NEIGHBOURHOOD EFFECTS IN ECONOMICS: A REVIEW

Dr Patricia Rice

University of Oxford

1. INTRODUCTION

There is a considerable body of literature extending across the social sciences relating to the role of social interaction in determining individual behaviour. It is suggested that individual decision-making across a wide range of activities is significantly influenced by the behaviour of those with whom they have frequent contact. As a result, the composition of the groups within which individuals work and live play an important role in determining individual outcomes and life chances. The range of behaviours considered is diverse – criminal activity; education participation; economic activity; smoking; sexual behaviour are among the activities for which individual choices are thought to be influenced by the degree to which other members of the their social network/neighbourhood engage in the activities.

The mechanisms by which social interaction modifies individual behaviour are seldom very clearly articulated in the literature. Much of the discussion is couched in terms of either (i) role model effects whereby the behaviour of the individual is influenced by characteristics and/or earlier behaviour of older members of the neighbourhood; or (ii) peer group effects which refer to contemporaneous behavioural influences among members of the neighbourhood and as such may be reciprocal. Role model and/or peer group effects give rise to imitative behaviour, either contemporaneous or across age cohorts. This imitative behaviour may be due to:

(i) psychological factors – the intrinsic desire to behave like others;

(ii) interdependences in constraints that individuals face, so that the costs – real or psychic - of a given behaviour depend on whether others behave in the same way. E.g. the opportunity costs of employment are higher and hence the costs of being unemployed are lower, when a large proportion of your social network are unemployed.

(iii) interdependencies in information transmission so that the behaviour of others provides the individual with information on the effects of such behaviour. E.g. the greater the proportion of the neighbourhood attending higher education, the better (less uncertain) the individual's estimate of the benefits of higher education.

These types of neighbourhood effects have attracted interest in that they provide a way of understanding the persistence of particular patterns of behaviour within communities. Role model effects - intertemporal social interactions – create the possibility of hysteresis in economic outcomes. If the decision to proceed to higher education is strongly and positively related to the percentage of the older age cohort attending higher education, then is possible that two neighbourhoods, one with high rate of participation in higher education and the other with a low rate of participation, converge to different rates of higher education participation in steady state. The existence of strong contemporaneous interdependencies - peer group effects – gives rise to the possibility of multiple equilibria. In each equilibrium, the behaviours are mutually reinforcing and individually rational, but the equilibria may not be equally desirable from a welfare perspective. The actual outcome that is observed may depend on a variety of factors: history, reactions to common shock etc.

The existence of neighbourhood effects has a number of policy implications. First, significant neighbourhood effects amplify the effects of changes to private incentives. Consider a policy such as Educational Maintenance Allowances designed to increase the incentives to continue in full-time further education. Significant neighbourhood effects imply that such a policy leads to higher participation rates among as non-recipients in the neighbourhood as well as direct recipients. Second, neighbourhood effects represent a straightforward example of a consumption externality. Individual pay-offs are directly affected by the choices of other individuals and these spillover effects are not fully reflected in market prices. Given this there is no presumption the equilibrium allocation of families across neighbourhoods is efficient, and policies designed to change the allocation of families neighbourhoods may increase welfare.

In what follows I provide an overview of neighbourhood effects within the economics literature. Throughout the next section I will use the term 'neighbourhood' to refer to a stable grouping of individuals within which there is a significant level of social interaction due to their close geographical/social proximity. This neighbourhood may take the form of a residential neighbourhood; classroom/school, workplace, or a social or ethnic grouping. A comprehensive survey is provided by Durlauf (2004) and this discussion draws heavily on this source; updated to include more recent literature relating to the UK context. I start by outlining the theoretical analysis of neighbourhood effects and the implications for individual-behaviour and for neighbourhood formation. I then consider some of the issues that arise in deriving robust estimates of neighbourhood effects before reviewing the main empirical findings.

2. THEORETICAL ANALYSIS OF NEIGHBOURHOOD EFFECTS

The theoretical work on neighbourhood effects focuses on two basic questions:

(a) How do the characteristics of neighbourhoods affect behaviour of individual members and in the aggregate?

(b) How does the existence of neighbourhood effects influence the membership of neighbourhoods?

