Diversity, inclusion and intersectionality: analysing education attainment level and ethnicity in responses to citizens' assembly recruitment processes

Authors: Nick Gill, Tom Lord
Acknowledgements: Gill Wyness, Graham Smith

Introduction

The purpose of this article is to interrogate the Sortition Foundation's current "gold-standard" recruitment methodology for Citizens' Assemblies. The OECD have released eleven "good practice principles for deliberative processes for public decision making" of which two are "inclusiveness" and "representativeness". We wish to test whether our recruitment methodology is optimal for these two measures.

Our investigation will involve the analysis of data about respondents to two recent recruitment processes for Citizens' Assemblies. The main theses we set out to test were the following:

- T1. Our current recruitment process results in a disproportionately high number of highly educated people among our respondents pool.
- T2. This disproportion is exacerbated if one considers the pool of respondents from minority ethnic groups.

If we consider "level of education" to form part of an individual's socio-economic status, then Thesis T1 is in line with what we expect - that socio-economic status has an impact on who registers. In recognition of this, we almost always collect data about educational attainment, or other proxy measures for socio-economic status, when we ask people to register their interest for an assembly: we perform our stratified selection using this data in order to ensure that the resulting group of assembly members has a representative profile of socio-economic statuses, thereby ensuring a diversity of voices in the room.

In this piece we wish to understand more precisely whether the use of different socioeconomic indicators is effective in ensuring a range of educational backgrounds among selected participants. In the discussion below we will focus on our use of the "Index of Multiple Deprivation" as a socioeconomic indicator. We are able to draw some conclusions about how this measure interacts with both education and ethnicity.

Thesis T2 presents more intrinsic problems: Our stratified selection is “blind to intersections”. So, for instance, if we collect data on ethnicity and on educational attainment, then our algorithm will ensure that the range of ethnicities in the room, and the range of educational attainment in the room, both match the population profile. The algorithm does not, however, ensure that within any given ethnicity there will be a range of educational attainments. So we may end up with, say, a group whose black members have, on average, higher educational attainment compared to (a) the black population; and/or (b) the non-black assembly members.

Such a scenario represents a suboptimal outcome of the recruitment process according to both of the OECD principles highlighted above: first, the assembly members are not optimally **representative** of the population from which they are drawn; second, **inclusiveness** has been compromised by the exclusion of specific marginalised groups (in the example above this would be black people with lower educational attainment). It is perhaps worth noting that there is significant literature establishing that representativeness and inclusiveness (or, more specifically, diversity) may not be entirely mutually compatible, so some compromise is often needed. We will not pursue this line of thought here: for our purposes it is enough to state that our aim is simply to optimise both insofar as that is possible.

We will see that the data we have available supports Thesis T1 very strongly and supports Thesis T2 less strongly. In both cases the data is, to this point, limited.

Our process

If you already understand how Sortition Foundation does recruitment, then skip this section!

Our standard process begins by sending out a large number of letters to randomly selected addresses. These letters give information about the event in question and ask people to register their interest, by phone or by filling in an online form. When someone registers they are required to share some information about themselves (e.g. date of birth, address). We also ask for a range of demographic information, dependent on the process and the context. Importantly for both of the events described below we asked people to share their ethnicity and their level of educational attainment.

When registrations close, we have a “pool of respondents”. We then run a selection algorithm which randomly selects the people who will make up the panel. This selection is *random* while also *stratified* meaning that the algorithm ensures that the profile of participants for the data we have collected reflects the overall population profile. So, for instance, the ethnicities of the selected assembly members will be proportionally representative of the population from which they are drawn.
We have discussed some of the minutiae around this process in an earlier article.

It is important to understand how the composition of the pool of respondents is related to the composition of the final group of selected panel participants. As just mentioned, we can use stratification to ensure that the profile of participants is representative of the population, even if the profile of the respondents is one. Note, though, that we can only choose a limited number of characteristics to stratify against.

Suppose, for instance, that we choose not to stratify against educational attainment and suppose, too, that our pool of respondents is skewed towards people with high levels of attainment compared to the population. Broadly speaking this leads to two potential outcomes for our final selected panel.

First, it is possible that, although we do not stratify against education, we do stratify against a different characteristic (some measure of socio-economic status, say) that correlates strongly with educational status in both the pool of respondents and the general population. Then our stratification against this measure will act as a proxy stratification for education and we can be confident that our final panel of respondents will have an educational profile which reflects that of the general population.

