



# Disadvantaged through discrimination? The role of employers in social stratification<sup>1</sup>

Michelle Jackson

---

## Abstract

Sociologists have consistently demonstrated that a rather strong association exists between an individual's social class origin and their social class destination, even after controlling for educational attainment. One explanation for this persisting association which is rarely addressed in research in social stratification and mobility is the extent to which class inequalities in access to advantaged class positions are due to discrimination by employers.

I set up a field experiment to test whether employers discriminate on the basis of class origin characteristics. I sent letters of job application for professional and managerial occupations to 2560 large UK companies, so as to compare the prospects of equally matched potential employees differing on a range of characteristics, some related to class of origin. The six treatment conditions in the experiment were: the name of the candidate, the type of school attended, the candidate's interests outside work, their sex, the university that they attended and their achieved degree class.

Results suggest that employers do pay attention to the class origin characteristics tested here, and that candidates with a name, school type and interests associated with the social elite are more likely to receive a reply to their application than candidates with the equivalent non-elite characteristics. However, the treatment conditions do not, on the whole, have significant effects on the employers' responses in and of themselves. Instead, employers appear to favour particular combinations of characteristics while penalising others.

**Keywords:** Discrimination; class origin; field experiment; social mobility; case-control analysis; recruitment

---

## Introduction

A central concern of social stratification and mobility research has been to document and explain the association between an individual's social class origin and their social class destination. Cross-national research consistently

---

shows that inequalities in relative mobility chances are found in most (and probably even all) modern industrial societies, and while these inequalities have declined in some societies over the past half-century, in no society have they disappeared altogether (Breen and Jonsson 2005; see also Breen 2004). In modern Britain, there has been a high degree of stability in relative mobility rates over time, such that a situation approximating constant social fluidity has prevailed (Goldthorpe and Mills 2008; Goldthorpe and Jackson 2007). This constancy embodies a strong degree of immobility towards the higher end of the class structure, and very low relative chances of mobility (in either direction) between these groups and those at the bottom.

In trying to explain differences in relative class mobility chances, sociologists have traditionally concentrated on individuals' educational attainment (see, for example, Blau and Duncan 1967; Halsey 1977), and, more recently, on their basic levels of ability and effort (Herrnstein and Murray 1994; Saunders 1997; Breen and Goldthorpe 1999). But while 'merit' variables can explain some part of the differences in relative mobility chances, even after controlling for ability, effort and educational attainment, there is still a significant association between an individual's class origin and their class destination. In summary, '... there are processes creating class inequalities in mobility chances that cannot be given a "meritocratic" legitimation of any kind that has so far been suggested' (Breen and Goldthorpe 1999: 18).

Class inequalities in mobility chances should be thought of as the consequence of a whole range of micro-level decisions by prospective employees and employers (Goldthorpe 1996; Jackson, Goldthorpe and Mills 2005; Jackson 2007). On one side, prospective employees make decisions about which jobs to apply for, in light of, for example, their educational attainments, skills, aptitudes, preferences and constraints. On the other side, employers make decisions about who to employ, in light of, for example, an understanding of the type of person appropriate for the position, preferences, and both legal and other constraints. In this paper I focus on the role of employers, examining decisions made by employers during the recruitment process which have the potential to impact on class inequalities in mobility chances. In particular, I ask why it is that at the same level of education, an individual from a privileged class background is advantaged over an individual from a less-privileged class background when applying for high-level occupations? An explanation which is rarely addressed in class mobility theory and research is discrimination: do employers for high-level occupations *discriminate* in favour of candidates from privileged class backgrounds?

The lack of attention given to discrimination as a potential mechanism for creating inequalities in access to high-level occupations is highlighted when we compare with research on race and ethnic inequalities. A substantial body of evidence demonstrates that individuals from racial or ethnic minorities are discriminated against when applying for jobs. Studies carried out over

the past forty years using experimental designs show that candidates of ethnic minority background are systematically discriminated against by employers (CRE 1980; CRE 1996; Daniel 1968; Esmail and Everington 1993; Firth 1981; Howitt and Owusu-Bempah 1990; McIntosh and Smith 1974; Newman 1978; Noon 1993; Bertrand and Mullainathan 2004). Employers have also been shown to discriminate on the basis of other characteristics, including sex, age, sexual orientation, disability (see Rodgers 2006 for an overview of research in this area) and possession of a criminal record (Pager 2003).

One reason why discrimination is given as an explanation for ethnic inequalities, for example, but not for inequalities in class mobility chances is perhaps that a potential employee's ethnicity is easily identified by reference to generally accepted criteria, and *can* thus form the basis of discriminatory judgments. But people do regularly make judgments about class background and act accordingly. Psychological research shows that people use cues to ascertain class background, and that social interaction can be influenced by class position (e.g. Sissons 1976; Evans 1988). Argyle argues that, 'Classes . . . function as social groups since members are attracted to one another because of their similarity of lifestyle, values and the rest, and they are aware of being members of the same group. They can recognize each other from speech and appearance, they conform to shared norms and they like outsiders less' (Argyle 1994: 289). Another example (far from academic, but illustrative) can be taken from *The Bluffer's Guide to British Class*, a tongue-in-cheek guide to negotiating the class structure:

It is clearly an enormous asset to be adept at class-spotting. The joy is that it can be done with a minimum of physical effort or bibliographical research: all you have to do is observe . . . What you must watch out for is the true trait coming out in spite of best endeavours and imposing its genuineness on the assumed mannerisms. For instance, a person may speak with a posh accent and even have a stiff upper lip or wear a deer-stalker, but the discerning [observer] will know that he or she is not the genuine article if the shoes are not made of leather. (Gammond 1990: 21)

As any stereotype may feed in to an employer's judgment of a potential employee, stereotypes associated with class background may lead to individuals from different class backgrounds being treated differently.

### ***Research questions***

In this paper, I test whether individuals from less privileged class backgrounds are disadvantaged when applying for high-level occupations *through* discrimination by employers. I use a research design which has been particularly favoured in studies of ethnic and racial discrimination in recruitment (cited

above): a field experiment, in which the researcher takes the role of a potential employee, making applications for a number of jobs. Typically, two or more letters of application are sent to a single employer. The letters are substantively the same, varying only in respect to the treatment condition (i.e. the variable of interest); thus all characteristics of an individual which might otherwise be deemed relevant to job performance (e.g. qualifications or experience) and all characteristics of the firm are controlled for. If, within a sample, one type of candidate receives better treatment on average than another type of candidate, the difference can be attributed to the treatment condition, and discrimination is said to have occurred.

