Influence of group members in multi-attribute utilities

Matthew Gilbert*, Simon French and
Jim Q. Smith

University of Warwick,
Coventry, CV4-7ES, UK
Email: m.g.gilbert@warwick.ac.uk
Email: simon.french@warwick.ac.uk
Email: j.q.smith@warwick.ac.uk
*Corresponding author

Abstract: This paper investigates measuring the influence of some group members on others in decision making. Being better able to identify potentially influential behaviour would be useful in supporting and subsequently auditing a decision. A new measure of the influence of individuals is given, which is analogous to the well-known Cook's distance used to identify influential data in regression. The theoretical properties of this measure are explored. A simple method to identify sub-groups within the group of decision makers is given. We investigate the efficiency of our new measures using large scale randomised studies. We use these measures to identify sub-groups of individuals with similar beliefs in a data set collected in a previous experiment.

Keywords: group decisions; multiple attributes; influence; cultural groups; common beliefs; Cook's distance; group utilities.

Reference to this paper should be made as follows: Gilbert, M., French, S. and Smith, J.Q. (2017) 'Influence of group members in multi-attribute utilities', *Int. J. Multicriteria Decision Making*, Vol. 7, No. 1, pp.1–28.

Biographical notes: Matthew Gilbert received his MSc of Mathematics and Statistics from the University of Warwick in 2013. He is currently in his final year of his PhD at the same university studying under Prof. French and Prof. Smith. He is funded through an EPSRC studentship.

Simon French is a Professor in the Department of Statistics at the University of Warwick and has had a long career in modelling risks and decisions in nuclear energy. He also worked extensively on risk communication and stakeholder engagement.

Jim Q. Smith has been a Professor of Statistics at the University of Warwick for about 20 years and published widely in statistics, decision analysis and artificial intelligence. He has a particular interest in rational and probabilistic decision making in complex environments including those with multiple agents.

1 Introduction

Working with a group to make decisions introduces several new interactions that are not present for single decision makers; for example influence and coalitions. Previous results show us voting systems cannot be entirely democratic; Arrow's impossibility theorem (Arrow, 1950), and work following from this, has shown the only hope of democratic group decisions is to use relative weights for each alternative (Keeney, 2002; Keeney and Winterfeldt, 2009; Nguyen et al., 2009). Also we know that every voting system is susceptible to manipulation (Gibbard, 1973; Satterthwaite, 1975), although we may be able to design a voting system such that it is too complex to feasibly manipulate (Procaccia, 2010). Despite these negative results, group decisions are an important part of society, where in most contexts a group decision rule might not exhibit unacceptable behaviour. However, the development of diagnostic technologies within the formal frameworks of decision making has received little attention.

We suggest methods that help identify individual influence over the group decision. These methods also help identify sub-groups or cultural groups that could be active within the group. An obvious family of diagnostics for this purpose exploits the long development within statistical regression to identify influential data. A natural choice is Cook's distance (Cook and Weisberg, 1982) and the coefficient of determination (Magee, 1990). To develop the methods demonstrated in this paper, we extended results such as these from regression into a utility setting. On a more experimental level such problems have already been considered in social choice and game theory (Marett and George, 2013).

The work in this paper had been carried out with two specific contributions in mind. The first is to provide an adaptation of Cook's distance from regression that can be applied to the group decision making setting. Specifically, to when a group utility function has been formed from several individual utility functions. This adaptation could then be used to identify potentially interesting or influential behaviour of individuals demonstrated within the group, to ensure that no group members are attempting to unfairly influence the group's utility function.

The second contribution is to provide an adaptation to the R^2 value from regression, which can be applied to the group decision making setting. This could be used to identify hidden coalitions that could have formed which could be attempting to control the group's utility function. In other scenarios (for instance e-participation surveys) it could be used to identify group members from similar background or cultural groups. These two methods have the overall aim of helping to provide more confidence in the other group members, and a more fair and stable platform for group discussions.

In Section 2 we motivate the need to identify interesting behaviour, and define contexts and notation. Section 3 explores the theoretical aspects and definitions associated with the influence measure given earlier, and draw the necessary links to regression to allow methods for sub-group selection to be developed, alongside giving a metric that can assess the success of the diagnostics. Section 4 applies these methods to large scale randomly generated groups, both with no influential group members and with a set of influential group members. Section 5 applies these diagnostics to a dataset from the selection of a nuclear waste disposal facility. Finally Section 6 explores some possible extensions and discusses some shortfalls and strengths of these methods, along with their practicality in real scenarios.

2 Context and motivation

2.1 Measures in different contexts

An influence or sub-group identification measure can have different meanings depending on the context (DeSanctis and Gallupe, 1987). For instance, group discussions may allow the possibility for an individual to influence preferences of other individuals, making detection of influence more difficult. In large-scale online elicitation group members may not be able to work together, but instead we may be able to detect similar cultural groups. We identify three specific types of context that could be of interest based upon the type of interactions between group members (although this list is not exhaustive).

- *Scenario 1:* face-to-face group meeting. Group members can see the other members whether it be in person or using video call software such as Skype, for example a board of directors.
- *Scenario 2:* 'chat room' decision making. Group members could be anonymous and discuss the situation without identifying other members, for example an online deliberation forum of local council.
- *Scenario 3:* non-discussion decision making. There is no formal discussion between group members with their only interaction being through elicitation of their individual beliefs, for example the CORWM elicitation (Phillips et al., 2006).

In Section 5 we explore a data set that falls under the Scenario 3 category which, as there is no formal discussion mechanism, causes the influence measure to pick up other interesting behaviour. For example, extremist views or the possibility that a group member did not understand the elicitation process. The common theme is that in these contexts it is worth investigating any individual that is flagged as 'potentially influential'. We say potential influence as we are only able to detect a chance of being influential. The influence value may also be high due to other reasons that should be explored on a case-by-case basis.

2.2 An example

To further motivate the need for such influence measures (and sub-group identification methods), we have provided a small example. Consider 2 large companies that are discussing a mutually beneficial business agreement alongside a single external representative of the area that will be affected by the agreement. The group is relatively small (two people from each company and the representative, totalling five people), and meets face-to-face (scenario 1 meeting) to discuss several options available for the agreement (and to produce a group utility functions through each individual's utilities). The external representative has openly stated their support of the agreement, citing the financial benefits the area would see from it, and so the meeting is expected to proceed quite smoothly. However, there is some concern that there may be ulterior motives at play during the decision making process. Ideally the group members would like a diagnostic tool that could be used during the group meetings, which could identify potentially influential (and strange) behaviour. If this type of behaviour could be identified for one or more of the individuals involved, then they could have grounds to launch an investigation into why this behaviour had been exhibited.

On the other hand, consider a similar deal being made between four companies who also meet face-to-face. In this circumstance it would be a concern of some of the attendants of the meeting that other companies may have agreed to sway the decision a certain way to better benefit their own interests (so a secret coalition had formed). In this case it would be very useful to have a tool that could attempt to identify the groups of people that hold suspiciously similar beliefs, in the hopes of identifying any such secret coalitions. Once identified, the other companies could then take appropriate actions to mitigate the decision against such 'unfair' activities. These are the types of tools we aim to lay the foundation for, by adapting methodologies developed in the area of regression into this utility setting.

2.3 Mathematical formulation

Suppose we have a group where each individual provides their own utility function, and the group agrees to combine these into a single linear utility function describing the beliefs of the group.

- a group comprised of N members, indexed by $i = 1,\ldots,N$, which can be described by s subgroups, where $1 \leq |s| \leq N$
- all utility functions describe the individual's preferences in terms of K attributes, $(A_j)_{j=1}^{K}$, where $a_j \leq A_j \leq b_j$ with $a_j, b_j \in \mathbb{R}$ for $j = 1,2,\ldots,K$
- an individual's preferences are described by a vectors of weights for the associated attributes, $\mathbf{u}_i^{(j)} = (u_i^{(1)}, u_i^{(2)},\ldots,u_i^{(K)})$ for $i = 1,\ldots,N$
- we assume that these weights are normalised such that $\sum_{j=1}^{K} u_i^{(j)} = 1$ and $u_i^{(j)} \geq 0$ for $j = 1,\ldots,K$ and $i = 1,\ldots,N$
- there exists a group decision operation, G, that maps from the N individual weight vectors to a single weight vector that describes the group's preferences, $\mathbf{u}_G$.

Definition 1: define the $(K - 1)$-simplex that each normalised linear utility function is contained in to be S: = $\{(x_1,\ldots,x_K)$: $\sum_{j=1}^{K} x_j = 1,\ x_j \geq 0$ for $j = \{1,\ldots,K\}$. Then define an operation G: $S^N \rightarrow S$ to be a *group decision operation*.

Definition 2: take the normalised linear utility function over K attributes: $\mathbf{u} = (u^{(1)}, u^{(2)},\ldots,u^{(K)})$. The value $u^{(j)}$ is defined as the *attribute weight* for attribute j, for $j = \{1,\ldots,K\}$.

