

DELTA POLL

Demographic Data Study: Analysis and Report for the Centre for Data Ethics and Innovation

Contents

Executive Summary.....	2
Aim	3
Method	3
Warm-Up Results	8
Conjoint 1: Third-party organisations.....	9
Differential opinions by respondent demographics and attitudes	14
<i>Comfort with sharing data</i>	14
<i>Generation</i>	14
<i>Disability status</i>	16
<i>Other demographic characteristics</i>	16
Conjoint 2: Predicting demographic data	17
Differential opinions by respondent demographics and attitudes	21
<i>Comfort with sharing data</i>	21
<i>Generation</i>	21
<i>Other demographic characteristics</i>	21
Open text responses	24
Word clouds.....	24
Structural topic models	26

DELTA POLL

Executive Summary

- △ When it comes to being asked to provide demographic information when applying for a job, generally speaking, UK adults are familiar with this practice and a large majority say that they are not concerned with this form of data sharing ([p. 8](#)).
- △ In the first conjoint experiment, respondents were presented with different scenarios where their demographic data is being collected by a third-party organisation in order to test for bias monitoring purposes. Out of all possible combinations, respondents feel most positively about sharing information on their age group to help an NHS hospital test whether their diagnostic tool successfully diagnoses patients from different demographic groups, where this data is then provided to a consumer rights organisations, is de-identified, and are only themselves able to make decisions about how it is used ([pp. 9-16](#)).
- △ Respondents feel least positive, however, about sharing information on their sexuality to a large technology company checking that their targeted advertising does not unfairly target certain groups of people, where their data is then provided to a large technology company third-party organisation, is personally identifiable, and where decisions about its usage are made by a dedicated group of ordinary citizens.
- △ Overall, the majority of respondents say that they would feel comfortable sharing their demographic data with a third-party organisation to help organisations check the fairness of their systems ([p. 12](#)).
- △ When asked to elaborate on any concerns they might have, respondents most commonly voiced concerns around their data falling into the wrong hands, the potential for identity theft or scams, or that the data might be used without their permission ([p. 25](#)).
- △ In the second conjoint experiment, respondents were instead presented with different scenarios where organisations make use of the demographic data they already have available in order to test the fairness of their systems. Out of the scenarios presented, respondents feel most positively towards an NHS hospital testing whether their diagnosis tool successfully diagnoses patients from different demographic backgrounds through using someone's postcode in order to predict their age group, where the organisation informs them privately about this ([pp. 17-21](#)).
- △ Conversely, respondents felt most negatively about a credit reference agency testing if they are providing credit scores fairly through using a photograph of someone's face to predict their disability status, where the organisation does not inform them privately about this and also does not publish information on their website.
- △ While a plurality of respondents say that they are comfortable with an organisation using information they already hold about them to predict their personal characteristics to help them test the fairness of their systems, this number is far lower than in the first conjoint, and those saying they would be neither comfortable or uncomfortable is significantly higher ([p. 20](#)).
- △ When asked to elaborate on any particular concerns they may have about this, commonly raised concerns include that they feel this would constitute an invasion of privacy, that this form of prediction would be inaccurate, unfair, or based on very limited data, and that the information may be leaked and end up in the wrong hands ([p. 26](#)).
- △ When asked to reflect on their overall experience of taking the survey, respondents most commonly said that they did not have encounter any issues, and that they found the subject matter to be interesting, compelling, and introduced them to new ideas. Other respondents, however, felt less positively, and said instead that they found it to be hard to complete, long and repetitive at times, and that it had peaked their level of concern in the subject matter ([p. 26](#)).

DELTA POLL

Aim

The aim of the survey is to analyse public attitudes towards the handling and usage of different forms of demographic data. Demographic information refers to personal characteristics such as an individual's gender, age, ethnicity, disability, occupation, and place of residence.

Organisations regularly use data to ensure that they provide fair outcomes for everyone. To do this, organisations may collect demographic information to test that the decisions they make about people from different backgrounds are fair. Attitudes towards two forms of data collection are investigated in two forced-choice conjoint survey experiments. The first experiment investigates attitudes towards third-party organisations who collect data on the behalf of organisations. The second experiment instead investigates attitudes towards organisations leveraging data they already have available in order to predict demographic information about their users.

Method

To address these questions, we fielded a survey to 4,046 adult respondents in the United Kingdom between 12th December 2022 and 3rd January 2023. The data has been weighted to be representative of the adult population as a whole.

The survey participants were asked a range of questions about their experience and levels of concern with providing information about their gender, ethnicity, or other demographic information when applying for a job. The main sections of the survey included two forced-choice conjoint survey experiments. The purpose of both experiments was to identify, if demographic data is being collected for bias monitoring purposes, in which forms do UK respondents feel most and least comfortable with.

