

Social class returns to higher education: comments on a paper by Bukodi and Goldthorpe with a response from the authors

Harvey Goldstein

University of Bristol

In a recent paper - Bukodi E and Goldthorpe JH. (2011) 'Social class returns to higher education: chances of access to the professional and managerial salariat for men in three British cohorts'. *Longitudinal and Life Course Studies*, 2, 185-201, 2011 - the authors adopt a particular methodological approach to the study of 'causality' in life course research. This highlights an important area of contention among social scientists. We have therefore, with the authors' agreement, taken this as an opportunity to open it up to a wider debate. In this issue of the journal we publish two commentaries, by Clarke, and by Legewie and Solga, that take issue with the Bukodi/Goldthorpe position on causality. The commentaries also question the relationship of the paper's contribution to the existing economics literature on 'returns to education' and the manner in which 'missing data values' are handled by the authors. See also a relevant tutorial published in the journal - Goldstein H. (2009) 'Handling attrition and non-response in longitudinal data'. *Longitudinal and Life Course Studies*, 1, 63-72, 2009 - that discusses methods for dealing with missing data. We also publish a response by Bukodi and Goldthorpe to these commentaries.

We are very grateful to the original authors as well as to those who contributed the commentaries, for taking the time to prepare their carefully argued positions. We hope that this kind of debate will form a regular feature of the journal and help to illuminate important methodological controversies, about which, of course, readers will form their own judgements.

Commentary by

Paul Clarke

Centre for Market & Public Organisation

University of Bristol

paul.clarke@bristol.ac.uk

The authors take a refreshing look at the well-worn subject of assessing returns from education. In doing so, they raise a number of interesting issues en route to justifying their methodological approach. The purpose of this commentary is not to criticise the authors' article, but instead to focus on some of the issues they raise that I believe are of wider interest to the readers of *Longitudinal and Life Course Studies*.

Causal and descriptive analyses

The first issues are those surrounding the authors' decision to eschew causal analysis in favour of a "descriptive" interpretation. To some extent, this is wise and understandable, because it is difficult to justify that a statistical analysis of observational data leads to estimates that can be

interpreted causally. On the other hand, a causal interpretation is desirable because it implies that the effect of education on salariat membership has been isolated from that of other pathways *not* driven by educational attainment. A causal interpretation is also necessary to influence policy makers, because such estimates approximate what we would expect to happen to people's salariat membership by improving their educational attainment.

A strictly descriptive interpretation provides no such information for understanding the social process or for policy: it merely states whether groups of individuals can be distinguished by their average outcomes. Perhaps the general tendency to play safe and avoid causality goes back to R.A. Fisher's remark that "statistics cannot prove causality", which is echoed by mathematical

statisticians like Philip Dawid. Dawid argues strongly that scientists should concern themselves only with “the effect of causes”, where the ‘cause’ itself is controlled by the experimenter. The alternative is described as looking for “the cause of effects”, namely, establishing which factors led to the observed outcomes taking those particular values. However, this is criticised as relying on “meta-physical” assumptions that cannot be tested on the study data (Dawid, 2000).

It is ironic that Fisher’s remark was made around the time that evidence of the link between smoking and lung cancer started to emerge. The success of epidemiologists in establishing this link has led today to a situation where most applied scientists would agree that such objections are too limiting for practice. Scientists ultimately seek definitive proof, but work towards it by gathering evidence to support their hypotheses as part of an on-going dialectic: some hypotheses will gain support as evidence accumulates, and others will fail. In this sense, the causal interpretation of any estimate has validity as far as one believes the initial hypothesis, and the extent to which competing explanations have been accounted for. As such, the authors’ analysis is somewhere between being descriptive and being causal, like most analyses carried out in quantitative social and biomedical research. After all, why would we be interested in reading a paper that reported a series of spurious correlations?

There is much mystique about “causal” analysis, but it is in fact based on a very simple idea. The most popular way of thinking about causality is based on the so-called Neyman-Rubin model, under which each individual has a set of potential outcomes; this is the approach widely used for the analysis of clinical trials in medical statistics. Under this model, an *individual* causal effect is the impact of a particular exposure on the potential outcomes of that individual. An example of an individual causal effect here is the difference between his/her (potential) salariat memberships after being exposed to different levels of education, say ‘high’ and ‘low’. The actual causal effect of an exposure on any individual is thus inherently unobservable, because it is a comparison between two alternative experimental scenarios for the same individual occurring at the same point in time and space; the *only* difference between these scenarios is in the level of educational attainment.ⁱ

However, the fact that such effects are unobservable does not prevent estimation of *average* causal effects, which is where randomisation comes in: we know there are no systematic differences between the two randomised groups prior to exposure, and so the difference between the two groups’ average responses (e.g. for low and high educational attainment) is due entirely to the exposure and estimates the average causal effect, namely, the average of the unobservable individual causal effects for the target population.