The second possibility is that education does not strongly correlate with any of the characteristics that we use for stratification. In this case, the *randomness* of our algorithm will naturally transmit any skewing in the pool of respondents into the final group of selected panel participants.

In what follows we choose to focus our attention on the profile of the pool of respondents, rather than that of the selected panel participants because, as we have just described, the latter is a direct consequence of the former. We choose this focus because, although our ultimate concern is that the group of selected panel participants is representative, the larger size of the pool of respondents means that any conclusions we draw about our recruitment process are likely to be more robust

The raw data

We will use data from the pools of respondents from the following recruitment processes:

- **D1.** The University College London Citizens' Assembly on Democracy in the UK. Recruitment took place in August 2021. SF sent 20,000 letters to addresses across the UK and, in response, 344 people aged 18 and over registered their interest. This assembly consisted of 75 people who considered the question "How should the UK's democracy work?"
- **D2.** The Coventry Art for the People Citizens' Assembly on Arts, Culture and Creativity. SF sent 15,000 letters to addresses around Coventry and, in response, 253 people aged 16 and over registered their interest. This assembly consisted of 50 people who considered the question "How will arts, culture and creativity shape a better future for Coventry?"

The relevant data for the two pools of respondents is as follows:

1. *Educational attainment*: The possible values here are "No qualification", "Level 1", "Level 2", "Level 3", Level 4 or above" and "Apprenticeship/ Other". The descriptions of these values can be found here. For both respondent pools there was some grouping of the data – the grouping was done in different ways as will be clear.

2. *Ethnicity*: In this case the discussion below focuses on three main ethnic groupings: [a] “white” encompassing, amongst other things, “white British”, “white Irish” and “white other”; [b] “black” encompassing “Black”, “African”, “Caribbean”, and “Black British”; and [c] “non-white” encompassing anyone who is not in group [a]. Note, in particular, that grouping [c] contains grouping [b] as a subset.
3. *IMD*: For data set D1 we also have information about the [Index of Multiple Deprivation](). This is a whole number between 1 and 10, calculated by the government for every postcode in the country. Number 1 indicates that the postcode is in the most deprived 10% of addresses in the country, while 10 indicates that the postcode is in the least deprived 10% of addresses in the country. IMD is also calculated for areas larger than individual postcodes (including [LSOAs]()).

We will also use population data from the 2011 census. This data is for England and Wales only. The most significant data set for our purposes was [Age by Ethnic group by Highest level of qualification](). Also useful was a data set connecting [IMD to ethnicity]().

Data analysis

The analysis we undertake here is quantitative although somewhat informal: we will present various pie charts in order to understand the distribution of certain properties in various groups of people. We will not undertake formal hypothesis testing at this stage.

The three populations

It is important to emphasise that the three data sets we use (D1, D2 and the 2011 census) are for different populations (respectively these are: the whole of the UK, the Coventry area, England and Wales).

An important, but untested, assumption in what follows is that the interaction between education, ethnicity and IMD is similar across the three populations. Although we do not have data for this assumption we can at least display the education attainment profile for the three groupings:

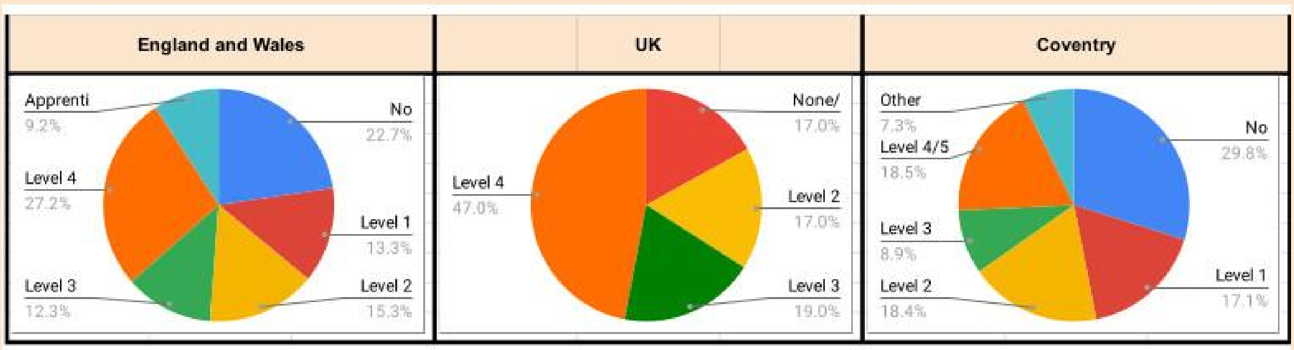

The education attainment profile for people in England and Wales, the UK and Coventry