I use a modified version of this experimental design, sending letters of job application for professional and managerial occupations to employers, so as to compare the prospects of equally matched potential employees who differ on a *range* of characteristics, some related to social class of origin. Professional and managerial occupations were chosen as the focus of the study as these are the jobs which carry the highest rewards, and can therefore be treated as being at the top of the British occupational structure. I sent unsolicited, or 'speculative' applications to employers, providing brief details of the 'candidate's' credentials and enquiring whether the company had any available positions.<sup>2</sup>

In an experimental design of this type, it is not realistic to believe that we can capture the nuances of the British class system to the same extent as social stratification and mobility researchers working with three- or four-digit occupational codes and employment status measures. Without using signals of class origin that are so strong that external validity would be compromised (for example, including the job titles of the candidate's parents in the application letter), it is difficult to imagine how a researcher could generate class origin measures which would communicate to an employer that a candidate has parents who would be classified in, say, NS-SEC Class 1 rather than NS-SEC Class 2. In addition, it is very difficult in a design of this type to isolate class from highly correlated social background characteristics relating to advantage and disadvantage. In particular, it is hard to see how the effect of social class could be considered independently from an effect of what sociologists would more properly call social status. In this paper, therefore, the conceptualization of class should be understood as being rather closer to the lay conception of class than to the social-scientific definition of class, in that it distinguishes between the socially and economically advantaged and the socially and economically disadvantaged.

There are two notable consequences of treating social class in this way. First, rather than trying to test the effect of relatively subtle differences in social position, I use a binary measure of class which distinguishes between individuals from privileged social backgrounds and those from less privileged backgrounds. Second, to emphasize the less rigorous conceptualization of class, I do not use the traditional nomenclature of stratification and mobility

research. Instead, I refer to those from privileged social backgrounds as the 'social elite' and to those from less privileged backgrounds as the 'non-elite'. The term 'elite' here refers to a group of people with superior social and economic status.<sup>3</sup>

The analyses will be guided by three research questions:

1. Can evidence of class discrimination in favour of the social elite be detected in this sample?
2. Is there evidence that the possession of *any individual* elite class characteristic confers a benefit to a candidate?
3. Are particular combinations of class- and non-class related characteristics more successful than others when it comes to receiving responses from employers?

## **Methodological issues**

### ***Sample***

In order to reach conclusions about discrimination in the UK labour market, a representative sample of companies is desirable. However, there are problems in defining the population of interest, given the lack of a list of all UK companies. In addition, there are three major constraints on the sample. First, the design of the experiment, described below, meant that a large sample size was necessary. Second, I aimed to sample the largest UK companies, as these companies have the largest amount to spend on recruitment, so are less likely to be negatively affected by dealing with fake applications. Furthermore, small companies are often avoided in studies of this type because of a concern that an increase in the number of applications received might lead employers to draw erroneous conclusions about the nature of the local labour market, thus negatively affecting both companies and real job applicants (McIntyre, Moberg and Posner 1980). And third, as the interest here is in examining the extent of discrimination faced when applying for high-level jobs, a company had to recruit for jobs at this level in order to be included in the sample. Each company also had to recruit for both professional *and* managerial occupations, as the letter combinations sent to each company were chosen at random (see below). As a result of these considerations, the sample of companies was gathered by taking the names and addresses of companies found in careers guides, information booklets and website directories. Any company named in these sources was included if it satisfied the criteria above. The total sample size is 2560 UK companies (the sample size being in part determined by the number of different combinations of treatment conditions, described below).<sup>4</sup>

**Measures**

There are six treatment conditions of interest in the experiment: sex, university attended, degree class, name, school type, and interests of the applicant. All conditions are binary, i.e. there are two values for each variable. It is important to note that not all treatment conditions are measures of class origin; only the name, school type and interests of the applicant are treated as indicators of class origin. Table I shows the treatment conditions, the values taken for each condition, and the signals which were used to communicate them to employers. Note that because two letters were sent to the same company, some treatment conditions required more than one signal for each value. An employer would be more than a little suspicious to receive two application letters from the same 'person', particularly if that person claimed an entirely different educational and personal history in the two letters.

I wanted the treatment conditions to approximate as closely as possible the real-world situation, both so that employers should believe that they were dealing with real job applicants, and so that the results could be treated as valid and generalizable. The aim was *not* to maximize the effects of the treatment conditions by using the most powerful manipulation that could be found. In choosing the signals for the treatment conditions, therefore, my aim was to use external information wherever it was available in order to maximize the external validity of the experiment. Where external information was not available and more subjective judgments were required, employers were consulted to check whether the treatment conditions signalled what they were supposed to signal.

**Table I:** *Description of the treatment conditions*

Treatment conditions	Values of the treatment condition	Examples of how values were signalled
Sex	Male Female	Mr . . . Ms . . .
University	High ranked university Low ranked university	Oxford/Cambridge Staffordshire/Sunderland
Degree class	High Lower	1 <sup>st</sup> Class 2 : 2
Name	Elite  Non-elite	Edward Acheson-Gray/Camilla Bevans-Brown Charles Bartle-Jones/Lucinda Ashford-Smith Gary Roberts/Stacey Brown Kevin Wilson/Donna Taylor
School type	Private school State school	Westminster Charterhouse Local High School
Interests	Elite interests  Non-elite interests	Skiing, tennis, operatic society/ Polo, skiing, yachting Snooker, football, darts/ Refereeing local football/netball matches, socializing, and playing the bass guitar/singing in a band formed with friends

The 'class' attributes were represented by the name, type of school attended and interests of the applicant:

- *Name*

Across the world, names operate as signals for a range of background characteristics, including most obviously, ethnicity, religion, sex and social standing (e.g. Lieberman 2000). First names are a particularly pure indicator of cultural background, in that, as Lieberman notes, '[i]t costs no more in dollars and cents to name a daughter *Lauren* or *Elizabeth* than it does to name her *Crystal* or *Tammy*' (2000: 24). In Britain, first names can be extremely strong signals of class origin; although some names may provide relatively ambiguous signals of class origin, many others are unambiguously class-linked. For example, a reviewer of Lieberman's *A Matter of Taste* (2000) writes, 'For English tastes, he is surprisingly scant on taste as formed by class. In this country, one has only to hear the three mellifluous syllables of Jeremy or Cressida to know where one is in social structure, and it is hard not to put down those baptisms to the mere preservation of superiority' (Inglis 2001; see Fox 2004 for a similar example). Surnames are similarly associated with class background: double-barrelled surnames are a hallmark of the elite classes. To find the names used in this experiment, I used names mentioned in the birth and death announcements of national newspapers, as well as the *Friends Reunited* website<sup>5</sup> (in order to minimize the possibility of using the name of an existing person, I always mixed first names with different surnames).

- *Type of school*

Private schools are overwhelmingly attended by the more advantaged classes (Halsey, Heath and Ridge 1980), making the distinction between private and state school education a marker of social background. The two private schools were taken from similar positions in a ranking of private and independent schools produced by *The Telegraph* newspaper. As students attend state schools close to where they live, the choice of the two state schools was constrained by the addresses to which replies should be sent (I do not provide the school names in order to protect the identity of the collaborators who allowed me to use their addresses). Out of the several state schools local to each address, the two chosen were very similar according to standard performance indicators (i.e. performance in public examinations and Ofsted scores).