The authors say on page 186 that: “This understanding of causation does, however, give rise to problems in that education is to a significant degree a matter of choice rather than simply a ‘treatment’ that is received [...] this choice is likely to be influenced by factors that may have their own direct effects on earnings.”

I see this as a comment on the historical relationship between causal effects and randomised experiments, but this link does not in itself rule out a causal approach. While prominent authors have also argued that a prerequisite for causal analysis is that a study can at least be *conceived* of as an experiment (e.g. Rubin 2008), the experiment does not have to involve either randomisation or a treatment-type intervention.

Rosenbaum and Rubin (1983) eloquently showed that the key to causal analysis is to model the exposure selection process.ⁱⁱ However, selection is a nebulous term that describes any form of non-constrained individual choice, assignment or institutional selection, and it does not have to involve randomisation. The temporal ordering of events means that the selection model can depend only on factors determined prior to selection. The difficult part is that, here and more widely, the mechanisms by which exposure is selected are rarely well understood. Theory and/or prior information must be used to identify the key pre-attainment factors most strongly associated with exposure selection, and one would not expect theory to identify all of these factors, or for all of these factors even to be measured.

The authors do include other variables in their analysis, but write:

“We include measures of ability and of social class origins not simply as controls but because they too are of substantive interest to us. We wish to know how these factors are associated with chances of

access to the salariat, considered both independently and in interaction with education” (page 187).

In other words, the authors acknowledge that the other variables are acting as controls, but that their interest also concerns the distinct effects these variables have on salariat membership.

As such, it seems straightforward to characterise their analysis as estimating the effect of educational attainment on salariat membership, while adjusting for some obvious alternative explanations. Hence, while this analysis cannot be regarded as an accurate representation of what an experiment would produce (if indeed such an experiment were possible), it could be a useful approximation based on the authors’ sound understanding of past evidence and theory.

Causality in economics

I am neither an economist nor an econometrician, but in my experience few of them view causality solely in terms of experiments and treatments. For instance, within the discipline there is a ‘structural’ modelling tradition based on models for the social and economic processes of interest. The rationale behind this approach is that, if these processes can be modelled satisfactorily, structural models can be used to predict the causal effects of intervening and changing how these processes operate.

Structural models are based on economic theory and can represent complex interactions between non-linear processes over time; but structural models can also be simple. For example, the simple linear regression model

$$y = \alpha + \beta x + e \quad (1)$$

has a structural interpretation. This model (and its generalisation to multiple regression) remains a staple of contemporary applied economics. The structural interpretation of the regression coefficient β here is as what happens to an *individual’s* outcome if we intervene and change x by one unit while holding e constant, where e represents the combined effect of all other factors affecting the outcome. It is clear from this definition that a structural model is simply a model of the potential outcomes, and thus causal.

In practice, the structural coefficient β cannot be estimated by fitting model (1) to the observed data if e and x are associated - that is, x is ‘endogenous’ - which is exactly what happens if

there are factors associated with the outcome that are also associated with exposure selection. Economists thus go to considerable lengths to develop explicit and plausible ‘identification strategies’ that enable estimation of ‘structural’ β . Other quantitative scientists also employ identification strategies for structural models in all but name, by including confounding variables in the model. One of the fundamental differences between economists and other social scientists, however, is that the former are sceptical of whether adjusting for observed confounders (so-called selection on observables) can ever satisfactorily adjust for confounding bias, which is why they will often say that analyses like this are not causal. I do not agree.

Instead, economists often seek identification by using an ‘instrumental variable’ that predicts exposure, but is independent of e and affects y only through its effect on x . The advantage of this approach is that the analyst does not have to identify and adjust for all of the confounding factors, and models like (1) can be analysed using, for example, two-stage least-squares. But choosing a valid instrumental variable is notoriously difficult, and often leads to supporting arguments that can test the limits of credulity. However, there is currently considerable debate within the discipline about whether researchers in the past were too lax in their choices of instruments, and the role of structural models in identifying causal effects (e.g. Heckman and Vytlacil 2007; Angrist and Pischke 2008).

Handling missing data

The authors use the missing data indicator method to handle the problem of missing covariates, in which the covariate is extended to include an extra ‘missing’ category (see tables 2 and 3). Unfortunately, there is a problem with this approach: even if the data are missing *completely* at random (i.e. in a way that has nothing to do with the variables in the analysis) then the resulting estimates will generally be biased.

Greenland and Finkle (1995) demonstrate why this approach is problematic for linear regression models, but exactly the same arguments apply for other generalised linear models. To see why, consider a situation where only the educational attainment variable has missing values. The missing data indicator method involves fitting the model

$$\text{logit}\{\Pr(y_{ic} = 1|x_{ic}, z_{ic})\} = \alpha + \beta x_{ic}^* + \gamma z_{ic}$$

where x_{ic}^* is a series of dummy variables representing the effect of educational attainment with a separate category for 'missing', and z_{ic} is a fully observed measure of ability.