Notes:

- Our population comparisons below will use the data for England and Wales [source]. This data is for those aged 16 and over. We use this population because it is the only one where we have an available breakdown of educational attainment according to ethnicity.
- Recall that we are seeking to test whether our pool of respondents, from data sets D1 and D2, tend to have higher educational qualifications than one would expect from the general population.
- The data for the UK is for those aged 19 to 64 and is grouped so that those with no qualification or with a Level 1 qualification are counted together [source]. This shows a generally higher level of attainment than for the England and Wales population. This potentially weakens confidence in the conclusions we draw below from our analysis of data set D1 (the UCL Democracy CA) – one might speculate that respondents in D1 have higher educational qualifications because they are drawn from the UK rather than from England and Wales. On the other hand the different age restriction undoubtedly accounts for some variation; in addition the inclusion of "apprenticeships" under the various levels will swell Levels 2 and 3 and will bring the two populations into line for these categories.
- The data for Coventry is for those aged 16 to 74 [source]. We have used the "urban" population of the Coventry and Warwickshire LEP as an approximation to the population of Coventry. Since this profile shows generally lower levels of educational qualification than for the England and Wales population, any divergence between the two would seem to strengthen confidence in the conclusions we draw below from our analysis of data set D2 (the Coventry Art for the People CA).

Education and ethnicity

Our principle line of inquiry concerns the connection between education and ethnicity for our respondent pool. First, the next set of charts give the population profile:

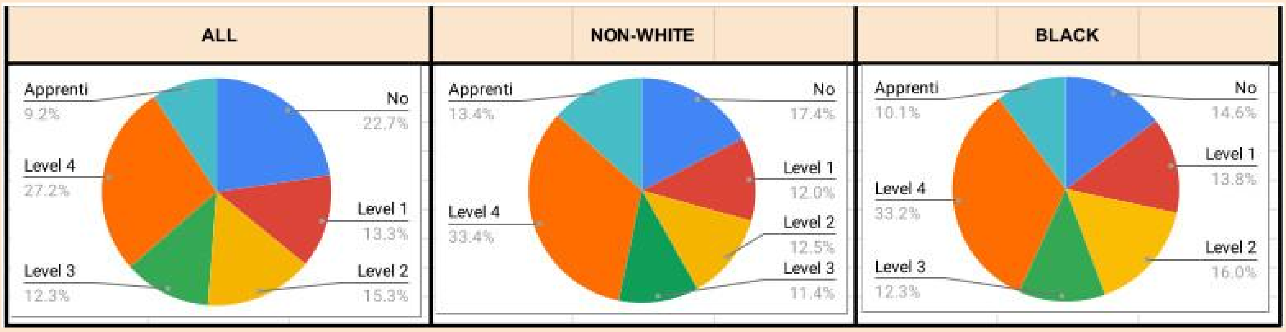

The educational attainment profile for people aged 16+ in England and Wales in 2011

Notes:
- The profile for Levels 1 and 3 is broadly similar for the three populations;
- The profile for the total percentage of people having Level 2 or an apprenticeship is broadly similar for the three populations;

- There is significant divergence at the two extremes of educational attainment: the non-white and black populations have a higher percentage of people who are qualified to Level 4 or above as compared to the overall population; and they have a lower percentage of people with no qualifications.

Now let us see the profiles for data set D1:

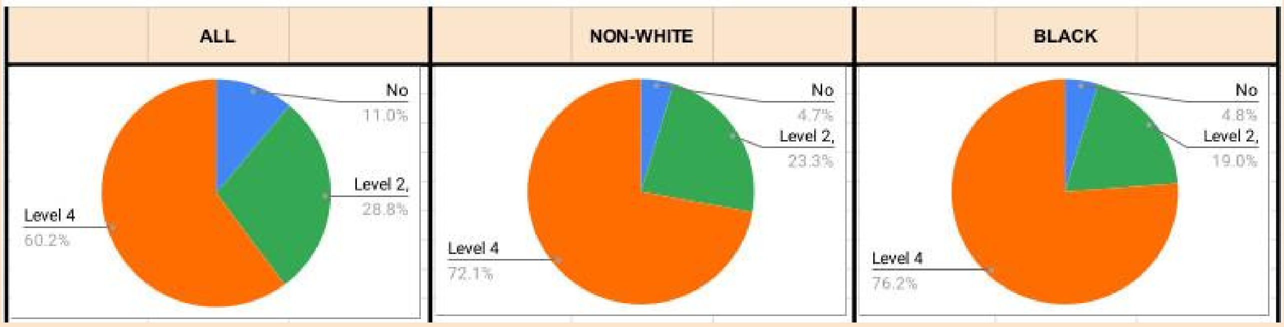

Education and ethnicity for the UCL Democracy Citizens' Assembly

Notes:

- We have some groupings here: the blue slice corresponds to "No qualification, Level 1"; the green corresponds to "Level 2, Level 3, Apprentice, Other"; the orange corresponds to "Level 4 and above".
- It is quite clear that the percentage of respondents who are highly educated is much greater, for all three categories, than for the corresponding population chart. (Note that the orange slices both here and above correspond to exactly the same education level.) A striking feature here is that, for all three groups, the percentage of respondents who are highly educated is around 2.2 times the percentage of the corresponding population who are highly educated. Thus ethnicity does not appear to affect response rates among the three populations of highly educated people.
- The situation for the least educated group here is noticeably different. First, for all three groups, the percentage of respondents with no qualification or a level 1 qualification is much lower than the percentage of the corresponding population who have this level of education. However, this decrease varies significantly across the three groupings: among the total population, the percentage of respondents with this lowest level of education is 0.31 times the percentage of the corresponding population at the same level; in contrast, the percentage of non-white (resp. black) respondents with the lowest level of education is 0.16 times (resp. 0.17 times) the percentage of the corresponding population at the same level. Thus, for this sample, a lower level of education is almost doubly effective in reducing responses from the non-white population as compared to the overall population.

Similarly data set D2:

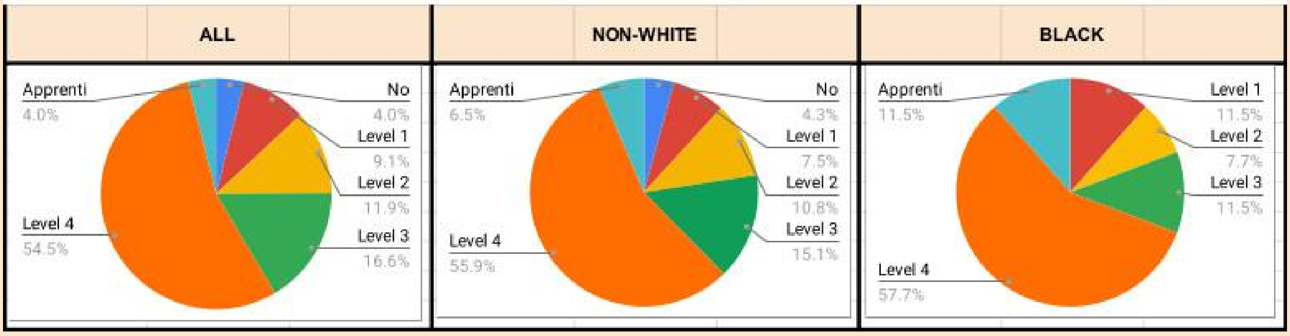

Education and ethnicity for the Coventry Art for the People CA

Notes:

- There are no groupings here, so we can compare these charts directly to the population charts (albeit keeping in mind that the populations are from different areas).
- Again we have very significant over-representation of highly educated people for all three groupings. This time the percentage of respondents who are highly educated is around 2.0 times the percentage of the population who are highly educated, but this multiplier drops to 1.7 when we consider non-white and black respondents.
- On the other hand for all three groups, the percentage of respondents with no qualification or a level 1 qualification is around 0.4 times the percentage of the corresponding population who have this level of education – so we see, again, that lower education levels result in a lower response rate but this time ethnicity does not appear to exacerbate this tendency. Note, though, that there are no respondents whatsoever who are both black and have no qualifications.

Education, ethnicity and IMD

The connection between ethnicity, education and IMD is interesting because we have in the past used education and IMD interchangeably as a measure of “socioeconomic status”. We wish to measure socioeconomic status because it is a way of understanding marginalisation in a population, and the inclusion of marginalised voices is crucial for any deliberative process.

