The Great British Class Fiasco: A Comment on Savage et al.: — Source link ⬈

Colin Mills

Institutions: University of Oxford

Related papers:

- A new model of social class : findings from the BBC's Great British Class Survey Experiment.

- Class Descriptors or Class Relations? Thoughts Towards a Critique of Savage et al.

- On social class, anno 2014

- Distinction: A Social Critique of the Judgement of Taste

- Thinking about Class

The Great British Class Fiasco: A Comment on Savage et al.

Colin Mills

University of Oxford

Abstract

Savage et al. (2013) claim they have produced a new model of the British class structure. They stress the innovative use of an internet survey, the BBC's Great British Class Survey (GBCS), but it plays no serious role in the generation of their class typology. What they do is a theory free (though Bourdieu inspired) data dredging exercise. What they derive is an arbitrary typology determined by a contingent fact – the size of their sample.

Corresponding author

Colin Mills, Nuffield College, Oxford OX1 1NF.

Email: colin.mills@sociology.ox.ac.uk

1. Introduction

This comment criticizes Savage et al. (2013) published in this journal. My argument is in five sections. Firstly I outline the standard view of typologies and compare it to Savage et al.'s. Secondly I deal with their data collection strategy and in particular the limit it imposes on their ability to make statistical inferences. Thirdly I point out the implications of their model selection strategy. Fourthly I reveal their inability to interpret their own results correctly and raise some additional questions about the operationalization of their concepts. Finally I discuss the substance of their typology and draw out some of the bizarre consequences that the authors must accept if they wish to claim consistency for their model.

2. On Typologies

Concepts are abstract entities. We use them to selectively organize experience. Typologies are closely related to concepts; they are the step we take in order to make abstract ideas more concrete. There can be as many typologies as there are points of view. We are free to fill a concept with whatever content we like and then derive a typology from it.

Typologies should be made for a purpose and we should ask both whether that purpose is worthwhile and whether the proposed typology is "fit for purpose". Though we are free to organize what we can learn of other people's circumstances with any conceptual scheme we like, there are logical constraints, and we must acknowledge what follows from doing things one way rather than another. I'll return to this point in section 6.

Some conceptual schemes are more useful than others, they, as Plato put it, cleave nature at its joints. Normally the purpose is given by the part that a concept (and its associated typology) plays in a theory that explains or predicts something and naturally this entails spelling out what it is that the concept is to play a part in explaining or predicting. From this it follows that a typology is a practical tool. Inventing a typology for its own sake is pointless.

Let me give two examples. In market research, market segmentation is a procedure that builds predictive models to help clients target products at people that are likely to buy them. In this kind of exercise all that matters is that good predictions are produced; the prediction model itself can be built in an entirely inductive manner. The categories that are produced, which may be an ad hoc combination of demographic, life-style and dispositional characteristics, are validated by their predictive value.

Now consider John Goldthorpe's most recent statement of what social class is about (Goldthorpe, 2000, 206-29), a view that underpins the construction of the National Statistics Socio-economic Classification (NS-SEC) (Rose & Pevalin, 2003). It focuses on selected aspects of the employment relationship and ignores others.

I've played a small part in empirically examining the correspondence between these things and the occupational categories that form part of the NS-SEC. In McGovern et al. (2007) together with my colleagues I argue that in principle there is no reason why "social class" indicators should not be measured at the level of the individual job and that if they were, the categories of the NS-SEC would be redundant.

In practice, this is unrealistic. No general purpose survey will collect the detailed information required for accurate measurement of the real variables of interest and thus it makes sense to chunk the world of employment up into occupations and then aggregate these into "classes" as long as we have evidence that these classes serve as adequate proxies for the underlying variables that practicalities dictate we can't measure directly. Goldthorpe can logically justify the "classification" of the underlying dimensions by saying: "I'm creating a tool that can be used in situations and with data where there are no direct measures of the things I'm interested in. By well validated groupings of occupations I can capture a substantial part of what I mean by 'social class'. "

Both the market researchers and Goldthorpe can justify their classifications by their objectives. Let's now ask whether Savage et al. can do the same. What do they expect their class typology to explain?

A close reading of the article itself gives no answer. The closest I can come to a clue is a statement in the *Guardian*'s 'Comment is free' section by the lead author: *"The concept of class matters…"* we are told *"...because we need a way of connecting accentuating economic inequalities to social and cultural differences which permeate our society. Rather than seeing our lifestyles and social networks as somehow separate from economic inequalities, there are overlaps that can work together to produce social advantage and disadvantage".*

That lifestyles, features of social networks, cultural differences and economic inequalities are correlated is not something that many would doubt but the demonstration of such a banal fact surely cannot be the principal objective of Savage et al.'s efforts. What is clear is that their

typology cannot have any role in *explaining* differences in the distribution of social and cultural capital for the simple reason that all of these things are built into the *definition* of their typology.

