



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

Journal:	<i>Sociology</i>
Manuscript ID:	Draft
Manuscript Type:	Article
Keywords:	Social Class, Measurement, Great British Class Survey

SCHOLARONE™
Manuscripts

view Only

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

Abstract

Savage et al. (2013) claim they have derived a new model of the British class structure. What they have actually done is something much less sociologically significant and more akin to the life-style profiling carried out by market research companies interested in market segmentation. They stress the innovative use of an internet survey and the large amount of information at their disposal. Their enthusiasm is shown to be ill-founded. Despite the hoopla surrounding the BBC's Great British Class Survey (GBCS), it turns out that it plays no serious role in the generation of Savage et al.'s class typology. What they have actually done is carried out a largely theory free (though Bourdieu inspired) data dredging exercise on a tiny amount of low quality data. What they discover is an arbitrary typology the main outlines of which are heavily influenced by a quite irrelevant contingent fact – the size of their sample.

1. Introduction

In this comment I am severely critical of an article by Savage et al. (2013) recently published in the pages of this journal. My contention is that 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. My argument is developed in six sections. Firstly I deplore a piece of terminological obscurantism, secondly, and more importantly, I outline what I believe is the standard view of the nature of typologies and compare the practical implications of that view to Savage et al.'s approach. Thirdly I deal with the implications of their data collection strategy and in particular the limitations it imposes on their ability to make valid statistical inferences. Fourthly I point out the uncomfortable implications of their model selection strategy. Fifthly I reveal their inability to read their own data 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 theoretical consistency for their model.

2. On the use of words

The word “capital” in conjunction with “cultural” and “social” has great obfuscatory potential and features prominently in Savage et al.’s article. The idea that people have cultural and social resources is perfectly sensible but the idea that they have cultural and social “capital” is not.¹ I can’t lend my cultural capital to another person or sell my social capital to them for their personal and exclusive use. While my economic capital can be liquidated my cultural and social resources cannot be alienated and a profit thus realized on them. My capacity to read, appreciate and enjoy Dostoyevsky in Russian may profit me, my children and Uncle Tom Cobley, but I can’t transfer my capacity to them (though I can encourage them to develop their own). Neither can I sell to them my love of fine wines or my taste for Bulgarian art house cinema. I can however, sell my Ferrari and give the proceeds to Bill Brewer thus increasing his economic capital and diminishing my own. Nothing is gained by appending the word “capital” to the words “cultural” or “social” except a veneer of chic profundity.

Generally speaking linguistic conventions that obscure important distinctions are bad tools for doing science and should not be perpetuated even when they have received the imprimatur of a guru figure. However, I fear in this case the battle for reason is lost and as I don’t want the substance of my comment to be buried under a dispute about the meaning of words I will, in what follows, defer to the now accepted usage while deploring the necessity of having to do so.

3. On Typologies

Concepts - like for instance social class - are abstract entities. We use them to selectively organize experience. This doesn’t make the experiences themselves (whatever they are) any less real for the people that have them. In everyday life people use a term like “social class” to describe who they are, what they are doing, how they have been treated and what they feel. However, social scientists have no obligation to use the same words in the same sense as lay actors or only to use concepts that such actors would be likely to use themselves.

1
2
3 Typologies or categorization schemes are closely related to concepts; they are the step we take in
4 order to make abstract ideas more concrete. There can be as many typologies as there are points
5 of view. We are free to fill a concept with whatever content we like and then derive a typology
6 from it. Things may get a little confusing when lay people use the same word to denote
7 something different from what we intend or when another scientist fills a familiar concept with
8 unfamiliar content. But we can live with this as long as everyone is clear about what they are
9 doing. The correct way to use the term “social class” has not been passed down on tablets of
10 stone following an encounter with a burning bush. Without an agreed convention all argument
11 about whether something is really a correct usage of “social class” is just so much wrangling
12 about words.
13
14
15
16
17
18
19
20
21
22

23 Typologies should be made for a purpose and we should ask both whether that purpose is worth-
24 while and whether the proposed typology is “fit for purpose”. We, as observers, are free to
25 organize what we can learn of other people’s experiences and circumstances with any conceptual
26 scheme we like, but, there are logical constraints in that we must acknowledge the consequences
27 of doing things one way rather than another. This point is of some consequence and I’ll return to
28 it below in section 7.
29
30
31
32
33
34