The idea that one can get an indication of educational attainment via knowledge of IMD score (and vice versa) is not entirely unreasonable: educational attainment of children and adults in an area forms part of the calculation of a postcode’s IMD score. However it is only a small part – 13.5% to be precise – and is combined with other factors including health, employment, crime and so on.

We need to understand the connection between a person’s educational attainment and the IMD score of their address. Our first chart demonstrates that, in England, **people who live in more deprived areas tend to have lower levels of educational attainment**.

The chart is drawn from data used in the calculation of IMD. The IMD score of an address combines a number of components in order to give an overall deprivation score. One component of this calculation is the “education, skills and training deprivation domain” which seeks to measure the educational deprivation of a population in a given area. So, specifically, to understand this chart:

1. England is split into 32844 areas known as LSOAs.
2. Each of these is assigned an overall IMD score and we have restricted our attention to the 9853 segments which are more deprived, i.e. in deciles 1-3.
3. The chart represents the educational decile for these areas. So, for instance 32.8% of these more deprived LSOAs are in the bottom decile (decile 1) for the specific measure of education; 27.3% are in decile 2 and so on.

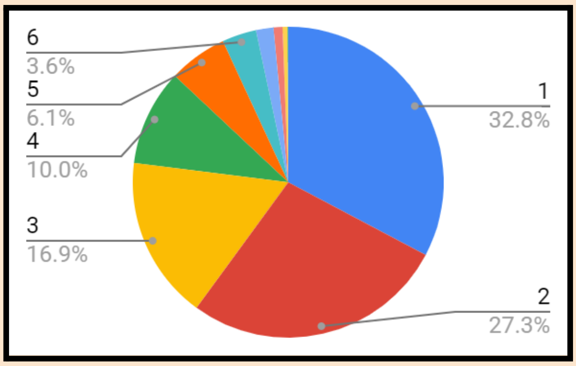

The “educational, skills and training deprivation decile” of people in England living in an IMD1 to IMD3 area

Notes: unsurprisingly, the educational deprivation score of an area strongly tracks the overall deprivation score.

The charts we present next pertain to data set D1. Again we look at the educational attainment of people from IMD1 to IMD3; this time, though, these people are from the respondents pool for this CA.

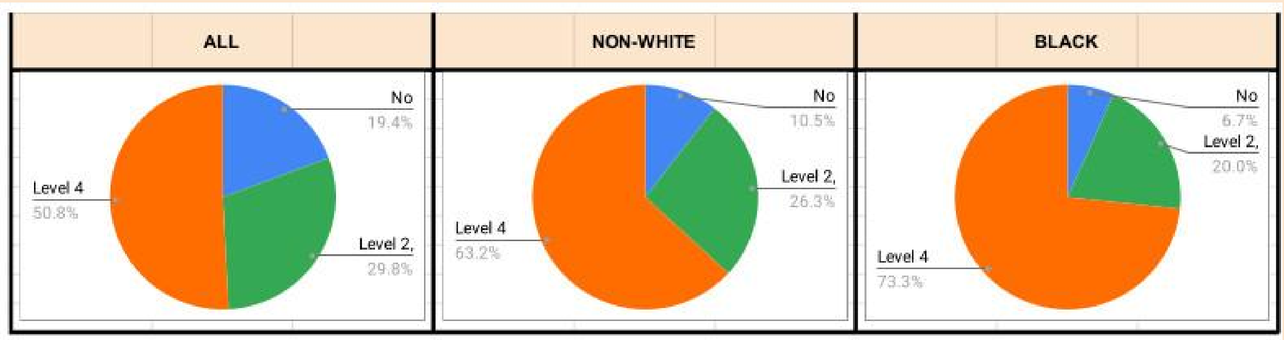

Educational attainment for respondents living in IMD1 to IMD3 areas for the UCL Democracy CA

Notes:

- We need to understand how the levels given here correspond to the educational decile given in the population chart. For instance, we saw earlier that 27.2% of the population in England and Wales are educated to Level 4 or above, so this places such people in education deciles 8,9 and 10 of the overall population.
- One sees immediately that the educational attainment for respondents living in more deprived areas is extremely divergent from the educational attainment of the overall population of more deprived areas – for all three groupings, more than 50% of respondents from more deprived areas are in education deciles 8,9 or 10 compared to around 5% of the overall population of more deprived areas. In short, **respondents who live in more deprived areas have high levels of educational attainment relative to the overall population of more deprived areas**.
- Indeed the overall education profile for respondents from IMD1-3 is very similar to the education profile we saw for the overall pool (see earlier charts). The restriction to IMD1-3 has resulted in a slight diminishing of the level of education compared with the overall respondents pool but the diminishing really is "slight". This is not consistent with what population data would suggest.