For the first conjoint experiment, respondents were presented with two different scenarios where their demographic data is being collected by a third-party organisation in order to test for bias monitoring purposes. The scenarios were manipulated on five dimensions. First, the *type of data* they asked to share (ethnicity; sex; age group; disability status; socio-economic status; sexuality). Second, the *use case* of the information they are asked to share (a car insurance company; a local authority; a large technology company; an employer; an NHS hospital; a bank; a government department – welfare benefits; a government department – state pensions; a credit reference agency). Third, the *third-party organisation* that the information is provided to (a large technology company; a technology start up; a regulator; a non-governmental organisation; a consumer rights organisation; a national statistics agency; a university; a private and secure database). Fourth, the *privacy and sharing* of the information (identifiable and shared with the use case organisation; identifiable but not shared with the use case organisation; de-identified and shared with the use case organisation). Fifth, who is responsible for the *control* of the data once collected (the use case organisation; a dedicated group of independent experts; a dedicated group of ordinary citizens; only you). Table 1 summaries the attributes and their levels.

DELTA POLL

Table 1. Attributes of conjoint 1 design

Attribute	Level number	Level words
Type of data	1	You are asked to share information on your ethnicity (for example, Black Caribbean; White British)
	2	You are asked to share information on your sex (for example, female; male)
	3	You are asked to share information on your age group (for example, 18-24 or 25-29).
	4	You are asked to share information on your disability status (for example, physical or mental health conditions)
	5	You are asked to share information on your socio-economic status (for example, type of occupation)
	6	You are asked to share information on your sexuality (for example, heterosexual or homosexual)
Use case	1	You are asked to provide this information to help a car insurance company to test if they are setting insurance prices fairly.
	2	You are asked to provide this information to help a local authority to test if they are allocating welfare benefits fairly.
	3	You are asked to provide this information to help a large technology company to check that their targeted advertising does not unfairly target certain groups of people.
	4	You are asked to provide this information to help an employer to test if they are reviewing job applicants' CVs fairly.
	5	You are asked to provide this information to help an NHS hospital to test whether their diagnostic tool successfully diagnoses patients from different demographic groups.
	6	You are asked to provide this information to help a bank to test if they are providing loans fairly.
	7	You are asked to provide this information to help a government department to test if they are allocating welfare benefits fairly.
	8	You are asked to provide this information to help a government department to test if they are administering state pensions fairly.
	9	You are asked to provide this information to help a credit reference agency to test if they are providing credit scores fairly.
Third-party organisation	1	You are asked to provide this information to a large technology company like Amazon Web Services, Microsoft or Oracle.
	2	You are asked to provide this information to a technology start up with expertise in data science and artificial intelligence.
	3	You are asked to provide this information to a regulator like the Equality and Human Rights Commission or the Information Commissioner's Office.
	4	You are asked to provide this information to a non-governmental organisation like Age UK or Scope.
	5	You are asked to provide this information to a consumer rights organisation like Citizens Advice or Which?.
	6	You are asked to provide this information to a national statistics agency like the Office for National Statistics.
	7	You are asked to provide this information to a university for research.

DELTA POLL

	8	You are asked to upload and store this information on a secure, private database on your phone or computer, which can only be accessed by other organisations with your permission.
Privacy / Sharing	1	This information is identifiable and will be shared with the organisation testing the fairness of their system.
	2	This information is identifiable but won't be shared with the organisation testing the fairness of their system.
	3	This information is de-identified and will be shared with the organisation testing the fairness of their system.
Control	1	Once you have provided consent for this information to be collected, decisions about how it is used will be made by [Organisation from use case]
	2	Once you have provided consent for this information to be collected, decisions about how it is used will be made by a dedicated group of independent experts.
	3	Once you have provided consent for this information to be collected, decisions about how it is used will be made by a dedicated group of ordinary citizens.
	4	Once you have provided consent for this information to be collected, decisions about how it is used will be made by only you.

Following an introduction screen describing the task, respondents were asked "If you had to choose, which of the following two options do you think you would prefer for an organisation sharing information with a third-party organisation to help organisations check the fairness of their systems?". For each scenario combination, all attribute levels were randomly assigned, with no restriction on level combination. We asked each respondent to complete the task five times, which provides us with 20,230 forced-choice combinations and 40,460 scenario-level ratings.

The outcome is a respondent's decision in the forced choice of their *scenario preference* (not selected; selected; don't know). We drop all *don't know* responses in the analysis below. The quantity of interest in the analysis below is the Average Marginal Component Effect (AMCE), which is the treatment effect of a particular scenario level (compared to an attribute baseline category) averaged over the joint distribution of all other attributes. Because each attribute was randomly assigned independent of the value of any other attributes, the AMCE can be estimated using a simple linear regression of the outcome variable. The unit of analysis is the scenario level, such that each respondent creates multiple scenario-level observations. We cluster standard errors at the respondent level.