- *Interests*

In this part of the application, the candidate tells an employer what they enjoy doing in their spare time. It is difficult to find external information about the interests that individuals from different class backgrounds claim to have when they apply for jobs. Instead, I had to fall back on stereotypes

of the interests of members of different classes, the stated interests of individuals listed on the *Friends Reunited* website, and findings from research on the association between class/status and interests in other countries (e.g. Bourdieu 1979). Different signals were chosen for men and women.

The non-class attributes were sex, university attended and degree class:

- *Sex*

Signalled by signing off the letter with 'Mr ...' or 'Ms ...' attached to a sex-specific name

- *University attended*

Designed to capture a distinction between high-ranked universities, which have high entrance requirements, and low-ranked universities, where the requirements are far lower. *The Times* newspaper produces a ranking of British universities each year based on entrance requirements, teaching and research quality, student-staff ratios and student satisfaction (amongst other things), and I compare the two top-ranked universities to two ranked in the lower half of the distribution (*The Times* (newspaper) 2006).

One feature of the British context is that universities are stratified by class, so that individuals from privileged backgrounds are overrepresented in high-ranked universities, while individuals from non-privileged backgrounds are overrepresented in lower-ranked universities. We therefore cannot rule out the hypothesis that the *university attended* measure also measures some aspects of social class.

- *Degree class*

British Honours degrees are given classes ranging from 1<sup>st</sup> class (the highest possible) to a 3<sup>rd</sup> (the lowest honours class). I compared a 1<sup>st</sup> class degree to a 2:2 (lower second class degree, one rank above a 3<sup>rd</sup>).

While the letters differed depending on the values of the treatment conditions chosen, they were substantively similar in other respects. All letters stated that the applicant was applying one year after completing university, that they had undertaken relevant work experience over the past year, and that the applicant had achieved high grades in their A-level examinations. Further, all applicants made enquires about job opportunities in the company, and asked for a company brochure and for further information about the company's graduate training programme. Applications were made for managerial and professional occupations.

Instead of sending two letters to the same company which differed only in relation to one treatment condition, I sent two letters which varied in respect to the treatment conditions chosen at random. A total of 64 different combinations can be derived from the six treatment conditions ( $2^6$ ), and each application letter combination was sent 80 times, giving a total of 5120 application

letters sent (i.e. two to each of the 2560 companies). Half of the applications were for managerial occupations, half for professional. The letter combinations were assigned to the companies at random, so that a company could potentially receive two letters of the same combination, or more likely, receive two letters of different combinations. While this design is a little more complex than the typical matched pairs approach used in many previous studies of discrimination, it allows for enormous flexibility of analysis, in that the effect of individual treatment conditions can be evaluated, as well as the effect of combinations of the treatment conditions.

Other features of the applications were also randomized in order to minimize the possibility of detection of the experiment by the companies. The two letters sent to each company were written on different coloured paper and envelopes, using different fonts, they were addressed to different (but equivalent) company personnel and were sent from and offered different contact addresses (both mail and email addresses were provided). In addition, two different letter styles were used (formal and less formal). Although the style differed, both letters were coherent and well-presented, with good grammar and spelling. I consulted with employers who see many job applications to pre-test the experiment materials and confirm that the two application letter styles were plausible, and that they were substantively similar while remaining different enough to avoid detection. A time lag of a month was introduced between the two letters sent to the same company (the first letter was sent at 'Date A', the second, a month later at 'Date B'). All of these variables designed to avoid detection were assigned randomly within each pair of applications sent to the same company.

### ***Ethical issues***

Research of this kind raises certain ethical questions due, primarily, to the degree of deception that is involved. The subjects of study – that is, the companies – are completely unaware of the purpose of the study, and, indeed, unaware that they are part of the study at all. Quite apart from any general concern that deception is ethically wrong, the time and money of companies involved has been unnecessarily wasted in dealing with the false applications.

It would be difficult to justify using this research design, complete with its deception, if it were possible to research discrimination using other, non-deceptive methods. However, the race discrimination literature provided strong evidence that what employers *say* they do when faced with candidates of different backgrounds is often quite different from what they *actually* do. Deception is an undesirable but necessary strategy for uncovering real world employment practices. In race discrimination research, experimental designs are justified on the basis that race discrimination is illegal as well as unfair. While discrimination on the basis of class origin is not (yet) illegal, it would

seem difficult to argue that individuals should not be protected from such discrimination, when they have as little choice over the class they are born into as over their race.

In designing the experiment I tried to minimize the negative consequences for the companies in the sample. As discussed above, I sampled large firms who are likely to have a large number of job applications each year, and thus the marginal cost of dealing with two more applications is likely to be small. Complete anonymity for the companies involved is assured. In short, the companies involved should suffer no appreciable negative consequences from being part of the research.<sup>6</sup>

### **Outcome measures**

Initially, when presented with a speculative application, a company chooses whether to reply or not reply to that application. If companies choose to reply to an applicant of one class background more frequently than to an applicant of the other class background, regardless of whether those replies are positive or negative, this provides us with an indication that these candidates are not treated equally (similar measures have been used previously in research on ethnic discrimination, e.g. Newman 1978; Howitt and Owusu-Bempah 1990). Although most would agree that a higher rate of replies to one type of candidate rather than another is hardly the most egregious form of discrimination that a society might be faced with, this measure is important because it provides an indication of how seriously a company treats applications from different candidates; if a company does not treat a candidate seriously, then there is no cost to ignoring that candidate's application. Overall, around 48 per cent of the applications received a reply (2330 replies from a total of 4896 applications – see the lower panel of Table II<sup>7</sup>). Within the replies, we can also

**Table II:** *Descriptive statistics for the whole sample and case-control sub-sample*

Company response at date A	Company response at date B		
	Reply	No reply	Total
Reply	834	353	1187
No reply	309	952	1261
Total	1143	1305	2448
		Whole sample	Case-control sub-sample
Reply rate		48%	50%
Range of reply rates for 64 combinations		30–61%	24–71%
Positive reply rate		0.88%	1.44%
Range of positive reply rates for 64 combinations		0–2.5%	0–7.1%
Total		4896	1324

distinguish between different types of responses, ranging from negative (i.e. a rejection), to positive (an interview offer, or request for an informal face-to-face meeting). We can therefore use the type of reply as a second outcome measure. I treat an interview offer or request for a face-to-face/telephone meeting as a positive response, and all other responses as negative. In the whole sample, 43 positive responses were received, which represents just under 1 per cent of all applications sent (just under 2 per cent of those that received a reply). The positive response rate is therefore low, but not peculiarly so given the high level of the jobs applied for and the mode of application.

### ***Case control analysis***

In the following section I use a case-control analysis. Case-control studies are rarely used in sociological research, but frequently used in epidemiology, for example in studies examining the determinants of cancer (Breslow and Day 1980; Breslow 1996; Keogh 2006). In designing a case-control study, the investigator selects cases because they exhibit an outcome of interest and then selects comparable controls which do not exhibit that outcome. Then, a comparison of cases with controls allows the investigator to identify possible risk factors in causing the outcome and to evaluate the relative importance of these risk factors.