Note that this group decision operation could be decided by the group, or could simply be taken to be an average function across each of the attribute weights (we assume the latter for the sake of consistency).

Notation: given a group decision operation G: $S^N \rightarrow S$. We state the *group attribute weights*, $\mathbf{u}_G$, to be;

$$\mathbf{u}_G = G(u)$$

Definition 3: define the group attribute weights when individual i^* is excluded, $\mathbf{u}_{G-1^*}$, as;

$$\mathbf{u}_{G_{-i*}} = G'(\mathbf{u}_{-i*})$$

where $G': S^{N-1} \rightarrow S$ is the group decision operation G which has been adjusted for $N - 1$ individuals and $\mathbf{u}_{-i*} = (\mathbf{u}_i)_{i=1, i \neq i*}^{N}$.

Note: we generally use G for all group decision operations, as the dimension of the co-domain should be clear from the context.

3 Influence measures

3.1 Influence in regression modelling

Our initial inspiration for the development of influence measures in group utilities came from strong links to standard regression modelling. See Sen and Srivastava (2011) for a summary of regression modelling. These links come from both from influential data and the set-up of the regression scenario. In particular the group decision operation in group utilities could be considered as analogous to the parameters in regression, given they perform the same rough function; combining a matrix of explanatory variable (or individual utilities) into a vector of response variables (or group utility function) that we are interested in. These links allowed us to consider how some of the common measures of influence in regression, such as Cook's distance (Cook and Weisberg, 1982), could extend into a utility setting.

3.2 Group decision operation

The success of the influence measure developed in this paper hinges on the group satisfying the properties outlined in Section 2.2. The most debatable property is that the group decides upon a group decision operation, so we consider this here.

Property 1: the group has agreed to use a group utility function to represent the beliefs of, and help make choices for the group.

Property 2: the group defines a group decision operation, G, that maps from N attribute weight vectors to a single weight vector for the group's preferences, u_G.

In particular we assume that the group uses a single group weight, $\mathbf{u}_G$, to describe the beliefs of the whole group. There are problems with this as the group utility function loses information about each individual's personal preferences, and the differences between the group members.

This may not be a problem in certain circumstances that call for consistency in public reporting. One example is an issue in public safety from the introduction of nuclear power to an area. The group overseeing public and environmental safety will meet to discuss the importance of different factors concerned with the issue, however when submitting reports to the public it should be clear that there was agreement within the group or the public will have little trust in the organising group. Using a single group utility function to describe the beliefs of the group allow for this consistency through being a clear statement of beliefs.

We have focussed only on contexts where using a group utility function has an advantage over less restrictive methods. This leads on to the problem of understanding what form the group decision operation can take. First we introduce two properties below, adapted from those seen in Arrow's view of a 'fair group' to ensure that the behaviour exhibited will be consistent and rational (Arrow, 1950). Note that with the introduction of highly correlated utility weights there is no analogous property to independence of irrelevant alternatives as no attributes can be irrelevant.

1. if all individuals in a group select their utility weight for attribute j such that $u_i^{(j)} > k$ for all $i = 1,\ldots,N$ and some constant $k \in (0, 1)$, then the group utility function must also have $u_G^{(j)} > k$
2. there is no individual i^* such that $u_{G_{i^*}} = u_G$ for any choice of utility functions for individuals $i \neq i^*$ and $i = 1,\ldots,N$.

Results from Keeney and Raiffa (1976) and Neumann and Morgenstern (1953) show that the form of decision operation in Assumption 3 below is the only type of group decision operation that satisfies similar conditions to Arrow's theorem. An important assumption to allow for this combination is that we have interpersonal comparisons of utilities. This is a key assumption in our context and should be applicable when the elicitation process is clear to all group members.

Property 3: for individual weight vectors $(\mathbf{u}_i)_{i=1}^N$ let $u_G^{(j)} = \sum_{i=1}^N w_i u_i^{(j)}$, where w_i is the weight assigned to individual i and $\sum_{i=1}^N w_i = 1.$

There are adaptations to this form that could be useful under certain contexts. For example requiring one attribute to have a minimum weight would cause the group decision operation to be similar to, but not quite, a linear combination as suggested. However it would still satisfy the conditions we have given above under the assumption that the group recognises that this type of minor dictatorship is required and adjustments are applied after the combination (e.g., government requirements).

Definition 4: individuals i_1 and i_2 are *permutable* if $C_{i_1}(\mathbf{u}, G) = C_{i_2}(\mathbf{u}^*, G)$, where $C_{i_1}(\mathbf{u}, G)$ is a measure of influence for individual i_1 over the group attribute weights $\mathbf{u}$ using the group decision operation G and $\mathbf{u}^*$ exchanges the attribute weights of i_1 and i_2 in $\mathbf{u}$. A group is permutable if all subsets of {1,…,N} are permutable.

The group decision operation we used was the arithmetic group mean where $w_i = N^{-1}$ for all $i = 1,\ldots,N$, which gives a simple average utility function. It must be highlighted that our choice of group decision operation is not a requirement on the suggested influence measure, but instead is being used for simplicity. This particular case allows for the permutability of the individuals seen below.

Observation 1: if $u_{i_1}^{(j)} = u_{i_2}^{(j)}$ for $j = 1,\ldots,K$, then $C_{i_1}(u, G) = C_{i_2}(u, G)$.

While this can be a useful property for the influence measure, it will often not be the case when the group decides to value the beliefs of one member more than another member (for instance the chairman of a committee) which could typically be the case in scenario 1.

Theorem 1: the only group decision operation of the form in Assumption 3 to exhibit permutability of individuals is when $w_i = N^{-1}$ for all $i = 1,\ldots,N$.

This property may be useful in a more general context where we assume that all group participants are to be treated equally, as sub-groups will be less dependent on having a high weight group member included. An example of this could be from scenario 3 where we may have large scale web-studies where participants do not know each other, and provide their individual utility function. In the case of a weighted group decision operation, we can draw similar conclusions with the added condition that the individuals have the same weight assigned to them.

Theorem 2: for individuals i_1 and i_2, $u_{i_1}^{(j)} = u_{i_2}^{(j)}$ and $w_{i_1} \in (0, 0.5)$. Then if $w_{i_2} > w_{i_1}$ we have that $C_{i_2}(u, G) > C_{i_1}(u, G)$.

3.3 Influence

3.3.1 Influence measure

The base influence measure that is given as $C_i(\mathbf{u}, G) = K^{-1}D(\mathbf{u}_G, \mathbf{u}_{G-i})$.. The function $D(.,.)$ is chosen from those given in Section 3.3.2 below. First however, note that we initially are considering influence with respect to the attribute weights as opposed to the final utility values of each individual. This is to allow us to answer both whether the influential could be influential, and how that individual has attempted to exert influence. Focussing only on the final utilities scores would provide us with the answer to the first of these questions, but it would require more work to understand how the individual had been attempting to exert influence. However we plan to perform further work on this to assess whether this extra information is useful, and to consider a more general but simple influence measure which would be easier to explain to groups. We initially need to address what is meant by influence more precisely.

Definition 5: given a distance measure D: $S \times S \rightarrow \mathbb{R} \geq 0$ and an influence measure $C_i(\mathbf{u}, G)$, then individual i is *influential at specified level* $\delta > 0$ if $C_i(\mathbf{u}, G) > \delta$.

Definition 5 gives a more precise understanding of an influential individual, and also gives groups freedom in how sensitive to manipulation they wish to be. Unfortunately we cannot appeal to asymptotics for a suitable δ level due to a limited number of both group members and attributes. The selection of a suitable δ value is explored in Sections 4 and 5.

Definition 6: if $C_i(u, G) = 0$, then individual i is *influentially irrelevant* to the decision for this specific set of utility functions and group decision operation.

The concept of an influentially irrelevant individual should be interpreted with care. We define this with respect to the decision and utilities given at the moment the influence measure is used. Any change in utilities from the group members could make a previously influentially irrelevant member relevant. Justification of a group member being influentially irrelevant in scenarios 1 and 2 is less likely due to the impact they could have on other utility functions. However the concept would be more applicable to scenario 3 where we can define extremely non-influential group members. To better understand this diagnostic we consider two methods of introducing influence based on increasing the individual's attribute weight of a single attribute (we can assume attribute 1 for individual 1 without loss of generality) by a fixed amount L. This value of L has bounds to ensure the attribute weights stay within [0, 1] after adjustment. The two methods differ in how we choose to balance this increase with the other attributes, as we still have the requirement for the attribute weights to sum to 1. This means that the reductions to the other attribute weights must sum to L. Define the utility function of individual 1 after influence is introduced to be $\mathbf{u}_{1,m*}$ where m is the case number from below, then we consider the two cases below;

- Case 1

$$\mathbf{u}_{1,1*} = \left(u_1^{(1)} + L, u_1^{(2)} - L, u_1^{(3)}, \ldots, u_1^{(K)}\right)$$

- Case 2

$$\mathbf{u}_{1,2*} = \left(u_1^{(1)} + L, u_1^{(2)} - \frac{L}{K-1}, u_1^{(3)} - \frac{L}{K-1}, \ldots, u_1^{(K)} - \frac{L}{K-1}\right)$$

Here we see that case 1 directly moves weight from one attribute to another, while case 2 increases the weight of attribute 1 at the expense of all other attributes equally. Our main interest was how the value of depends on the case used. For this we compared $D(\mathbf{u}_G, \mathbf{u}_{G,1*})$ and $D(\mathbf{u}_G, \mathbf{u}_{G,2*})$, where $\mathbf{u}_{G,m*}$ is the group utility function when considering case m.

