Equality of Opportunity and Kernel Density Estimation:

an Application to Intergenerational Mobility*

Donal O'Neill, Olive Sweetman and Dirk Van de gaer

Abbreviated Title: **Equality of Opportunity, Kernel Density and Intergenerational Mobility**

Economics Dept., NUI Maynooth, Maynooth, Co. Kildare, Ireland

Phone: 353-1-7083728
Fax: 353-1-7083934

E-Mail: olive.sweetman@may.ie; donal.oneil@may.ie; dirk.vandegaer@may.ie

* We are grateful to Walter Bossert, Tom Fomby, Marc Fleurbaey, Ramses Abul Naga and Erik Schokkaert for their helpful comments. We would also like to thank seminar participants at NUI Galway, NUI Maynooth, the annual conference of the Irish Economic Association (Westport, March 1999), the Applied Econometrics Society meetings (Pau, May 1999), the Conference on Intergenerational Transfers, Taxes and the Distribution of Wealth (Uppsala, June 1999) and the Econometric Society Summer meetings (Madison, June 1999). The usual disclaimer applies.

1. Introduction

When economists talk about equality, they typically have equality of outcomes, like welfare or income in mind. However such a view of equality is not entirely satisfactory and theories of equality of opportunity have been developed and proposed as an alternative[1]. To highlight the distinction between equality of outcome and equality of opportunity, consider two young adults, David and Dick. They are identical twins who grew up in the same family, went to the same high school and had the same circle of friends. After high school, they have to decide what to do next. Their parents are sufficiently wealthy and open minded, so that David and Dick can do anything that they want. Both decide to go into auto mechanics. David works hard, while Dick is lazy. As a consequence, David becomes very rich, while Dick remains poor. Assuming that egalitarians care for people's incomes, should this inequality of income be a source of concern? For egalitarians that believe in equality of outcome the answer is yes, for egalitarians that believe in equality of opportunity the answer is no. The reason is that both had the same set of opportunities. David decides to work hard, while Dick chooses to put in no effort. Equality of opportunity holds individuals accountable for their choices.

Next, consider two different young adults. Charlene's parents are rich, while Christine's parents are poor. Both of them would like to become a surgeon. The training to become a doctor takes a lot of time, and Christine's parents cannot afford to send their daughter to college so she becomes a nurse instead. They both work equally hard in their chosen professions. Surgeons have much higher lifetime earnings than nurses. Consequently, Charlene's lifetime earnings will be much higher than Christine's. In this case, outcome egalitarians and opportunity egalitarians will agree that this inequality is undesirable.

If we focus on equality of outcome then any inequality is undesirable. Equality of opportunity, on the other hand, holds people responsible for some differences in outcomes, like those that are the result of genuine choice, but not for others, like those that result from social background, racial or sexual discrimination. The latter are called non-responsibility characteristics. This has one immediate consequence for the informational requirements of the empirical analysis of inequality. If we are interested in inequality of income, we only need data on the distribution of incomes. If attention is focused on inequality of opportunity for income these data will not suffice. We then have to identify and measure variables or characteristics for which people are responsible, and variables for which they aren't responsible, in order to calculate an individual's opportunity set.

In the next section we further examine the concept of equality of opportunity by defining the opportunity sets of individuals conditional on characteristics for which they are not responsible. In Section 3 we describe the nonparametric methodology we use to estimate the opportunity sets of individuals while in Section 4 we use this methodology to examine the intergenerational mobility between father's income (non-responsible characteristic, x) and son's income (outcome, y). We conclude in Section 5.

2. Equality of Opportunity

The second example given in the introduction implicitly raises the issue of the evaluation of opportunity sets. Charlene has clearly superior opportunities: Christine's choice set is probably a subset of Charlene's, but can we somehow quantify this difference? The answer to this question is non-trivial, e.g., Sen (1985)[2], and many different answers have been given, see Bossert et al. (1999). A related issue focuses on redistribution among individuals. For instance, we may want to compensate people for the influence of the non-responsibility characteristics, yet preserve the influence of responsibility characteristics -see, e.g., Fleurbaey (1995c) for a discussion that puts the issue of responsibility into a general egalitarian perspective. In example one, all of the non-responsibility characteristics of David and Dick are equal so we would want to hold Dick accountable for being lazy. With equality of opportunity, individuals cannot be held responsible for being born into a poor family and so in example two, Christine should be compensated for this fact.[3]

Our paper provides a non-parametric procedure for estimating opportunity sets for those interested in evaluating differences in these sets. With this approach, we can also analyse the impact of responsibility and non-responsibility characteristics on outcomes. Let the opportunity set of a person, S_x, be determined by a vector of non-responsibility characteristics, x. The elements of S_x are the relevant outcomes. A person's non-responsibility characteristics determine the opportunity set out of which she can chose, while the responsibility characteristic determines which outcome she chooses from that set. Different people that are of the same type may choose different options out of their opportunity set, and therefore obtain a different relevant outcome, $z=y[p,x]$. One can think of z as utility or lifetime income. p represents the variable that causes differences in the outcome for which we do not want to compensate. p can be thought of as a taste parameter inducing people to work harder. We will assume that p is a scalar. The first branch of the literature referred to above tries to order (distributions of) opportunity sets, S_x, the second branch of the literature imposes requirements on $y[p,x]$. To examine either of these issues, we need to be able to identify the opportunity sets.

The first step in doing this concerns the identification of p. We will assume that the distribution function of p is continuous[4]. Roemer (1993, 1996, 1998) proposes a procedure that can be motivated by the combination of two assumptions[5]:

SINC (Strictly Increasing): $y[p,x]$ is a strictly increasing function of p.