A case-control approach is the most appropriate for analysis of these data, as the central aim is to determine whether or not companies are selective in their response to applications from 'candidates' of differing class background. The logic underlying the choice of a case-control approach for this study can be laid out in the following way. As companies receive two letters in this experiment, three different response combinations can be observed: a company replies to both applications, it replies to neither, or it replies to just one of the applications. If a company replies to both of the letters that it receives, or alternatively neither of the letters that it receives, then this gives us (potentially) important information about the company itself, but by definition, no information about a company's selective response to the two different letters.<sup>8</sup> Analysing a sample including companies which reply to both or neither applications would dilute the overall effect of the treatment conditions because for these companies the effect must by definition be zero (as they are not selective). In other words, the very large variance in response behaviour between companies would result in a loss of sensitivity in measuring the effects of interest. It is only in the mixed cases, where a company replies to one candidate but not the other, that we have unambiguous information on how selective a company is in replying to applications. Consequently, to assess the effect of the treatment conditions on the outcomes considered here, it is necessary to restrict the sample to those companies which respond selectively: i.e. a matched case-control sub-sample. Analysis of the case-control

sub-sample can show whether companies are selective in their response to different types of candidates, and if they are, the analysis can also show the magnitude of the difference in responses to different types of candidates.

I match cases to controls on the basis of the first outcome, reply. That is, I limit the sample to contain companies that replied to just one of the applications that they received. Table II shows a summary of the number of companies who replied at Date A against the number of companies who replied at Date B. To construct the case-control sub-sample, I select the 353 companies who replied at Date A but not Date B, and the 309 companies who replied at Date B but not Date A, providing an overall sample size of 662 companies (i.e. 1324 applications). Given the nature of the case-control design, the response rate is fixed at 50 per cent, while the positive response rate is around 1.5 per cent.

## **Results**

### ***Class discrimination***

I begin by asking whether evidence of class discrimination can be found, by testing whether candidates with all three elite characteristics are treated more favourably than candidates with all three non-elite characteristics. Of the 64 different letter combinations, eight combinations included candidates with all three elite characteristics, and eight included candidates with all three non-elite characteristics. Comparing the percentage of replies received by these candidates, in the case control sub-sample those with all three elite characteristics ( $n = 155$ ) received a reply 58 per cent of the time, while those with all three non-elite characteristics ( $n = 184$ ) received a reply only 44 per cent of the time.<sup>9</sup> A Chi-square test shows that this difference is significant at the 0.01 level, indicating that we can reject the null hypothesis of no discrimination. When it comes to receiving a positive reply, a significant difference cannot be detected (which may partly be due to the relatively small number of cases of the relevant combinations): 2.2 per cent of letters from non-elite candidates received a positive reply, compared with 2.6 per cent of letters from elite candidates.

On the basis of this evidence, therefore, discrimination can be detected in the propensity of employers to reply to applications, but not in the propensity of employers to reply positively.

### ***The role of individual class characteristics***

I now move on to ask whether there is evidence that possession of any elite characteristic confers a benefit in and of itself, by examining the percentage of replies and positive replies for each of the treatment conditions.

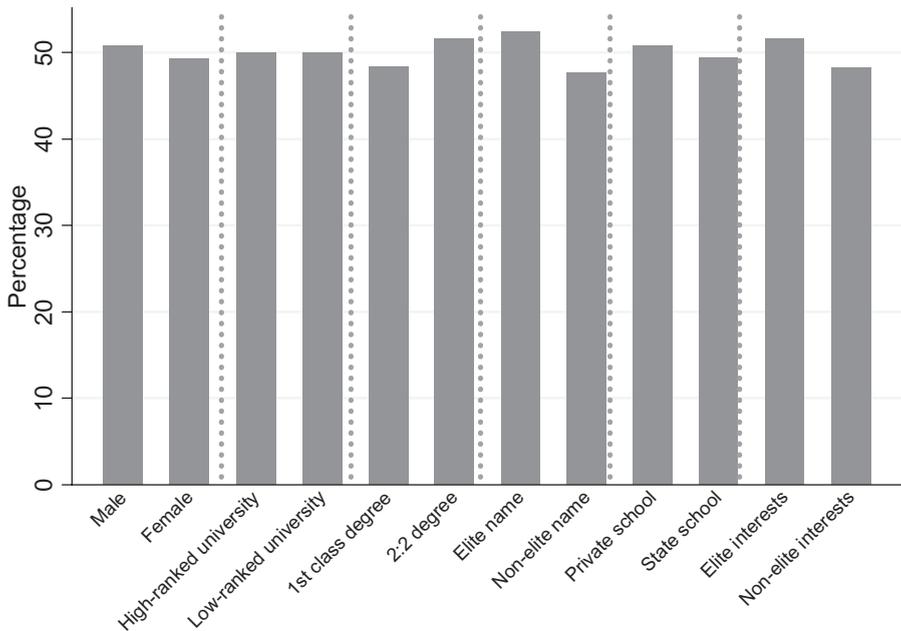
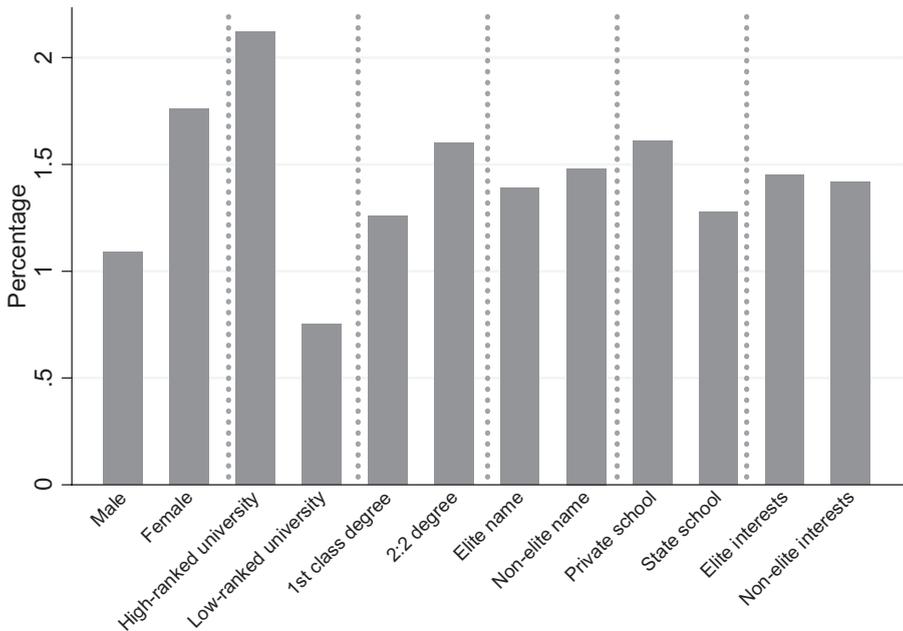
**Figure I:** Percentage of replies by treatment conditions

Figure I shows the percentage of replies received for each category of the treatment conditions, and it is clear that the response rate is rather flat, hovering around the 50 per cent level. For each treatment condition, the response rate is almost identical for the two categories, and where there are differences, these are very small. The one treatment condition for which a small difference in response rates can be found is *name*: letters sent from a person with an elite name are more likely to receive a reply than those sent from someone with a non-elite name. A Chi-square test shows that this finding does not quite achieve statistical significance, with a p-value of 0.09.