35 Some conceptual schemes are simply of more practical use than others in the sense that they, as
36 Linnaeus put it, paraphrasing Plato, cleave nature at its joints. Not everyone need have the same
37 purpose and cleave the same joints, but a clear purpose they must have. Normally that purpose is
38 given by the part that a concept (and its associated typology) plays in a theory that is meant to
39 explain or predict something and naturally this entails spelling out clearly what it is that the
40 concept is to play a part in explaining or predicting. From this it follows that a typology is a
41 practical tool. Inventing a typology for its own sake is a pointless exercise and can have no
42 scientific value.
43
44
45
46
47
48
49
50

51 Let me give two examples, one short the other longer, to make my line of reasoning more
52 concrete. In commercial market research market segmentation is a procedure that builds
53 predictive models to help clients target products at people that are likely to buy them. In this kind
54 of exercise all that matters is that good predictions are produced; the prediction model itself can
55
56
57
58
59
60

1
2
3 be built in an entirely inductive manner. The categories that are produced, which may be an ad
4 hoc combination of all sorts of demographic, life-style and dispositional characteristics, are
5 validated by their predictive value. If it works that is good enough and nobody has to pretend that
6 the categories consist of real social groups.
7
8
9

10
11
12 Now consider John Goldthorpe's most recent statement of what social class is about (Goldthorpe,
13 2000, 206-29) a view that underpins the construction of the National Statistics Socio-economic
14 Classification (NS-SEC) (Rose & Pevalin, 2003) and the European Socio-economic
15 Classification - ESec - (Rose & Harrison, 2010). It focuses our attention on some selected
16 aspects of the employment relationship - output/input monitoring difficulty and human asset
17 specificity – and ignores others. What precisely these terms mean need not concern us here.
18
19
20
21
22
23

24
25 I've played a small part in empirically examining the correspondence between these things and
26 the occupational categories that form part of the NS-SEC. In Author A ref (2007) together with
27 my colleagues I go so far as to argue that in principle there is no reason why “social class”
28 indicators should not be measured at the level of the individual job and that if they were, then, in
29 a sense, the categories of the NS-SEC would be redundant.
30
31
32
33
34

35
36 In practice, of course, this is unrealistic. No general purpose survey will collect the detailed
37 information required for accurate measurement of the real variables of interest and thus it makes
38 perfect sense to chunk the world of employment up into occupations and then aggregate these
39 into "classes" as long as we have evidence that these classes serve as adequate proxies for the
40 underlying variables that practicalities dictate we can't measure directly. This is a matter of
41 pragmatics disciplined by evidence. Goldthorpe can logically justify the "classification" of the
42 underlying dimensions by saying: “I'm creating a tool that can be used in situations and with data
43 where there are no direct measures of the things I'm interested in. By well validated groupings of
44 occupations I can capture a substantial part of what I mean by “social class” - ie the employment
45 relations relevant aspects of jobs”.
46
47
48
49
50
51
52
53

54
55 Goldthorpe had a clear reason for creating his original class categories. He wanted to describe
56 patterns of mass social class mobility in Britain and therefore he made distinctions that were
57
58
59
60

1
2
3 sensible given that purpose. When that purpose changed - for example to describing cross-
4 national patterns of mass social class mobility - further distinctions became necessary (farmers
5 are not that numerous in Britain but they are and were in Poland). Hence the evolution from the
6 original class schema to the EGP schema and its variants
7
8
9

10
11
12 Both the market researchers and Goldthorpe have a justification for their classifications in terms
13 of their objectives. Let's now ask whether what Savage et al. do can be justified along similar
14 lines? What do they expect their class typology to explain? How would they expect to use it?
15
16
17

18
19 These turn out to be very difficult questions to answer either from a close reading of the article
20 itself or of the various statements that Savage and his colleagues have subsequently made. The
21 closest I can come to a clue is a statement in *The Guardian's* 'Comment is free' section by the
22 lead author: "*The concept of class matters...*" we are told "...because we need a way of
23 connecting accentuating economic inequalities to social and cultural differences which permeate
24 our society. Rather than seeing our lifestyles and social networks as somehow separate from
25 economic inequalities, there are overlaps that can work together to produce social advantage
26 and disadvantage".
27
28
29
30
31
32
33

34
35 That lifestyles, features of social networks, cultural differences and economic inequalities are
36 correlated is not something that many would doubt, though surprisingly Savage et al. provide
37 rather little straightforward evidence of their degree of inter-correlation. But the demonstration
38 of such a banal fact surely cannot be the principal objective of Savage et al.'s efforts.
39
40
41
42

43
44 What is clear is that their typology cannot have any role in *explaining* differences in the
45 distribution of social and cultural capital for the simple reason that all of these things are built
46 into the *definition* of their typology. If they want to examine and possibly explain the variation
47 in the extent, strength and depth of a population's social network ties - a perfectly reasonable
48 thing to want to do - a tool that already is partly defined in terms of the explanandum, will not
49 help them to discover any new knowledge. The same is true if they want to explain variation in
50 life-style, cultural consumption or cultural participation.
51
52
53
54
55
56
57
58
59
60