For the second conjoint experiment, respondents were presented with two different scenarios where proxy data that an organisation already holds is being used to predict their demographic information. The scenarios were manipulated on five dimensions. First, the *use case* for assessing the fairness of a system (a car insurance company; a local authority; a technology company; an employer; an NHS hospital; a bank; a government department - welfare benefits; a government department - state benefits; a credit reference agency). Second, the *proxy data* that will be used (postcode; full name; a photograph of your face; social media use). Third, the *demographic characteristic* that the organisation is trying to predict (ethnicity; sex; age group; disability status; socio-economic status; sexuality). Fourth, the *privacy* of the demographic characteristics (identify an individual personally; identify characteristics but not identifying personally; identify groups of people). Fifth, the *transparency* of the organisations decisions (privately inform; privately inform and

DELTA POLL

publish information on website; not inform and not publish information on website). Table 2 summaries the attributes and their levels.

Table 2. Attributes of conjoint 2 design

Attribute	Level number	Level words
Use case	1	A car insurance company wants to check whether they are setting insurance prices fairly.
	2	A local authority wants to check whether they are allocating welfare benefits fairly.
	3	A technology company wants to check whether their advertising does not unfairly target certain groups of people.
	4	An employer wants to check whether they are reviewing job applicants' CVs fairly.
	5	An NHS hospital wants to test whether their diagnostic tool successfully diagnoses patients from different demographic groups.
	6	A bank wants to test if they are providing loans fairly.
	7	A government department wants to test if they are allocating welfare benefits fairly.
	8	A government department wants to test if they are administering state pensions fairly.
	9	A credit reference agency wants to test if they are providing credit scores fairly.
Proxy data	1	They will use your postcode...
	2	They will use your full name...
	3	They will use a photograph of your face...
	4	They will use information based on your social media use (for example, accounts you follow on Instagram or Twitter)...
Demographic characteristic	1	... to predict your ethnicity (for example, Black Caribbean; White British)
	2	... to predict your sex (for example, female; male).
	3	...to predict your age group (for example, 18-24 or 25-29).
	4	...to predict your disability status (for example, physical or mental health conditions).
	5	... to predict your socio-economic status (for example, type of occupation)
	6	...to predict your sexuality (for example, heterosexual or homosexual).
Privacy	1	The organisation will make predictions about your own demographic characteristics that can be used to identify you personally.
	2	The organisation will make predictions about your own demographic characteristics, but apply technology that prevents them from identifying you personally.
	3	The organisation will make predictions about the demographic characteristics of groups of people, but not about any individual specifically.

DELTA POLL

Transparency	1	The organisation will inform you privately that they will predict demographic information about you personally.
	2	The organisation will inform you privately that they will predict demographic information about you personally <i>and</i> publish information on their website informing members of the public that they are predicting demographic information to help them monitor for bias.
	3	The organisation will not inform you that they will predict personal information about you personally, and will not publish any information on their website informing members of the public that they are predicting demographic information to help them monitor for bias.

Following an introduction screen describing the task, respondents were asked “If you had to choose, which of the following two options do you think you would prefer for an organisation using information they already hold about you to predict your personal characteristics?”. For each scenario combination, all attribute levels were randomly assigned, with no restriction on level combination. We asked each respondent to complete the task three times, which provides us with 12,138 forced-choice combinations and 24,276 scenario-level ratings.


The outcome is a respondent’s decision in the forced choice of their *scenario preference* (not selected; selected; don’t know). We drop all don’t know responses in the analysis below. As with conjoint 1, our quantity of interest is the Average Marginal Component Effect (AMCE). The unit of analysis is the scenario level, such that each respondent creates multiple scenario level observations. We cluster standard errors at the respondent level.

DELTA POLL

Warm-Up Results

Figure 1 shows the results for CDEI1. Generally speaking, UK adults are familiar with providing demographic information, such as their gender or ethnicity, when applying for jobs. Just short of three in four (72%) say that they have done this, compared to one in five (20%) who have not. A further 7% of respondents say that they don't know.


Figure 1. *Have you ever had to fill in a form where you were asked to provide information about your gender, ethnicity, or other demographic information when applying for a job? (CDEI1)*



Yes No Don't know

Figure 2 shows the results for CDEI2. When asked instead about their level of concern when providing demographic information when applying for a job, more than three quarters (77%) of respondents say that they *would not be concerned*, compared to fewer than one in five (17%) who *would be concerned*. Around one in twenty (6%) respondents instead say that they don't know.