The problem here is that the group with missing values (i.e. those people for whom x_{ic}^* equals the missing category) comprises people who have different (but missing) educational attainments, and so the effect of ability (z_{ic}) cannot be estimated by holding educational attainment fixed, which is essential for estimating the regression coefficients of any generalised linear model. Only if educational

attainment and ability are completely independent of each other will the indicator method work, and this rarely holds in practice.

It is clearly tempting to use the missing data indicator method because it is simple and enables data from the incomplete cases to be kept in the analysis. However, it is difficult to justify a method that is biased even if data are missing completely at random. If one does not wish to drop incomplete cases, then methods for multiple imputation should always be preferred (e.g. Royston 2005; Carpenter et al 2011).

References

- Angrist JD and Pischke JS. (2008) *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.
- Dawid AP. (2000) Causal inference without counterfactuals (with discussion). *Journal of the American Statistical Association*, 95, 407-448. (<http://www.jstor.org/stable/2669377>)
- Carpenter J, Goldstein H and Kenward MG. (2011) REALCOM-IMPUTE software for multilevel multiple imputation with mixed response types. *Journal of Statistical Software* 43, 5. (www.jstatsoft.org)
- Greenland S and Finkle WD. (1995) A critical look at methods for handling missing data in epidemiologic regression analyses. *American Journal of Epidemiology*, 142, 1255-1264. (<http://aje.oxfordjournals.org/content/142/12/1255.abstract>)
- Heckman JJ and Vytlacil EJ. (2007) Causal models, structural models, and econometric policy evaluation. *Handbook of Econometrics*, 6, 4779-4874. ([http://dx.doi.org/10.1016/S1573-4412\(07\)06070-9](http://dx.doi.org/10.1016/S1573-4412(07)06070-9))
- Robins JM, Greenland S and Hu FC. (1999) Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome (with discussion). *Journal of the American Statistical Association*, 94, 687-712. (<http://www.jstor.org/stable/2669978>)
- Rosenbaum PR and Rubin DB. (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41-55. (<http://dx.doi.org/10.1093/biomet/70.1.41>)
- Royston P. (2005) Multiple imputation of missing values: an update of ice. *Stata Journal*, 5, 527-536. (http://www.stata-journal.com/article.html?article=st0067_2)
- Rubin DB. (2008) For objective causal inference, design trumps analysis. *Annals of Applied Statistics*, 2, 808-840. (<http://dx.doi.org/10.1214/08-AOAS187>)

Endnotes

ⁱ At least, individual causal effects are unobservable without further assumptions about the equivalence of the experimental conditions at different points in time; for example, cross-over trials in pharmaceutical research.

ⁱⁱ Both Rosenbaum and Rubin (1983) and Rubin (2008) actually advocate the use of propensity scores for causal analysis. I do not go into this approach here, which is required primarily to avoid unnecessary modelling assumptions, but I refer to the reader to either of these papers for further details.

Commentary by Joscha Legewie

Columbia University
Department of Sociology

Heike Solga

Head of Research Unit 'Skill Formation and
Labour Markets'
Social Science Research Center Berlin (WZB)

Introduction

Bukodi and Goldthorpe (2011) present an interesting analysis of the changing relation between education and social class position across three British cohorts from 1946, 1958, and 1970. In the introduction, the authors define their objectives: First, they intend to make a sociological contribution to the debate on returns to education by focusing on social class rather than on earnings. Second, they believe their focus on social class will help enrich the theoretical debate on returns to education by highlighting social relations and employment contracts as central aspects of economic life. Third, they dismiss the focus on causality, replacing it with what they call a “descriptive” approach—that is, examining the effects of education, ability, and social class origin in combination (and without controlling for confounding variables).

In our commentary, we argue that Bukodi and Goldthorpe’s paper is more of a *critique* than an *extension* of the (economic) returns to education literature. By dismissing causality, the paper essentially fails to contribute to the literature on educational returns because the concept of *returns* to education is inherently causal. On the other hand, we believe Bukodi and Goldthorpe make an important, albeit implicit, contribution by challenging the underlying ideas in the returns to education literature, namely that of free markets and that of education solely as an indicator of productivity.

Returns to education as “treatment” or “choice”

Bukodi and Goldthorpe analyze the social class returns to higher education using a “descriptive” approach based on logistic regressions and marginal estimates of parameters. They argue that returns to education from a causal perspective “give rise to problems, in that education is to a significant degree a matter of choice rather than simply a ‘treatment’ that is received, and also, in that this choice is likely to be influenced by factors that may have their own direct effects on earnings” (Bukodi and Goldthorpe 2011, 186). Whereas the authors rightly point out that education is also a matter of choice and influenced by

other factors, they fail to acknowledge that economists would wholeheartedly agree. In particular, the mainstream economists would say that education is a “choice” based on calculations about human capital investments and returns to these investments. Sociologists would, rather, argue that this choice is limited, given the fact that educational opportunities are socially structured and unequally distributed. The term “treatment” does not imply that there is no choice. It rather points at the question of randomization in experimental studies and causal inference. This leads us to the second part of our comment.