This assumption is quite natural if a higher value for p is thought of as inducing more effort. Let $F_z^*[z|x]$ and $F_p^*[p|x]$ denote the cumulative distribution function of z conditional on x, and the cumulative distribution of p conditional on x, respectively. SINC implies that

$$F_z^*[y[p',x]|x]=F_p^*[p'|x] \qquad (1)$$

In words: an individual's level of effort is smaller than the α-th percentile in the distribution of the effort within his type if and only if his outcome is below the α-th percentile in the distribution of outcome within his type.

IND (Independence): $F_p^*[p'|x]$ is independent of x.

Again this assumption is quite acceptable since it is difficult to imagine in what sense people could be held responsible for p if the independence axiom does not hold[6].

Together (1) and IND imply[7] Roemer's identification axiom:

RIA (Roemer's Identification Axiom): $F_z^*[y[p',x_1]|x_1] = F_z^*[y[p'',x_2]|x_2] \Rightarrow p'=p''$

RIA says that, if two people of different types are at the same percentile of the distribution of the outcome within their type, then they have the same value for p. Let $\pi= F_z^*[z|x]$ denote the cumulative distribution of the outcome (z) given the level of non-responsibility (x). Assume that this is a strictly increasing function of z. $F_z^{*-1}[\pi|x]$

expresses the income level obtained by an individual of type x and who was at the 100*π-th percentile of the cumulative distribution of outcome within his type. Looking at $F_z^{*-1}[\pi|x]$ is, because of RIA, equivalent to looking at $y[p,x]$.

$F_z^{*-1}[\pi|x]$ visualises the dependence of the outcome on responsibility and non-responsibility characteristics. It allows us to draw income as a function of $\pi \in [0,1]$ for different values of x. The opportunity set for a particular type x is given by the outcomes someone of type x can obtain by varying his responsibility characteristic p or π. Consequently, we find the opportunity set of individuals of type x as

$$S_x = \{(z,\pi) \in (R^+ \times [0,1]) \,|\, z = F_z^{*-1}[\pi|x]\} \tag{2}$$

where R^+ is the set of non-negative real numbers. If $F_z[z|x]$ was known we could depict the opportunity sets for different types of individuals, as suggested by Fleurbaey, described in Bossert et al (1999) and found in Roemer (1998, p.75). Visual inspection of these sets will provide some idea about the degree of inequality of the opportunities offered to different types and the extent to which people who have the same tastes or expended level of effort obtain different outcomes.[8]

In the literature it is usually assumed that society somehow determines the elements of x. It can be expected that x is a multi-dimensional variable including things like race, sex, parental background and innate ability. The data requirements for the calculation of these conditional opportunity sets are vast. Therefore we illustrate the above framework in the context of intergenerational mobility[9]. In this context, z is the son's income and x, the vector of non-responsibility characteristics, is parental income. Conditioning only on father's income might make the independence assumption less plausible. However, scholars often look at data on intergenerational mobility as providing an indication of the extent of inequality of opportunity in society. The above analysis makes explicit that such an interpretation can only be maintained if some kind of independence assumption holds.

Roemer (1998) and Betts and Roemer (1998) propose to construct opportunity sets by running a regression of π on the outcome variable. That procedure has two obvious disadvantages. First, the approach is parametric so a functional form has to be imposed. Second, the procedure is only operational if the number of types is finite, otherwise π cannot be identified. Alternatively it has been suggested that intergenerational transition matrices can be used to obtain a rough approximation of the opportunity sets open to children from different descent – see Van de gaer, Schokkaert and Martinez (1998). Transition matrices are non-parametric, but are based on a discrete approximation of the conditional distribution. Both parents' incomes and children's incomes have to be divided into classes, called bins. As a consequence, transition matrices are sensitive to the choice of bins. As noted in Quah (1996), this approach can distort dynamics when the underlying observations are not discrete. To overcome these problems, we propose using kernel density estimation techniques to estimate the entire conditional distribution of children's income which avoids the need for arbitrary discretization of the data.

3. Bivariate Kernel Density Estimation

Denoting father's income by x and son's income by z, the distribution of z given x can be written as:

$$f[z|x] = \frac{f[z,x]}{f_x[x]} \tag{3}$$

where $f_x[x]$ is the marginal distribution of father's income and $f[z,x]$ is the joint distribution of z and x . To estimate this distribution we replace both the numerator and denominator of this expression with non-parametric estimates. The marginal distribution of father's income is estimated using adaptive kernel density techniques for univariate distributions. That is:

$$\hat{f}_{xA}[x] = \frac{1}{nh_x}\sum_{i=1}^{n}\frac{1}{w_i}K\left[\frac{x-x_i}{w_i h_x}\right] \qquad (4)$$

where h_x is the bandwidth for father's income and $K[.]$ is the kernel[10]. The subscript A on the estimator indicates that the adaptive procedure is used.

The joint distribution of fathers' and sons' income (the numerator of (3)) is estimated as

$$\hat{f}_A[z,x] = \frac{1}{nh_z h_x}\sum_{i=1}^{n}\frac{1}{w_i^{\,2}}K\left[\frac{z-z_i}{w_i h_z}\right]K\left[\frac{x-x_i}{w_i h_x}\right] \qquad (5)$$

The adaptive kernel estimator adjusts the window width so that the window is narrower where the density is high and wider where the density is low, thus retaining detail where data are plentiful and reducing noise where data are sparse.