The story is somewhat similar for the positive reply measure, shown in Figure II, although there is apparently more variation than in the previous figure. Once again, the positive response rate is rather flat, and Chi-square tests indicate that the only significant relationship detected between the treatment conditions and whether a positive reply is received is for the *university attended* treatment condition. Candidates who attended a high-ranked university are almost three times more likely to receive a positive reply than not, compared to candidates who attended a low-ranked university ( $\chi^2 = 4.38$  for 1 d.f.,  $p = 0.04$ ). The positive response rate is rather similar for all three measures of class background, and the small differences shown in Figure II are statistically insignificant.

**Figure II:** *Percentage of positive replies by treatment conditions*

### *The role of combinations of characteristics*

It appears that while some treatment conditions are associated with different response and positive response rates, on the whole, there is little relationship between the individual treatment conditions and the outcome measures. But in the previous analysis, there was a significant difference between the response rate for candidates with all three elite characteristics and the response rate for candidates with all three non-elite characteristics. Clearly, while the individual treatment conditions do not generate different response rates on the whole, particular *combinations* of characteristics are more successful in eliciting a response from employers than others. The design of the experiment makes it possible to ask which combinations of treatment conditions are most successful.

In order to examine the different rates of success of different combinations of the treatment conditions, I estimate a latent class model. This is an exploratory technique in which the latent variable is a function of the six treatment conditions, and the response variable is linked to the latent variable according to a set of classification probabilities (see Breen 2000: 377–8, for a similar use of latent class analysis (LCA)). The results of the LCA can be used to describe patterns of association between the treatment conditions and the response variable, and to point towards hypotheses for further investigation. The response variable in the LCA is a three-category variable: the letter receives no reply, a negative reply or a positive reply.<sup>10</sup>

One of the principal assumptions of LCA is that, within each latent class, conditional independence exists between variables. However, in some circumstances it is appropriate to relax the conditional independence assumption, allowing for local dependencies between variables (Hagenaars 1988).<sup>11</sup> In the analysis presented here, I allowed for local dependence between the three treatment conditions measuring class background. There are both theoretical and methodological reasons for this. Theoretically, it is reasonable to expect that three measures of the same concept should exhibit local dependence, even within the same latent class. But it also turns out that a model which does not include local dependencies does not have a global solution – at least three separate local solutions were found in repeated runs of the analysis. Investigation of the residuals from these local solutions suggested that local dependencies between the three class measures might be a problem. When the conditional independence assumption was relaxed for these measures, a stable solution was found, which improved on the fit of the previous local solutions.

Under conventional criteria, the independence model postulating one latent class in fact fits the data ( $p = 0.92$ ), which is unsurprising given that the relationships between the individual treatment conditions and the outcome measures are on the whole insignificant. However, the model fit significantly improves for each latent class added, and the best fit is achieved for a model with three latent classes, which fits the data extremely well ( $p = 0.99$ ). The model and fit statistics are presented in Table III.<sup>12</sup> The three latent classes have different profiles on the response variable, and in light of this, can be labelled as follows. The first class can be seen as comprising the ‘acknowledged, with some success’ in the experiment: letters in this class have a 55 per cent chance of receiving a negative response from employers, and a 4 per cent chance of receiving a positive response. The second class comprises the ‘acknowledged, but unsuccessful’: letters in this class are less likely to receive a positive reply (a 2 per cent chance), but are still likely to receive a negative response (a 56 per cent chance).<sup>13</sup> The final class comprises the ‘failures’, where letters are extremely unlikely to receive a response (only 25 per cent of letters in this class receive a reply), and the positive response rate is the lowest at 1 per cent.

The three latent classes are also characterized by three different patterns of conditional probabilities summarizing the distribution of treatment conditions within the classes. The conditional probabilities show the probability of a candidate having, say, an elite or a non-elite name, *conditional on* being found in each of the three latent classes. If, within a particular latent class, the conditional probability of having an elite name was 50 per cent (and therefore the corresponding probability of having a non-elite name was also 50 per cent), this would mean that conditional on being found in that latent class, a ‘candidate’ is just as likely to have an elite name as a non-elite name. Differences which move away from this 50-50 distribution are therefore of interest. The

**Table III:** Results of latent class analysis, allowing for local dependence between name, school and interests

		Acknowledged, with some success	Acknowledged, but unsuccessful	Failures
Size of latent class		0.25	0.52	0.23
Response	No reply	0.41	0.42	0.75
	Negative reply	0.55	0.56	0.24
	Positive reply	0.04	0.02	0.01
Sex	Male	0.22	0.66	0.39
	Female	0.78	0.34	0.61
University	High-ranked university	0.90	0.30	0.50
	Low-ranked university	0.10	0.70	0.50
Degree	1 <sup>st</sup>	0.65	0.33	0.63
	2:2	0.35	0.67	0.37
Name	Elite name	0.59	0.51	0.34
	Non-elite name	0.41	0.49	0.66
School	Private school	0.50	0.48	0.41
	State school	0.50	0.52	0.59
Interests	Elite interests	0.38	0.63	0.44
	Non-elite interests	0.62	0.37	0.56
<b>Model fit statistics</b>				
1 class	$G^2 = 104.8, 126 \text{ d.f.}, p = 0.92$			
2 classes	$G^2 = 82.4, 113 \text{ d.f.}, p = 0.99$			
<b>3 classes</b>	<b><math>G^2 = 50.1, 100 \text{ d.f.}, p = 0.99</math></b>			
4 classes	$G^2 = 40.2, 87 \text{ d.f.}, p = 0.99$			

conditional probabilities for all six of the treatment conditions can be used to build a picture of the type of ‘candidate’ who might be found within each of the latent classes, as they describe the different combinations of attributes that are associated with particular response patterns.

Conditional on being found in the ‘acknowledged, with some success’ class, applicants were most likely to be female, with a 1<sup>st</sup> class degree from a high-ranked university, with an elite name, but non-elite interests. We might argue that this class can be taken as reflecting selection on the basis of meritocratic characteristics, since conditional on being in this class, a candidate is highly likely to have attended a high-ranked university (with 90 per cent probability) and is also likely to have achieved a 1<sup>st</sup> class degree (with 65 per cent probability), and this is associated with a higher response and positive response rate. However, the meritocratic interpretation depends on an acceptance that attendance at a high-ranked university is seen to be merit-based. If, in contrast, employers treat a degree from Oxford or Cambridge as a signal both of the merit of the candidate and of their class background, the interpretation of this latent class as being meritocratic is a little more questionable. Furthermore, conditional on membership in this class, candidates have a higher probability of having an elite rather than non-elite name, although at the same time they are more likely to have non-elite rather than elite interests.