Returns to education as an inherently causal concept

The concept of returns to education is an inherently causal concept (Brand and Xie 2010). At its core is a counterfactual question about the differences or “gains” with respect to some outcome (usually earnings, or social class position in this case) under different levels of education: “How much more would Emma earn had she received three more years of education?” Any confounding variables—that is, variables that are causally prior to education and systematically related to both education and the outcome—undermine this concept. In particular, it becomes meaningless to speak of *returns* to education if the relation between education and the outcome can be attributed to confounding variables and not to education itself. Parents’ aspirations for their children, for example, might influence both the child’s education and his or her later occupational or social class position. As a consequence, part of the relation between education and social class might be due to the confounding influence of parents’ educational aspirations so that the effect is biased upwards.¹ While potentially interesting, this upward bias in the effect of education on social class challenges the interpretation of the relation as *returns* to education: Does it make sense to talk about *returns* to education if the actual reason for the effect is not education itself but some other confounding variable, such as parental aspirations for their children? Thus a

descriptive approach, as favoured by the authors, does not capture *returns to education*, if in the end a person's likelihood of obtaining a higher class position or higher earnings is not a consequence of their higher educational attainment, but of a different, causally prior attribute.

In addition, economists are generally willing to include factors such as cognitive ability, if available, as a proxy for individuals' competences. We have to acknowledge that in the 'returns to education' literature, education and educational degrees are taken as proxies of an individual's productivity, skill endowment, and (train)ability. As soon as economists have more detailed information (such as competence or ability measures), they do use this information in their analyses (see Hanushek et al 2010). The research on returns to early life course investments in human capital (or education, respectively) by Nobel Prize winner James Heckman, for example, shows that even non-cognitive skills (such as personality traits) are treated as part of human capital (see Heckman and Yona Rubinstein 2001). Accordingly, including ability measures is not an extension of the returns to education literature for the simple reason that economists are happy to investigate the relative impact of the cognitive and non-cognitive dimensions of human capital.

“Social class” as a (sociological) critique of the returns to education literature

By framing their paper—incorrectly, in our view—as an “extension” of the returns to education literature, Bukodi and Goldthorpe undermine the real contribution of their analysis to this line of research. We believe their true contribution is not simply to add a dependent variable (social class position instead of earnings) but to offer a lively *critique* of the returns to education literature. In particular, the authors' emphasis on the social class context of skill formation, and their sociological differentiation between education as certificate and education as development of cognitive abilities, undermine the underlying ideas in the returns to education literature, namely that of free markets and that of education solely as an indicator of productivity. Such a critique is neither at odds with the authors' actual research questions (p 187, as opposed to the way the paper is framed in the introduction) nor with their empirical analysis. In particular, we see an implicit, twofold critique as

the actual contribution of Bukodi and Goldthorpe's paper—a contribution, however, that the authors fail to make explicit.

First, the authors emphasize the importance of social origin and cognitive ability as additional, possibly interconnected factors for labour market outcomes. Examining these three factors in combination (education, class origin, and cognitive ability) and across three birth cohorts, brings the authors closer to a sociological perspective on skill formation (Müller and Jacob 2008). It points to the necessity to distinguish different dimensions of “education”—such as Bourdieu's or Collins's distinction between education as certificate (as institutionalized cultural capital or credential) or education as development of cognitive competences (as incorporated cultural capital or skills) (Bourdieu 1982; Collins 1971). Relating their findings to these sociologically informed concepts of skill formation would have yielded a much richer sociological contribution to the returns to education literature, one that would have raised the question: returns to what—to education as certificate or to education as cognitive competences? As the first dimension is often overlooked by economists, this framing would have allowed a pointed critique of the dominance of human capital theory, with its emphasis on competence and skills, in the research on returns to education.

Second, the authors promote social class simply as an additional outcome in a literature allegedly focused on earnings: “(...) there is no reason, from either an academic or a policy point of view, why *only* earnings returns should be considered. In this paper we focus instead on social class returns” (Bukodi and Goldthorpe 2011, 186). The true contribution, however, would have been to emphasize the importance of social class as a relational concept, thereby challenging the underlying idea of free markets in the returns to education literature. In particular, with their focus on the “chances of *access* to the professional and managerial salariat” (ibid), the authors question rather than extend the concept of “individual returns”. They point at the structural conditions that inhibit a direct translation of education into earnings through processes such as social closure and other mediating factors related to the class concept. “When we hear from all sides the demand for an introduction of regular curricula and special examinations, the reason behind it is, of course, not

COMMENT AND DEBATE

a suddenly awakened ‘Thirst for education’ but the desire for restricting the supply for these positions and their monopolization by the owners of educational certificates. Today the ‘examination’ is the universal means of this “monopolization” (Weber 1994, 130). In this way, Sørensen (2000) criticized the human capital literature by framing the returns to education process as a rent-seeking process—that is, “as returns on assets that are in fixed supply because single owners of the asset [social classes] control the supply of those assets so that the supply will not respond to an increase in price” (Sørensen 2000, 1525, our addition). More recently, Weeden (2002) used social closure processes at the occupational level to partially account for earnings inequalities between occupations (which, in Weber’s sense, are means of social class closure, connected to educational certificates, differences in skill requirements, and employment contracts). With their focus on social classes, Bukodi and Goldthorpe implicitly highlight the structural conditions that shape the relation between education and earnings through processes such as social closure.