Figure 2. *Generally speaking, do you think you would or would not be concerned about providing information about your gender, ethnicity, or other demographic information when applying for a job? (CDEI2)*



Would not be concerned Would be concerned Don't know

DELTA POLL

Conjoint 1: Third-party organisations

Having assessed initial experience with and comfort towards providing demographic data in job applications, we turn to the results from the first conjoint. The point estimates and confidence intervals are presented in Figure 3 and the full results in Table 3. The figure presents the results from a model where the outcome is whether a scenario was chosen or not. The vertical line at zero represents no preference either way and any point estimates to the right of the line indicate that any given level is preferred and, conversely, any point estimates to the left of the line suggest that it is not. Further, any points with confidence intervals that include the value of 0 suggest there is no statistically significant difference.


The baseline category for each attribute is indicated in brackets, and the corresponding point estimates for each level should be interpreted as preferences compared to the baseline, holding constant all other attributes. For example, the point related to *age group* for “Type of data” should be compared to *ethnicity*. The point estimate is positive and statistically significant at the 95% confidence level. As such, UK respondents prefer to share information on their age group as opposed to their ethnicity.

- △ Following this logic for the remaining point estimates, we find that, compared to sharing information about *ethnicity*, respondents significantly prefer to share information about their *age group*, however, are more negative towards sharing their *disability status*, *sexuality*, and *socio-economic status*.
- △ There is no statistically significant difference between sharing information on their sex or their *ethnicity*.
- △ Turning to use case, compared to being asked to provide the information to help *a car insurance company to test if they are setting insurance prices fairly*, respondents significantly prefer to provide information to *a government department to test if they are allocating welfare benefits fairly* or *administering state pensions fairly*, *a local authority to test if they are allocating welfare benefits fairly*, *an employer to test if they are reviewing job applicants' CVs fairly*, and *an NHS hospital to test whether their diagnostic tool successfully diagnoses patients from different demographic groups*. Out of each of the use cases, the largest positive effect is for *an NHS hospital*, meaning that out of all of the possible options respondents feel most positively towards their data being used for this purpose.
- △ There is no statistically significant difference between *a car insurance company* and *a bank to test if they are allocating loans fairly*, *a credit reference agency to test if they are providing credit scores fairly*, and *a large technology company to check that their targeted advertising does not unfairly target certain groups of people*.
- △ Turning to the third party organisation respondents are asked to provide the information to, compared to *a large technology company*, respondents significantly prefer each of the other third party organisations. The largest positive effect is for *a consumer rights organisation like Citizens Advice or Which?*.
- △ The baseline for privacy and sharing is that the information is *identifiable and will be shared with the organisation testing the fairness of the system*, compared to this, respondents significantly prefer the information to be *identifiable and not shared with the organisation testing the fairness of their systems* and *de-identified and will be shared with the organisation testing the fairness of their system*.
- △ Finally, for control of the data, compared to control being with the organisation given for the use case, there is no statistically significant difference between *a dedicated group of independent experts*, however respondents significantly prefer for the decisions to be made

DELTA POLL

by themselves only and, conversely, are least positive towards control and decisions being made by *a dedicated group of ordinary citizens*.

Figure 3. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 1 scenario attribute level compared to the baseline level of the attribute, standard errors clustered by respondent. Bars show 95% confidence intervals. (CDEI_S1_1)*



Overall, across the various attributes, UK respondents feel most positively about sharing information on their *age group* to help *an NHS hospital* to test whether their diagnostic tool successfully diagnoses patients from different demographic groups, where this data is then provided to a *consumer rights organisation*, is *de-identified* and will be shared with the organisation testing the fairness of their system, and decisions about how it will be used are made by only the person themselves.

Conversely, UK respondents feel least comfortable with being asked to share information on their *sexuality* to a *large technology company like Amazon Web Services, Microsoft or Oracle* where the information is *identifiable and shared with the organisation testing the fairness of the system*, where once consent has been provided the decisions about how it is used will be made by *a dedicated group of ordinary citizens*. Although the point estimate for *a large technology company checking that their targeted advertising does not unfairly target certain groups of people* is the most negative, it is not significantly different from the second least popular option.

DELTA POLL

Don't know responses, which were selected by 18% of respondents, are excluded from the analysis described above. We find usage of don't know responses varies by respondent characteristics where older respondents make greater usage of don't know responses than younger ones: 12% of 18-44, 21% of 45-64, and 24% of those aged 65 and over.