Theorem 3: assume we introduce influence to individual 1 on attribute 1 as in Case 1 and Case 2 above, if we also assume that $u_G^{(2)} < u_G^{(j)}$ and $u_1^{(j)} > 0$ for $j \neq = 2$ and $j \in \{1,\ldots,K\}$ and choose L such that all adjusted utilities are in (0, 1) then we have;

$$D(\mathbf{u}_G, \mathbf{u}_{G,1*}) > D(\mathbf{u}_G, \mathbf{u}_{G,2*})$$

This theorem shows that transferring weight from one attribute to another is easier to detect when the attribute being reduced is small relative to the other utilities. Note that these calculations were done for the KL-divergence, although all distance measures behave similarly. The non-negativity condition is for simplicity. This result extends trivially for zero utility weights.

The more interesting property is that the case which produces a higher value for the influence measure depends on the size of $u_G^{(2)}$ relative to the other utilities $u_G^{(j)}$ for $j = 3,\ldots,K$. When $u_G^{(2)}$ is relatively large then case 2 gives a larger value, while when $u_G^{(2)}$ is relatively small the opposite is true. The behaviour of the influence measure between these extremes seems relatively smooth, although the details have yet to be specified (in particular for equality between the cases). This behaviour comes from the fact that when $u_G^{(2)}$ is relatively large, the other utility weights change by a relatively larger proportion than $u_G^{(2)}$ due to all utilities being reduced by $\frac{L}{K-1}$. This shows that the influence measure responds more to relative changes in utilities rather than absolute changes, which is what we would hope. If we had changed case 2 to reduce each utility weight relative to the original weight for each attribute, it is likely that the influence measure would perform similarly in both cases.

3.3.2 Distance measures

The distance measures considered for use in C_i are given in Table 1 from our initial tests, all behave similarly. The first three are common measures of distance between probability distributions, while the cosine divergence was adapted from Chung et al (1989).

Table 1 Definitions of the four distance measures we have considered

Distance measure	*Definition*
Jeffreys' distance	$D(w,v)=\sum_{j=1}^{K}(\sqrt{w^{(j)}}-\sqrt{v^{(j)}})^2$
KL-divergence	$D(w,v)=\sum_{j=1}^{K}\left(w^{(j)}\log\frac{w^{(j)}}{v^{(j)}}\right)$
J-divergence	$D(w,v)=\sum_{j=1}^{K}\left((w^{(j)}-v^{(j)})\log\frac{w^{(j)}}{v^{(j)}}\right)$
Cosine divergence	$D(w,v)=\frac{1}{2}\left[1-\sum_{j=1}^{K}f(w^{(j)}-v^{(j)})\right]$ where $f(w^{(j)},v^{(j)})=\begin{cases}(w^{(j)}v^{(j)})^{\frac{1}{2}}\cos\left(\alpha\log_2\frac{w^{(j)}}{v^{(j)}}\right) & \text{for } \left\lvert\alpha\log_2\frac{w^{(j)}}{v^{(j)}}\right\rvert<\frac{\pi}{2}\\ 0 & \text{otherwise}\end{cases}$

Note: Both w and v represent vectors of length K, where w^j and v^j are the j^{th} element of each vector.

3.3.3 Sub-groups

The method used to assess the likelihood of a sub-group collaborating was adapted from the widely used R^2 value in regression. This was called the adjusted R_s^2 value which is given below for a specific sub-group s.

$$R_s^2=1-\frac{\sum_{i\in s}\sum_{j=1}^{K}\left(u_i^{(j)}-u_{G_s}^{(j)}\right)^2}{\sum_{i\in s}\sum_{j=1}^{K}\left(u_i^{(j)}-u_G^{(j)}\right)^2}$$

The adjusted R^2 value measures how different a sub-groups beliefs are to the whole group by re-evaluating a group utility function using only the members of that sub-group. Similarly to R^2 in the regression sense, we expect groups with values close to 1 to be potentially working together. This preliminary definition of R_s^2 has provided promising results when applied to generated data and to our data set. However for particularly large groups (for instance we considered interactions between 46 people as a single group in our data set) using this method to consider all possible subgroup possibilities is very inefficient. As such we would only recommend using this method for smaller group sizes, or evaluating specific sets of sub-groups in larger group settings.

3.3.4 Assessment of success

For our influence measure, we assessed success by generating the individual utility functions for large numbers of groups with and without influential individuals to compare the true and false positive rates. A true positive in this setting is when an individual that was constructed to be influential is flagged as being potentially influential. A false positive however is when an individual that was not constructed to be influential is flagged as potentially influential. The true positive and false positive rates are therefore the percentage of cases which falls under each definition respectively. Traditionally in statistics we fix one of these rates and maximise (or minimise depending on the rate) the other, however in this situation both of these rates are very important to success and the trade-off between them should be considered (Lashner, 2006).

For the sub-group identification method we first checked that we can identify a sub-group designed to work together in a larger group. We then applied the method all non-trivial sub-groups (sub-group size between 2 and $N - 1$) of newly generated groups where there were no sub-groups specifically designed to be working together. We counted the number of flagged sub-groups of each size. This was repeated for newly generated groups where a sub-group designed to work together had been included and contrasted with results from when no sub-groups were specifically included. Finally we applied these methods to our dataset and analysed whether the reported sub-groups and potentially influential individuals were sensible in the given context.

4 Numerical studies

4.1 Single influential individual

Our first experiment took a set of simulated utility functions, and manipulated the utilities of one individual on a single attribute so that we would expect them to be flagged as potentially influential. We then compared this with simulated utility functions where no group members should be particularly influential. Unless otherwise specified we assume ten group members providing utility functions on eight attributes.

To create an individual's utility function, a randomised 'base' utility function, $\mathbf{u}_b$ was sampled to represent the average preferences of the group. This base utility function was taken and a random deviation was applied to each attribute weight to reflect the individual's personal preferences. The utility function was then renormalised such that the attribute weights sum to 1. The algorithm used is given below;

Step 1 Sample the base utility function through $\mathbf{u}_b^{(j)} \sim U[0,1]$, for $j = 1,\ldots,K$.

Step 2 Renormalise the base utility function such that $\sum_{j=1}^{K} \mathbf{u}_b^{(j)} 1$.

Step 3 Sample each individual's utility function from the base utility function through

$$\mathbf{u}_i^{(j)} \sim N\left(\mathbf{u}_b^{(j)}, \frac{\mathbf{u}_b^{(j)}}{2K}\right), \text{ for } i = 1,\ldots,N \text{ and } j = 1,\ldots,K.$$

Step 4 Renormalise all individual utility functions such that $\sum_{j=1}^{K} \mathbf{u}_i^{(j)} 1$ for $i = 1,\ldots,N$.

All individuals in the group derive their personal utility functions from the same base utility function $\mathbf{u}_b$ so that the attribute weights remain relatively consistent. This was not a requirement for any of the methods; however it mirrors a real situation more closely than when there is no dependence. It also makes it simpler to introduce the concept of influence into the group.

The randomisation around the base utility function was sampled from a normal distribution with parameters $\mathbf{u}_b^{(j)}$ and $\mathbf{u}_b^{(j)}(2K)^{-1}$ to ensure that the groups beliefs would stay centred around the base utility function, but with enough variance to demonstrate individuals having different beliefs. The variance was dependent on $\mathbf{u}_b^{(j)}$ to keep utility weights close to the base weight, relative to the value of the weight itself. For instance a weight of 0.05 should have less variance than a weight of 0.4. The variance was divided by 2 K to account for the effect large changes would have on other utility weights at the normalisation step. Note that if the attribute weight became negative, it was set to 0 to suggest the individual does not believe that attribute is important.

When an influential individual is introduced into the group, an independent selection of m attributes will become inflated relative to their initial value (assume m = 1 unless otherwise stated). This adjustment will be independently repeated for every influential individual introduced to the group. When we introduce an influential individual, we follow steps 5 and 6 below to introduce influential individuals 1 to n (assumed to be independent of each other).