These are important aspects of the social class concept that are relevant for the debate on returns to

education and not merely the fact “that individuals’ earnings come from the *jobs* that are offered by employers” (Bukodi and Goldthorpe 2011, 186). For this argument, economic segmentation theory would be sufficient, attributing earning differentials to differences in individuals’ marginal productivity, based on job content, skill requirements, working conditions, and work contract (e.g. in terms of employment duration). What is more important to social class is the fact that neither access to higher education nor access to jobs—even when recruiting via educational credentials—are solely based on individuals’ productivity (as the returns to education literature suggests). Instead, both returns and individual productivity depend on access, and access is shaped by social closure and practices such as “opportunity hoarding” —that is “categorically unequal access to valued outcomes (...), e.g. how to sort students, whom to hire etc.” (Tilly 1998, 11 and 15). Although the authors briefly touch on this argument in their introduction (p 186), they fail to embed their work in this literature and do not explicate the implications for the returns to education literature. Understanding Bukodi and Goldthorpe’s contribution in these terms merits a truly sociological critique of the literature on returns to education.

References

- Bourdieu P. (1987) *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press, Cambridge.
- Brand JE and Xie Y. (2010) Who Benefits Most from College? Evidence for Negative Selection in Heterogeneous Economic Returns to Higher Education.” *American Sociological Review* 75, 273-302.
- Bukodi E and Goldthorpe JH. (2011) Social class returns to higher education: chances of access to the professional and managerial salariat for men in three British birth cohorts. *Longitudinal and Life Course Studies* 2,185–201.
- Collins R. (1979) *The credential society: An historical sociology of education and stratification*. Academic Press, New York.
- Hanushek EA, Ludgar W, Peterson P, Baker E and McGaw B. (2010) Education and Economic Growth. *International Encyclopedia of Education*. Elsevier, Oxford, 245-252.
- Heckman JJ and Rubinstein Y. (2001) The Importance of Noncognitive Skills: Lessons from the GED Testing Program. *The American Economic Review* 91, 145-149.
- Müller W and Jacob M. (2008) Qualifications and the Returns to Training Across the Life Course. In Mayer KU and Solga H. eds. *Skill Formation: Interdisciplinary and Cross-National Perspectives*. Pp. 126-172. Cambridge University Press, Cambridge.
- Sørensen AB. (2000) Symposium on class analysis: Toward a sounder basis for class analysis. *American Journal of Sociology* 105, 1523-1558.
- Weber M. (1994) The Rationalization of Education and Training. In Grusky DB. ed. *Social Stratification: Class, Race, and Gender in Sociological Perspective*. Pp. 129-131. Westview, Boulder, CO.
- Weeden KA. (2002) Why Do Some Occupations Pay More Than Others? Social Closure and Earnings Inequality in the United States. *The American Journal of Sociology* 108, 55-101.
- Tilly C. (1998) *Durable Inequality*. University of California Press, Berkeley, CA.

Endnote

ⁱ Note that this example refers to parental aspirations that are prior to the child’s education.

Response by **Erzsébet Bukodi**

Department of Social Policy and Intervention
and Nuffield College, Oxford

John H Goldthorpe

Department of Quantitative Social Science
Institute of Education, London

Causes, Classes and Cases

The commentaries on our paper by Clarke and by Legewie and Solga, which we appreciate, concern three main topics: causality, social class as a restriction on the supply side of labour markets, and missing data. Our response is structured accordingly.

Causes

Our general position in this regard is that of ‘causal pluralism’: i.e. we do not believe that there is any one ‘right’ characterisation of causality, but rather that different approaches to causality will be more or less appropriate to different substantive areas of inquiry and to the problems that arise within them (cf. Cartwright 2007). We would not therefore wish to accept what Clarke refers to as the Neyman-Rubin approach, in which a causal effect is understood as the impact of exposure to some intervention or treatment, as in some way capturing the essentials of the matter. On grounds set out at greater length elsewhere (Goldthorpe 2001) we would regard this approach as one that is in fact largely *inappropriate* to the social sciences, and not only because it derives from experimental research of a kind rarely possible for social scientists; two further reasons can be advanced.

First, as some of the more sophisticated proponents of the approach do indeed recognise, it has difficulty in dealing with cases where causation could be thought to inhere in individuals’ goals, beliefs, reasoning and decisions. Thus, to take an example from Holland (1986), ‘She did well on the exam because she was coached by her teacher’ is a claim that can be accommodated to an understanding of causality as the effect of a treatment; but not the claim ‘She did well on the exam because she studied for it’. In this latter case, as Holland accepts, a causal account involving the student’s goal of doing well, her - correct - belief about the means of achieving this goal, and her reasoned decision to take this route, is scarcely compatible with the idea of her being exposed to a

treatment analogously, say, to a patient being given a drug.