Table 3. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 1 scenario attribute level compared to the baseline level of the attribute (standard errors, clustered at the respondent level, in parentheses)*

Attribute	Level	Estimate
Type of data (Baseline: Ethnicity)	Sex	0.031 (0.158)
	Age group	0.032* (0.015)
	Disability status	-0.051** (0.017)
	Socio-economic status	-0.062*** (0.016)
	Sexuality	-0.097*** (0.015)
Use case (Baseline: A car insurance company)	A local authority	0.056* (0.018)
	A large technology company	-0.021 (0.018)
	An employer	0.058** (0.018)
	An NHS hospital	0.132*** (0.019)
	A bank	-0.001 (0.018)
	A government department (welfare benefits)	0.041* (0.019)
	A government department (state pensions)	0.060*** (0.018)
	A credit reference agency	-0.018 (0.019)
Third party organisation (Baseline: A large technology company)	A technology start up	0.065*** (0.017)
	A regulator	0.109** (0.018)
	A non-governmental organisation	0.107*** (0.018)
	A consumer rights	0.103*** (0.018)
	A national statistics agency	0.131*** (0.017)
	A university	0.124*** (0.017)
	A secure and private database	0.054** (0.018)

DELTA POLL

Privacy/Sharing (Baseline: Identifiable and shared)	Identifiable and not shared	0.041 (0.012)
	De-identifiable and shared	0.077 (0.012)
Control (Baseline: The use case)	Dedicated independent experts	-0.003 (0.013)
	Dedicated ordinary citizens	-0.033 (0.014)
	Only you	0.082 (0.013)

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Number of observations: 33,448

Number of respondents: 3,769

Figure 4 shows the results for CDEI_S1_2. Around two thirds (65%) of respondents say that they would be comfortable sharing information about their gender, ethnicity, or other demographic information with a third-party organisation. A further approaching a quarter (23%) say that they would be *neither comfortable nor uncomfortable*, while only one in ten (10%) say that they would be uncomfortable. As with the pilot, we randomised the question ordering such that half of respondents would be asked their comfort levels before the conjoint and half after. We find that question order did not significantly influence responses.

Figure 4. *Taking everything into account, how comfortable or uncomfortable do you think you would feel sharing information about your gender, ethnicity, or other demographic information with a third-party organisation to help organisations check the fairness of their services? (CDEI_S1_2)*





Figure 5. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 1 scenario attribute level compared to the baseline level of the attribute for those who are comfortable and uncomfortable about sharing their data (CDEI_S1_1)*



DELTA POLL

Differential opinions by respondent demographics and attitudes

Comfort with sharing data

Do responses to the conjoint differ by those who say that they are comfortable or uncomfortable in response to CDEI_S1_2? In figure 5, we present the results separately for these respondents with those who say that they are either *very comfortable* or *fairly comfortable* in the left hand panel and those who say that they are either *fairly uncomfortable* or *very uncomfortable* in the right hand panel.

For the vast majority of attributes and levels, there is no statistically significant difference between the two sets of results. In a number of places, the point estimates are in the same direction but the results are not statistically significantly different which is in large part due to the wider confidence intervals due to the smaller number of respondents who said that they would be uncomfortable.

However, for control of the data there are some intuitive differences: while there are no significant differences between the baseline – the *use case* – and a *dedicated group of independent experts* or a *dedicated group of ordinary citizens* for those who would be comfortable sharing their data, this is not case for those who would be uncomfortable. Here, we see that for those who would be uncomfortable sharing their data, they significantly prefer the *use case* to have control over their data compared to both *independent experts* and *ordinary citizens*.


Generation

We also assess whether responses differ by generation. In figure 6, we present the results for Millennials in the left hand panel, Generation X in the middle panel, and Baby Boomers in the right hand panel. While for the most part the results are broadly similar, there are a number of differences. We summarise these difference below.

- △ Type of data: Compared to *ethnicity*, Millennials feel significantly more negatively towards data on *disability status* and *socio-economic status* being collected, both effects that are not found for Generation X or Baby Boomers. Likewise, Baby Boomers are significantly more positive towards information on their *sex* being collected, where for Millennials and Generation harbour no preference.
- △ Use case: Perhaps intuitively, compared to a *car insurance company*, Baby Boomers feel more positively towards a *government department to test if they are administering state pensions fairly*. This effect does not hold for the other generations. While there are significant positive coefficients for an *NHS hospital* for all generations, the size of the effect is largest for Baby Boomers.
- △ Third party organisation: Compared to being asked to provide their data to a *large technology company*, Millennials prefer for it to be stored on a *secure and private database*, where Generation X and Baby Boomers hold no preference.
- △ Privacy and sharing: Both Generation X and Baby Boomers prefer their data to be *de-identified and shared* and *identifiable but not shared* compared to *identifiable and shared*, however Millennials do not.
- △ Control: Compared to the use case, Baby Boomers are significantly more negative towards *ordinary citizens* making decisions about how the information is used. Millennials and Generation X, however, do not have a preference.

DELTA POLL

Figure 6. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 1 scenario attribute level compared to the baseline level of the attribute for Millennials, Generation X, and Baby Boomers (CDEI_S1_1)*



DELTA POLL

Disability status

Beyond differences in preferences by comfort level and generation, we also discovered some intuitive differences in response patterns by other respondent demographics. Respondents who report that their day-to-day activities are *limited a lot* because of a health problem or disability harbour significantly more positive attitudes towards sharing information on their disability status being shared than those who activities are either *limited a little* or *not limited*.