1
2
3 To carry out an empirical investigation to demonstrate that economic, social and cultural capital
4 are correlated with each other seems to me to be setting the intellectual bar a little too low. Who
5 has ever claimed that they are not? We might also ask ourselves what is revealed by relabeling
6 the perfectly well understood distinctions of income, wealth, cultural preference and social
7 connectivity as class?
8
9

10 11 12 13 14 4. On sample selection 15 16

17
18 Savage et al. are refreshingly candid in mentioning the flawed nature of their data, but they are
19 less forthcoming about the implications of these flaws. In fact mostly they just plough on
20 regardless. In their Panglossian world the data are allowed to be unsuitable for the task at hand
21 but nevertheless substantive conclusions are breezily drawn, as if, by wanting it to be so, all
22 difficulties can magically be made to disappear.
23
24
25
26

27
28 The GBCS is nothing more than a self-completion questionnaire administered over the internet.
29 By their very nature self-completion questionnaires don't admit sophistication. Respondents are
30 left to their own devices to make what they can of the questions. The great attraction of using the
31 internet is that large amounts of data can be collected cheaply, but this comes at the cost of
32 losing control over who is actually providing the answers. All a respondent needs is a valid email
33 address. If they have more than one there is nothing, other than boredom, to prevent them from
34 completing the questionnaire on multiple occasions. And if they do, there will be hidden non-
35 independence between observations.
36
37
38
39
40
41
42

43
44 Internet surveys which have as their target the general population are known to be subject to self-
45 selection bias. Savage et al. don't try to hide this and in fact surprisingly demonstrate its
46 consequences for the GBCS by benchmarking against the very NS-SEC typology they elsewhere
47 affect to disdain. Managers and professionals are vastly over-represented while the working-class
48 are vastly under-represented. There is only a fifth as many routine and technical craft workers in
49 the GBCS sample as there should be and only a third as many routine workers. Applying post-
50 stratification weights derived from an existing large and reliable survey such as the Quarterly
51 Labour Force Survey (QLFS) is not a viable option because this can only correct for selection
52
53
54
55
56
57
58
59
60

1
2
3 related to variables with a known (or reliably estimable) distribution in the population of interest.
4
5 Some variables of relevance – primarily the standard demographic face sheet variables – are
6
7 available in the QLFS, others – the distribution of cultural and social capital – are not, yet are
8
9 obviously “mission critical” for Savage et al. Thus an additional source of information about the
10
11 observable sources of response bias is needed. Enter the GfK survey.

12
13 Despite having over 160,000 responses to the GBCS the information that Savage et al. use to
14
15 derive their empirical claims comes from a small quota sample (the so called GfK sample) of
16
17 1026 respondents. This is a surprisingly flimsy source upon which to build their grand inductive
18
19 edifice but without it the GBCS data is, as they realize, essentially useless. In fact it turns out that
20
21 even with it the GBCS is less valuable than they would lead us to believe, but I’ll return to this
22
23 point below.

24
25 Given that the quota sample is the bedrock of the whole enterprise, and the GBCS is for
26
27 analytical purposes set aside, we might reasonably expect to be told something substantial about
28
29 it. What quota controls were applied? Were they interlocking? What was the target population?
30
31 Which weights, if any, were applied and where did they come from? All we are in fact told is
32
33 that that both GfK and the authors are satisfied that "the demographics are nationally
34
35 representative". It’s a curious sort of science that substitutes the firmness of the beliefs of the
36
37 investigators for publically verifiable information.

38
39 I could go on to outline the well-known problems of making coherent inferences from sample
40
41 data selected by a non-probability method. This is an old story, though no less important for that,
42
43 but I will leave it for others to take it up.² It is sufficient to say that with a probability sample
44
45 there is a coherent procedure for making sample to population inferences including, most
46
47 importantly, inferences about uncertainty: with a quota sample there is none, unless one is
48
49 prepared to assert that the "correct" data generating model has been identified. I doubt that
50
51 Savage et al. will wish to claim the latter - it would, for one thing, be inconsistent with their
52
53 general philosophy of data analysis which, in as far as I understand it, seems to favour more
54
55 "exploratory" approaches.