2.1 NEIGHBOURHOOD EFFECTS AND INDIVIDUAL DECISION-MAKING:

Role model and peer group effects on individual behaviour may be captured by assuming that the individual's behavioural choices depend on neighbourhood factors in two distinct ways. The first is through a set of predetermined characteristics of neighbourhood members referred to as contextural effects (Manski, 1993[#]). Examples of contextural variables are income levels, occupational group, education, ethnic group. The second neighbourhood influence comes through the individual's beliefs about the behavioural choices of other members of the neighbourhood (referred to by Manski as endogenous effects).

More formally, consider I individuals who are members of a common neighbourhood, n . Each individual makes a choice *y(i)* from some set of all possible behaviours. Individual choices are assumed to maximise a pay-off (utility) function which depends on both individual-specific and neighbourhood-specific factors as follows:

- a set off individual-specific characteristics: e.g. age, gender, income;
- a set of predetermined characteristics of the other members of the neighbourhood (contextural effects): ethnicity, income, education, social class;

- the (subjective) beliefs of individual i about the contemporaneous choices *y(n, -i)* of other members of their neighbourhood (endogenous effects).

The benchmark model assumes that the individual's beliefs are rational in the sense that in equilibrium, subjective beliefs about behaviour are consistent with actual behaviour. The model is greatly simplified by assuming that the pay-off functions depends only on the expected value (i.e. the average) of the choices of other neighbourhood members (e.g. Brock and Durlauf, 2001[#]). More complex models consider the distribution of behaviours within the neighbourhood. For example, we may consider that the effect on individual i of the contemporaneous behaviour of others in her neighbourhood varies with proximity or frequency of contact.

The implications of neighbourhood effects for individual outcomes depend critically on what is assumed about the shape of the pay-off function with respect to the endogenous effects. In general, it is assumed that individual pay-offs exhibit "complementarity" – individual i's pay-off from a higher level of *y(i)* increases with the choice levels *y(n, -i)* of others in the neighbourhood. In other words the marginal benefit of *y(i)* is greater, the higher is the average choice level of other neighbourhood members. This assumption creates an incentive for members of a neighbourhood to behave similarly. When complementarities are strong enough, they can give rise to multiple equilibria such that two neighbourhoods with identical compositions can exhibit different behaviours.

There is a body of work which boorows heavily from the statistical mechanics literature in physics to model different forms of neighbourhood structures in which pay-off interdependencies vary with proximity (see Arthur, Durlauf and Lane (eds), 1997[#]). In these models, individuals are considered to be arrayed in some space in which distances between individuals may be defined, and the endogenous effects on individual pay-offs vary with these distances. These models suggest that the properties of outcomes are not very sensitive to the particular form of interaction structure - different interaction structures result in broadly similar aggregate behaviours.

Finally, there is a growing literature analysing social interactions explicitly in terms of information flows across social networks. A social network is a collection of active bilateral relationships and the pay-off function for any individual is assumed to depend on the behaviour of those members of the network whom she is has a direct bilateral relationship. Calvo-Armengol and Jackson (2004) use this framework to analyse peer effects on education and employment choices; Ioannides and Datcher Loury (2004) analyse job information networks.

2.2 Equilibrium Neighbourhood Configurations:

Given the presence of neighbourhood effects in individual pay-off functions, location decisions depend on beliefs about the equilibrium composition of neighbourhoods. What are the implications of this for neighbourhood formation? More specifically, how are individuals with different attributes allocated across neighbourhoods in equilibrium? It is this allocation that determines the extent to which neighbourhoods vary with respect to their observable characteristics (i.e. contextural effects). The most common characteristic studied in the literature is income, although other attributes – education, ethnicity - have been considered.

Most of the interest has been on the extent to which neighbourhoods are stratified with respect to income or some other characteristic, where stratification with respect to characteristic x implies that range of values of x within neighbourhoods do not overlap. In general neighbourhood effects per se are not sufficient to produce to stratification. Equilibrium stratification with respect to a characteristic x (e.g. education) requires that an individual's willingness to pay for a higher average x in the neighbourhood is positively correlated with the individual's level of x. In other words, more highly educated individuals are willing to pay more to live in neighbourhoods with a higher average education. In general, this condition will hold where there is strong complementarity in preferences.

The results on the existence and stability of stratified equilibria appear robust to a wide range of alternative specifications in terms of neighbourhood structure - number of neighbourhoods; size of neighbourhoods; the degree of heterogeneity of individuals; the mechanisms determining neighbourhood membership. However, stratification of neighbourhoods does depend critically on the assumption that individuals/ families are distinguished along a single dimension – income or education or race.