Secondly, as Clarke himself acknowledges, the Neyman-Rubin approach is concerned with the effects of causes - with, say, ‘average treatment effects’ - and not with the causes of effects. This focus may be entirely appropriate in many areas of applied research including, in the social sciences, policy evaluation studies. But it does at the same time seem excessively restrictive. A central scientific concern has always been with the causes of effects. As Popper has remarked (e.g. 1972, 115), in all the sciences *it is effects that raise the problem*: i.e. that constitute that which is to be explained and for which causal explanations are sought.

This last point is one directly relevant to our concern with description. In our view, the primary role to be played by statistics in the social sciences, or at least in sociology, *is* a descriptive one. It lies in establishing, and in determining the precise form of, social regularities - of a probabilistic kind - within specific populations. These regularities we would see, following Popper, as constituting the *explananda* of sociology. Clarke suggests that causal rather than descriptive interpretations of statistical results are desirable because they imply the isolation of different effects. But what we would stress is that a concern with isolating effects and checking for spurious relationships, is just as relevant in analyses that aim at description, as in those with causal ambitions. Thus, in our paper, our main aim was to establish patterns of association - hopefully, non-spurious - between individuals’ social origins, their cognitive ability, their educational qualifications and their chances of access to the professional and managerial salariat, and the way in which these patterns have changed over time.¹

In a social science context, we would in general favour an interpretation of the results of regression analyses as having only descriptive force - rather than regression being seen as a means of moving, as Freedman has put it (1997), ‘from association to

causation'. Clarke observes that a descriptive interpretation of regression means that one can speak only about how different groups of individuals, as defined by the independent variables of the analysis, are distinguished - i.e. about how they vary - in relation to the outcome variable. This we would accept - apart from the implication that description is then in some way an inferior task. We would, rather, see the accurate description of social regularities or, in other words, the task of getting the *explananda* right, as one that is important in itself and *prior to* that of trying to explain why the regularities established are as they are.

A causal interpretation of regression would appear warranted insofar as coefficients can be taken as having, in Freedman's words (1997, 117), 'a life of their own', outside of the data from which they are estimated. But this seems far less likely to be the case in the social than in the physical sciences, and primarily because of less developed theory. For example, in quantifying Hooke's Law (Freedman 2010, 11-15), the coefficient obtained from regressing the extension of a spring on its load could be regarded as capturing an inherent property of that spring, or type of spring, and one that itself causes the data from which the estimate is made to be as they are. But do regression coefficients for the earnings returns to education have such a life of their own? There seems no compelling evidence for believing so - for believing that the underlying theory is strong enough for 'structural β s', in Clarke's phrase, to be identified; and in turn we would doubt the validity - and, so far as policy is concerned, the wisdom - of taking such coefficients as a basis for counterfactual propositions.

In the approach to causality that we would ourselves wish to follow, what is important is not determining the effects of causes - i.e. of treatments - but rather gaining an understanding of the causal processes, or mechanisms, that generate established effects. As Cox has observed (1992, 297), it is a 'major limitation' of the Neyman-Rubin approach that 'no explicit notion of an underlying process' is introduced - no notion of a process 'at an observational level that is deeper than that involved in the data under immediate analysis'.

Thus, to take the case of smoking and lung cancer to which Clarke refers, epidemiologists did, as he states, establish a robust association between the two - which could be, and was, taken as highly

suggestive of a causal relation. But the actual demonstration of a causal, generative process came only with the isolation of known carcinogens in cigarette smoke, histopathological evidence from the bronchial epithelium of smokers with lung cancer, and so on. An understanding of causal processes in this sense could in fact be regarded as a *prerequisite* for any attempt at producing structural models - for knowing, in the first place, the proper variables to include. As Duncan has pointed out (1975, 152), there were no structural models for the epidemiology of malaria 'until the true agent and vector of the disease were identified, although there were plenty of correlations between prevalence of the disease and environmental conditions'ⁱⁱ.

In sociology, the established effects are the empirical social regularities that descriptive statistical analyses can demonstrate, and what is then required is some account of how these regularities are actually produced. Such causal accounts cannot be cranked out from statistical analysis itself but will, we believe, have in general to be given in terms of some theory of social action: that is, in terms of the goals, beliefs, reasoning and decisions of individuals acting within particular structures of opportunity and constraint. That is to say, they will be accounts entailing precisely the understanding of causality that the Neyman-Rubin approach is not designed to handle. Statistics are of course likely to come again into play in testing the validity of such accounts, although in this case the relevant data could be quite different - at, say, a more 'micro' level - from those from which the regularities constituting the *explananda* were derived.

We will elaborate further on the foregoing in the course of our discussion of the second of the topics we take up.