56
57 So Savage et al. have a problem. They have a mountain of highly self-selected poor quality data
58
59 sitting on top of a mole hill of (slightly) better quality data. Ultimately they want to use the latter
60

1
2
3 to make a sensible calibration of the former, but they have to do this without any basis for
4 assessing the uncertainty surrounding the numbers they estimate.
5
6

7
8 To take just one example of the difficulties this creates, much fuss is made about the discovery
9 of a social class category they call the “elite”. They tell us that this is about 6 per cent of the
10 population. If the data were from a simple random sample a rough 95% confidence interval
11 would be ± 1.5 per cent. But that is foolishly optimistic and takes no account of any other source
12 of variance inflating error. To be on the safe side one could easily double it. So the “elite” could,
13 on the basis of these data, be 3 per cent of the population, which, though still an implausibly
14 large number, is moving in the right direction, but it could, on the other hand, be 9 per cent of the
15 population. Nine per cent is so large that Savage et al. might want to think again about the
16 appropriation of the “elite” label.³ My point is very simple. Savage et al. have no strong basis for
17 favouring either of these numbers (or anything in between) other than wishful thinking.
18
19
20
21
22
23
24
25

26
27 The seven social classes that Savage et al. purport to find fall out of the application of a
28 particular statistical model – a so called latent profile model – to the 1026 responses in the quota
29 sample plus 1 case consisting of the 161,400 responses to the GBCS each given a weight of
30 1/161400. This means that the parameter estimates that determine how respondents are allocated
31 to social classes are derived to all intents and purposes entirely from the small quota sample and
32 that the influence of the GBCS is negligible. At least the quota sample rescues, if only at a
33 rhetorical level, a slightly desperate situation. After all it would have been embarrassing for both
34 the BBC and the academics to have collected all that data and not used it.
35
36
37
38
39
40
41

42
43 So far so good: GBCS data can be allocated to the latent class categories on the basis of a model
44 that is based almost entirely on the information contained in the GfK survey. There will be a
45 considerable amount of uncertainty in doing so because the parameters of the latent profile
46 model will not be well determined, but let’s suspend our disbelief for there is another serious
47 snag.
48
49
50
51

52
53 Even if we accept that the model enables us to allocate individuals in the GBCS data to the right
54 classes, it does not, as in fact Savage et al. appear to concede, allow the relationship between the
55 class categories and any variable in the GBCS external to the latent profile model to be estimated
56
57
58
59
60

1
2
3 without bias. This turns out to be important, for this is in fact what Savage et al. go on to do, for
4 instance when they tabulate data on occupations conditional on class membership, or when they
5 map the density of class membership by geographical region.
6
7

8
9
10 Once one thinks about it the problem is obvious. Calibration of the GBCS by the GfK only
11 works for variables that have been observed in both surveys but it cannot make an appropriate
12 adjustment for selection into the GBCS on the basis of unobservables.
13
14

15
16
17 Consider the following set-up. Imagine a population in which we are able to calculate scores on a
18 latent variable by combining a suitable statistical model with information on three observed
19 variables, say economic capital, cultural capital and social capital. Let's call this variable X . We
20 also have information on two other variables. The first combines all the factors that influence the
21 propensity of an individual to respond to internet based surveys. For the sake of concreteness
22 let's call it "curiosity" or C for short. The second is a variable of interest that is not included in
23 our latent variable model, let's call it Y , for example a simple 1/0 dichotomy indicating whether
24 or not an individual is a university graduate .
25
26
27
28
29
30
31

32
33 In the population it is straightforward to calculate $P(Y=1|X)$ the probability that Y equals 1
34 conditional on X . Now assume the following simple selection rule: there is a threshold T such
35 that if $C > T$ an individual will respond to an internet survey and if $C \leq T$ an individual will not
36 respond. Further assume (for simplicity) that X , C and T are all positively correlated. When we
37 carry out an internet survey what we observe is $P(Y=1|X, C > T)$ – the probability that Y equals 1
38 conditional on X and that C is greater than the threshold value - and this will not, except in
39 uninteresting cases, be equal to $P(Y=1|X)$. The size of the difference will depend on the strength
40 of the inter-correlation between X , C and Y and the point at which we set T .
41
42
43
44
45
46
47
48

49 We know that there is strong self-selection into the GBCS on the basis of observable
50 characteristics and it is completely implausible to assume that there is no self-selection on the
51 basis of unobserved "curiosity". Even if we correctly assign a GBCS respondent on the basis of
52 the GfK data to be a member of the "traditional working class" they would, on average, be a
53 highly unusual member of this class. They would have a "curiosity" score that was higher than
54
55
56
57
58
59
60

1
2
3 the average in their class and because C is correlated with Y they would also have higher than
4 average propensity to be a graduate than a non-graduate.
5
6
7