The second class, which I labelled the ‘acknowledged, but unsuccessful’ above, can be seen as comprising candidates who are not very successful in achieving a positive response, but are treated as serious applicants in that they receive plenty of responses overall. Candidates in this class are most likely to be males who have achieved a 2:2 degree from a low-ranked university. Neither the name of the applicant nor the school attended distinguish members of this class, but conditional on membership, these candidates are likely to have elite interests (with 63 per cent probability).

The final class identified in the analysis comprises the ‘failures’; those candidates who are highly unlikely to receive any response to their application, and even more unlikely to receive a positive response. A distinctive feature of this class is that, conditional on membership, the candidates are likely to have all three non-elite attributes. That is, they have around a 60 per cent chance of having a non-elite name, having attended a state school and having non-elite interests. In addition, they are likely to be female with a first class degree, although they are no more likely to have attended a high-ranked university than a low-ranked one. This profile suggests somewhat non-meritocratic selection, with candidates who exhibit the group of non-elite characteristics doing particularly badly, despite the fact that candidates in this latent class are high-achievers.

### **Summary and discussion**

I have addressed three main research questions in this paper. First, I asked whether evidence of class discrimination in favour of the social elite could be detected in this sample. Results showed that candidates with all three elite characteristics receive more replies than candidates with all three non-elite characteristics, but that the propensity to receive a positive reply is not significantly affected by social background. Second, I asked whether possession of *any individual* elite class characteristic conferred a benefit to a candidate. Other than a small (and significant at the  $p < 0.1$  level) positive effect of holding an elite rather than non-elite name on receiving a reply, elite class characteristics did not influence the propensity to receive a reply or to receive a positive reply, in and of themselves. Finally, I asked whether particular combinations of class- and non-class related characteristics were more successful than others in the applications. The LCA identified a latent class of ‘failures’, made up of candidates who were likely to exhibit all three non-elite characteristics, and two other latent classes with different patterns of conditional probabilities attached to them. Notably, the ‘successes’ were not defined by candidates exhibiting all three elite characteristics, as we might expect if employers were making a simple class-based judgment.

These findings point to a rather complex picture of how employers use information about social background when selecting potential employees. To

**Table IV:** *Crosstabulation of outcome of application against consistency of class measures*

Bundles of attributes	Outcome of application			Total
	No reply	Negative reply	Positive reply	
All non-elite	56.0	41.9	2.2	184
2 non-elite, 1 elite	49.3	49.7	1.1	471
1 non-elite, 2 elite	51.0	47.9	1.2	514
All elite	41.9	55.5	2.6	155
Total	662	643	19	1324

employers reading the application letters, the 64 different combinations of the treatment conditions represent different types of candidates, and it is important not to forget the social and psychological context of the experiment when interpreting the statistical findings. The expectations that employers have about which characteristics ‘match’ other characteristics, which characteristics ‘clash’ with other characteristics, and which combinations of characteristics unambiguously classify an individual as a certain ‘type’ of person, will all feed in to the propensity of employers to look favourably (or otherwise) upon applications that they receive.

Psychological research on stereotypes suggests that consistency is important in order to allow people to comfortably categorize others as one type or another. Information which is inconsistent with stereotypes stands out and requires more cognitive processing than information which is consistent (e.g. Hastie and Kumar 1979; Stern et al. 1984; Sherman, Macrae and Bodenhausen 2000; although see also Heider et al. 2007 for a discussion of the conditions under which inconsistent information requires more cognitive processing). In the case of the applications considered here, employers can be more certain of a candidate’s background if the class signals are consistent with one another: they can be more certain that a candidate is from an elite background if they have all three elite background characteristics, and, conversely, more certain that a candidate is not from an elite background if they have all three non-elite characteristics. On the other hand, combinations of elite and non-elite characteristics could be seen as inconsistent by employers, increasing the uncertainty about the social background of the applicant. Inconsistent combinations might also raise questions about the candidate – maybe they are embellishing the truth, or maybe they just sound a bit odd (for example, Gary Roberts, who attended the local comprehensive school yet lists polo, skiing and yachting as hobbies). It is impossible to test these hypotheses with the data available, but it is possible to check whether there is any evidence for the ‘consistent’ combinations being more successful than the ‘inconsistent’ ones.

Table IV shows a crosstabulation of the three category outcome measure (used in the latent class analysis) against a simple measure of consistency; a four category variable which ranges from candidates who have all three

non-elite characteristics to candidates who have all three elite characteristics. The most successful group has all three elite characteristics, and the odds of receiving a negative reply rather than no reply at all are 1.8 times greater for this group than for the group with all three non-elite characteristics (this difference is highly significant:  $p = 0.010$ ). This clear difference suggests that where class signals are consistent, employer discrimination can be detected. The 'inconsistent' combinations are less likely to receive a negative reply than the elite group, but more likely than the non-elite group. But when it comes to receiving a positive response, the non-elite group performs rather similarly to the elite group, and the inconsistent combinations are less successful. Unfortunately this finding does not provide strong evidence for either of the hypotheses discussed above. Instead, it suggests that consistency is important in some circumstances (receiving a positive response), but that even when the bundle of characteristics is consistent, full bundles of non-elite characteristics are less successful than full bundles of elite characteristics.

Rather than a straightforward consistency explanation, a more sophisticated extension might be to argue that some signals are stronger than others, so that a strong signal of an elite background, such as name, is not necessarily perceived as being inconsistent with the other non-elite characteristics. Indeed, the presence of a strong elite signal changes the meaning of the other signals in the application. So two candidates might claim to have an interest in football, darts and socializing, but the social meaning of that statement will differ depending on whether the candidate is named Edward Acheson-Gray or Gary Roberts. To borrow a phrase from Becker (1966), we might argue that some signals act as a 'master status', determining the meaning of other signals in the application.

In fact, an explanation of this type fits rather well with the results of the LCA. If we once again consider the 'failure' class identified in the analysis, we can note that conditional on being in this class candidates are likely to have all three non-elite characteristics – there is no elite characteristic present to suggest that the candidate is anything other than of non-elite background. In contrast, the classes comprising the 'acknowledged, with some success' and the 'acknowledged, but unsuccessful' both encompass a strong elite background signal. Candidates in the 'acknowledged, with some success' class are likely to have an elite name, and while they have non-elite interests, the name acts as the master status. In the 'acknowledged, but unsuccessful' class, candidates are found to have elite interests, but no other distinctive social background profile, suggesting again that the interests act as a master status. Although the lack of effects of holding any *individual* elite characteristic on the outcome measures might be interpreted to mean that no class treatment condition acts as a master status without fail, this could again simply point to the importance of combinations of class and non-class characteristics.