Step 5 Sample each influential individual's utility functions independently from their original utility function using $\tilde{\mathbf{u}}_i^{(l)} \sim N\left(\mathbf{u}_i^{(l)} + \frac{2}{3}\mathbf{u}_i^{(l)}, \frac{\mathbf{u}_b^{(l)}}{K}\right)$ where $l \sim U[\{1,\ldots,K\}]$, for $i = 1,\ldots,n$.

Step 6 Renormalise the influential individual utility functions such that $\sum_{j=1}^{K} \tilde{\mathbf{u}}_i^{(j)} = 1$ for $i = 1,\ldots,n$.

The parameters of the normal deviation were chosen as they will usually place the attribute value far enough away from the rest of the group, but also not far enough such that it is obvious the individual is trying to be influential. While we have only presented results for when n = 1, our influence measure works for larger numbers of influential individuals. In this case we would consider success of result on the individual level, rather than on the whole group. For instance, if we introduce two influential individuals, and only one of the individuals is flagged then only one true positive value would be recorded. The undetected influential individual would still be accounted for in the false negative rate (opposite of the true positive rate, these two rates must sum to 1) of the influence measure.

4.1.1 Scaling

Using this algorithm for generating group utility functions and influential individuals, tests were done with 1,000 randomly generated groups, of differing sizes to investigate what δ levels are suitable for each of the distance measures. The distance measures were found to need some slight scaling to account for group size, attribute number (for all distance measures) and α value (for just the cosine divergence). Figure 1 shows an example of the robustness to group size after this scaling for Jeffrey's distance. This robustness can also be seen when scaling for attribute number and α value.

Figure 1 Set of box plots showing the distribution of scaled influence values for an influential individual at group sizes 4, 8, 12, 16 and 20 using Jeffrey's distance

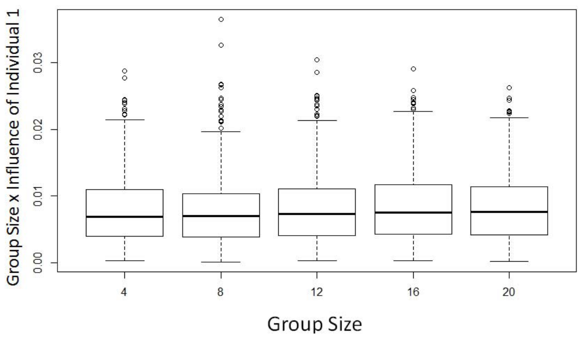

Our tests suggested that it would be simpler to include attribute number scaling directly on the δ value we compare the influence measure to, rather than the influence measure itself. To do this we need a base value δ_b which can be used to calculate the δ level for different numbers of attributes. Thus we define δ to be $\delta = \frac{\delta_b}{K+1}$. We have provided an initial suggestion of δ_b in Section 4.1.3 when using the cosine divergence. Note that the cosine divergence displayed slightly different patterns to the other distance measures, and needed to be scaled by N^2 and $(K + 1)^2$ instead of N and $(K + 1)$ like the other distance measures. The scaled form of the influence measure for each set of distance measures is given below where scaling for attribute number is used for the δ value.

$$C_i(\mathbf{u}, G) = \frac{N^2 D(\mathbf{u}_G, \mathbf{u}_{G_{-i}})}{K\alpha^2} \text{ for cosine divergence.}$$

$$C_i(\mathbf{u}, G) = \frac{ND(\mathbf{u}_G, \mathbf{u}_{G_{-i}})}{K} \text{ for other distance measures.}$$

4.1.2 Selection of δ

With our influence measures producing similar values no matter the group or attribute set-up, we can move our attention to the selection of a suitable δ level for each distance measure. Recall in Section 3.3.4 we reviewed the definitions of true and false positive results, and the need to balance these rates in our scenario. We aim to maximise the true positive rate, while minimising the false positive rate. However the optimal trade-off between these two rates is dependent on the group's preferences. We provide a value of δ that we think is suitable for most cases in 4.1.3. This δ value corresponds to a 93.4% true positive rate, and between 2.7% and 4.4% false positive rate dependent on the distance measure chosen. This value can then be used, along with the scaling options suggested in Section 4.1.1 to calculate δ_b, which allows generalisation to other attribute numbers for the same true and false positive rates.

4.1.3 Results

The groups we use for analysis during this section were randomly generated according to the process seen in Section 4. Each group had ten members, with eight attributes and a single influential individuals being introduced where appropriate. In other words, $N = 10$, $K = 8$ and $n = 1$ or 0 for influential and non-influential groups respectively.

Figure 2 Box plots comparing distribution of 25,000 scaled influence values for an influential and non-influential individual in a groups of size 10 using Jeffrey's distance

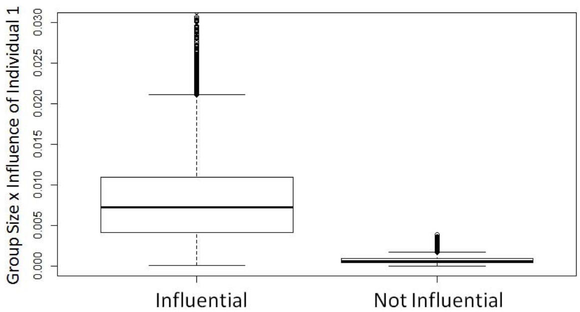

The results for Jeffrey's distance can be seen in Figure 2 which shows a clear difference between the influenced and non-influenced groups, however larger values of the non-influenced group are close to the lower quartile of the influenced group. This means we cannot choose a perfect δ value. So we need to decide on a δ value with the right balance of true and false positives.

Table 2 Comparisons of true positives against false positives from two sets of 25,000 simulated groups for each distance measure, and associated δ values as column headings

Jeffrey's distance	0.0072	0.009	0.0108	0.0126	0.0144
Influential positives (%)	97.96	96.94	95.86	94.72	93.42
Non-influential positives (%)	32.59	20.38	12.31	7.44	4.38
KL-divergence	0.0216	0.0243	0.027	0.0297	0.0324
Influential positives (%)	95.88	95.00	94.10	93.14	92.23
Non-influential positives (%)	12.31	8.48	5.77	3.79	2.44
J-divergence	0.0432	0.0486	0.054	0.0594	0.0648
Influential positives (%)	95.86	94.97	94.06	93.09	92.19
Non-influential positives (%)	12.31	8.49	5.76	3.79	2.46
Cosine divergence (α=1)	0.03888	0.04374	0.0486	0.05346	0.05832
Influential positives (%)	96.12	95.25	94.34	93.38	92.40
Non-influential positives (%)	10.15	6.74	4.24	2.68	1.73

Table 2 shows the true and false positive rates for different δ values for each of the distance measures we had considered. By comparing a true positive rate of roughly 93.4% with the associated false positive rate, Jeffrey's distance performs the worst while KL-divergence and J-divergence performed similarly to each other. Cosine divergence performed the best as it minimised the false positive rate the most, improving it by over

1% from the KL-divergence. We can calculate δ_b from these values. For example for cosine divergence, if we select $\delta = 0.05346$ then we know $\delta_b = 4.33026$.

Table 3 Comparisons of true positives against false positives from sets of 1,000 simulated groups for the cosine divergence with different values of α at $\delta = 0.05346$

α value	0.5	0.9	1	1.1	1.5	2
Influential positives (%)	97.3	91.5	92.9	92.2	91.4	93.9
Non-influential positives (%)	5.8	2.6	2.1	3.4	1.8	1.4

Due to the best performance coming from the cosine divergence we may ask if we can improve the results by considering different values of α. We must be careful as increasing α too much will make the influence measure too concentrated, and so we only consider values of α up to 2. Table 3 shows we can see a slight downward trend for the false positive rate as we increase α. We reran the tests with the larger sample size of 25,000 groups for just $\alpha = 2$ to explore the consistency of results. The false positive rate was 1.76% and the true positive rate was 92.61% which was only a slight improvement over when $\delta = 0.05832$, and so we will proceed by keeping $\alpha = 1$.

4.2 Sub-groups

The method proposed to identify possible sub-groups is the adjusted R^2 value. We have considered two group types; a 'standard' group with individual utility functions derived from steps 1 to 4 seen in Section 4.1, and an 'altered' group where we have adjusted the attribute weights of a known sub-set of the members to reflect collaboration. This ensures we know which sub-group has been specifically designed to collaborate at this stage. When we introduce a collaborating sub-group, there are two variables to be defined. The number of attributes the subgroup is trying to influence is m, and the size of the sub-group is n. To introduce a sub-group to the group, we followed the steps below;

Step 1 Create a set of individual utility functions for the group as in steps 1 to 4 in Section 4.1.

Step 2 Produce a mean utility weights vector; $\bar{\mathbf{u}} = \frac{\sum_{i=1}^{N} \mathbf{u}_i^{(j)}}{N}$ for $j = 1,\ldots,K$.