8
9 The magnitude of the bias induced by selection into the GBSC on the basis of unobservables is
10 unknown and, I believe, not estimable from these data. It would however be heroic to assume it
11 didn't exist and doubly heroic to assume that it wasn't compounded by a similar selection bias
12 affecting inclusion in the GfK quota sample.
13
14
15
16

17
18 It's conceivable that many readers of this journal might not appreciate the full implications of the
19 preceding argument. To put it in a nutshell, it is highly likely that any comparison between
20 Savage et al.'s class categories with respect to variables contained in the GBCS but not used to
21 derive the class categories – such as for instance educational attainment, type of university
22 attended, geographical location and so forth - will be biased to a degree that is likely to be
23 substantively significant.
24
25
26
27
28

29 30 5. On model selection 31 32

33 We are told remarkably little by Savage et al. about the discovery of their seven social classes.
34 They blandly state that the Bayesian Information Coefficient (BIC) is minimized when seven
35 latent classes are assumed. That may be good enough for the majority of *Sociology's* readers
36 most of whom will have to trust that the authors know what they are doing. But should we have
37 to place our trust in them?
38
39
40
41
42

43
44 We are not told for instance how much worse things become when six or eight latent classes are
45 assumed, or the extent to which a six or eight class solution resembles the seven class solution.
46 In fact we are given no sense whatsoever of model uncertainty or indeed why the gospel
47 according to BIC is to be preferred to any other model selection criterion.⁴ In fact there is a dark
48 secret hidden away here. The selection of seven classes is mostly determined by the small sample
49 size of the GfK survey and without any well-articulated theoretical grounds for distinguishing
50 the classes it is difficult to see how it could be otherwise.
51
52
53
54
55
56
57
58
59
60

1
2
3 In this respect it's just as well that Savage et al., to all intents and purposes, discard the GBCS
4 data. If they had estimated their latent profile model with 161,400 cases the logic of their *modus*
5 *operandi* would have forced them to conclude that rather than there being seven classes there
6 were many more. The moral of the story is this: if you have no purpose in mind behind the
7 application of a data reduction technique or no well-articulated expectations, then the data will
8 lead you by the nose. But "the data", or at least the amount of data, is completely arbitrary and so
9 will be your conclusions. The Columbus like discovery of seven social classes is simply a
10 consequence of an earlier decision about how much data to collect. This does not look like
11 cleaving nature at its joints.
12
13
14
15
16
17
18
19

20
21 Perhaps you are thinking this is unduly dismissive. Surely, I hear you say, Savage et al. are
22 pursuing an open-minded inductive strategy rather than a closed minded deductive approach.
23 This simply won't do. Let's agree that a sensible model selection strategy must be based not
24 only on formal statistical criteria but also on whether the model makes sense in terms of the
25 configuration of variables that go into it; but it is just too facile to pass off unconstrained post-
26 hoc interpretations as though they are facts of nature. This is particularly true when one allows
27 oneself, as Savage et al. do, to go on an inductive fishing expedition that encompasses variables
28 quite external to the latent profile model itself.
29
30
31
32
33
34
35
36

37 There is nothing wrong with using external variables to validate a typology once you have made
38 it, but peeking at them while making inductive decisions about whether there are six, seven, eight
39 or seventy-seven latent classes is having your cake and eating it. It's an old saw, but none the
40 less true that any sociologist unable to come up with a satisfying interpretation of an arbitrary
41 number of latent classes should probably be considering a career change. We shouldn't confuse
42 the sociologist's rationalization with reality.
43
44
45
46
47

48 6. On Measurement

49
50 I want to turn now to the building blocks of the Savage et al. class model – some of the basic
51 questions answered by the respondents (to the GfK survey). But before I do I want to first slay a
52 particular dragon.
53
54
55
56
57
58
59
60

1
2
3 In the course of discussing Goldthorpe's class schema Savage et al. say "...the schema has been
4 shown to be of less use in explicating the wider cultural and social activities and identities ...
5 which do not appear to be closely linked to people's class position, as defined by the Goldthorpe
6 class schema..."
7
8
9

10
11 This is an interesting claim, but I wonder how it sits with the evidence they set before our eyes.
12 Savage et al.'s Figure 1 (pp 227) is a bi-plot derived from a multiple correspondence analysis of
13 a set of categorical measures of the extent to which respondents take part in various, broadly
14 defined, cultural activities – what they would call “cultural capital”.
15
16
17

18
19 If we follow a line from NW to SE we can discern a dimension that puts high-brow cultural taste
20 at the NW pole (people visit stately homes, go to the theatre, museums and galleries, dislike fast
21 food etc) and low brow taste in the SE corner. If we look at their Figure 2 (pp 228) which
22 projects so called "supplementary points" onto Figure 1 what do we see?
23
24
25