2.3 Equilibrium Neighbourhood Allocations, efficiency and Policy.

As noted in the Introduction, neighbourhood effects represent a straightforward example of a consumption externality. Individual pay-offs are directly affected by the choices of other individuals and these spillover effects are not fully reflected in market prices. Given this there is no presumption the equilibrium allocation of families across neighbourhoods is efficient. Equilibrium allocations may have a level of stratification that is too high or too low depending on the precise shape of the pay-off functions. Benabou (1996)[#] highlights the importance of the trade-off between complementarity in preferences and diminishing returns in the pay-off function. Consider a simple case where the individual pay-off to an activity – say education - depends only on their own choice level *y(i)* and the average choice level of the other members of the neighbourhood, *y(n,-i)*. Complementarity - the marginal benefit of *y(i)* is greater, the higher is the average choice level in the neighbourhood *y(n,-i)* – implies efficiency gains to stratification. However, diminishing returns – the marginal benefit of *y(n,-i)* declines as *y(n,-i)* increases – favours a more even distribution of *y(n,-i*) across neighbourhoods. Essentially the balance between these two factors determines whether net efficiency gain to reallocating individuals across neighbourhoods to produce a more even distribution of average *y.* Finally it should be noted that the above argument relates to efficiency only and does not take account of equity arguments for such reallocations.

3. ESTIMATION OF NEIGHBOURHOODS EFFECTS:

Evidence on neighbourhood effects comes from two main sources: estimation of econometric models using observational data, and a small number of quasi-experiments consisting of government interventions into the residential choices of individuals and families. Before reviewing the main findings we review the measurement and econometric issues that arise in estimation.

3.1 Definition of Neighbourhood

There is little systematic analysis of what constitutes a 'neighbourhood' for these purposes. Most empirical studies define neighbourhoods in terms of physical space although it is recognised that two individuals in the same physical space may experience very different social interactions depending on education, class ethnicity etc. Aizer and Currie (2002)[#] allow residential proximity and ethnicity to jointly define relevant 'neighbourhoods' Conley and Topa (2002)[#] considered alternative measures of proximity based on physical distance; travel time distance, racial and ethnic distance; occupational distance. They conclude that ethnic distance is the most salient dimension along which neighbourhoods exhibit correlation.

For the vast majority of empirical studies, the neighbourhood structure is determined by the constraints imposed by the available data. Most empirical work is restricted to considering a single pre-defined neighbourhood structure, often based on administrative units – local education authority, postcode sector, census tract, classroom, school. Very often the scale considered is much greater than that suggested by theories of social interaction, and the scope for testing the robustness of estimates to the definition of neighbourhood are very limited. Two recent UK studies, Buck (2001) and Bolster et al (2007), have investigated the appropriate spatial scale for residential neighbourhood effects using data from the British Household Panel Studies (BHPS). They conclude that in general

area effects decline with spatial scale, with the strongest effects at the level of the Census enumeration district (approx. 500 residents); much smaller than the scale considered in the majority of studies. Further, conditioning on neighbourhood effects at the very local level, there is no evidence of significant neighbourhood effects at larger spatial scale; in other words no evidence of spillover effects across neighbourhoods.

3.2 ECONOMETRICS ISSUES

Consider a basic regression model incorporating neighbourhood effects based on the model described in 2.1.

$$y_i = \beta X_i + \alpha Z_{n(i)} + \lambda m_{n(i)} + \varepsilon_i \qquad (1)$$

Where y_i is the outcome of interest for individual i, for example the level of educational attainment; X_i is a vector of r observable personal characteristics of individual i; $Z_{n(i)}$ is a vector of s observable characteristics of individual i's neighbourhood n; and $m_{n(i)}$ denotes the expected value of y for members of the neighbourhood n(i). The parameters α and λ capture the contextural and endogenous neighbourhood effects respectively.

This model is referred to in the literature as the linear-in-means model. The endogenous effects are determined by the expected (average) behaviour in the neighbourhood only. This assumption is appropriate in larger groups where individuals take the behaviour of other members as given. In small groups, additional complications can arise from the possibility of strategic interaction whereby i anticipates the effects of her behaviour on other members of the group. This model forms the basis for much of the empirical literature for the US and the UK.