The local window factors used are given by:

$$w_i = \left(\frac{\tilde{f}_g}{\hat{f}_k[z_i, x_i]}\right)^{1/2} \qquad (6)$$

where $\hat{f}_k[z,x]$ is the fixed window kernel estimate of $f[z,x]$ and $\tilde{f}_g$ is the geometric mean of $\hat{f}_k[z,x]$. Implementation of this method thus involves a two-step estimation strategy (see Silverman (1986) and Fox and Long (1990)). First $\hat{f}_k[z,x]$ is estimated using a fixed window width. This initial window width was calculated using Scott's optimal bandwidth for the bivariate normal density (Scott (1991), page 152)[11]. Next this density is used to calculate the weights involved in constructing the final density estimates in equations (4) and (5). We follow Trede (1998) in assuming that the kernel is multiplicative. This simplifies the expression for the cumulative conditional distribution. However, as noted by Trede, more general specifications are possible.[12]

The conditional distribution is estimated by replacing the terms in equation (3) with our estimates derived from (6) and (4)[13]. One can think of this density as providing an estimate of the transition matrix as the number of cells in the matrix tends to infinity. In doing so the kernel approach avoids the arbitrary discretization inherent in traditional transition matrix estimation.

In much of our analysis we will be interested in the cumulative distribution of son's income conditional on fathers income. Extending the approach discussed above we estimate this directly as

$$\hat{F}^{*}_{zA}\left[z|x\right]=\frac{\sum_{i=1}^{n}\frac{1}{w_i}K\left[\frac{x-x_i}{w_i h_x}\right]G\left[\frac{z-z_i}{w_i h_z}\right]}{\sum_{i=1}^{n}\frac{1}{w_i}K\left[\frac{x-x_i}{w_i h_x}\right]} \tag{7}$$

where $G(z)=\int_{-\infty}^{z}K(t)dt$ is the cdf of the kernel function. The opportunity set for children of parents with income level x_i can then be estimated as

$$\hat{S}_{x_i}=\left\{(z,\pi)\in\left(R^{+}\times[0,1]\right)\middle|z=\hat{F}^{*-1}_{zA}\left[\pi|x_i\right]\right\} \tag{8}$$

4. Empirical Results

As an illustration of the methodology discussed in sections 2 and 3, we now compute and estimate the sets S_{x_i} in the context of intergenerational mobility. In part a) we begin by deriving the opportunity sets for the case where fathers' and sons' incomes are known to be joint lognormally distributed. In part b) we create an artificial data set by drawing observations from a joint lognormal distribution. We use these observations to estimate the opportunity sets using the methodology developed in section 3. We then compare our estimates to the true opportunity sets derived in part a). Finally in part c) we apply our procedure to actual data on father and son's income taken from the National Longitudinal Surveys.

a) Opportunity Sets when Incomes are Jointly Lognormally Distributed

To visualise opportunity sets we must first specify a distribution. In this sub-section we assume that the incomes of sons and fathers are jointly lognormally distributed.[14] Once this distribution is specified we can find (numerically) the distribution of z (sons' incomes) given x (father's income). To derive the opportunity sets we first fix a value for father's income. Conditional on this value we plot son's income against the percentile points of the conditional distribution. Examples of these opportunity sets are given in Figure 1, where the y-axis measures the son's income relative to the mean of son's incomes and the x-axis represents the percentile points of the conditional distribution.

FIGURE 1 ABOUT HERE

Panel (a) shows the opportunity sets when incomes are not correlated. Since the distribution of son's income conditional on father's income does not depend on x, the opportunity sets are the same for all levels of father's income. With this specification there is no inequality of opportunity. To obtain average income, one has to have a level of effort equal to the 60th percentile in the effort distribution, irrespective of father's income.

Panel (b) shows the opportunity sets when incomes are correlated, with the correlation equal to .5. In this case a son's opportunities will depend on the level of his father's income and so the opportunity sets will differ. We plot 3 such opportunity sets corresponding to 'poor', 'average' and 'rich' fathers. For the purpose of drawing these graphs a poor father is defined as the individual at the 25th percentile of the father's income distribution. Average income is defined to be the median of the father's distribution and rich fathers are defined to be those at the 75th percentile of the distribution. It is clearly advantageous to have a rich father. Given a particular amount of effort, a son whose father is rich always has a higher level of income than one whose father is poor. The underlying income differences can be quite substantial. If you work at a median level of effort and your father was rich, your income will be 40 per cent higher than the level of income obtained by someone who put in the same amount of effort, but who was unfortunate to have a poor father. To obtain the average level of income, sons of poor fathers have to work much harder than those with rich fathers. Such differences are indicative of the existence of a substantial amount of inequality of opportunity.

b) Nonparametric Estimation of Opportunity Sets with Simulated Lognormal Data

Non-parametric estimators are traditionally data intensive, see, e.g., Scott (1992). However in most studies of intergenerational mobility the number of observations are quite small.[15] We generate an artificial data set, with 300 observations from a joint lognormal distribution and parameters equal to those used in deriving Figure 1b). The estimated opportunity sets from these data along with the true sets are given in Figure 2.[16] The solid line represents the true opportunity set while the dashed line represents the estimated set using the nonparametric method described in Section 3 above.

FIGURE 2 ABOUT HERE

Figure 2 suggests that even with small samples our estimator is capable of providing estimated opportunity sets, which are similar to those generating the data.

c) Intergenerational Mobility: NLS Data

This section of the paper illustrates our approach by applying it to an analysis of intergenerational mobility in the U.S. Concerns about the intergenerational transmission of well-being arise from the belief that such a process poses a significant barrier to achieving equality of opportunity. However the studies examining intergenerational mobility to date (for example Becker and Tomes (1979), Solon (1992) and Zimmerman (1992)) have tended to emphasize the correlation between father and son's income. Such an analysis is useful in that it determines the extent to which children of poor families are themselves poor. However, as discussed in the earlier sections, knowing only the correlation between incomes is not sufficient to examine equality of opportunity. To do this we must estimate the opportunity sets facing individuals from different parental backgrounds. We do this by estimating equation (8) using data from the National Longitudinal Surveys (NLS).