26
27 What we see is that the categories of the conventional NS-SECs which Savage et al. believe are
28 not closely linked to "cultural and social activities and identities" seem to fall from top to bottom
29 roughly in a NW to SE fashion. What this means is that they are indeed predictive of the first
30 cultural dimension of their Figure 1.
31
32
33

34
35 Savage et al.'s initial belief is simply wrong but as if this isn't bad enough they then ignore the
36 obvious implication. We are told that one of the motivations for the project is the fact that
37 conventional class categories are inadequate for the sociological analysis of the sorts of cultural
38 tastes and practices they are interested in. But when their own data show that they are mistaken,
39 rather than abandoning their initial premise they just plow on regardless and apparently hope that
40 no one will notice! Whether this is the result of *folie de grandeur* or lack of authorial
41 coordination we will, I suppose, never learn unless they care to tell us. However, the least we can
42 say is that they have in this respect been badly served by the journal's referees who should have
43 saved them from this particular embarrassment.
44
45
46
47
48
49
50

51
52 What of the second dimension in Figure 1? This is apparently a great discovery and given the
53 impressive label "emerging cultural capital" though in what sense it is emerging and what it is
54 emerging from is never vouchsafed. This dimension runs from SW to NE and once you grasp the
55 labels it is obvious that it distinguishes the sorts of things that young people like doing from the
56
57
58
59
60

1
2
3 sorts of things that older people dislike or don't do.
4
5 This interpretation is confirmed by inspection of Figure 2 which shows that the supplementary
6
7 age category points line up in the same direction.
8

9
10 I don't doubt that younger people do and enjoy things that older people don't but what has this to
11
12 do with social class? Do people's tastes and activities change as they get older? Yes, they do. Do
13
14 people's tastes depend on when they were born? Yes, I believe that is true too. So what
15
16 dimension 2 actually represents is a mix of life-cycle and cohort effects. Again we have
17
18 revelation by relabeling. Nobody denies that the young and old are different. It's a bold, and
19
20 indeed novel, claim however to assert that they are different social classes. I'll return to the
21
22 implications of this absurdity in section 7.

23
24 I'll only deal briefly with the way in which Savage et al. attempt to measure economic capital.
25
26 They are to be commended for at least acknowledging that wealth and not just income is
27
28 important. Unfortunately they seriously underestimate the difficulty of collecting reliable wealth
29
30 data. Respondents to the GBCS are asked to provide information about their savings, including
31
32 the value of their pensions. In one sense this is a very sensible thing to do. Many people have a
33
34 considerable proportion of their assets locked up in a pension plan. However expecting them to
35
36 provide a sensible estimate of the value of their pension is a stretch too far.

37
38 Before we even consider the basic conceptual question of whether their assets are actually liquid
39
40 (try as a 30 year old to borrow money against future pension income), we have to confront the
41
42 difficulty of calculating the present discounted value. Amongst other things this will involve
43
44 choosing a suitable discount rate and making a guess about how long one is going to survive
45
46 after reaching pensionable age. In principle there is no difficulty in arriving at a sensible number:
47
48 in practice it strains credibility to believe that many people will have bothered to do this whilst
49
50 completing the GBCS.

51
52 Let's moving on to social capital. This is supposed to be about connectivity - the amount of
53
54 contact that you have with other people and the social range of those contacts. It's a perfectly
55
56 respectable idea, but I don't see that it has any straightforward connection to the idea of social
57
58 class.
59
60

1
2
3 If you are a disadvantaged person I can see that having a broader rather than a narrower social
4 network may bring you some advantages. After all if you are a domestic-cleaner and become
5 unemployed and you know an investment banker they might employ you to clean their house or
6 look after their kids. It's not obvious that if you are an investment banker the crude economic
7 advantage of having a socially wider rather than a socially narrower set of acquaintances is going
8 to be that great. After all if you don't know someone who would love to look after your kids or
9 clean your house you can always find somebody from an agency or ask your next-door
10 neighbour who they use.
11
12

13
14 In practice Savage et al., while paying lip service to the notion of network diversity, just ignore it
15 and actually only use a measure of the average occupational status score of the respondent's
16 social contacts and a count of what they call the "mean number of social contacts reported". The
17 latter designation is, strictly speaking, inaccurate. It is in fact just a count of the affirmative
18 answers to a sequence of questions that ask: do you know socially anyone who is a X? where X
19 is any one of 34 different occupations. So someone who knows a large number of stockbrokers
20 and nobody else would get a score of 1, as would somebody who has but a single friend in the
21 world. This is not a very sensible way to estimate the size of an ego-centric network.
22
23