IDENTIFICATION

For the moment, we assume that the classical linear regression model assumption $E(\varepsilon_i \mid X_i , Z_{n(i)}) = 0$ is satisfied. Assuming beliefs are rational then averaging over the *y* in the neighbourhood and taking expectations in equation (1) we have

$$m_{n(i)} = (\beta X_{n(i)} + \alpha Z_{n(i)})/(1-\lambda) \qquad (2)$$

where $X_{n(i)}$ denotes the average values of the X variables in the neighbourhood. If we substitute (2) into (1) then we obtain a reduced-form estimating equation that depends only on observables

$$y_i = \beta X_i + aZ_{n(i)} + cX_{n(i)} + \varepsilon_i$$

$$where \quad a = \alpha/(1-\lambda) \quad and \quad c = \lambda\beta/(1-\lambda) \qquad (3)$$

The identification problem comes down to whether unique estimates of the structural parameters α, β and λ can be derived from estimates of the reduced-form coefficients a, b and c. At first sight this appears straightforward; if r=1, the model is exactly identified, otherwise there is over-identification. However, this ignores the possibility of overlap between the individual observable characteristics (X) and the observable neighbourhood characteristics (Z). We might expect these to include elements in common – income, ethnic group etc. Suppose that they include same set of characteristics so $X_{n(i)} = Z_{n(i)}$ then identification fails and we cannot disentangle α and λ. This is what is referred to in the literature as the 'reflection problem' – it is not possible to distinguish the direct effect of the variables Z on individual behaviour from their indirect effects as reflected through individual i's beliefs i.e. $m_{n(i)}$.

In so far as policy interventions are based on influencing behaviour by changing the $Z_{n(i)}$, it is important for policy evaluation that we are able to distinguish the direct and indirect effects. In general, identification in the 'linear in means' model requires that there is at least one individual characteristic (X) that affects outcomes but is not included as a neighbourhood contextural variable (Z).

Most empirical studies ignore the distinction between contextural and endogenous effects, and simply work with an estimating equation which includes a set of individual-specific observable characteristics and a set of neighbourhood-specific observable characteristics. This raises a number of issues for interpretation of parameter estimates unless it is the case that no endogenous effects ($\lambda = 0$). The reduced-form estimating equation (3) implies that all the individual variables should be included also at the neighbourhood level. If this is not the case and endogenous neighbourhood effects are significant ($\lambda \neq 0$) then the estimates of (3) will be subject to omitted variable bias. Further, suppose there is a variable, say education, that exerts an influence at both the individual level (included in the X) and as a contextural effect (included in the Z). In such cases, the estimated coefficient on the neighbourhood education variable in the reduced form equation (3) is equal to (a+c) and therefore is not a 'neighbourhood effect' but a combination of contextural effects (α), endogenous effects (λ), and individual effects (β).

It is worth noting that the identification problem arises because of linearity of model – $m_{n\,(i)}$ is a linear function of the X and Z. The problem does not arise if the endogenous effects are determined by moments of the distribution of the y other than just the expected value. Moreover, in general the 'reflection' problem is avoided in non-linear models such as discrete choice models – probit, multinomial logit - or in duration type models.

SELECTION

As we discussed above, neighbourhood composition is endogenously determined through individual location decisions. Neighbourhood choice depends in part on individual unobserved heterogeneity ε_i and hence, ε_i and $Z_{n(i)}$ are correlated and $E(\varepsilon_i \mid X_i , Z_{n(i)}) \neq 0$. This selection bias implies that OLS estimates of the parameters of (3) are biased and inconsistent.

There are two standard remedies to the selection problem. The first is to obtain consistent parameter estimates using the a selection-correction procedure as proposed by Heckman and others. This requires that we model neighbourhood selection explicitly (see for example Ioannides and Zabel). The usual concerns regarding these techniques apply here – namely that the selection-correction involves strong parametric assumptions on the distributions of unobservables that cannot be tested.

An alternative approach is to use an Instrumental Variables method. This requires that we identify variables that are highly correlated with the observables characteristics $Z_{n(i)}$ but uncorrelated with the unobserved individual heterogeneity. As always the difficulty is in identifying valid instruments ad avoiding the problems of weak instruments.