The NLS data contain information on four cohorts of individuals, each consisting of approximately 5,000 members. These individuals were initially surveyed in 1966 and then surveyed repeatedly throughout the 1970's and early 1980's. The sampling design used by the Census Bureau made it possible for any given NLS household to include respondents in more than one of the cohorts. This allows us to match some children from the young cohorts with information obtained on their parents in the older cohort surveys. In total 1,039 older men resided with a son who was a respondent in the young men's cohort. When more than one match was possible we follow Zimmerman (1992) and use the match involving the eldest son. We are left with 876 father-son pairs from the "mature men" and "young men" cohorts using data through 1981.

In order to obtain a more reliable measure of permanent income for fathers, we average fathers' income over the years 1965, 1966 and 1968. Furthermore, only observations for which both the father and son were fully employed were retained in the sample.[17] To measure income we use data on wages and salaries expressed in 1981 prices using the CPIU historical index. When individuals with missing wage data were excluded from the sample we were left with a sample of 204. Summary statistics for this sample are provided in Table 1.[18]

TABLE 1 ABOUT HERE

The first step in implementing our approach is to estimate the density of son's income conditional on father's income. To make our results transparent we express both incomes relative to their respective means. The estimated conditional densities are given in figure 3, while the contours of these densities are given in figure 4.

FIGURE 3 ABOUT HERE

FIGURE 4 ABOUT HERE

The contours in figure 4 highlight the positive relationship between fathers' and sons' incomes. To the extent that son's income is independent of father's income we would expect the contours of the conditional density to be horizontal lines. However, the estimated contours are clearly upward sloping, indicating a positive relationship between fathers' and sons' incomes. [19] This is even more apparent when we restrict attention to fathers with less than twice the mean income.[20]

There is some evidence that the peaks at the lower end of the distribution are higher than any other parts of the distribution, which would indicate less mobility from lower income ranges. However, given the small sample sizes we should exercise caution in interpreting small differences in the estimated density.

While this analysis complements the traditional measures of intergenerational mobility the goal in our paper is to use these estimates to visualise the opportunity sets facing individuals from different parental backgrounds.[21] These are presented in Figure 5.

FIGURE 5 ABOUT HERE

As before we present the opportunity sets for children from poor, average and rich families. Looking at income differences among children with the same parental background shows the impact of variations in characteristics for which children may be held responsible (example effort). Thus we see that even among poor fathers children who exert substantial effort can wind up with income levels greater than the overall average. Fixing the percentile of the distribution on the other hand and comparing outcomes across parental backgrounds provides information on the degree of inequality of opportunity. For instance we see that among children who choose to exert little effort (say the 10th percentile of the respective distributions) children from rich families can expect to earn 56% more ([.67/.43]-1) than children of poor families. Alternatively one can see that in order to reach average income children from poor families have to exert a level of effort which would place them at the 70th percentile of their distribution. To reach the same level of income children from rich families need only reach the 40th percentile of their distribution. These differences are substantial and highlight the existence of substantial inequality of opportunity.

4. Conclusion

The goal of this paper has been to develop a non-parametric procedure for estimating the opportunity sets available to individuals of different types. While the framework can be applied in any context, the application which we have focused on concerns the impact of parental income on a child’s opportunities. Our results indicate significant inequality of opportunity across family backgrounds. Children from poor families can expect to do significantly worse than those from rich families, even where both children exert the same level of responsibility.

The development of a procedure for estimating opportunity sets is the first step in conducting an empirical analysis of equality of opportunity. In section 2, we briefly surveyed the theoretical literature attempting to quantify differences in opportunity sets. To date only crude estimates of opportunity sets have been available for this analysis. Our approach provides a robust non-parametric means of estimating the opportunity sets needed to carry this analysis further. An advantage of the non-parametric approach developed in our paper is that it is amenable to the introduction of controlling regressors and thus suited to estimating the counterfactual distributions needed to study the effect of policy measures on the distribution of opportunities. We are currently working on this issue. In our view the interaction of existing theoretical literature on equality of opportunity with a non-parametric estimation strategy, capable of carrying out counterfactual experiments, represents a significant development in applying social choice analysis.

References

Becker, G., and N. Tomes (1979), "An Equilibrium Theory of the Distribution of Income and Intergenerational Inequality." *Journal of Political Economy*, 87, 1153-89.

Betts, J.R. and J, Roemer (1998), "Equalizing Opportunity Through Educational Finance Reform." Mimeo.

Bossert, W. (1995), "Redistribution Mechanisms Based on Individual Characteristics." *Mathematical Social Sciences*, 29**,** 1-17.

Bossert, W., and M. Fleurbaey (1996), "Redistribution and Compensation." *Social Choice and Welfare*, 13, 343-355.

Bossert, W, Fleurbaey, M and D. Van de gaer (1999), " Responsibility, Talent and Compensation: a Second Best Analysis." *Review of Economic Design*, 4, 35-55.

Bossert, W., P.K. Pattanaik, and Y. Xu (1994), "Ranking Opportunity Sets: an Axiomatic Approach." *Journal of Economic Theory*, 63, 326-345.

Corak, M. and A. Heisz (1999), " The Intergenerational Earnings and Income Mobility of Canadian Men." *Journal of Human Resources*, 34, 504-533.