24
25 There is actually a very simple way to get an idea of the relative size of the ego-centric networks
26 in a population. You give people a list of names taken from the page of a telephone directory or
27 suchlike and you ask them how many people they know who are called Smith, Jones, Johnson
28 etc.
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45

46 7. On social classes

47
48 Taken together I think the criticisms I make above are sufficient to destroy the credibility of
49 Savage et al.'s intellectual project. However, I'm aware that many sociologists are unpersuaded
50 by methodological arguments alone so in this section I am going to invite a further suspension of
51 disbelief. Let's pretend that there are convincing replies to all of the points I've made so far and
52
53
54
55
56
57
58
59
60

1
2
3 that we can therefore take Savage et al.'s seven classes seriously. Let's do that and see where it
4 leads.
5
6

7
8 Difficulties immediately arise when Savage et al. attempt to interpret the seven social classes
9 they discover (Tables 5 & 6, pp 230). Normally one would do this in terms of the profiles of the
10 observed variables associated with each of the latent classes. But looking at Table 6, it becomes
11 apparent that this is not straightforward. There are differences between the classes but which
12 differences are important and which are trivial? We are given no criteria, either statistical or
13 theoretical for deciding and of the two it is the lack of theoretical criteria that is of greatest
14 importance.
15
16
17
18
19

20
21
22 In fact the interpretation and almost certainly the labels applied to the classes actually come from
23 information quite extraneous to the latent profile analysis itself (Table 7, pp 231). Bizarrely one
24 of the key ingredients is the percentage classified in or stemming from one of the occupationally
25 based NS-SEC categories that Savage et al. claim are inadequate for their purposes.
26
27
28
29

30
31 Age seems to play a rather prominent part in the construction of Savage et al.'s classes as is
32 made clear for instance by examining their Figure 2 which clearly shows that age is an important
33 correlate of "emerging" cultural capital. While this result is consistent with their methodological
34 approach and would be unobjectionable if their objective was market segmentation, it is more
35 difficult to see how this can be reconciled with more conventional understandings of social class.
36 Usually classes have been thought of as collectivities, or at least groupings, within which
37 individuals could potentially spend the whole of their lives.⁵ However Savage et al.'s inductive
38 method allows classes to be distinguished by tastes and activities that have a rather strong age
39 gradient attributable either to life-cycle or cohort specific causes. The implication is that
40 individuals can either "grow out of" their class, perhaps because they become too old to get
41 down to the gym or even more strangely become stuck because of the cultural tastes they
42 acquired during early adulthood.
43
44
45
46
47
48
49
50
51
52

53
54 Table 7 is in fact quite useful to a critic for it reveals the odd nature of the classes that are
55 discovered. Take for instance the "traditional working class". The average age of this group is
56 sixty-six! They have modest incomes, are low on "emerging cultural capital" yet thirty per cent
57
58
59
60

1
2
3 of them have, or had, jobs that would be classified as professional or managerial (how working-
4 class is that?). If the average age is sixty-six then a large proportion of this group are actually
5 pensioners. Pensioners are an important interest group, but what analytical insight do we achieve
6 by calling them a social class?
7
8
9

10
11 Life-cycle plays a role in distinguishing what Savage et al. term the "elite" and the "established
12 middle class". In terms of the variables tossed into their statistical food blender the latter differ
13 from the former principally by having smaller incomes, less savings and living in cheaper
14 houses. In terms of social and cultural capital they don't differ at all. When we look at Table 7
15 what we find is that the "established middle class" are on average eleven years younger than the
16 elite and other differences are quite marginal. A substantial part of the difference between these
17 two classes could quite plausibly be attributed to life-cycle stage plus having a stake in the
18 South-East's housing market.
19
20
21
22
23
24
25
26

27
28 Life-cycle influences also play a role in the identification of the "emergent service workers".
29 These are young people – their average age is thirty-four – who go to gigs, join a gym and watch
30 sport. They have only a small amount of savings and little invested in home ownership, neither
31 of which is surprising given their low average age. If you were to say that this is a life-style
32 group, say young single people or young couples without children then there would be little to
33 object to: but a social class? It's a strange sort of social class that people will grow out of simply
34 by ageing, getting married and having kids.
35
36
37
38
39
40
41