UNOBSERVED HETEROGENEITY – INDIVIDUAL, FAMILY AND NEIGHBOURHOOD EFFECTS

The selection issue is one aspect of a more general problem in these models of controlling for unobservable effects at the individual, family and neighbourhood level that are likely to be correlated with the observables variables in the estimation equation, which results in bias and inconsistent parameter estimates.

A number of solutions have been proposed depending on data availability. For example, individual panel data for two or more points in time when the individual is resident in the same location allows one to control for unobserved fixed effects at the individual and neighbourhood levels..

In a similar manner, data on siblings allow one to control for family fixed effects. Consider the basic estimating equation (3).

$$y_{s,j} = \beta X_{s,j} + \alpha Z_{n(j)} + cX_{n(j)} + \varepsilon_j + u_{s,j}$$

Where s denotes sibling and j denotes family and ε_j denotes the unobserved family fixed effects. By considering the differences in outcomes between siblings, we can control for the family fixed effects. It shouldn't be assumed that in controlling for family fixed effects, we resolve the problems associated with the selection of families into neighbourhoods. This is the case only if it can be assumed neighbourhood selection has the same impact on all family members, in other words if it is a fixed effect. However, this is not likely to be the case if siblings are raised in the neighbourhood at different times.

Finally sibling data has been used for non-parametric analysis of the relative importance of family and neighbourhood influences on individual outcomes. This is based on a comparison of the covariance of outcomes for siblings in same neighbourhood is compared with the covariance of outcomes for a pair of unrelated individuals in the same neighbourhoods. This approach has the advantage of avoiding the need to make strong parametric assumptions that cannot be tested, but it relies on the assumption that differential treatment of siblings with families is fully accounted for by random error term.

4. EMPIRICAL FINDINGS

4.1 ECONOMETRICS STUDIES WITH OBSERVATIONAL DATA ON INDIVIDUALS

In what follows, I summarize the main findings from an extensive US and UK literature. For a more detailed summary of the results see Table 2 of Durlauf (2004). These studies are largely based on the 'linear-in-means' model described above, but vary considerably with respect to the individual outcomes considered, definitions of the neighbourhoods, the neighbourhood variables used and estimation methods.

(a) The majority of empirical studies report statistically significant neighbourhood effects although this may reflect in part a publication bias against negative results. In general, neighbourhood effects are found to be relatively small in magnitude when compared to the influence of individual/family factors.

This finding is confirmed by the results of correlation studies using sibling data to compare the relative importance of family background and neighbourhood on outcomes. Duncan et al (2001)# study correlations in education achievement and delinquent behaviour between siblings, friends, classmates and neighbours and find that sibling influences are much stronger than the others, and friendship effects dominate those for school or neighbours.

(b) The range of neighbourhood variables selected for the empirical analysis is very wide, and the choice is seldom motivated by theory. It is difficult to relate the empirical findings to specific microeconomic mechanisms and consequently neighbourhood effects remain something of a black box.

(c) There is relatively little systematic analysis of robustness of findings to model specification. Where this has been investigated, there is evidence that the magnitude and significance of neighbourhood effects is sensitive to choice of individual level controls (e.g.Ginther Haverman and Wolfe 2000#, McCulloch, 2001).

(d) Few studies address the distinction between endogenous and contextural effects and the implications for estimation. The majority of studies ignore the issue and simply work with an estimating equation which includes a set of individual-specific observable characteristics and a set of neighbourhood-specific observable characteristics. As explained in (3.2), this poses problems for the interpretation of the reduced form estimates unless, it is the case that there are no endogenous neighbourhood effects, only contextural effects. The small number of studies that have addressed the distinction between contextural and endogenous effects have found that endogenous effects are statistically significant (e.g. Ioannides; 2002; Ioannides and Zabel, 2002[#]). Moreover, when proper account is taken of endogenous neighbourhood effects, the results on contextural neighbourhood effects are more mixed – statistically significant in some cases, not in others.

(e) The evidence on the robustness of neighbourhood effects to controlling for unobserved heterogeneity at the neighbourhood or family level is mixed. Using sibling data to investigate neighbourhood effects on education outcomes, Aaronson (1998[#]) found neighbourhood effects remained significant having controlled for family fixed effects; a similar study by Plotnick and Hoffman (1999[#]) found this not to be the case.