Cowell, F.A. (1995), *Measuring Inequality*. Harvester Wheatsheaf, Hemel Hempstead, 2nd edn.

Fleurbaey, M. (1995a), "The Requisites of Equal Opportunity." In W.A. Barnett, H. Moulin, M. Salles and N.J. Schofield (eds), *Social Choice, Welfare and Ethics*, Cambridge University Press, Cambridge, 37-53.

Fleurbaey, M. (1995b), “Three Solutions for the Compensation Problem.” *Journal of Economic Theory*, 65, 505-521.

Fleurbaey, M. (1995c), “Equality and Responsibility.” *European Economic Review*, 39, 683-689.

Fleurbaey, M. (1995d), “Equal Opportunity or Equal Social Outcome?” *Economics and Philosophy*, 11, 25-55.

Fleurbaey, M. (1998), “Equality Among Responsible Individuals.” In J.-F. Laslier, M. Fleurbaey, N. Gravel and A. Trannoy (eds), *Freedom in Economics: New Perspectives in Normative Analysis*, Routledge, London, 206-234.

Fox, J and J.S. Long (1990) (eds), *Modern Methods of Data Analysis*. Sage, California.

Gravel, N. (1998), “Ranking Opportunity Sets on the Basis of their Freedom of Choice and their Ability to Satisfy Preferences: a Difficulty.” *Social Choice and Welfare,* 15, 371-382.

Iturbe-Ormaetxe, I. (1997), “Redistribution and Individual Characteristics.” *Review of Economic Design*, 3, 45-55.

Kranich, L. (1996), “Equitable Opportunities: an Axiomatic Approach.” *Journal of Economic Theory*, 71, 131-147.

Minicozzi, A. (1997), “Nonparametric Analysis of Intergenerational Income Mobility.” Doctoral Dissertation, University of Wisconsin Madison.

Ok, E. A. and L. Kranich (1998), “The Measurement of Opportunity Inequality: a Cardinality-Based Approach.” *Social Choice and Welfare*, 15, 263-287.

Pattanaik, P. K. and Y. Xu (1990), “On Ranking Opportunity Sets in terms of Freedom of Choice.” *Recherches Economiques de Louvain*, 56, 383-390.

Pen, J. (1980), "A Parade of Dwafs (and a few Giants). " In A.B. Atkinson (ed.), *Wealth, Income and Inequality*, Oxford University Press, Oxford, 47-55.

Quah, D. (1996), "Empirics for Growth and Distribution: Stratification, Polarization and Convergence Clubs." LSE Discussion paper.

Roemer, J. (1993), "A Pragmatic Theory of Responsibility for the Egalitarian Social Planner." *Philosophy and Public Affairs*, 22, 146-166.

Roemer, J. (1996), *Theories of Distributive Justice.* Harvard University Press, Cambridge.

Roemer, J. (1998), *Equality of Opportunity.* Harvard University Press, Cambridge.

Scott, D. (1992), *Multivariate Density Estimation.* Wiley, New York

Sen, A. (1985), *Commodities and Capabilities.* North Holland, Amsterdam.

Silverman, B.W. (1986), *Density Estimation for Statistics and Data Analysis.* Chapman and Hall, London.

Solon, G. (1992), "Intergenerational Income Mobility in the United States." *American Economic Review*, 82, 392-408.

Trede, M. (1998), "Making mobility visible: a Graphical Device." *Economic Letters*, 59, 77-82.

Van de gaer, D. (1993), Investment in Human Capital and Equality of Opportunity, doctoral dissertation, KULeuven, Leuven.

Van de gaer, D, E.Schokkaert, and M. Martinez (1998), "Measuring Intergenerational Mobility and Equality of Opportunity." Economics Department Working Papers Series N78/05/98, NUI Maynooth.

Zimmerman, D. (1992), "Regression Toward Mediocrity in Economic Stature." *American Economic Review*, 82, pp 409-429

Table 1: Summary Statistics for NLS Sample

Variable Name	*Mean (Std)*
Fathers' Log Earnings 1965	9.87 (.58)
Fathers' Log Earnings 1966	9.91 (.52)
Fathers' Log Earnings 1968	9.95 (.56)
Fathers' Age 1966	49.7 (3.7)
Sons' Log Earnings 1981	9.90 (.50)
Sons' Age 1981	32.73 (2.80)
Beta[a]	.41 (.06)
N	204

[a] This is the regression estimate when sons' income in 1981 is regressed on the three year average of fathers' income in 1981 dollars.

Figure 1: Opportunity Sets of Sons when Fathers' and Sons' Incomes are Jointly Lognormally Distributed