Step 3 Sample each individual's utility function (in the sub-group) from the mean utility weights vector through $\mathbf{u}_i^{(j)} \sim N\left(\frac{3}{2}\bar{\mathbf{u}}^{(j)}, \frac{\bar{\mathbf{u}}_b^{(j)}}{2K}\right)$, for $i = 1,\ldots,n$ and $j = 1,\ldots,m$.

Step 4 Renormalise individual utility functions (from the sub-group) such that $\sum_{j=1}^{N} \mathbf{u}_i^{(j)} = 1$ for $i = 1,\ldots,n$.

This produces a sub-group of size n that have roughly similar beliefs which fall outside the average of the group, allowing for personal preferences within the sub-group (hence the relatively tight variance). We specifically move people away from the group as this is the situation we are more interested in, as having a sub-group that is just as average of the rest of the group is not very informative.

4.2.1 Scaling

Figure 3 shows the effect of scaling the adjusted R^2 value to the power $\frac{1}{n}$, where n is the sub-group size. While there are slightly lower distributions for lower and higher sub-group sizes (2, 3, 8 and 9), the adjusted R^2 value remains consistently high enough to provide reliable detection. This can be seen from comparing to groups with no sub-group specifically introduced, seen below in Figure 4.

Figure 3 Box plots comparing the distributions of scaled adjusted R^2 values for 100 groups of size 10 considering six attributes when we introduce sub-groups of various sizes

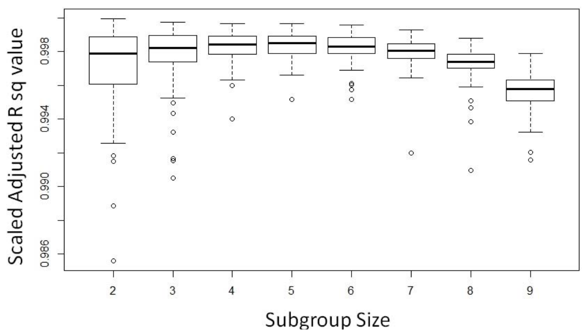

Figure 4 Box plots comparing the distributions of scaled adjusted R^2 values for groups of size 10 considering six attributes when we introduce sub-groups of various sizes (left) or when we do not introduce sub-groups (right)

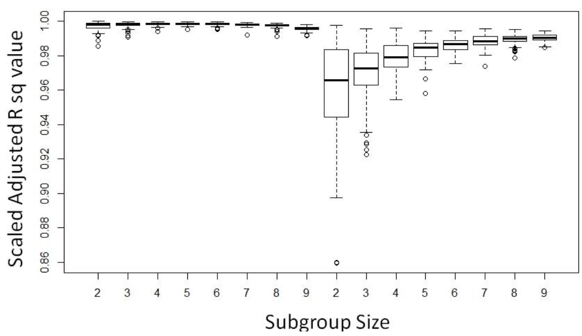

While Figures 4 and 3 show that our proposed scaling method is not perfect due to the more extreme group sizes, the scaling appears to be sufficient to consistently identify collaborators, while keeping the false positive rate low. To account for this scaling factor, we redefine the adjusted R^2 accordingly, as shown below.

$$R_s^2 = \left(1 - \frac{\sum_{i \in s} \sum_{j=1}^{K} \left(u_i^{(j)} - u_{G_s}^{(j)}\right)^2}{\sum_{i \in s} \sum_{j=1}^{K} \left(u_i^{(j)} - u_G^{(j)}\right)^2}\right)^{\frac{1}{n}}$$

We can see that when we look at groups where we have forcibly introduced a sub-group (the left side of box-plots) it gives a far higher value than when looking at a group with no sub-groups, which is quite promising. Unfortunately the scaling has had the opposite effect on the sample with no sub-groups, but despite this the difference between the two sets of groups is still clear.

4.2.2 Results

We can see from our initial explorations that the adjusted R^2 value can identify sub-groups relatively consistently when we introduce them and we know which sub-group we are looking for. However this is unlikely to be the case in practice if members of the group are concerned about any coalitions that have been formed that they are unaware of. Because of this we looked at all possible subgroups of 100 generated groups that have each had a sub-group included of a specific size, and record how many sub-groups are discovered for differing group sizes and sub-group sizes. The main scenarios we considered are given in Table 4.

Table 4 Summary of the sub-group tests performed

Test number	*1*	*2*	*3*	*4*	*5*	*6*
Number of group members (N)	8	8	12	12	16	16
Number of attributes (K)	6	6	10	10	10	10
Sub-group size (n)	3	0	4	0	6	0
Number of sub-group attributes (m)	3	0	5	0	3	0

The six tests shown in Table 4 were repeated 100 times, and all sub-groups were checked each time to see if the adjusted R^2 value was greater than 0.995 (which was a preliminary value chosen from initial tests). We then counted the number of sub-groups that passed this check for each possible sub-group size, and these results are given below. Keep in mind that these are the number of reported sub-groups for all 100 simulated groups, not just a single group, and so 300 reported sub-groups is an average of three sub-groups reported per group.

The results in Table 5 demonstrate a relatively low false positive rate as, for all groups with no sub-group forcibly introduced; the average number of significant sub-groups was less than 1. We can also see that there may be some difficulty in detecting smaller sub-groups, as on average only half of the subgroups of size 3 in test 1 were reported as significant. This behaviour could be expected when recalling Figure 3 and is something that should be kept in mind when exploring a set of data for sub-groups.

Table 5 Number of significant sub-groups of each size

Sub-group size	*2*	*3*	*4*	*5*	*6*	*7*
Test 1 (sub-group)	170	50	4	1	0	0
Test 2 (no sub-group)	22	2	0	0	0	0

Sub-group size	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*
Test 3 (sub-group)	1,213	1,283	1,420	1,456	1,029	418	71	0	0	0
Test 4 (no sub-group)	35	1	0	0	0	0	0	0	0	0

Sub-group size	*2*	*3*	*4*	*5*	*6*	*7*	*8*
Test 5 (sub-group)	2,814	5,418	9,482	1,3883	1,4208	9,595	3,967
Test 6 (no sub-group)	61	9	3	1	0	0	0
Sub-group size	*9*	*10*	*11*	*12*	*13*	*14*	*15*
Test 5 (sub-group)	910	93	0	0	0	0	0
Test 6 (no sub-group)	0	0	0	0	0	0	0

It is from tests 3 and 5 we can start to see the power this sub-group identification method. In both cases there is a noticeable peak around the sub-group size that was introduced. The reason that the average number of reported significant sub-groups is relatively larger (for example 14.2 for test 3 and 142 for test 5) is that it is counting slight changes to the sub-group as well, so for instance including three of the four sub-group members we introduced, and one other non-sub-group member that has somewhat close beliefs. This means the best way to use this method is to look for sets of common individuals in the reported sub-groups. This allows us to deal with the larger number of report significant sub-groups.

5 Case study

5.1 Dataset

The data we use is taken from E.C. Atherton's (1999) thesis which explored the temporal issues in decision making, and we used the data gathered for her first experiment 'the hypothetical decision'. Participants in a web experiment were asked to give attribute weights (using an online tool) on the importance of several attributes over several eras of different lengths on the construction of types of facility for nuclear waste disposal. Five attributes were considered (construction, health, environment, operating and accidents) over four time periods (immediate, 0–100 years, 100–500 years and 500–1,000 years). In other words, each participant provided 4 utility functions, one for each time period. The participants were then asked to weight the importance of each of the time periods so that their four utility functions could be combined into a single utility function representing all four attributes over all four time periods. For more information on the design of this experiment, see Atherton and French (1998). The distribution of utility weights of each of these five attributes can be seen below in Figure 5 and an example utility function of a participant can be seen in Table 6.

Figure 5 Box plots comparing the distributions of attribute weights assigned to each of the attributes for all participants (see online version for colours)

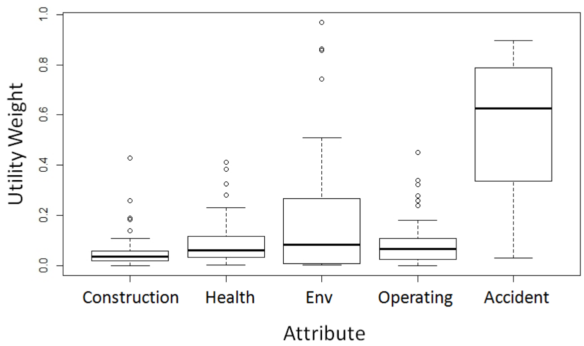

Table 6 An example of the normalised utility function for subject 14 (five decimal places)

Attribute	*Construction*	*Health*	*Environment*	*Operating*	*Accidents*
Participant 14	0.09625	0.05919	0.06978	0.11068	0.66410

Forty-six participants were included in the study where most participants were between 20 and 30 (31 total), with the other participants being spread out between 30 and 50. Geographically, 33 participants were based in the UK, 8 from the EU, and the remaining five participants were non-EU. This shows that the biggest differences tended to be between the importance of 'accidents' or 'environment'. Keep in mind that increasing the weight of one attribute necessarily decreases the weight of the other attributes due to the normalisation condition. Thus participants with extreme values in one attribute will have associated extreme weights for other attributes. In general it seems participants moved weight from 'accidents' to the attributes they felt were particularly important which accounts for the large variation of accident attribute weights. The mean and standard deviations for the attribute weights of all participants can be seen in Table 7.