To place the findings of this study in the context of research on other forms of discrimination, it might be helpful to consider the parallels with research on

racial and ethnic discrimination. The clearest commonality between the two types of study is the finding that employers do respond to stimuli such as names, rather than simply replying to or ignoring all application letters that they receive. They also appear to distinguish between different types of candidates when deciding whether to reply or not reply, and whether to reply positively or negatively. Where the findings differ from research on racial and ethnic discrimination is in the lack of a clear effect of altering a single characteristic. Time after time, research on racial and ethnic discrimination demonstrates that employers treat minority candidates differently from candidates of other groups. But *experiments* examining the extent of racial and ethnic discrimination differ in an important way from the experiment reported here: with some exceptions (e.g. Pager 2003; Bertrand and Mullainathan 2004), most experiments of this kind employ matched pairs designs with a single treatment condition. Thus, the effect of changing a *range* of characteristics on employers' propensities to discriminate is not established. One interesting question is how far the findings of racial and ethnic discrimination studies would change if more than one indicator of ethnic background was chosen – perhaps 'incongruous' combinations would here too reduce the chances of finding a main effect of discrimination? Conversely, would a matched pairs design have produced a simpler class discrimination effect than the findings of this study?

The results of this experiment suggest that class background characteristics *are* relevant to the recruitment process for high-level occupations, and that furthermore, some evidence of discrimination can be found. It is clear that in order to understand *how* employers use information about class background characteristics when recruiting, and in particular, the social and psychological processes which are at play, further research employing other methods is needed.

(Date accepted: May 2009)

## Notes

1. For help in conducting the experiment and for useful comments and advice, I thank Richard Breen, David Cox, Geoffrey Evans, Diego Gambetta, John Goldthorpe, David Grusky, Gwilym Hughes, Alexandra Jackson, Susan Jackson, William Jackson, Ruth Keogh, Steffen Lauritzen, Laurence Lessard-Phillips, Graham Pike, Vicki Pike, Martín Sanchez-Jankowski, and participants in the RC28 Conference in Nijmegen and in a sociology seminar at Reading University. I

would also like to express my gratitude to the referees, who provided extremely helpful comments on an earlier version of this article. I am grateful to the British Academy for funding the research project (SG-38534).

2. One advantage of testing discrimination through speculative applications is that these applications are not typically subject to the highly regulated and standardized process that applications in response to

advertisements are. We might therefore expect to obtain a purer measure of an employer's tendency to discriminate at the beginning of the recruitment process by examining applications of this type.

3. This definition of 'elite' is distinct from the definition provided by 'elite theorists', for whom 'elite' is synonymous with 'ruling class'.

4. The sample permits conclusions drawn from this experiment to be generalized to the population of relatively large UK companies. It is quite possible that smaller companies do not exhibit the same discriminatory behaviour as the large companies in this sample; one might imagine that, unconstrained by the bureaucratic recruitment procedures characteristic of very large companies, smaller companies have more opportunity to discriminate. This hypothesis, however, can ultimately only be tested by including smaller companies in a field experiment, which seems difficult to justify on ethical grounds.

5. Newspapers were used to gather names as different newspapers are associated with different readerships. For example, *The Telegraph* is seen to be a newspaper for the privileged, while *The Mirror* is seen to be a 'lower class' newspaper (see e.g. Chan and Goldthorpe 2007: 1110, who show that those in Higher Professional occupations read broadsheet newspapers at a far higher rate than General Labourers, and in contrast, General Labourers read 'Red-Top Tabloids' at a far higher rate than Higher Professionals). The *Friends Reunited* website is a place where adults can make contact with old school friends. Every school in Britain is listed, with lists of ex-pupils provided for each school (by leaving year). I used lists of pupils at schools with very different class compositions to find names associated with privileged and non-privileged individuals. See <http://www.friendsreunited.co.uk> (most recently accessed May 2009).

6. The project was submitted to a departmental ethics committee for consideration, and subsequently received approval.

7. Although 5120 letters were sent to 2560 companies, 112 companies had to be excluded from the sample because they had closed down by the time the letters were sent, or the address was incorrect.

8. There are many reasons why a company might reply to both candidates; these might include a company policy to acknowledge all material received by post, a highly standardized recruitment system, or indeed a genuine interest in both applications received. In contrast, if a company replies to neither candidate, this could be due to a policy to never consider speculative applications, a policy to never read the application letters that it receives, or alternatively a lack of interest in either candidate.

9. As this is a case-control sub-sample, it is unwise to take between-treatment differences in the absolute percentages of replies and positive replies received to the applications as an indication of the level of discrimination within the whole population. Instead, interest is in whether some companies are selective in their response to different types of candidates.

10. As the number of positive responses received was rather low, a number of cells in the crosstabulation of response by treatment conditions had no cases in them. A small constant of 0.2 was therefore added to every cell in the table to allow for a stable solution.

11. Hagenaaers writes that, 'Sometimes the constraints imposed by conditional or local independence are unrealistic. Test-retest effects, response consistency effects, or, in general, correlated response errors and omitted variables may produce relationships between the manifest variables in addition to the relations produced by the latent variables' (1988: 380).

12. One reasonable question might be whether it is appropriate to continue with the LCA, given the good fit of the independence model. While the independence model does fit, I would argue that an independence model does not accurately reflect the complicated structure of the data, as evidenced through the very different response

rates for the 64 letter combinations. Importantly, the model fit is significantly improved for a model postulating three latent classes: in comparison to a model with one latent class, the three class model shows a highly significant reduction in  $G^2$  of 54.7 for 26 d.f. ( $p < 0.01$ ). Furthermore, probabilities derived from the independence model can be shown to less accurately reflect real patterns in the data than those derived from the three class model (further details available from author).

13. Note that although the first two latent classes look to be rather similar in terms of

their chances of receiving a negative reply, what distinguishes them is the chance of receiving a positive reply. Although the difference between receiving 4 per cent and 2 per cent positive replies appears to be small, the positive reply rate is very low overall, and in relative terms this is an important difference (i.e. those in the first latent class are twice as likely to receive a positive reply as those in the second). A model postulating two latent classes has a significantly worse fit than a model postulating three latent classes (the three class model shows a significant reduction in  $G^2$  of 32.3 for 13 d.f. ( $p < 0.01$ )).