(f) The robustness of neighbourhood effects to controlling for self selection appears to depend on method used. Among the studies that use IV to attempt to control for selection bias, many – but not all - report IV estimates of neighbourhood effects that are far smaller in magnitude than the OLS counterparts, and in some cases, no longer statistically significant (e.g. Evans at al, 1992[#]). Few studies that have explicitly corrected for selection using Heckman-type methods, but one such, Ioannides and Zabel (2002)[#] , found strong evidence of neighbourhoods effects on housing demand having controlled for selection.

4.2 STUDIES USING (QUASI) EXPERIMENTAL DATA

Given the many interrelated econometric issues that arise in estimating neighbourhood effects from observational data, it is not surprising that considerable attention has been give to the findings of the small number of (quasi) experimental studies involving direct government interventions in the residential decisions of individuals.

The Gautreaux programme of 1967 assigned of poor families in Chicago public housing in poor neighbourhoods to (i) alternative public housing in Chicago and (ii) suburban communities outside Chicago. Rosenbaum (1995[#]) reports substantially better socio-economic outcomes for those living in the suburban communities. The differences are found to be particularly pronounced with respect to outcomes for children; for example 54 percent of the children of families who moved to the suburbs attended college, as compared with a figure of 21 percent for those who remained in public housing in the city.

The Moving to Opportunity (MTO) programme was implemented in five US cities – Baltimore, Boston, Chicago, Los Angeles and New York – from 1994. This programme provides housing subsidies to a randomly selected group of families. Families within the subsidized group are randomly assigned to one of two types of support: (i) housing subsidies and professional support to move to neighbourhoods where poverty rates are below 10% (the experimental group); (ii) unrestricted housing vouchers (section 8 group). A comparison of outcomes for the experimental and section 8 groups provide the impact of the neighbourhood reallocation on behaviour, as distinct from the income effect of a housing subsidy.

A number of studies have evaluated the effects of MTO across the different cities. On balance, these find evidence of gains for children in both groups – experimental and section 8 groups - particularly for girls. Katz et al (2001[#]) study of Boston families report a substantial reduction in behavioural

problems for children in both groups, with children in experimental group exhibiting fewer health problems. However, more recent follow-up studies (Kling et al 2005, 2007) found evidence that social and behavioural problems were higher among young males in the experimental group. Similar results on improvement in behavioural problems were found in the case of NY families but no evidence on differences in health outcomes. Ludwig et al (2001)[#] report evidence of a reduced incidence of juvenile crime for among Baltimore families. In contrast to child outcomes, the evidence a significant impact on adult outcomes is slight. Some evidence of improved mental health outcomes for women. However, Kling et al (2004) could find no significant effects on economic outcomes – adult employment, earnings, welfare receipt.

As Durlauf (2004) discusses, the evidence from quasi-experimental data is not without its problems, particular with regard to selection issues. In the case of the Gautreaux programme, the basis on which families were relocated is unclear, and the sample studied includes only families that moved to the suburbs and remained. Similarly with the MTO programme, only families who thought they might benefit volunteered, and moreover, not all the families in the programme made use of the vouchers. The lack of evidence of substantial differences in outcomes between the section 8 and the experimental groups in the MTO programme suggest that the income effects of voucher eligibility is as important for outcomes as the effects of changing neighbourhood. The MTO programme provides little evidence on what aspects of neighbourhoods are responsible for improvements in outcomes. For example, the improvement in health outcomes may be due to better quality housing rather than anything to do with the neighbourhood. Finally there is the important question of the general equilibrium effects of policies of this type. Moving large numbers of poor families to affluent neighbourhoods may be expected to induce general equilibrium effects in terms of location decisions of other families etc. Thus it is far from clear that the observed effects of the MTO programme would be replicated if implemented on a large scale.

4.3 Recent UK Studies

A number of recent studies for the UK have combined longitudinal data from the BHPS with census data on neighbourhood characteristics - Buck (2001); McCulloch (2001), Bolster et al (2007). The first two of these – Buck and McCulloch – examine the relationship between a ward-level index of deprivation and a range of economic, health and social outcomes for individuals. These studies treat the BHPS sample as a pooled-cross section and little attempt is made to systematically address the issues unobserved heterogeneity and selection bias. Both studies find that the inclusion of individual and household observable characteristics attenuates the neighbourhood effects.