3. On sample selection

The GBCS is a self-completion questionnaire administered over the internet. Internet surveys are known to be subject to self-selection bias. Savage et al. acknowledge this and demonstrate its consequences for the GBCS by benchmarking against the NS-SEC typology they elsewhere affect to disdain. Managers and professionals are vastly over-represented. Applying post-stratification weights is not a viable option because this only corrects for selection on variables with a known (or reliably estimable) distribution in the population of interest. Some variables of relevance – primarily the standard demographic face sheet variables – are available, others – the distribution of cultural and social capital – are not, yet are obviously "mission critical". Thus an additional source of information about the observable sources of response bias is needed. Enter the GfK survey.

Despite having over 160,000 responses to the GBCS the information that Savage et al. base their empirical claims on is a quota sample (the so called GfK sample) of 1026 respondents. This is a flimsy source upon which to build their inductive edifice but without it the GBCS data is essentially useless. In fact it turns out that even with it the GBCS is less valuable than they lead us to believe.

Given that the quota sample is the bedrock of the whole enterprise, and the GBCS is for analytical purposes set aside, we should be told something substantial about it. All we are in fact told is that that both GfK and the authors are satisfied that "the demographics are nationally representative".

I could go on to outline the well-known problems of making inferences from sample data selected by a non-probability method but I will leave this for others to take up. Suffice to say that with a probability sample there is a coherent procedure for making sample to population inferences including, most importantly, inferences about uncertainty: with a quota sample there is none, unless one is prepared to assert that the "correct" data generating model has been

identified. Savage et al. have a mountain of highly self-selected poor quality data sitting on top of a mole hill of (slightly) better quality data. Ultimately they want to use the latter to make a sensible calibration of the former, but they have to do this without any basis for assessing the uncertainty surrounding the numbers they estimate.

To take just one example of the difficulties this creates, much fuss is made about the discovery of a social class category they call the "elite". They tell us that this is about 6 per cent of the population. If the data were from a simple random sample a rough 95% confidence interval would be ± 1.5 per cent. But that is foolishly optimistic and takes no account of any other source of variance inflating error. One could easily double this number. So the "elite" could, on the basis of these data, be 3 per cent of the population, but it could also be 9 per cent of the population. Nine per cent is so large that Savage et al. might want to think again about the appropriation of the "elite" label. They have no basis, other than wishful thinking, for favouring either number (or anything in between).

The seven social classes that Savage et al. find fall out of a so called latent profile model estimated on the GfK quota sample plus 1 case consisting of the 161,400 responses to the GBCS each given a weight of 1/161400. This means that the parameter estimates that determine how respondents are allocated to social classes are derived almost entirely from the quota sample. The GBCS respondents are then allocated to the latent class categories on the basis of the model. There will be a considerable amount of uncertainty in doing so because the parameters of the latent profile model will not be well determined, but let's suspend our disbelief for there is another snag.

Even if we accept that the model enables us to allocate individuals in the GBCS data to the right classes, it does not allow the relationship between the class categories and any variable in the GBCS external to the latent profile model to be estimated without bias. This is important, for it is in fact what Savage et al. go on to do, for instance when they tabulate data on occupations conditional on class membership, or when they map the density of class membership by geographical region. Calibration of the GBCS by the GfK only works for variables that have been observed in both surveys but it cannot make an appropriate adjustment for selection into the GBCS on the basis of unobservables.

Imagine a population in which we are able to calculate scores on a latent variable by combining a suitable statistical model with information on three observed variables, say economic capital, cultural capital and social capital. Let's call this variable X. We also have information on two other variables. The first combines all the factors that influence the propensity of an individual to respond to internet based surveys. For the sake of concreteness let's call it "curiosity" or C for short. The second is a variable of interest that is not included in our latent variable model, let's call it Y, for example a simple 1/0 dichotomy indicating whether or not an individual is a university graduate.

In the population it is straightforward to calculate $P(Y=1|X)$ the probability that Y equals 1 conditional on X. Now assume the following simple selection rule: there is a threshold T such that if $C > T$ an individual will respond to an internet survey and if $C <= T$ an individual will not respond. Further assume that X, C and T are all positively correlated. When we carry out an internet survey what we observe is $P(Y=1|X,C>T)$ – the probability that Y equals 1 conditional on X and that C is greater than the threshold value - and this will not, except in uninteresting cases, be equal to $P(Y=1|X)$. The size of the difference will depend on the strength of the inter-correlation between X, C and Y and the point at which we set T.