Table 7 Summary statistics of the dataset

Attribute	*Construction*	*Health*	*Environment*	*Operating*	*Accidents*
Mean	0.05695	0.09422	0.18848	0.09631	0.56493
Standard deviation	0.07716	0.09668	0.25064	0.10008	0.26303

Note: The group utility function is assumed to be the mean here.

There are first some differences to our simulations that need to be considered. First, the number of participants is larger than any of the simulated studies but the scaling should compensate for this. Second, as the participants did not collaborate, the utility functions are not as consistent as we would expect. Thirdly, zero utility weights will be replaced by a small value (1×10^{-9}) to avoid numerical issues. Finally when considering sub-groups, we are unable to explore every possible sub-group due to the size of the power set (246), and while this is a more open problem, we only explored sub-groups of size 6 (roughly ten million sub-groups).

We also need to be clear about what the influence means in the given context. In this case the participants were not involved in group discussions, and so had no chance to influence the views of the other participants so all influence they wish to exert is contained in their utility function. Therefore we might expect 'potentially influential' participants to be one of two things; they may not understand the utility elicitation method as there was no human interaction for when participants may have had problems, or they may have an extremist view about a certain issue and so inflate their attribute weights for this issue. Using group utilities here also makes sense due to the significance of the issue, as any recommendation put forward for nuclear waste disposal should have a clear source.

5.2 Results

For the analysis of these results we be used the cosine divergence, with $\alpha = 1$, as it showed better performance than the other distance measures during our simulations. As there are five attributes, we considered an individual as potentially influential if their adapted Cook's distance is greater than $\delta = 0.12029$. Figure 7 shows why the participants were flagged in Figure 6; they are all far from the group's average. In particular all individuals that weight environment highly in Figure 5 were flagged as potentially influential. Also we can see that the accident weight is generally reduced when other attributes are weighted more heavily.

Figure 6 Scatter plot of the adapted Cook's distance values for all participants, with a δ line (red) (see online version for colours)

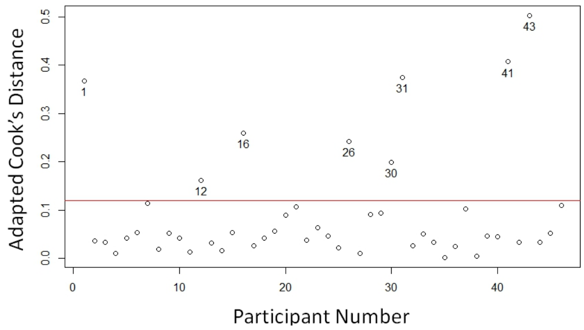

Figure 7 The utility weights for participants 1 (blue), 12 (red), 16 (green) and 26 (purple) in the top plot and 30 (blue), 31 (red), 41 (green) and 43 (purple) in the bottom plot with the group utility function (black) (see online version for colours)

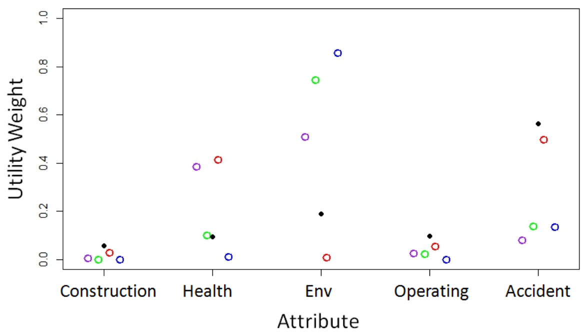

Figure 7 The utility weights for participants 1 (blue), 12 (red), 16 (green) and 26 (purple) in the top plot and 30 (blue), 31 (red), 41 (green) and 43 (purple) in the bottom plot with the group utility function (black) (continued) (see online version for colours)

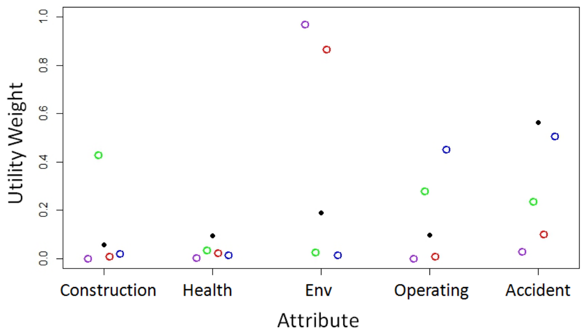

Table 8 Information about each of the participants that were flagged as potentially influential

Participant	*Age*	*Country of origin*	*Job title*	*Sex*
1	23	UK	Student	M
12	28	Belgium	Researcher	M
16	20	UK	Student	F
26	42	UK	PhD student	M
30	45	Australia	Assistant director	M
31	26	UK	PhD student	F
41	32	UK	Research fellow	M
43	28	Norway	Post-doctoral	M

Table 8 shows several expected patterns. For instance, the younger people demonstrated more concern for the environment, or the Assistant Director that weighting operating cost very highly. This gives us some confidence that the results fit the context. Unexpected values could also be useful, for example we could investigate participant 41 to see if they understood the utility elicitation process as they weighted construction costs very highly, when it was usually a very minor cost in the long term (even when participants weight immediate effects highly).

Table 9 Sub-groups of size 6 which have an adjusted R^2 value of over 0.998

Sub-group number																		
1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	9	9
9	9	9	9	9	9	9	9	9	9	9	9	15	15	15	15	18	15	18
15	15	15	15	15	15	15	18	18	18	18	19	18	18	18	18	19	18	19
18	18	18	18	18	19	19	19	19	22	39	39	19	19	19	39	39	19	39
19	19	22	39	39	39	39	39	39	39	40	40	22	39	39	40	40	39	40
39	40	39	40	42	40	42	40	42	40	42	42	39	40	42	42	42	40	42

Note: Each column is a reported sub-group, with the elements being the participant number.

While we identified the students with high environmental weights, we did not consider other students. To answer this question we can appeal to our subgroup identification method. We only considered sub-groups of size 6 due to computational limitations. Comparing the adjusted R^2 values to our tolerance of 0.995, we found 1,215 reported sub-groups. This number is still large (which is a by-product of having a 46 participant group), however for now we consider those with particularly large adjusted R^2 values (over 0.998). This then the returns 19 sub-groups given in Table 9.

From Table 9 we see the same participants are being included in many of these groups. In particular, participants 6, 9, 15, 18, 19, 39 and 40 all appear in at least 12 of these sub-groups, so if we look at this specific sub-group we see an adjusted R^2 value of 0.998483, which is also very high. Comparing the elicited utilities from these participants we can see all of them weighted accidents very highly (between 0.83 to 0.9), and so we have been successful in identifying a group that shares similar beliefs. We can also see a very strong trend demographically with these participants, as they were all students between the ages of 20 and 29. This shows that between applying our influence measure, and a simple application of our sub-group identification, we have been able to identify two clearly distinct groups within the younger participants from just 11 of the participants.

6 Discussion and extensions

6.1 Diagnostics performance and usage

In Sections 4 and 5 we explored the behaviour of our proposed methods in both large-scale simulated studies and a case study concerning nuclear waste disposal. The data set that was used proved a good test for the robustness of our methods to group size, and the results were consistent with what we may expect. Our simulated results for influence detection were promising and quite useful in selection of the most suitable distance measure. The cosine divergence was chosen due to its empirical performance, giving the lowest false positive rate (2.7%) coupled with just as strong a true positive rate (93.4%) as any of the other distance measures.

When applying the influence measure to our case study, we identified some interesting behaviour showing the influence measure works in larger scenarios. If our influence measure had been applied at the time of the study, the online elicitation process could have been improved as there was some evidence that at least one participant did not understand the elicitation. Also expected behaviours were identified, such as the younger participants generally giving more extreme environment attribute weights. This could have been politically driven or influenced by parties such as Greenpeace. The influence measure could have been used to flag these participants for a follow-up study. It was also interesting to see that participants would usually increase one weight at the expense of just one other weight (usually accidents).

6.1.1 Scenario 1: boardroom style meetings

This is the scenario where our proposed diagnostics would likely have the least impact. This is due to individuals being more likely to influence the group by changing the beliefs of the other group members. Therefore an individual's influence can more easily be divided amongst the whole group's individual utility functions. We would expect our influence measure to be less useful here due to a lower true-positive rate, although it can be used to detect group members that misreport their beliefs. The sub-group identification could be more useful as it can check specific sub-groups the group may be worried about while also detecting sub-groups that remain undetected. Due to the difficulty of influencing a group decision individually, it is likely group members would form sub-groups to increase their influence over the group.