Classes

As we understand them, Legewie and Solga have two main concerns about our paper. First, they see it as unfortunate that we decline to make causal claims because we can then do little to extend economists' analyses of returns to education; and second, they think that we also miss an opportunity to develop a critique of economists' analyses that is implicit in our paper, relating to the part played by social class in imposing restrictions on the supply side of labour markets.

On the first issue, we have already set out why we believe that the description - as accurate as possible and as sophisticated as necessary - of what is to be explained should precede attempts at explaining it; and, further, why we believe that, at least in sociology, such explanatory attempts require the specification of causal processes in terms of individuals' actions rather than the estimation of treatment effects. We would certainly be responsive to suggestions of the kind Legewie and Solga make of further variables that might be included in our analyses. Thus, in current work (Bukodi and Goldthorpe, 2011) we develop the concept of social origins, in relation to children's educational attainment, so as to include parental social status and parental education as well as parental class; and in future research we plan to include measures of various of children's non-cognitive attributes. But all this is with the aim, in the first place, of arriving at better description rather than that of producing causal accounts of the kind in which it is variables rather than individuals that do the acting. So far as extending economists' analyses is concerned, our main concern was in fact quite different to what Legewie and Solga would have it be. It was simply to argue that the relation between education and individuals' class positions could be of greater interest than that between their education and their (current) earnings in view of the evidence that class position is a more comprehensive indicator of individuals' economic situation and life-chances.

On the second issue, we would agree with Legewie and Solga that social class imposes restrictions on the supply side of labour markets, and to a greater extent than tends to be recognised in at least mainstream economics. But it was not our aim to enter into this contested terrain in the paper under discussion. We intend to do so in future - but then, we have to say, on different lines to those that Legewie and Solga would have us follow, and for reasons that derive directly from our position as set out in the previous section.

Legewie and Solga see the main bases for a critique of economists' analyses of the earnings returns to education as lying in the existing sociological literature on 'credentialism', 'social closure' and 'segmented' labour markets. We are less impressed by what has been achieved in pursuing these concepts. We find two other literatures more promising, and primarily because

they do comprise attempts at accounts of causal processes that could generate established empirical regularities.

First, in the case of inequalities in educational attainment associated with individuals' class origins, we would note research (Erikson *et al.*, 2005; Jackson *et al.*, 2007; Jackson and Jonsson, eds., forthcoming) that seeks to distinguish between 'primary' effects, as reflected in class differences in actual educational performance, and 'secondary' effects as reflected in class differences in educational *choices*, given performance (these latter effects being usually discussed, and often underestimated, by economists under the rubric of 'credit constraints'). Relevant models of educational choice, hypothesising underlying causal processes, are then available (Erikson and Jonsson, 1996; Breen and Goldthorpe, 1997) and have been quite widely subject to empirical test with, so far, encouraging results (Goldthorpe, 2007, vol. 2, ch. 4).

Second, as regards the association between individuals' educational attainment and the jobs, forms of employment relations and thus class positions into which they enter, we would be more appreciative than Legewie and Solga of what certain economists have proposed as alternatives, or at least complements, to the dominant human capital theory. Again, possible underlying causal processes have been at least outlined, and appear, from a sociological standpoint, more realistic and comprehensive than those invoked or implied by human capital theory.

For example, signalling and screening theory (cf. Weiss, 1995) suggests that rather than, or as well as, education enhancing productivity, it serves to indicate it and allows employers to identify it in potential employees - with the possibility then being raised of *other* attributes of individuals associated with their social origins being similarly used in signalling and screening processes. Job competition theory (Thurow, 1976) proposes that job returns to education are prior to earnings returns and, further, that in job competition what matters is not how much education an individual has but rather how much *relative to others* - so that families in more advantaged classes are likely always to engage in 'defensive' expenditure in order to maintain their children's educational edge. And 'incentive enhancing preference theory' (Bowles, Gintis and Osborne, 2001a,b) sees

education increasing individuals' productivity not simply through the transmission of knowledge and skills but through schools and colleges reinforcing - or in some cases perhaps offsetting - processes of family socialisation so as to endow individuals with values, norms and preferences of a kind that make them more responsive to the incentives and sanctions that employers can deploy.

In a research project on which we are about to begin (together with Heather Joshi and Jouni Kuha) we will aim to test further the causal claims deriving from such theories - with a view to their refinement and development - against the data of the British cohort studies. These will allow us to follow men and women from their social origins, through their educational careers and then through the job, occupational and class trajectories that they follow up to their middle years.

Cases

Clarke raises the problem of missing data, and expresses doubts about the way we handle it - i.e. through the 'missing data indicator' method. We recognise the importance of the problem but we would further observe that with the data-sets of the cohort studies that we use, attrition - i.e. the permanent loss of cohort members over successive sweeps of the studies - could be reckoned as a more serious source of missing data than that of item non-response on which Clarke appears to focus. Statisticians have given a good deal of attention to identifying the nature and extent of biases that might result from attrition; but far more work would still seem needed, on the lines of that of Hawkes and Plewis (2006), aimed at modelling attrition as a basis for some form of corrective weighting.

