $g(y_{01}, \dots, y_{0T})$ and $g(y_{11}, \dots, y_{1T})$, but not the joint densities over time over both outcomes.

Now suppose that Y_{0t} and Y_{1t} are both generated by a common factor f (e.g. ability, motivation) so that

$$\begin{aligned} Y_{0t} &= \mu_{0t} + \alpha_{0t}f + \varepsilon_{0t}, & t = 1, \dots, T, \\ Y_{1t} &= \mu_{1t} + \alpha_{1t}f + \varepsilon_{1t}, & t = 1, \dots, T, \end{aligned}$$

⁸ This is the version of identification at infinity discussed in Heckman (1990).

where the ε_{0t} and ε_{1t} are mutually independent of each other, f , and all other $\varepsilon_{1t'}$, $\varepsilon_{0t''}$, $t \neq t', t''$.⁹ All of these error components are assumed to have mean zero. A common factor generates both potential outcomes. If we can recover the distribution of the common factor, we can compute the joint distribution of counterfactuals up to some signs for the covariances.

Within each regime we can compute the following covariances:

$$\begin{aligned} \text{Cov}(Y_{0t}, Y_{0t'}) &= \alpha_{0t} \alpha_{0t'} \sigma_f^2, & t \neq t', t, t' = 1, \dots, T, & \text{ for } Z \in \mathbf{Z}_0, \\ \text{Cov}(Y_{1t}, Y_{1t'}) &= \alpha_{1t} \alpha_{1t'} \sigma_f^2, & t \neq t', t, t' = 1, \dots, T, & \text{ for } Z \in \mathbf{Z}_1. \end{aligned}$$

For concreteness suppose $T = 3$, so we have three panel wage observations. Then

$$\begin{aligned} \text{Cov}(Y_{01}, Y_{02}) &= \alpha_{01} \alpha_{02} \sigma_f^2, \\ \text{Cov}(Y_{01}, Y_{03}) &= \alpha_{01} \alpha_{03} \sigma_f^2, & \text{ for } Z \in \mathbf{Z}_0, \\ \text{Cov}(Y_{02}, Y_{03}) &= \alpha_{02} \alpha_{03} \sigma_f^2, \end{aligned}$$

and

$$\begin{aligned} \text{Cov}(Y_{11}, Y_{12}) &= \alpha_{11} \alpha_{12} \sigma_f^2, \\ \text{Cov}(Y_{11}, Y_{13}) &= \alpha_{11} \alpha_{13} \sigma_f^2, & \text{ for } Z \in \mathbf{Z}_1, \\ \text{Cov}(Y_{12}, Y_{13}) &= \alpha_{12} \alpha_{13} \sigma_f^2. \end{aligned}$$

If we assume $\alpha_{01} = 1$ or $\sigma_f^2 = 1$, we can identify all of the rest of the factor loadings. The proof is straightforward. Assume $\alpha_{03} \neq 0$. Then

$$\frac{\text{Cov}(Y_{01}, Y_{02})}{\text{Cov}(Y_{01}, Y_{03})} = \frac{\alpha_{02}}{\alpha_{03}}.$$

Given $\sigma_f^2 = 1$, we can use

⁹ The means may depend on the covariates.

$$\text{Cov}(Y_{02}, Y_{03}) = \alpha_{02} \alpha_{03}$$

to obtain

$$(\alpha_{03})^2 = \frac{\text{Cov}(Y_{02}, Y_{03})\text{Cov}(Y_{01}, Y_{03})}{\text{Cov}(Y_{01}, Y_{02})}$$

and we can identify α_{03} up to sign and hence can identify α_{02} and α_{01} up to sign. If we normalize $\alpha_{01} = 1$, we can identify α_{02} , α_{03} up to sign and σ_f^2 . Since the sign of f is unknown, the sign of the factor loadings is unknown.

With this information in hand, we can identify the variances of the uniquenesses, ε_{0t} , ε_{1t} of the outcomes:

$$\begin{aligned} \text{Var}(\varepsilon_{0t}) &= \text{Var}(Y_{0t}) - \alpha_{0t}^2 \sigma_f^2 & t = 1, \dots, T, \\ \text{Var}(\varepsilon_{1t}) &= \text{Var}(Y_{1t}) - \alpha_{1t}^2 \sigma_f^2 & t = 1, \dots, T. \end{aligned}$$

Suppose that $f, \varepsilon_{0t}, \varepsilon_{1t}, t = 1, \dots, T$ are normally distributed. Then from the information just presented obtained from the subsamples associated with Z_0 and Z_1 we can identify the density of f and hence the joint density of $(Y_{01}, Y_{11}, \dots, Y_{0T}, Y_{1T})$. Using the outcome data within schooling choices we can identify the distribution of f and hence estimate the joint distribution of schooling choices across potential outcomes provided that we fix the sign of the factor loadings.

In our companion paper we present two methods for resolving the ambiguity regarding the sign of the covariances. The first method explicitly models the choice process and uses the covariance between choices and outcomes to pin down the sign of the factor loadings and the covariances of the potential outcomes *across* schooling levels. The second method uses an indicator of the factor (e.g., an ability test) to resolve the sign problem. Both approaches rely on the same basic idea of using the covariances of Y_{1t} and Y_{0t} with a common third variable to identify the sign of the factor loading. The second approach is easier to motivate and we do so here.

Suppose that we have access to one ability test for each person. Measured ability is

$$A = \mu_A(X) + \beta f + \varepsilon_A,$$

where $\mu_A(X)$ is the mean of ability, X are the covariates predicting ability, and ε_A is mutually independent of $(\varepsilon_{01}, \dots, \varepsilon_{0T}, \varepsilon_{11}, \dots, \varepsilon_{1T})$ and the f .

We can compute

$$\text{Cov}(A, Y_{0t}) = \beta \alpha_{0t} \sigma_f^2 \quad t = 1, \dots, T$$

for persons who do not attend school (e.g., do not attend college) and

$$\text{Cov}(A, Y_{1t'}) = \beta \alpha_{1t'} \sigma_f^2 \quad t' = 1, \dots, T$$

for persons who attend school (e.g., college). Assuming $\beta \neq 0$ and $\alpha_{0t}, \alpha_{1t'} \neq 0$ we can identify

$$\frac{\text{Cov}(A, Y_{1t'})}{\text{Cov}(A, Y_{0t})} = \frac{\alpha_{1t'}}{\alpha_{0t}} \quad t = 1, \dots, T, t' = 1, \dots, T \text{ and } t \neq t'$$

and hence we fix the sign of $\text{Cov}(Y_{0t}, Y_{1t'})$ for all t and t' . This resolves the ambiguity regarding the sign of the covariances.

In our companion papers we show that we can obtain this joint density without a normality assumption for f or $\varepsilon_{0t}, \varepsilon_{1t}, t = 1, \dots, T$. We extend our analysis to allow for vector f so there may be many factors, not just one. We show that it is possible to nonparametrically identify the joint density of potential outcomes provided that the number of panel data wage measurements is large, in a sense we make

precise in our companion papers, relative to the number of factors.¹⁰ We do not need to invoke “identification at infinity”, i.e. we can dispense with the requirement that there are subsets of Z where there is no selection. We also consider a model with multiple discrete choices (schooling levels) instead of just two. With these counterfactual distributions determined, we can identify the impact of social policy on the distributions of outcomes and returns.

¹⁰ In our companion papers, we show how indicators of f can be used to supplement, or replace, panel data. This type of identification is familiar to users of LISREL (see Jöreskog and Sörbom, 1979).