Recruitment from more deprived areas

We have seen, so far, that

- The number of lower educated people in our respondents pool is disproportionately small;
- The number of lower educated people from more deprived areas in our respondents pool is disproportionately small.

To complete the picture we would like to know if the number of people from more deprived areas in our respondents pool is disproportionately small. Were this to be the case, then we might speculate that this is contributing to the disproportionately small number of lower educated people.

In fact this is not the case. The next charts are for data set D1 again and they give the IMD score of respondents:

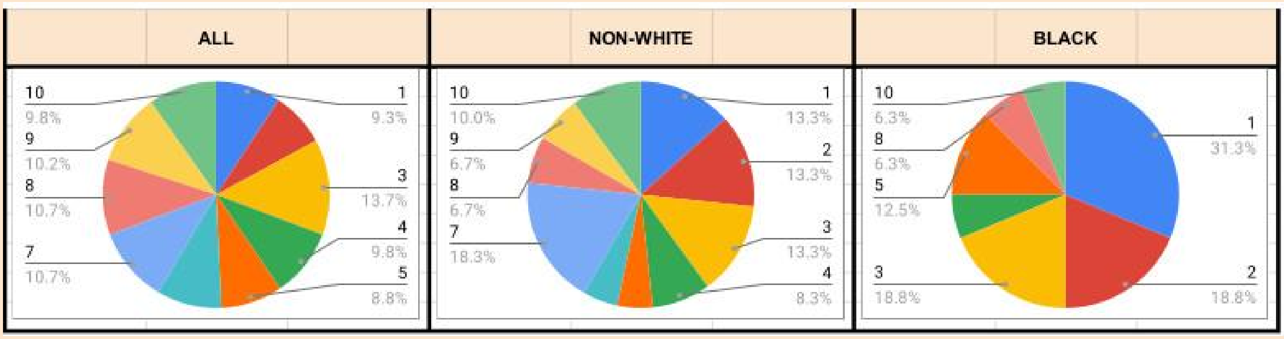

The IMD score for respondents to the UCL Democracy CA

In order to draw accurate conclusions from this data we need a note about our process. We have shown in previous investigations of our data that we receive proportionally fewer responses to invitations from addresses in more deprived areas. To compensate for this, when randomly selecting addresses to receive invitations, our process proceeds in two stages: first, we select 80% of required addresses at random from the whole area; we then select 20% of required addresses at random from those that have an IMD score of at most 3. Thus we send a disproportionate number of invitations to more deprived areas (note that some of the initial 80% of all addresses will also be in more deprived areas).

With that in mind, some notes:

- The overall IMD profile of the pool of respondents is remarkably close to the population profile: we would expect 10% from each decile and we get something pretty close to that – our 20% overselection from IMD1-3 seems to provide the right correction needed here.
- The IMD profile of non-white and black respondents is weighted towards more deprived areas as, anecdotally, one might expect (see below). This is particularly clear for black respondents, 67% of whom are from more deprived areas as opposed to 33% of the total respondents. This suggests that our recruitment process is successfully recruiting respondents from minority ethnic groups living in more deprived areas.
- Thus, while our recruitment process results in proportionally fewer respondents with lower educational attainment than would expect from the population, there are no such issues with regard to areas of deprivation: the process does result in a proportional number of respondents from more deprived areas and this remains true when we restrict to the non-white and black populations.

To confirm our anecdotal suspicions mentioned in the second bullet above, it would be useful to have a breakdown of the IMD profile of the non-white and black populations of the UK. We have been unable to source a complete breakdown however partial information is available. This tells us that 9.9% of the population of England live in IMD1, 9% of the white population of England live in IMD1, 15% of the non-white population of England live in IMD1 and 15.2% of the black population of England live in IMD1. This mirrors the trend we see in the charts above (indeed one might conclude that we are being excessively successful in obtaining black respondents from IMD1… however the small sample size mandates caution in drawing such a conclusion).