Against this background, the 'missing data indicator' method then has obvious attractions in helping to keep up the number of individual cases represented in each cohort. Nonetheless, we see the force of the arguments that Clarke sets out, and we have in fact come to accept that the potential dangers of the method in introducing bias into estimates do outweigh its advantages. In more recent research (e.g. Bukodi and Goldthorpe 2011) we have therefore reverted to the 'complete cases' procedure - i.e. that of including in analyses only those cases in which there is no missing data on the set of variables of interest. We may, though, add that we have repeated the main analysis of the paper under discussion using only complete cases - with a reduction of the total N of about 25% - and that the results obtained, available on request, do not differ in any consequential way from those originally reported using missing data indicator variables (Table 2): one marginally significant effect now becomes non-significant. Is more to be learnt from this than that we just got lucky?

Clarke proposes that if one does not wish to lose incomplete cases, then methods of multiple imputation (MI) should be applied. Until recently, it could be said that MI has been far more often recommended by statisticians than practised by social researchers, and chiefly because of the apparently high start-up costs involved. However, software developments now make its use more attractive. In the research project referred to above, we plan to use MI and, in the early stages of the project, to treat various substantive issues through analyses based both on complete cases and on MI. Some methodological interest may again attach to the comparison of the results obtained.

References

- Bowles S, Gintis H and Osborne M. (2001a) Incentive-enhancing preferences: personality, behavior and earnings. *American Economic Review*, 91, 155-8.
- Bowles S, Gintis H and Osborne M. (2001b) The determinants of earnings: a behavioural approach. *Journal of Economic Literature*, 39, 1137-1176.
- Breen R and Goldthorpe JH. (1997) Explaining educational differentials: towards a formal rational action theory. *Rationality and Society*, 9, 275-305.
- Bukodi E and Goldthorpe JH. (2011) Decomposing social origins: the effects of parents' class, status and education on the educational attainment of their children. Typescript, Nuffield College Oxford.
- Cartwright N. (2007) *Hunting Causes and Using Them*. Cambridge University Press, Cambridge.
- Cox DR. (1992) Causality: some statistical aspects. *Journal of the Royal Statistical Society, Series A*, 155, 291-301.
- Duncan OD. (1975) *Introduction to Structural Equation Models*. Academic Press, New York.
- Erikson R and Jonsson JO. (1997) Explaining class inequalities in education: the Swedish test case. In R Erikson and JO Jonsson. eds. *Can Education be Equalized?* Westview Press, Boulder, CO.
- Erikson R, Goldthorpe JH, Jackson M, Yaish M and Cox DR. (2005) On class differentials in educational attainment. *Proceedings of the National Academy of Sciences*, 102, 9730-3.
- Freedman DA. (1997) From association to causation via regression. In SR McKim and SP Turner. eds. *Causality in Crisis?* Notre Dame University Press, Notre Dame, IN.
- Freedman DA. (2007) *Statistical Models and Causal Inference*. Cambridge University Press, Cambridge.
- Goldthorpe JH. (2001) Causation, statistics and sociology. *European Sociological Review* 17, 1-20. (Reprinted in Goldthorpe, 2007)
- Goldthorpe JH. (2007) *On Sociology* (2nd ed. 2 vols.), Stanford University Press, Stanford, CA.
- Hawkes D and Plewis I. (2006) Modelling non-response in the National Child Development Study. *Journal of the Royal Statistical Society, Series A*, 169, 479-91.
- Holland P. (1986) Statistics and causal inference. *Journal of the American Statistical Association*, 81, 945-60.
- Jackson M, Erikson R, Goldthorpe JH and Yaish M. (2007) Primary and secondary effects in class differentials in educational attainment. *Acta Sociologica*, 50, 211-29.
- Jackson M and Jonsson JO. eds. (forthcoming) *Determined to Succeed? Performance, Choice and Education*.
- Popper KR. (1972) *Objective Knowledge*. Oxford University Press, Oxford.
- Thurow LC. (1976) *Generating Inequality*. Macmillan, London.
- Weiss A. (1995) Human capital vs. signalling explanations of wages. *Journal of Economic Perspectives*, 9, 133-54.
- Xie Y. (2007) Otis Dudley Duncan's legacy: the demographic approach to quantitative reasoning in social science. *Research in Social Stratification and Mobility*, 25, 141-56.

Endnotes

ⁱ Legewie and Solga claim that in taking over the concept of 'returns' to education from economists we imply causality. But what matters here, in our view, is usage, not the supposedly 'inherent' meanings of words; and we see no difficulty in talking about the 'returns' of one kind or another that are *associated with* education, while leaving open questions of causal process (see further below).

ⁱⁱ Our view that in sociology the results of regression analyses are in general better interpreted in descriptive rather than causal terms does in fact go back to Duncan. For a revealing discussion of his position, see Xie (2007) and the distinction made between Gaussian and Galtonian conceptions of regression.