6.1.2 Scenario 2: chat room meetings

In this scenario we would expect group members to exaggerate some of their beliefs (particularly when they can remain anonymous), which our influence measure can detect so that the truth of their reported utilities can be investigated. Also group members will generally be less malleable to other group members because of less trust in this setting. However, group members may be more open to working with other group members to manipulate the outcome due to the ease of hiding these discussions from other group members. Our sub-group identification diagnostic should allow us to detect a selection of subgroups that could be working together to allow the group to use this information to re-assess the truth of the group utilities.

6.1.3 Scenario 3: non-discussion decision making

In this scenario we have no discussion between group members, meaning individuals influence the group utilities only through their own beliefs. We would expect exaggeration of group member's beliefs to be more extreme than in scenario 2 as people may feel they need to 'pull' the group towards their beliefs. Our influence measure can detect these extreme views (as we saw in the case study), so that follow up studies or slight adaptations to the decision process could be made. Particularly high influence values could also correspond to individuals not understanding the system of utility elicitation well (which they likely would have had more help with in the other scenarios), and so it could be used as a diagnostic to test for how good the implemented elicitation process is.

The sub-group identification method can be useful for detecting group members with similar backgrounds or affiliations (for example political interests on environment like in our case study), and so could be useful for helping to identify demographics that tend to share the same beliefs. This could be useful in deciding how representative the sample was of the general beliefs of the target population of the decision (in our case study, the sample is not very representative of the general UK population due to the majority of participants being students in their 20s).

6.2 Extensions

We considered three main extensions to the methods developed in this paper. First we can explore adaptations to the one-out restriction that the current method uses, as real scenarios could have complex interactions between group members. We may be able to draw on more links to regression, for example stepwise forward selection in variable selection (see Harrell, 2001). Using this it may be possible to select the 'most important' group members to the decision. Exploration can also be done into effects of removing multiple group members, and whether the order and time of removals makes a significant impact on the group utility function. This could also help us understand how 'missing data' affects the process.

We could also use the developed diagnostic tests to attempt to identify deceptive individuals. This is a difficult problem as we need to use behavioural observations alongside the influence values for each individual. For example we may expect that group members with much influence over the group utility function that interacts more than normal with the group could be a deceiver. This would be useful extension for scenarios 1 and 2. We should also keep in mind that false positives in this case could be far more common than for the influence measure, and so confidence in the diagnostic tools would drop substantially.

Finally we could turn our focus to the detection of unwanted influence in a scenario where we expect group members to influence other member's beliefs instead of misreporting their own. We use the term 'unwanted' influence here because is both expected and accepted by the group. The problem comes when group members do not realise they have been influenced. We may try to elicit each group member's utility function before any discussions to give an estimate of their 'prior' utility function. We may then need to draw on some more behavioural observations to estimate how malleable each individual is, although group hierarchy could also be used. This could then be used alongside the utility function elicited after the meeting (posterior) to lessen the effect of the meeting. However keep in mind there is a large difference between an individual's initial beliefs being corrected by the group and being influenced by a member within the group, despite both these cases giving similar outcomes.

7 Conclusions

Both the influence measure and sub-group identification methods developed show some promise, although work still needs to be done to allow them to be more tailored to specific scenarios. The influence measure can be particularly useful for identification of issues with an elicitation process in a large scale study, and also in identifying individuals with more extreme beliefs. This could clearly be seen when exploring the data for the choice of nuclear waste disposal, where more extreme beliefs are identified with ease, and also where an example subgroup was identified with very similar beliefs, even when we could not explore all possible sub-groups due to the scale of the problem. In smaller scenarios, the influence measure needs to be adapted to account for group members attempting to influence the decision through other member's utilities. We hope that once some of these suggested improvements have been implemented, that the proposed diagnostics will be a strong addition to a group's arsenal of diagnostic tests.

References

Arrow, K.J. (1950) 'A difficulty in the concept of social welfare', *The Journal of Political Economy*, Vol. 58, No. 4, pp.328–346.

Atherton, E.C. (1999) *Valuing the Future: Issues in Modelling Intertemporal Decisions*, PhD thesis, University of Manchester.

Atherton, E.C. and French, S. (1998) 'Valuing the future: a MADA example involving nuclear waste storage', *Journal of Multi-Criteria Decision Analysis*, Vol. 7, No. 6, pp.304–321.

Chung, J.K., Kannappan, P.L., Ng, C.T. and Sahoo, P.K. (1989) 'Measures of distance between probability distributions', *Journal of Mathematical Analysis and Applications*, Vol. 138, No. 1, pp.280–292.

Cook, R.D. and Weisberg, S. (1982) *Residuals and Influence in Regression*, Chapman and Hall, New York.

DeSanctis, G. and Gallupe, R.B. (1987) 'A foundation for the study of group decision support systems', *Management Science*, May, Vol. 33, No. 5, pp.589–609.

Gibbard, A. (1973) 'Manipulation of voting schemes: a general result', *Econometrica*, Vol. 41, No. 4, pp.587–601.

Harrell, F.E. (2001) *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, Springer-Verlag, New York.

Keeney, R.L. (2002) 'Common mistakes in making value trade-offs', *Operations Research*, Vol. 50, No. 6, pp.935–945.

Keeney, R.L. and Raiffa, H. (1976) *Decisions with Multiple Objectives*, John Wiley and Sons, New York.

Keeney, R.L. and Winterfeldt, D. (2009) *Practical Value Models*, Published Articles and Papers, Paper 36.

Lashner, B.A. (2006) 'Sensitivity-specificity trade-off for capsule endoscopy in IBD: is it worth it?', *The American Journal of Gastroenterology*, Vol. 101, No. 5, pp.965–966.

Magee, L. (1990) 'R^2 measures based on Wald and likelihood ratio joint significance tests', *The American Statistician*, Vol. 44, No. 3, pp.250–253.

Marett, K. and George, J.F. (2013) 'Barriers to deceiving other group members in virtual settings', *Group Decisions and Negotiations*, Vol. 22, No. 1, pp.89–115.

Neumann, J.V. and Morgenstern, O. (1953) *Theory of Games and Economic Behavior*, 3rd ed., Princeton University Press, Princeton.

Nguyen, H.T., Kosheleva, O. and Kreinovich, V. (2009) 'Decision making beyond Arrow's impossibility theorem, with the analysis of effects of collusion and mutual attraction', *International Journal of Intelligent Systems*, Vol. 24, No. 1, pp.27–47.

Phillips, D., Egan, M. and Airoldi, M. (Eds.) (2006) *CoRWM: MCDA Decision Conference*, Catalyze Limited.

Procaccia, A.D. (2010) 'Can approximation circumvent Gibbard-Satterthwaite?', *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*, pp.836–841.

Satterthwaite, M. (1975) 'Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions', *Journal of Economic Theory*, Vol. 10, pp.187–217.

Sen, A. and Srivastava, M. (2011) *Regression Analysis – Theory, Methods and Applications*, 4th ed., Springer-Verlag, New York.

Proofs

1 *Theorem 1*

First consider the permutability of any two individuals i_1 and i_2 and assume that $\mathbf{u}^*$ is another set of utilities where $\mathbf{u}_{i_1} = \mathbf{u}_{i_2^*}$, $\mathbf{u}_{i_2} - \mathbf{u}_{i_1^*}$ and $\mathbf{u}_i = \mathbf{u}_{i^*}$ for $i = 1,\ldots,N$ with $i \neq = i_1$, i_2.

From the permutability of i_1 and i_2 it should not matter where one of the individuals are within the group before being removed, so we have that;

$$D\left(\mathbf{u}_G, \mathbf{u}_{G-i_1}\right) = D\left(\mathbf{u}_{G^*}, \mathbf{u}_{G-i_1^*}\right) \tag{1.1}$$

where $\mathbf{u}_G = \sum_{i=1}^{N} w_i \mathbf{u}_i$ and $\mathbf{u}_{G-1_1} = \sum_{i=1,\, i\neq i_1}^{N} \tilde{w}_i \mathbf{u}_i$ with w_i and $\tilde{w}_i$ being the weight and adjusted weight (different group size) assigned to individual i respectively.

1.2 *Case 1: equal weights*

Assume $w_i = N^{-1}$ and $\tilde{w}_i = (N-1)^{-1}$. Then from (1.1) we require;

$$\begin{aligned} &D\left(N^{-1}\sum_{i=1}^{N} \mathbf{u}_i, (N-1)^{-1}\sum_{i=1, i\neq i_1}^{N} \mathbf{u}_i\right) \\ &\quad = D\left(N^{-1}\sum_{i=1}^{N} \mathbf{u}_{i^*}, (N-1)^{-1}\sum_{i=1, i\neq i_2}^{N} \mathbf{u}_{i^*}\right) \end{aligned} \tag{1.2}$$

Considering each summand separately we know from our definition of $\mathbf{u}^*$ that $\mathbf{u}_i = \mathbf{u}_{i^*}$ holds for $i \neq = i_1$, i_2, so for (1.2) to hold we need that $\mathbf{u}_{i_2} = \mathbf{u}_{i_1^*}$ which is true from the definition of *.