Bolster et al (2007) address the question of how the characteristics of an individual's current neighbourhood impact on their income trajectories over one, five and ten year time horizon. By focusing on changes in income, rather than the level, this study mitigates many of the estimation problems associated with sample selection and other sources of unobserved heterogeneity. The fact that neighbourhoods are poor because poor households locate within them implies a positive correlation between the average income level of the neighbourhood and the income level of the individual household, but it has no implications for the relationship between the average income level of the neighbourhood and changes in household income. Bolster et al examine the robustness of the estimates to varying the spatial scale of neighbourhoods and to possible measurement errors. Their overall conclusion is that there is little robust evidence that neighbourhood disadvantage has a significant impact on subsequent income growth.

A second strand of recent UK literature focuses on classroom or school effects on pupil achievement. A number of these studies make use of the rich source of longitudinal provided by UK schools administrative data which allows controls for unobservable individual, family and school fixed effects (see Gibbons and Telhaj (2008) for a recent example). The conclusions here are similar to above; namely, the effect of the peer group on the achievement of an individual pupil is small and in many cases not statistically significant.

5. CONCLUSIONS

In this short paper I have summarized the main issues that arise in identifying and estimating area or neighbourhood effects on individual outcomes. The robust estimation of neighbourhood effects from observational studies presents the researcher with a number of interrelated problems arising from issues of identification, selection and unobserved heterogeneity. The approaches taken to resolving these estimation issues are diverse, and yield mixed results. However, on balance, the evidence from this source indicates that neighbourhood effects, where statistically significant, are not large in magnitude in comparison with the impact of individual and family/household characteristics. This conclusion is supported by the evidence from the small number of (quasi) experimental studies. On the basis of this body of evidence, it is difficult to make a strong case for area-based anti-poverty policies in preference to policies focused on individuals or families. Area-based policies such as regeneration schemes benefit individuals and households through the provision of jobs and income, but the spill-over effects on households in the neighbourhood not directly affected are expected to be small.

The review of the empirical literature highlights a number of issues that merit further attention. The first is the question of the definition and measurement of neighbourhood for these purposes. Most of the empirical literature is based on pre-defined geographical units, often at a much larger scale than suggested by considerations of social interaction. The few studies that have investigated the question of spatial scale suggest that the magnitude of neighbourhood effects decline with scale and this suggests that further work at the very local level may yield more robust estimates.

Much of the evidence base is drawn from estimation of reduced-form 'linear-in-means' type models that assume that neighbourhood effects are the same for all members of the group and may be captured by considering the average value of some set of observables neighbourhood characteristics. Such models have obvious advantages in terms of availability of the necessary data and estimation but they fail to capture the essence of social interaction in which pay-off interdependences vary with the nature and frequency of contact. This points to the need for further empirical research based on more structural models that allow for a degree of heterogeneity in neighbourhood effects

REFERENCES

Full details of references marked # may be found in Durlauf (2004)

Bolster, A, Burgess, S., Johnston, R., Jones, K., Propper, C., and R. Sarker (2007), "Neighbourhoods, households and income dynamics: a semi-paramteric investigation of neighbourhood effects", Journal of Economic Geography, 7, pp 1-38.

Buck, N. (2001) "Identifying neighbourhood effects on social exclusion". Urban Studies, 38: 2251–2275.

Calvo-Armengol (2006) "Social networks and labour market outcomes" *Els Opuscles del CREI* , 17, January.

Durlauf, S. (2004) Neighborhood effects. Handbook of Regional and Urban Economics. In J. V. Henderson and J.-F. Thisse (eds) vol. 4, pp 2175 -2242.

Gibbins, S and S, Telhaj, "Peers and achievement in England's secondary schools" SERC Discussion Paper 1, July 2008.

Kling, J. R., Liebman, J., Katz, L., Sanbonmatsu, L. (2004) Moving to opportunity and tranquility: neighborhood effects on adult economic self-sufficiency and health from a randomized housing voucher experiment. Princeton IRS Working Paper 481.

Liebman, J, Katz, L. Kling, J. (2004) Beyond treatment effects: estimating the relationship between neighbourhood poverty and individual outcomes in the MTO Experiment Working Paper 493, Industrial Relations Section, Princeton.

McCulloch, A. (2001) Ward-level deprivation and individual social and economic outcomes in the British household panel study. Environment and Planning A, 33: 667–684.

Dr Patricia Rice
Department of Economics
University of Oxford

19th September 2008