42 And this is the point. A fundamental flaw with Savage et al.'s typology is that they fail to take its
43 content seriously enough. If they did they would have to confront consequences which should
44 give them cause to rethink their approach. For example to bring up one obvious absurdity, they
45 would have to accept that people could change their social class at will simply by changing their
46 cultural preferences. By giving up Mills & Boon and taking up Dostoyevsky or turning off the
47 Beatles and turning on to Beethoven they could alter their class position. Of course people do
48 change class both inter and intra-generationally, but if it was as easy as this the solution to
49 Britain's so called "social mobility problem" would have been spotted long ago.
50
51
52
53
54
55
56
57
58
59
60

1
2
3 7. Conclusions.
4
5

6 I want to make it easy for Savage et al. to reply to my criticisms so I will end with a series of
7 questions that put as simply as I can manage the points I think they need to address. What is your
8 typology meant to explain? Why should we have confidence in a typology built on the basis of
9 such a small amount of data? What will you do when your method is applied to a larger amount
10 of data and you discover, as you undoubtedly will, a larger number of classes? What use can the
11 GBCS (as opposed to the GfK) data actually be put to? Do you accept that your own data shows
12 that cultural consumption is related to conventional measures of social class? Do you really
13 believe that changing one's social class can be a matter of getting out of bed and making a
14 serious effort to like Brahms or to attract a few more Facebook friends?
15
16
17
18
19
20
21
22

23 These are serious questions and I hope Savage et al. will take them seriously. If they do then I,
24 and more importantly the readers of this journal, will learn something from their answers.
25
26

27 References
28
29
30
31

32 Goldthorpe, J H (2000) *On Sociology*, Oxford, Oxford University Press.
33
34

35 Goldthorpe, J H (2007) "Cultural Capital": Some Critical Observations. *Sociologica*, 2, 2.
36 Available at <http://www.sociologica.mulino.it/journal/issue/index/Issue/Journal:ISSUE:2>
37
38

39 Kuha, J (2004) AIC and BIC: Comparisons of Assumptions and Performance, *Sociological*
40 *Methods and Research*, 33 (2) 188-229.
41
42
43

44 Rose, D and Pevalin, D J (eds.) (2003) *A Researcher's Guide to the National Statistics Socio-*
45 *economic Classification*, London, Sage Publications.
46
47

48 Rose, D and Harrison, E (2010) *Social Class in Europe: An Introduction to the European Socio-*
49 *economic Classification*, London, Routledge.
50
51

52 Savage, M Warde, A and Devine, F (2007) Comment on John Goldthorpe, *Sociologica*, 2,2.
53 Available at <http://www.sociologica.mulino.it/journal/issue/index/Issue/Journal:ISSUE:2>
54
55
56
57
58
59
60

1
2
3 Savage, M. (2013) Concerned about the BBC's Class Calculator? Let me explain, *The Guardian*,
4 *Comment if free*, 10th April. [http://www.guardian.co.uk/commentisfree/2013/apr/10/bbc-class-](http://www.guardian.co.uk/commentisfree/2013/apr/10/bbc-class-calculator)
5 [calculator](http://www.guardian.co.uk/commentisfree/2013/apr/10/bbc-class-calculator).
6
7

8
9 Savage, M., Devine, F., Cunningham, N., Taylor, M., Li, Y., Hjellbrekke, J., Le Roux, B.,
10 Friedman, S. and Miles, A. (2013) A New Model of Social Class? Findings from the BBC's
11 Great Class Survey Experiment, *Sociology*, 47 (2) 219-250.
12
13

14
15 Weakliem, D. (1999) A Critique of the Bayesian Information Criterion for Model Selection,
16 *Sociological Methods and Research*, 27 (3) 359-97.
17
18
19
20
21
22
23
24
25
26
27

28
29 ¹ On the incoherence of the concept of cultural capital see Goldthorpe(2007) and the reply by Savage, Warde and
30 Devine (2007) in the same issue.
31

32 ² One reason I don't want to dwell on the non-probability nature of the GfK sample is that it is obvious that Savage
33 et al. were in a tough spot when the GBCS failed to deliver data fit for purpose. They were forced to make the best
34 of a bad job without the resources to do much better. Though this explains how the sorry state of affairs came about,
35 it doesn't amount to a licence to carry on as if everything in the garden is rosy and it can't be part of a defence of
36 the scientific weaknesses of their paper. To understand is to forgive may be a reasonable motto in some parts of life,
37 but this doesn't wash in science.
38
39

40 ³ Nine per cent is in fact not that much less than the roughly eleven per cent of the population conventionally placed
41 in the highest category of the NS-SEC – higher managerial, administrative and professional occupations.
42

43 ⁴ On the technical issues involved see Kuha (2004) and Weakliem (1999)
44

45 ⁵ I'm grateful to John Goldthorpe for this point.
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60