Other demographic characteristics

We also investigated whether responses, particularly in relation to the *type of data* collected, varied by respondent sex, ethnicity, or social grade. We find that they do not.

DELTA POLL


Conjoint 2: Predicting demographic data

In figure 7 and table 4 we present the results for the second conjoint. The analysis and interpretation of the results remain the same as the first conjoint.

- △ For use case, the baseline is *a car insurance company wants to check whether they are setting insurance prices fairly*. Compared to the car insurance company, respondents significantly prefer *a government department testing if they are allocating welfare benefits fairly, a government department testing if they are administering state pensions fairly, and an NHS hospital testing whether their diagnosis tool successfully diagnoses patients from different demographic backgrounds*.
- △ There is no statistically significant difference between the car insurance company and the remaining use cases: *a bank testing if they are providing loans fairly, a credit reference agency testing if they are providing credit scores fairly, a local authority checking whether they are allocating welfare benefits fairly, a technology company checking whether their advertising does not unfairly target certain groups of people, and an employer checking whether they are reviewing job applicants' CVs fairly*.
- △ For the proxy data used, the baseline is the use of a person's *postcode*. The use of postcode data is significantly preferred to the use of *a photograph of your face, social media usage, and someone's full name*.
- △ Turning to the demographic characteristics that the organisation is trying to predict, the baseline category is *ethnicity*. Respondents significantly prefer the use of proxy data to predict their *age group*, and, conversely, are less positive towards the prediction of their *disability status*.
- △ There are no statistically significant differences between the prediction of ethnicity compared to the remaining levels: *sex, sexuality, and socio-economic status*.
- △ For the privacy of the information, the baseline category is that the organisation will *make predictions about your own demographic characteristics that can be used to identify you personally*. There are no statistically significant differences between the baseline and the other two levels: the organisation will *make predictions about your own demographic characteristics, but apply technology that prevents them from identifying you personally* and *make predictions about the demographic characteristics of groups of people, but not about any individual specifically*.
- △ Finally, for the transparency of the process, the baseline category is that the organisation will *inform you privately that they will predict demographic information about you personally*. Respondents significantly prefer this to both *not inform you that they will predict personal information about you personally, and will not publish any information on their website informing members of the public that they are predicting demographic information to help them monitor for bias* and *inform you privately that they will predict demographic information about you personally and publish information on their website informing members of the public that they are predicting demographic information to help them monitor for bias*.

DELTA POLL

Figure 7. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 2 scenario attribute level compared to the baseline level of the attribute, standard errors clustered by respondent. Bars show 95% confidence intervals. (CDEI_S2_1)*



In sum, respondents feel most positively towards *an NHS hospital* testing whether their diagnosis tool successfully diagnoses patients from different demographic backgrounds using someone's *postcode* to predict their *age group* and the organisation will *inform you privately that they will predict demographic information about you personally*. Focusing only on the point estimates, respondents feel most positively towards the organisation making *predictions about your own demographic characteristics, but apply technology that prevents them from identifying you personally*, however this difference is not statistically significant.

DELTA POLL

Conversely, respondents feel least positively about *a photograph of their face* being used to predict their *disability status* and the organisation will *not inform you that they will predict personal information about you personally, and will not publish any information on their website informing members of the public that they are predicting demographic information to help them monitor for bias*. Again, focusing only on the point estimates, respondents feel least positively about *a credit reference agency testing if they are providing credit scores fairly* and the organisation making *predictions about your own demographic characteristics that can be used to identify you personally* but these are not statistically distinct from the second least preferred levels.

Don't know responses, which are given by 24% of respondents, are excluded from the analysis above. We find usage of don't know responses again varies by respondent characteristics where older respondents make greater usage of don't know responses than younger ones: 16% of 18-44, 31% of 45-64, and 32% of those aged 65 and over.

Table 4. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 2 scenario attribute level compared to the baseline level of the attribute (standard errors, clustered at the respondent level, in parentheses)*

Attribute	Level	Estimate
Use case (Baseline: A car insurance company)	A local authority	0.039 (0.025)
	A technology company	-0.007 (0.024)
	An employer	0.024 (0.024)
	An NHS hospital	0.115*** (0.026)
	A bank	-0.008 (0.025)
	A government department (welfare benefits)	0.062* (0.025)
	A government department (state pensions)	0.076** (0.025)
	A credit reference agency	-0.027 (0.025)
Proxy data (Baseline: Your postcode)	Your full name	-0.078*** (0.016)
	A photograph of your face	-0.257*** (0.018)
	Social media use	-0.184*** (0.018)
Demographic characteristics (Baseline: Ethnicity)	Sex	-0.012 (0.022)
	Age group	0.053* (0.021)
	Disability status	-0.050* (0.021)
	Socio-economic status	-0.017