We know that there is strong self-selection into the GBCS on the basis of observable characteristics and it is implausible to assume that there is no self-selection on the basis of unobserved "curiosity". Even if we correctly assign a GBCS respondent on the basis of the GfK to be a member of the "traditional working class" they would, on average have a "curiosity" score that was higher than the average in their class and because C is correlated with Y they would also have a higher than average propensity to be a graduate than a non-graduate.

The magnitude of the bias induced by selection into the GBSC on the basis of unobservables is unknown and, I believe, not estimable from these data. It would however be heroic to assume it didn't exist and doubly heroic to assume that it wasn't compounded by a similar selection bias affecting inclusion in the GfK quota sample.

Many readers of this journal might not appreciate the implications of the preceding argument. To put it in a nutshell, it is highly likely that any comparison between Savage et al.'s class categories with respect to variables contained in the GBCS but not used to derive the class categories – such as for instance educational attainment, type of university attended, geographical location and so forth - will be biased to a degree that may be substantively significant.

4. On model selection

We are told little by Savage et al. about the discovery of their seven social classes. They state that the Bayesian Information Coefficient is minimized when seven latent classes are assumed. We are not told how much worse things are when six or eight latent classes are assumed, or the extent to which these solutions resemble the seven class solution. We are given no sense whatsoever of model uncertainty. In fact there is a dark secret. The selection of seven classes is a function of the sample size of the GfK survey and without well-articulated theoretical grounds for distinguishing the classes it could not be otherwise. It is just as well that Savage et al discard the GBCS data. If they had estimated their latent profile model with 161,400 cases they would conclude that there were rather more than seven classes. The number of classes is a consequence of a decision about how much data to collect. This does not look like cleaving nature at its joints.

Savage et al. will no doubt claim that they are pursuing an open-minded inductive strategy but this simply won't do. A sensible model selection strategy must be based not only on formal statistical criteria but also on whether the model makes sense in terms of the configuration of variables that go into it; but it is just too facile to pass off unconstrained post-hoc interpretations as though they are facts of nature. This is particularly true when one allows oneself, as Savage et al. do, to go on a fishing expedition that encompasses variables external to the latent profile model itself. It is permissible to use external variables to validate a typology once you have made it, but peeking at them while making decisions about whether there are six, seven, eight or seventy-seven latent classes is having your cake and eating it.

5. On Measurement

In the course of discussing Goldthorpe's class schema Savage et al. say *"...the schema has been shown to be of less use in explicating the wider cultural and social activities and identities ... which do not appear to be closely linked to people's class position, as defined by the Goldthorpe class schema..."*

This is a bold claim, but it is not supported by Savage et al.'s own Figure 1 (pp 227). Following a line from NW to SE is a dimension that puts high-brow cultural taste at the NW pole (people visit stately homes, go to the theatre, museums and galleries etc) and low-brow taste in the SE corner. Their Figure 2 (pp 228) which projects "supplementary points" into the space of Figure 1 shows that the categories of the conventional NS-SECs fall from top to bottom roughly in a NW to SE fashion. This means that they are indeed predictive of the first cultural dimension.

Savage et al.'s initial belief is wrong. We are told that one of the motivations for the project is that conventional class categories are inadequate for the sociological analysis of the sorts of cultural tastes and practices they are interested in. But when their own data show that they are mistaken, rather than abandoning their initial premise they carry on regardless.

What of the second dimension in Figure 1? This is given the label "emerging cultural capital" though in what sense it is emerging and what it is emerging from we are never told. It runs from SW to NE and it distinguishes the sorts of things that young people like doing from the sorts of things that older people dislike or don't do. There is no doubt that younger people do and enjoy things that older people don't but what has this to do with social class? Do people's tastes and activities change as they get older? Yes. Do people's tastes depend on when they were born? Yes. What dimension 2 represents is a mixture of life-cycle and cohort differences. Nobody denies that the young and old are different. It's a novel claim however to assert that they are different social classes. I'll return to the implications of this absurdity in section 6.

Savage et al. are to be commended for acknowledging that wealth and not just income is important to their concerns, but they underestimate the difficulty of collecting reliable wealth data. Respondents to the GBCS are asked to provide information about savings and pensions. However expecting respondent's to provide an accurate estimate of the value of their pension is a stretch too far.

Before we even consider whether their assets are actually liquid (try as a 30 year old to borrow money against future pension income), we have to confront the difficulty respondents will have in calculating the asset's present discounted value. This will involve choosing a suitable discount rate and making a guess about how long they are going to survive after reaching pensionable age. In principle there is no difficulty in arriving at a sensible number: in practice it strains credibility to assume that many people will have bothered to do it.