(a) The Opportunity Sets of Sons when Incomes are Uncorrelated

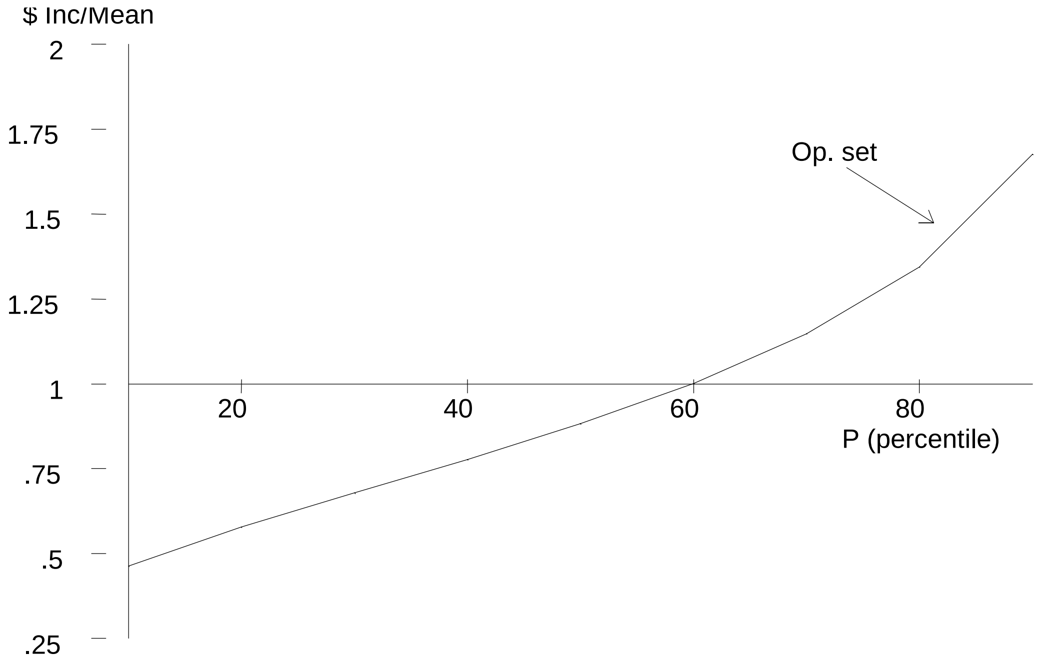

(b) The Opportunity Sets when the Correlation Coefficient=0.5

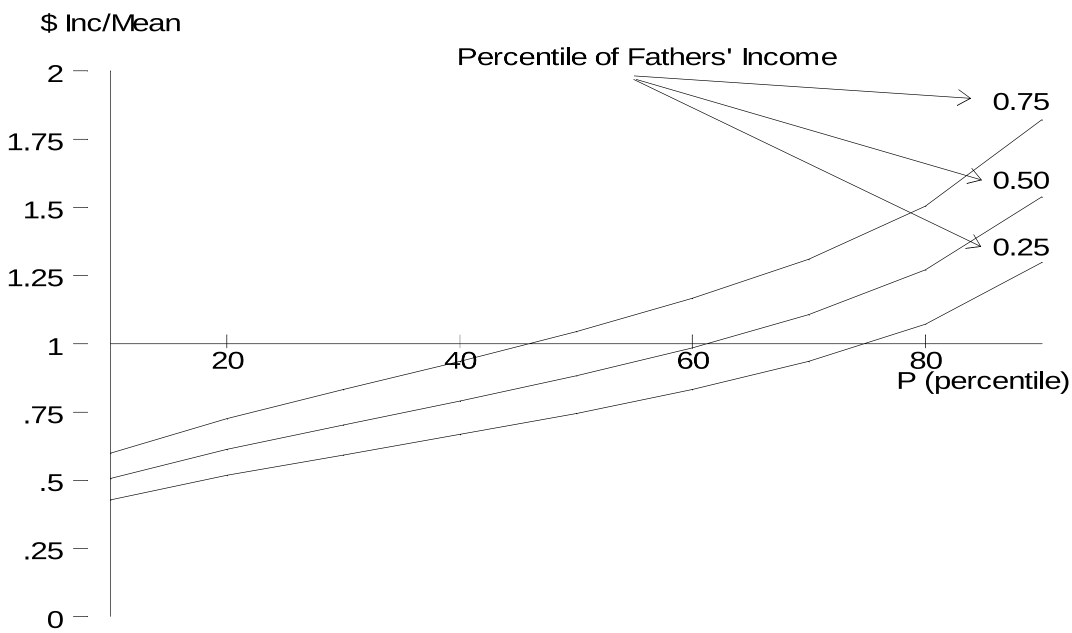

Figure 2: True and Estimated Opportunity Sets when Fathers' and Sons' Incomes are Jointly Lognormally Distributed.

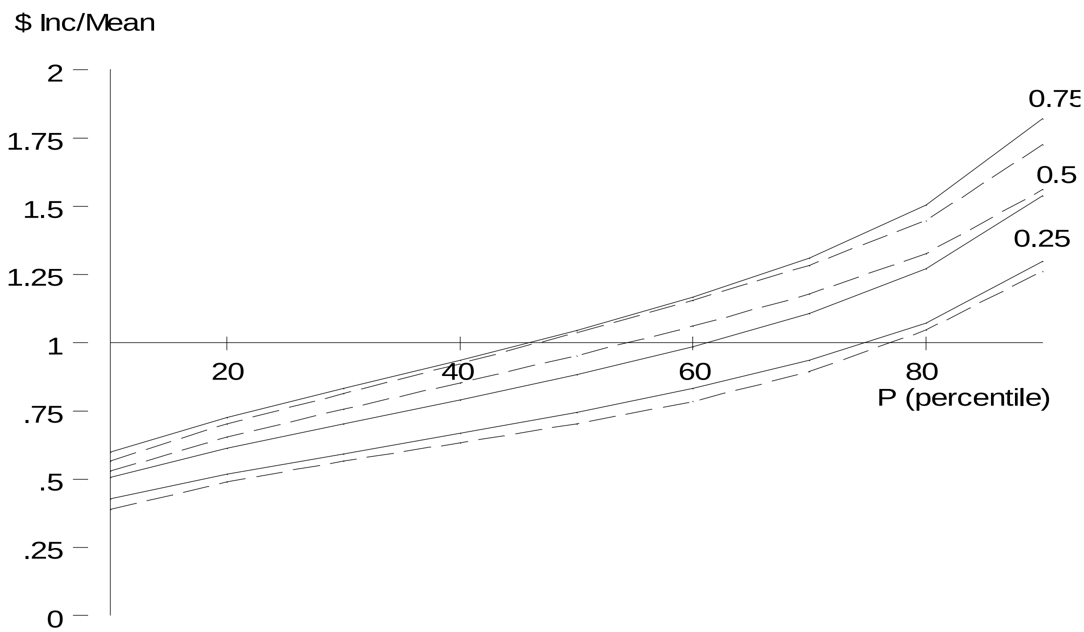

Figure 3: Estimated Conditional Densities of Sons' Income given Fathers' Income using NLS data.