DELTA POLL

		(0.021)
	Sexuality	-0.040
		(0.021)
Privacy (Baseline: Make predictions about your own characteristics that can identify you personally)	Predictions about your own characteristics but prevented from personal identification	0.008 (0.016)
	Predictions about groups of people but not specific individuals	0.028 (0.016)
Transparency (Baseline: Inform you privately)	Inform you privately and publish information on website	-0.053*** (0.016)
	Not inform you privately and publish information on website	-0.087*** (0.016)

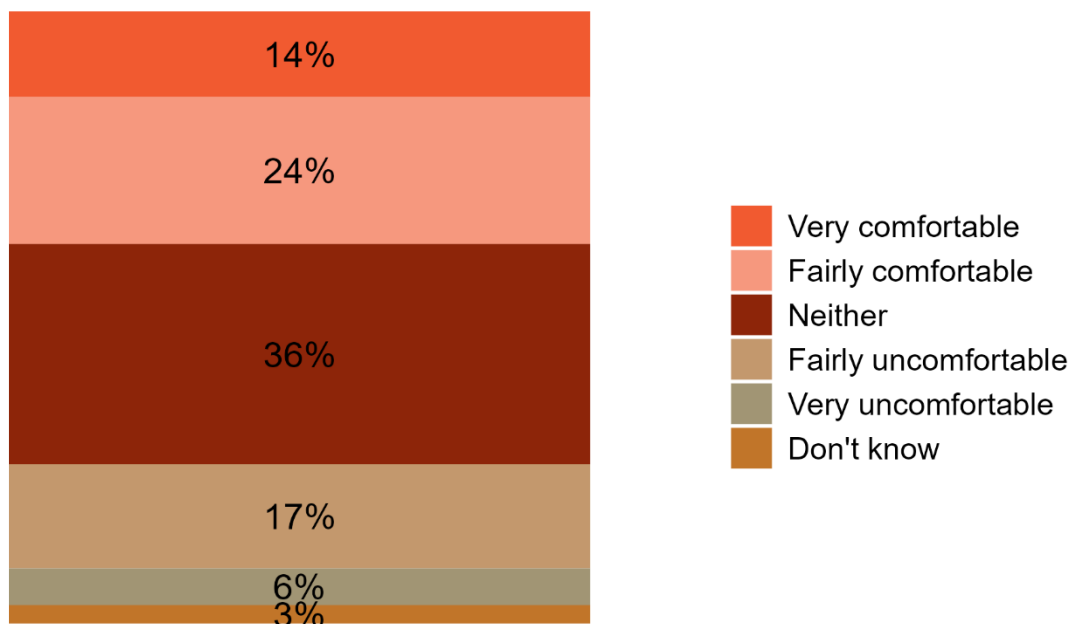
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Number of observations: 18,314

Number of respondents: 3,539

Figure 8 shows the results for CDEI_S2_2. Compared to CDEI_S_1_2, fewer respondents say that they would be comfortable with this form of demographic data collection and usage. Just over a third (36%) of respondents say that they would be comfortable, the same proportion (26%) say that they would be *neither comfortable nor uncomfortable*. A larger proportion of respondents, approaching one in five (23%), also say that they would be uncomfortable with an organisation using information they already hold to predict personal characteristics. As with the pilot, we randomised the question ordering such that half of respondents would be asked their comfort levels before the conjoint and half after. Again, we find that question order did not significantly influence responses.

Figure 8. *Taking everything into account, how comfortable or uncomfortable do you think you would feel if an organisation used information they already hold about you to predict your personal characteristics to help them test the fairness of their services? (CDEI_S2_2)*



DELTA POLL

Differential opinions by respondent demographics and attitudes

Comfort with sharing data

As with the first conjoint, we assess whether there are differential response patterns by those who say that they are comfortable or uncomfortable with an organisation using information they already hold to predict their personal characteristics. The results are presented in figure 9.

Across the board the responses by those who are comfortable and those who are uncomfortable are similar. In a number of cases, however, where differences are statistically significant for those that are comfortable, they are not for those who are uncomfortable. This is clearest for the *use case*, where there are no statistically significant differences for those that are uncomfortable, where there are for those that are comfortable. Beyond this the two sets of results are very similar to one another.

Generation

As with the first conjoint, we again assess whether responses differ by generation. In figure 10, we present the results for Millennials in the left hand panel, Generation X in the middle panel, and Baby Boomers in the right hand panel. While response patterns are similar in many ways across generations, there are a number of differences.