## Bibliography

- Argyle, M.** 1994 *The Psychology of Social Class*, London: Routledge.
- Becker, H.S.** 1966 *Outsiders. Studies in the Sociology of Deviance*, New York: Free Press.
- Bertrand, M. and Mullainathan, S.** 2004 'Are Emily and Greg More Employable than Lakisha and Jamal?' *American Economic Review* 94(4): 991–1013.
- Blau, P.M. and Duncan, O.D.** 1967 *The American Occupational Structure*, New York: Wiley.
- Bourdieu, P.** 1979 *Distinction. A Social Critique of the Judgement of Taste*, London: Routledge&Kegan Paul.
- Breen, R.** 2000 'Why is Support for Extreme Parties Underestimated by Surveys? A Latent Class Analysis', *British Journal of Political Science* 30(2): 375–82.
- Breen, R.** (ed.) 2004 *Social Mobility in Europe*, Oxford: Oxford University Press.
- Breen, R. and Goldthorpe, J.H.** 1999 'Class Inequality and Meritocracy: A Critique of Saunders and an Alternative Analysis', *British Journal of Sociology* 50(1): 1–27.
- Breen, R. and Jonsson, J.O.** 2005 'Inequality of Opportunity in Comparative Perspective: Recent Research on Educational Attainment and Social Mobility', *Annual Review of Sociology* 31: 223–43.
- Breslow, N.E.** 1996 'Statistics in Epidemiology: The Case-Control Study', *Journal of the American Statistical Association* 91(433): 14–28.
- Breslow, N.E. and Day, N.E.** 1980 *Statistical Methods in Cancer Research. Volume 1: The Analysis of Case-Control Studies*, Lyon: International Agency for Research on Cancer.
- Chan, T.W. and Goldthorpe, J.H.** 2007 'Social Status and Newspaper Readership', *American Journal of Sociology* 112(4): 1095–134.
- CRE** 1980 *Half a Chance? A Report on Job Discrimination Against Young Blacks in Nottingham*, Nottingham: Commission for Racial Equality.
- CRE** 1996 *We Regret to Inform You ...* London: Commission for Racial Equality.
- Daniel, W.W.** 1968 *Racial Discrimination in England: Based on the P.E.P. Report*, Harmondsworth: Penguin.
- Esmail, A. and Everington, S.** 1993 'Racial Discrimination against Doctors from Ethnic Minorities', *British Medical Journal* 306: 691–92.
- Evans, G.** 1988 *Causal Explanation, Social Class and Perceived Efficacy*, D.Phil. Thesis, University of Oxford.
- Firth, M.** 1981 'Racial Discrimination in the British Labour Market', *Industrial and Labor Relations Review* 34(2): 265–72.

- Fox, K.** 2004 *Watching the English. The Hidden Rules of English Behaviour*, London: Hodder and Stoughton.
- Gammond, P.** 1990 *The Bluffer's Guide to British Class: Bluff Your Way in British Class*, Horsham: Ravette.
- Goldthorpe, J.H.** 1996 'Problems of "Meritocracy"' in R. Erikson and J.O. Jonsson (eds.), *Can Education be Equalized? The Swedish Case in Comparative Perspective*, Boulder, Colorado: Westview Press.
- Goldthorpe, J.H. and Jackson, M.** 2007 'Intergenerational Class Mobility in Contemporary Britain: Political Concerns and Empirical Findings', *British Journal of Sociology* 58(4): 525–46.
- Goldthorpe, J.H. and Mills, C.** 2008 'Trends in Intergenerational Class Mobility in Modern Britain: Evidence from National Surveys, 1972–2005', *National Institute Economic Review* 205(1): 83–100.
- Hagenaars, J.A.** 1988 'Latent Structure Models with Direct Effects Between Indicators', *Sociological Methods and Research* 16(3): 379–405.
- Halsey, A.H.** 1977 'Towards Meritocracy? The Case of Britain' in J. Karabel and A.H. Halsey (eds) *Power and Ideology in Education*, New York: Oxford University Press.
- Halsey, A.H., Heath, A.F. and Ridge, J.** 1980 *Origins and Destinations: Family, Class and Education in Modern Britain*, Oxford: Clarendon Press.
- Hastie, R. and Kumar, P.** 1979 'Person Memory: Personality Traits as Organizing Principles in Memory for Behaviors', *Journal of Personality and Social Psychology* 37: 25–38.
- Heider, J.D., Scherer, C.R., Skowronski, J.J., Wood, S.E., Edlund, J.E. and Hartnett, J.L.** 2007 'Trait Expectancies and Stereotype Expectancies have the Same Effect on Person Memory', *Journal of Experimental Social Psychology* 43: 265–72.
- Herrnstein, R.J. and Murray, C.A.** 1994 *The Bell Curve: Intelligence and Class Structure in American Life*, New York: Free Press.
- Howitt, D. and Owusu-Bempah, J.** 1990 'The Pragmatics of Institutional Racism: Beyond Words', *Human Relations* 9: 885–99.
- Inglis, F.** 2001 'Cressida Reveals All Before You Meet Her', *Times Higher Education Supplement*, 8<sup>th</sup> June.
- Jackson, M.** 2007 'How Far Merit Selection? Social Stratification and the Labour Market', *British Journal of Sociology* 58(3): 367–90.
- Jackson, M., Goldthorpe, J.H. and Mills, C.** 2005 'Education, Employers and Class Mobility', *Research in Social Stratification and Mobility* 23: 3–34.
- Keogh, R.H.** 2006 *Comparing Methods of Analysis for Matched Case-Control Studies*, D.Phil. Thesis, University of Oxford.
- Lieberson, S.** 2000 *A Matter of Taste: How Names, Fashions and Culture Change*, New Haven: Yale University Press.
- McIntosh, N. and Smith, D.J.** 1974 *The Extent of Racial Discrimination*, London: P.E.P.
- McIntyre, S., Moberg, D. and Posner, B.** 1980 'Discrimination in Recruitment: An Empirical Analysis: Comment', *Industrial and Labor Relations Review* 33(4): 543–47.
- Newman, J.M.** 1978 'Discrimination in Recruitment: An Empirical Analysis', *Industrial and Labor Relations Review* 32(1): 15–23.
- Noon, M.** 1993 'Racial Discrimination in Speculative Application: Evidence from the UK's Top 100 Firms', *Human Resource Management Journal* 3(4): 35–47.
- Pager, D.** 2003 'The Mark of a Criminal Record', *American Journal of Sociology* 108(5): 937–75.
- Rodgers, W.M.** 2006 *Handbook on the Economics of Discrimination*, Cheltenham, UK: Elgar.
- Saunders, P.** 1997 'Social Mobility in Britain: An Empirical Evaluation of Two Competing Explanations', *Sociology* 31: 261–88.
- Sherman, J.W., Macrae, C.N. and Bodenhausen, G.V.** 2000 'Attention and Stereotyping: Cognitive Constraints on the Construction of Meaningful Social Impressions', *European Review of Social Psychology* 11(1): 145–75.
- Sissons, M.** 1976 *Social Class and Social Interaction*, D.Phil. Thesis, University of Oxford.

**Stern, L.D., Marrs, S., Millar, M.G. and Cole, E.** 1984 'Processing Time and the Recall of Inconsistent and Consistent Behaviours of Individuals and Groups', *Journal of Personality and Social Psychology* 47: 253–62.

**The Times (newspaper)** 2006 *The Times Good University Guide: The Definitive Guide to Britain's Universities*, London: Times Books.