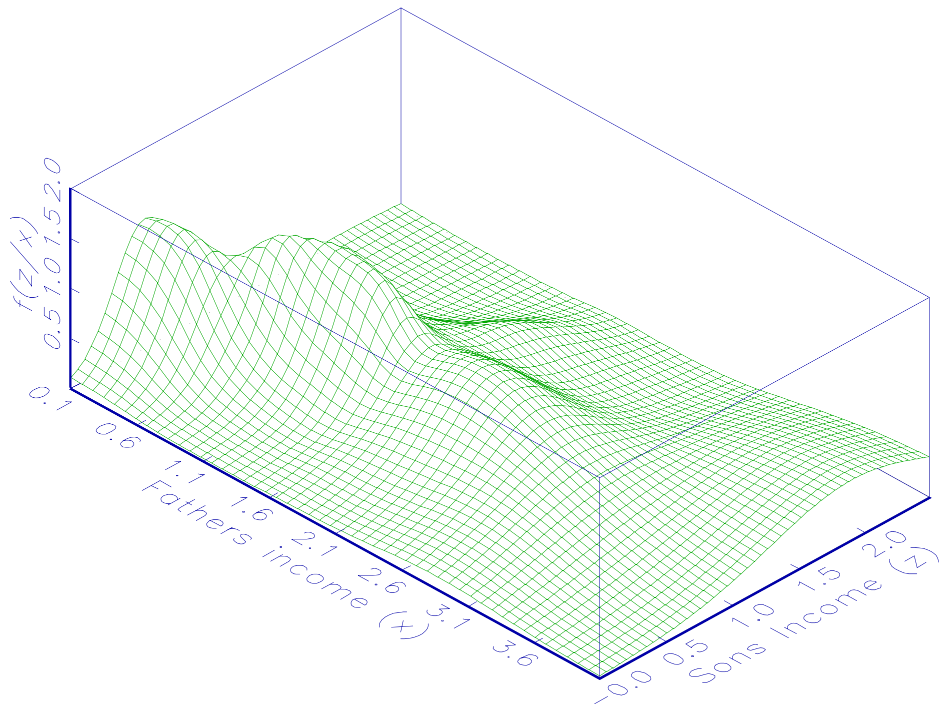

Figure 4: Estimated Contours of the Conditional Distributions using NLS data.