6.On social classes

In this section I invite a further suspension of disbelief. Let's pretend that there are convincing replies to all of the points I've made so far and we take Savage et al.'s seven classes seriously.

Difficulties immediately arise when Savage et al. interpret the seven social classes they discover (Tables 5 & 6, pp 230). Normally one would do this in terms of the profiles of the observed variables associated with each of the latent classes. But looking at Table 6, it becomes apparent that this is not straightforward. There are differences between the classes but which differences are important?

In fact the interpretation and certainly the labels applied to the classes actually come from information quite extraneous to the latent profile analysis itself (Table 7, pp 231). Bizarrely one of the key ingredients is the percentage classified in the NS-SEC categories that Savage et al. claim are inadequate.

Age also seems to play a prominent part in the construction of the classes. While this would be unobjectionable if the objective was market segmentation, it is difficult to see how it can be reconciled with conventional understandings of social class. Usually classes are thought of as groups within which individuals could potentially spend the whole of their lives. However Savage et al.'s inductive method allows classes to be distinguished by tastes and activities that have a strong age gradient. The implication is that individuals can either grow out of their class,

because they become too old to get down to the gym or become stuck because of the cultural tastes they acquired during early adulthood.

Table 7 is a critic's gold-mine. Take for instance the "traditional working class". The average age of this group is sixty-six! They have modest incomes, are low on "emerging cultural capital" yet thirty per cent have, or had, jobs that would be classified as professional or managerial. A large proportion of this group are pensioners. Pensioners are an important interest group, but what analytical insight do we get from calling them a social class?

Life-cycle plays a role in distinguishing what Savage et al. term the "elite" and the "established middle class". The latter differ from the former principally by having smaller incomes, less savings and living in cheaper houses. In terms of social and cultural capital they don't differ at all. When we look at Table 7 we find that the "established middle class" are on average eleven years younger than the elite and other differences are quite marginal. A substantial part of the difference could plausibly be attributed to life-cycle stage and having a stake in the South-East's housing market. Again we must ask: what is gained by relabeling age and now geography as "class"?

Life-cycle also plays a role in the identification of the "emergent service workers". These are young people who go to gigs, join a gym and watch sport. They have few savings and little invested in home ownership. If you were to say that this is a life-style group, say young single people or young childless couples then there would be little to object to: but a social class? It's a strange sort of social class that people will grow out of simply by ageing, getting married and having kids.

In short Savage et al. fail to take the content of their typology seriously. If they did they would have to confront consequences which should give them cause to rethink their approach. For example they would have to accept that people could change their social class at will simply by turning off the Beatles and turning on to Beethoven. Of course people do change class but if it was as easy as this the solution to Britain's so called "social mobility problem" would have been spotted long ago.

7. Conclusions.

My conclusion is that for the reasons I outline above and for others that space limitations prevent me from mentioning the GBCS is a fiasco. It is so theoretically and methodologically flawed that it can contribute little of value to our understanding of the structure of systematic social inequality in the UK.

I want to make it easy for Savage et al. to reply to my criticisms so I will end with the questions I think they need to address: What is your typology meant to explain? Why should we have confidence in a typology built on the basis of such a small amount of data? What will you do when your method is applied to a larger amount of data and you discover, as you undoubtedly will, a larger number of classes? What use can the GBCS (as opposed to the GfK) data actually be put to? Do you accept that your data shows that cultural consumption is related to conventional measures of social class? What is gained by relabeling age-cohort and life-cycle stage as "social class"? And finally, do you really believe that changing one's social class can be a matter of getting out of bed and making a serious effort to like Brahms or to attract a few more Facebook friends?

References

Goldthorpe, J H (2000) *On Sociology*, Oxford, Oxford University Press.

McGovern, P, Hill, S, Mills, C and White, M (2007) *Market Class and Employment*, Oxford, Oxford University Press.

Rose, D and Pevalin, D J (eds.) (2003) *A Researcher's Guide to the National Statistics Socio-economic Classification*, London, Sage Publications.

Savage, M. (2013) Concerned about the BBC's Class Calculator? Let me explain, *The Guardian, Comment if free*, 10th April. http://www.guardian.co.uk/commentisfree/2013/apr/10/bbc-class-calculator.

Savage, M., Devine, F., Cunningham, N., Taylor, M., Li, Y., Hjellbrekke, J., Le Roux, B., Friedman, S. and Miles, A. (2013) A New Model of Social Class? Findings from the BBC's Great Class Survey Experiment, *Sociology*, 47 (2) 219-250.

Colin Mills is a Reader in the Department of Sociology, University of Oxford and a Professorial Fellow of Nuffield College.