Conclusions and implications

Let us list a number of conclusions, our confidence in these conclusions, and some potential responses:

The pool of respondents tends to over-represent highly educated people and under-represent people with lower educational attainment. There is **strong** evidence for this conclusion. It suggests that, if we continue to recruit respondents using current methods, then we need to stratify across indicators that correlate in the respondents pool with educational attainment. If we stratify on such indicators, then it is clear that the resulting assemblies will better represent the wider population..

The skewing mentioned above is exacerbated in some ethnic (and gender) groupings. There is **some** evidence for this conclusion but our results here are not definitive. It seems possible that over representation of well-educated people is not heavily affected by ethnicity; on the other hand the under representation of lower educated people may be exacerbated in the non-white population.

Stratifying, as described in the previous point, will hopefully ameliorate some of this skewing but, since the algorithm does not explicitly calculate intersections, there is a potential weakness in our process here.

A technical solution would be to amalgamate categories where we suspect skewing in the respondents pool: rather than having two categories (*ethnicity* and *education*) we have a single category for both, in which we have values like (*white no qualifications, white Level 1,* through to *black Level 4 and above).* This is problematic because it results in an explosion of the number of values, and makes the stratification process difficult.

A second solution, of a rather similar kind, would be to use some kind of scoring system around marginalisation that picks up key intersections between (say) ethnicity, education and the deprivation score of one's address. The design of such a scoring system would require significant data and experiment.

A more promising line of work would be to seek to adjust the part of our process devoted to recruiting respondents. If we focus, first, on achieving the *inclusion* objective of our recruitment process, then we must recognise that certain marginalised voices **are** potentially under-represented in the assemblies we select: namely those of ethnic minorities with lower educational attainment.

This suggests that we should try and recruit respondents from these under-represented groups in a different way. If we can ensure that such marginalised voices are well-represented in the respondents' pool, then our selection algorithm will naturally ensure that these voices will be represented in the assemblies we select.

The benefit of a hybrid recruitment strategy to achieving inclusion of marginalised voices has already been noted in the literature. Our strategy is to identify the groups we are potentially missing through data analysis of the type given above and then to intentionally seek out representatives who are willing to join the respondents pool. Practically, we have started to investigate the possibility of combining our standard postal recruitment with some targeted door knocking.

The IMD profile of the pool of respondents is generally in line with the population, even when considering ethnicity. There is **some** evidence for this conclusion… But we would need to analyse more processes to be sure. The 20% over-invitation in more deprived areas seems to result in a good overall profile. It also seems to ensure a reasonable proportion of respondents from minority ethnic backgrounds who live in more deprived areas

– this is a portion of the population who are marginalised and so it is important that their voice is heard.

Using IMD as a category for stratification does not address skewing in educational attainment. There is **strong** evidence for this conclusion: the IMD profile and the education profile of the pool of respondents are not aligned. At a population level there is significant correlation between overall deprivation and educational deprivation, however this correlation does not carry over to the pool of respondents. It would appear that stratifying through IMD alone will, in general, result in a group who are more highly educated than the general population.

More evidence is needed. This conclusion is indisputable. Our sample is too small to be definitive and further analysis is needed. Indeed the 11th principle of good practice given by the OECD is “evaluation”: every deliberative process should be analysed and evaluated by participants and organisers to ensure that best practice is maintained.

Specific areas that we would like to investigate further include:

- *The efficacy of using other socio-economic indicators for stratification purposes*. We have demonstrated above that the use of IMD does not interact well with educational attainment, however maybe other indicators (e.g. tenure/ benefits/ income) would work better (perhaps in combination, as per the scoring system suggestion above).
- *Individual wealth versus IMD score*: Another obvious measure of socio-economic status is household income. It would be interesting to ask all of the questions we have asked above for education, but for household income. In particular we would like to know whether the respondents we recruit from areas with a lower IMD score are themselves subject to financial insecurity. We would like to know if we should stratify across household income, or if a wide spread of responses across IMD scores is sufficient to ensure representation along this measure. Such an investigation would require new data as we do not customarily collect information about household income when we recruit.
- *The granularity of our educational analysis*. We have not distinguished between educational attainment above Level 4. Given that we have over-representation in this category it would seem important to know whether, for instance, a disproportionate number of respondents have a degree (Level 6 or higher). If this were the case, then the imperative to correct for this through stratification would be even stronger than it already is.