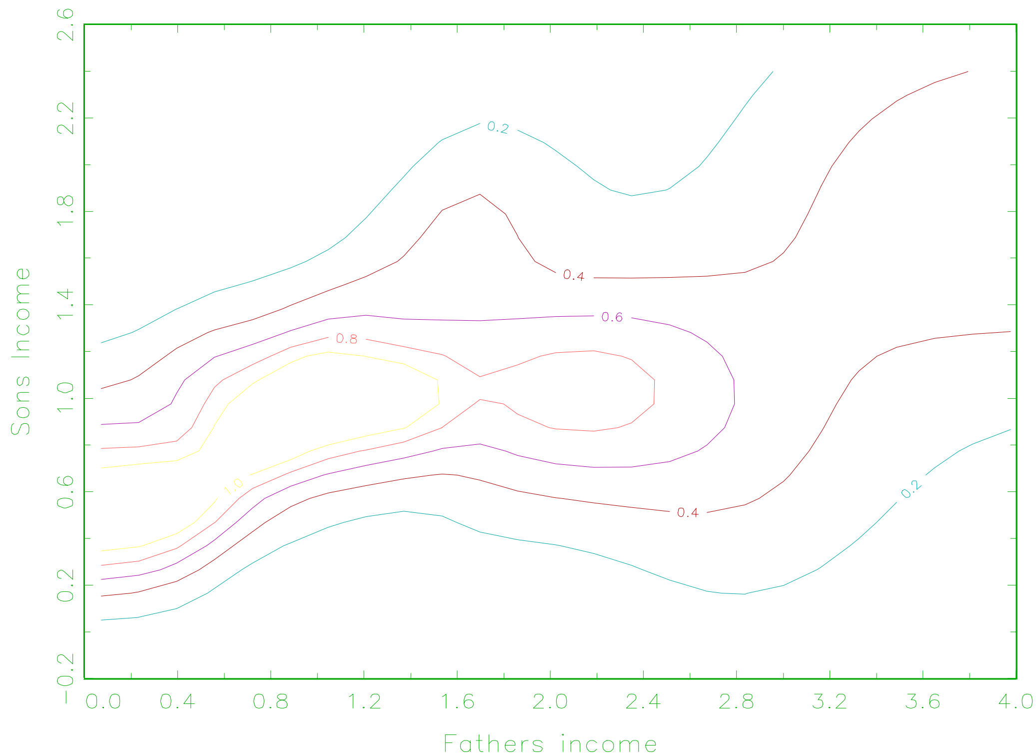

Figure 5: Estimated Opportunity Sets from the NLS data

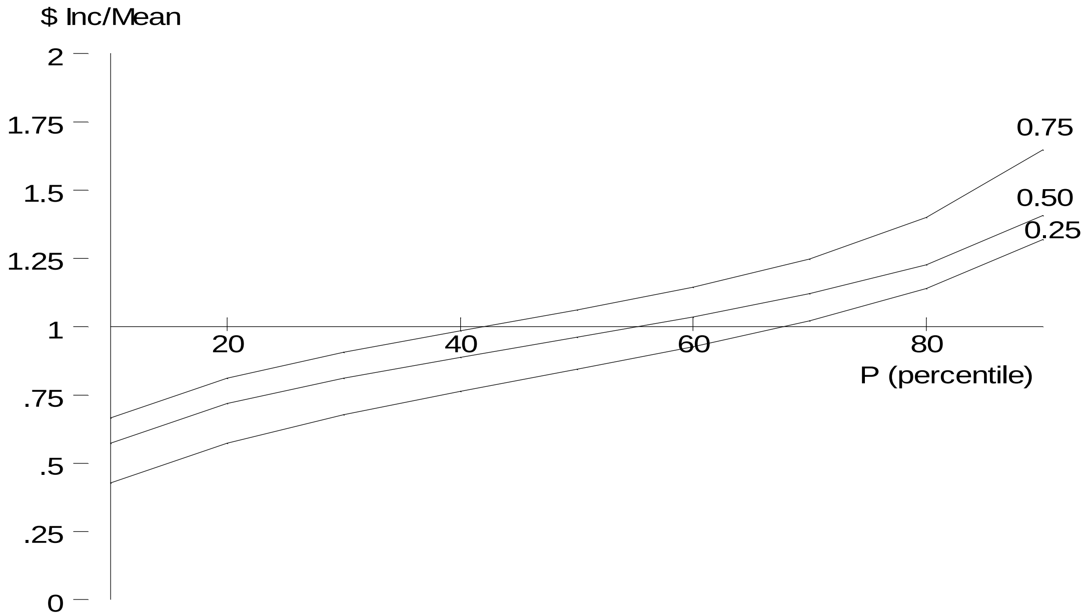

[1] There is a vast philosophical literature about these issues. This literature is critically assessed in Fleurbaey (1995d). See also Roemer (1996) for an overview.

[2] See Pattanaik and Xu (1990), Bossert, Pattanaik and Xu (1994), Kranich (1996), Gravel (1998) and Ok and Kranich (1998) for related discussions.

[3] Fleurbaey (1995a, b) shows that in general it is not possible to simultaneously compensate for the influence of non-responsibility characteristics while preserving the influence of responsibility characteristics. Bossert (1995) comes to the same conclusion when income is the relevant outcome. Bossert and Fleurbaey (1996) and Iturbe- Ormaetxe (1997) discuss these issues further. Roemer (1993) and Van de gaer (1993) formulate a general objective to capture some of the above intuitions. Bossert et al (1999) discuss the links between the redistribution literature and the literature on social choice.

[4] The framework can be easily adapted to deal with a discrete distribution of p.

[5] See also Fleurbaey (1998, p.221). He discusses how the procedure might work when some of the x variables are observable.

[6] The independence axiom is the formal representation of Roemer's assumption of charity: "What I call the assumption of charity says that, within any type, that distribution *(referring to the distribution of effort)* would be the same, were we able to factor out the (different) circumstances which define types." (Roemer, 1998, p.15- the text in italics added for clarity)

[7] Note that the procedure for drawing opportunity sets is valid under wider circumstances than SINC and IND. These are sufficient conditions only. Consider the case where $y[p,x]$ is non-decreasing in p for some types but not strictly increasing over the entire support of p. The conditional cumulative distributions of outcomes for these types will then have a flat segment. Hence their inverse does not exist. We can still draw the opportunity sets of these types, however. The resulting opportunity set will be horizontal over part of the support [0,1] of π.

[8] Alternatively, the graph of S_x can be interpreted as a conditional Pen Parade, see, e.g., Pen (1980) or Cowell (1995).

[9] Van de gaer, Schokkaert and Martinez (1998) evaluate to what extent traditional measures of intergenerational mobility conform to ideas of equality of opportunity.

[10] Throughout our analysis a Gaussian kernel is assumed.

[11] We experimented with a range of bandwidths and obtained similar results to those presented in the paper.

[12] We have also carried out our estimation using more general bivariate normal kernels. The estimates obtained using these kernels were similar to those presented in the paper.

[13] A similar approach has been used by Quah (1996) to analyse income convergence across countries and by Trede (1998) to examine income mobility over time for an individual.

[14] We calibrate the distributions such that the mean of the log of sons' and fathers' incomes and the standard deviation of the log of their incomes are equal to those in our NLS sample.

[15] An exception to this is the work by Corak and Heisz (1999) who have approximately 400,000 observations in their data.

[16] The computer programmes used to estimate these opportunity sets were written for Stata and are available from the authors upon request.

[17] We follow Zimmerman and define fully employed as those working on average 30 hours a week for at least 30 weeks a year.

[18] The correlation between the log of fathers and sons income for this data set is .41, which is similar to those obtained by Zimmerman using average NLS data.

[19] The fact that the slope of the contours are less than 1 is consistent with the pattern of regression towards the mean found in previous studies.

[20] In our sample, 96% of the observations fall in this range. The contours estimated outside that range are determined with a low degree of precision.

[21] We could also use our non-parametric estimates of the conditional density to test the linearity assumption prevalent in previous work on intergenerational mobility. We do not carry out such an analysis in this paper, choosing instead to concentrate on estimating the opportunity sets. For a related discussion on non-linearities in the intergenerational mobility process see Minicozzi (1997) and Corak and Heisz (1999).