- ▲ Use case: Compared to *a car insurance company*, Baby Boomers significantly prefer *a government department* testing that they are fairly providing *state pensions* and *welfare benefits* where both Millennials and Generation X have no preference. Likewise, both Generation X and Baby Boomers prefer *an NHS hospital*, where Millennials again have no preference.
- ▲ Proxy data: Millennials, compared to the use of their *postcode* being used to predict their demographic information, do not feel significantly more negatively towards the use of their *full name*, where Generation and Baby Boomers do.
- ▲ Demographic characteristics: Compared to their *ethnicity* being predicted, Millennials feel more positive towards the prediction of their *age group* while Generation X and Boomers do not. Baby Boomers, however, are more negative towards the prediction of their *disability status* and *sexuality* where Millennials and Generation X have no preference.

Other demographic characteristics

We assessed whether preferences, in particular for the *demographic characteristics* that are being predicted, vary by respondent disability status, ethnicity, sex, or social grade. We find no variation in preferences among these demographic groups


DELTA POLL

Figure 9. Estimated Average Marginal Component Effects (AMCEs) of each conjoint 2 scenario attribute level compared to the baseline level of the attribute for those who are concerned and unconcerned about having their data predicted (CDEI_S2_1)



DELTA POLL

Figure 10. *Estimated Average Marginal Component Effects (AMCEs) of each conjoint 2 scenario attribute level compared to the baseline level of the attribute for Millennials, Generation X, and Baby Boomers*



DELTA POLL

Open text responses

In addition to the conjoint results, the survey also included a number of open text responses. We present the results of two forms of analysis of these responses here: word clouds and structural topic models.

Word clouds

First, to investigate in brief respondents' responses to the open text responses, we create word clouds for CDEI_S1_3, CDEI_S2_3, and PCheck3. Prior to estimating these, we removed all English language stop words (e.g., I, me, my) and punctuation from the texts. Figure 11 presents the top collocations for CDEI_S1_3, figure 12 presents the same information for CDEI_S2_3, and figure 13 for PCheck3. In each figure, the most frequently a word is used the larger it appears and the dark the colour it appears in.

Figure 11. Word cloud for CDEI_S1_3 "In your own words, what are your main concerns, if any, about sharing information about your gender, ethnicity, or other demographic information with a third-party organisation to help organisations check the fairness of their services?"



DELTA POLL

Figure 12. Word cloud for CDEI_S2_3 "In your own words, what are your main concerns, if any, about sharing information about an organisation using information they already hold about you to predict your personal characteristics to help them test the fairness of their services?"



Figure 13. Word cloud for PCheck3 "Thank you for taking part. In your own words, please say how you felt completing this survey".



DELTA POLL

Structural topic models

Second, to delve deeper into the topics that respondents raised, we also estimated statistical topic models. To do so, we estimated three correlated topic models (CTM) for the three open text response questions in the data: CDEI_S1_3, CDEI_S2_3, and PCheck3. The CTM is an unsupervised learning approach which assumes that the frequency with which words co-occur within different responses provides information about the topics that feature in these responses. The key output of a topic model is a matrix of topic proportions that measures the degree to which each response in the data features each of the estimated topics.

For CDEI_S1_3, "In your own words, what are your main concerns, if any, about sharing information about your gender, ethnicity, or other demographic information with a third-party organisation to help organisations check the fairness of their services?", 34% of respondents said that they *don't know* and did not provide a response. Of the remaining 66% of respondents who provided a response, we find that the most commonly raised issues are as follows:

- △ Concerns around data falling into the wrong hands;
- △ Concerns around the potential of identity theft or scams;
- △ Concerns that systems are not secure and will be sold on to further parties;
- △ Concerns around data breaches;
- △ Concerns that the data will be used without their permission;
- △ Concerns that third parties are not sufficiently trustworthy;
- △ Concerns that the data will not be securely stored and could be hacked;
- △ No concerns or worries expressed

Following the same analysis strategy for CDEI_S2_3, "In your own words, what are your main concerns, if any, about sharing information about an organisation using information they already hold about you to predict your personal characteristics to help them test the fairness of their services?", we find that 44% of respondents said that they *don't know* and did not provide a response. Of the remaining 56% of respondents who provided a response, we find that the most commonly raised issues are:

- △ Concerns around prediction feeling intrusive or like an invasion of privacy;
- △ Concerns that this sort of prediction is inaccurate or unfair;
- △ Concerns that organisations are using personal data without consent;
- △ Concerns that predictions rely on bad, unlimited, or incorrect information;
- △ Concerns that the information is leaked, ends up in the wrong hands, or is used by scammers;
- △ Concern about photos being made public;
- △ No concerns or worries expressed

Finally, for PCheck3, "Thank you for taking part. In your own words, please say how you felt completing this survey.", we find that 5% of respondents opted not to provide a response. Of the remaining 95%, the most commonly raised themes are:

- △ Thought that the survey was interesting and compelling;
- △ Thought that it was confusing and hard to complete;
- △ Thought that it dealt with interesting topics;
- △ That it was thought provoking and introduced new ideas;
- △ Thought that it was long and repetitive