1.2 *Case 2: unequal weights*

We proceed through proof by contradiction. Assume that $w_{i_1} \neq w_{i_2}$. By the permutability of i_1 and i_2 we know that;

$$D\left(\sum_{i=1}^{N} w_i \mathbf{u}_i, \sum_{i=1,\, i\neq i_1}^{N} \tilde{w}_i \mathbf{u}_i\right) = D\left(\sum_{i=1}^{N} w_i \mathbf{u}_{i^*}, \sum_{i=1,\, i\neq i_2}^{N} \tilde{w}_i \mathbf{u}_{i^*}\right) \tag{1.3}$$

Similarly to in (1.2) we know that each summand for $i \neq i_1$, i_2 must be equal, and so we only need to show that $w_{i_2}\mathbf{u}_{i_2} = w_{i_1}\mathbf{u}_{i_1^*}$. We know that $\mathbf{u}_{i_2} = \mathbf{u}_{i_1^*}$. from the definition of $\mathbf{u}^*$, however our assumption was that $w_{i_1} \neq w_{i_2}$ and so we arrive at a contradiction. Therefore we cannot have permutability with unequal weights.

2 *Theorem 2*

Assume for two individuals i_1 and i_2 we have that $w_{i_2} > w_{i_1}$ and $\mathbf{u}_{i_1} = \mathbf{u}_{i_2}$. We want to show that;

$$C_{i_2}(\mathbf{u}, G) > C_{i_1}(\mathbf{u}, G) \tag{2.1}$$

Due to the construction of $C_i(., .)$ we can remove any normalisation terms and the problem is simplified to showing;

$$D\left(\mathbf{u}_G, \mathbf{u}_{G-i_2}\right) > D\left(\mathbf{u}_G, \mathbf{u}_{G-i_1}\right) \tag{2.2}$$

As $D(., .)$ is our distance measure and is increasing in divergent arguments, so it suffices to show that $\mathbf{u}_{G-i_2}$ has moved further than $\mathbf{u}_{G-i1}$. Note that we mean moves further in all attributes here, as both i_1 and i_2 have the same utility functions and that $\tilde{w}_i = w_i \frac{N}{N-1}$ denotes the weight adjusted for group size.

$$\mathbf{u}_{G-i_1} = \sum_{i=1,\, i\neq i_1}^{N} \tilde{w}_i \mathbf{u}_i = \sum_{i=1,\, i\neq i_1, i_2}^{N} \tilde{w}_i \mathbf{u}_i + \tilde{w}_{i_2} \mathbf{u}_{i_2}$$

$$\mathbf{u}_{G-i_2} = \sum_{i=1,\, i\neq i_2}^{N} \tilde{w}_i \mathbf{u}_i = \sum_{i=1,\, i\neq i_1, i_2}^{N} \tilde{w}_i \mathbf{u}_i + \tilde{w}_{i_1} \mathbf{u}_{i_1}$$

We know that $\tilde{w}_{i_2} u_{i_2} > \tilde{w}_{i_1} u_{i_1}$ from our assumptions that $w_{i_2} > w_{i_1}$ and $\mathbf{u}_{i_1} = \mathbf{u}_{i_2}$ implying that $\mathbf{u}_{G-i_2}$ has moved further from the group utility $\mathbf{u}_G$ than $\mathbf{u}_{G-i1}$, and so necessarily we have that (2.2) holds which implies (2.1) holds.

3 *Theorem 3*

3.1 *Case 1*

Suppose $\mathbf{u}_1 = (u_1^{(1)}, u_1^{(2)}, \ldots, u_1^{(K)})$ with $u_i^{(j)} > 0$ for $j = 1,\ldots,K$ and that $u_G^{(2)} < u_G^{(l)}$ for $l = 3,\ldots,K$. We wish to influence individual 1's utility for attribute 1 at the expense of attribute 2. Set $\mathbf{u}_{1,1*} = (u_1^{(1)} + L, u_1^{(2)} - L, \ldots, u_1^{(K)})$, where $L < \min(1 - u_1^{(1)}, u_1^{(2)})$. Then we have the group utility functions given by;

$$\mathbf{u}_G = \sum N^{-1} u_i$$

$$\mathbf{u}_{G,1*} = \sum_{i=1}^{N} N^{-1} u_i^{(1)} + N^{-1} u_{1,1*}.$$

Clearly the group utilities for attributes $3,\ldots,K$ will remain the same, so we only need to consider attributes 1 and 2. For attribute 1 we have that;

$$u_{G,1*}^{(1)} = \sum_{i=2}^{N} N^{-1} u_i^{(1)} + N^{-1}\left(u_i^{(1)} + L\right).$$

From this expression we can see that $u_{G,1*}^{(1)} = u_G^{(1)} + \frac{L}{N}$ and in a similar way we can derive that $u_{G,1*}^{(2)} = u_G^{(2)} - \frac{L}{N}$. Now consider the impact this has on the KL-divergences when comparing the two group utilities.

$$D\left(\mathbf{u}_G,\mathbf{u}_{G,1*}\right)=u_G^{(1)}\log\left(\frac{u_G^{(1)}}{u_G^{(1)}+\frac{L}{N}}\right)+u_G^{(2)}\log\left(\frac{u_G^{(2)}}{u_G^{(2)}-\frac{L}{N}}\right)$$

We now take the derivative of this with respect to L to give us a basis to compare against case 2;

$$\frac{dD\left(\mathbf{u}_G,\mathbf{u}_{G,1*}\right)}{dL}=-\frac{u_G^{(1)}}{\left(Nu_G^{(1)}+L\right)}+\frac{u_G^{(2)}}{\left(Nu_G^{(2)}-L\right)} \tag{3.1}$$

It is worth noting that we can see that the distance measure is increasing in L as we might expect.

3.2 Case 2

Now we repeat these same steps for Case 2, where influence is introduced such that $\mathbf{u}_{1,2*}=\left(u_1^{(1)}+L,u_1^{(2)}-\frac{L}{K-1},\ldots,u_1^{(K)}-\frac{L}{K-1}\right)$, where L is defined similarly as before. Redoing the same calculations we get that $u_{G,2*}^{(1)}=u_G^{(1)}+\frac{L}{N}$ as before and $u_{G,2*}^{(j)}=u_G^{(j)}-\frac{L}{N(K-1)}$ for j = 2,3,…,K. If we then consider the impact this has on the KL-divergence, we have;

$$D\left(\mathbf{u}_G,\mathbf{u}_{G,2*}\right)=u_G^{(1)}\log\left(\frac{u_G^{(1)}}{u_G^{(1)}+\frac{L}{N}}\right)+\sum\nolimits_{j=2}^{K}u_G^{(j)}\log\left(\frac{u_G^{(j)}}{u_G^{(j)}-\frac{L}{N(K-1)}}\right)$$

$$\frac{dD\left(\mathbf{u}_G,\mathbf{u}_{G,2*}\right)}{dL}=-\frac{u_G^{(1)}}{\left(Nu_G^{(1)}+L\right)}+\sum\nolimits_{j=2}^{K}\frac{u_G^{(j)}}{\left(N(K-1)u_G^{(j)}-L\right)}$$

3.3 Comparison of the cases

The first point we note when comparing equations (3.1) and (3.2) is that first term (relating to attribute 1) is the same for both cases, and so we can ignore this term for now for the purposes of comparison of the two expressions. We aim to show that;

$$\frac{u_G^{(2)}}{\left(Nu_G^{(2)}-L\right)}>\sum\nolimits_{j=2}^{K}\frac{u_G^{(j)}}{\left(N(K-1)u_G^{(j)}-L\right)}$$

Let $u_G^{-}=\min(u_G^{(2)},u_G^{(3)},\ldots,u_G^{(K)}=u_G^{(2)})$. Then when comparing each element of the summand in (3.2) to $u_G^{(-)}$, we know that;

$$\frac{u_G^{(j)}}{\left(N(K-1)u_G^{(j)}-L\right)} \le \frac{u_G^{(-)}}{\left(N(K-1)u_G^{(-)}-L\right)} \text{ for all } j = 2,3,\ldots,K.$$

As this holds for all elements of the summation, we have that;

$$\sum_{j=2}^{K} \frac{u_G^{(j)}}{\left(N(K-1)u_G^{(j)}-L\right)} \le \sum_{j=2}^{K} \frac{u_G^{-}}{\left(N(K-1)u_G^{-}-L\right)} = \frac{(K-1)u_G^{(-)}}{\left(N(K-1)u_G^{(1)}-L\right)}$$

$$< \frac{(K-1)u_G^{(-)}}{\left(N(K-1)u_G^{(1)}-(K-1)L\right)} = \frac{u_G^{(2)}}{\left(Nu_G^{(2)}-L\right)}$$