Reliability Updating in Decision Making Problems -

PhD Transfer Report.

Donnacha Bolger

(under the supervision of Dr. Brett Houlding)

October 2013.

Abstract

This report attempts to provide a framework by which individuals who must make decisions may incorporate the opinions and beliefs of other relevant participants into their decision making process. A variety of novel techniques by which this can be done are presented, as well as attempts to justify their usage, and compare their respective performances. We also consider a more complex framework, in which a collection of individuals must determine which decision out of a collection of alternatives is deemed to be optimal for the group as a whole, by combining their collective beliefs and utility functions. A constant and consistent terminology has been used throughout this report, with a glossary of the relevant notation included in Appendix A.

Contents

1 Introduction

Decision making is an act that is carried out by millions of individuals (and groups) all around the world on a daily basis. The consequences of these decisions will greatly vary in their severity - clearly a decision pertaining to whether or not to invade a foreign territory is far more important then one pertaining to which brand of milk to buy. Nevertheless, in both cases the question may be asked of "how sound is the decision making process being implemented"? It appears that decisions, both of a crucial and of a non-crucial nature, are often made in an *ad hoc* and, as we shall return to in more depth later, inconsistent manner, with no consideration of the introduction of a formal framework for this procedure. Yet attempts to derive such a framework, whereby users can choose a course of action that is deemed to be in some sense optimal to them, have been ongoing for centuries.

Section 2 of this report provides an extensive overview of the subject of decision theory, beginning from its most fundamental components, and building upon these to consider increasingly complicated situations. Initially we address problems of a single decision maker (DM) before extending to the setting of group problems, and the consequences of Arrow's Impossibility Theorem. This section concludes with a brief review of recent developments in the area of combining expert judgements, which will be a crucial stepping stone in what follows.

The original contributions of this report begin in Section 3, in which several novel approaches are developed for dealing with a specific type of decision problem, which we believe little effort has been made to consider before. We focus primarily on subjective reliability updating, and how a DM may assess the accuracy of information received from various diverse sources. An illustrated example demonstrates the strengths and weaknesses of these approaches, and compares their performance against the pre-existing methodology. Section 4 attempts to extend the work of the previous section to a group setting, and while this is still a work in progress, we belief that some interesting theory has been developed. Lastly we conclude in Section 5 with a discussion of some of the many further areas for research that are open to us, extending our own developments into new and challenging directions.

2 Review

The two pillars on which modern day statistical decision making are built are the concepts of probability and utility, to which we shall give substantial discussion to in what follows. Continuing on from the introduction of these two topics (as well as comments on the interesting relationship between them) we will discuss the idea of maximising expected utility, as well as the objections raised and alteratives proposed to this approach.

2.1 Probability

Probability is a concept then is often misinterpreted, and commonly misdefined. Below we will give a rigourous, and indeed axiomatic, treatment of what a valid probability actually is. First and foremost, probability is a method which is used to quantify uncertainty about some unknown event of interest. In a decision making framework we can see that uncertainty is something that will frequently be faced by a DM who is unsure of the exact consequences that will occur as a result of the decisions she makes. Probability is used to provide some description to the unsureness of a DM about the environment which she inhabits. Kolmogorov (1950) provided a strict axiomatic framework for its use, which is still applied today. The three axioms of Kolmogorov are given below, where Θ denotes the event space of possible events.

- If $\theta \in \Theta$ then $P(\theta) \in \mathbb{R}$ and $P(\theta) \geq 0$, *i.e.*, for all $\theta \in \Theta$ there is a non-negative probability of θ occuring.
- There is some universal event $\Theta^* \subseteq \Theta$ such that $P(\Theta^*) = 1$.
- If $\theta_1, \theta_2, \ldots, \in \Theta$ are a collection of countable mutually exclusive events (by which we mean

that no two events can occur simultaneously) then

$$P(\theta_1 \cup \theta_2 \cup \ldots) = \sum_{i=1}^{\infty} P(\theta_i)$$

By countable we mean that the set is either of finite size or it is countably infinite (see Cantor, 1874 for more details).

From these rudimentary axioms some other important properties of probability can be derived, for instance:

- Monotonicity: If $\Theta_1 \subseteq \Theta_2$ for $\Theta_1, \Theta_2 \subseteq \Theta$ then $P(\Theta_1) \leq P(\Theta_2)$.
- Null set: $P(\varnothing) = 0$.
- Bounds of Probability: For all $\theta \in \Theta$, $0 \leq P(\theta) \leq 1$, *i.e.*, a probability must lie in the closed $[0, 1]$ interval.

Any collection of probabilities failing to meet the three Kolmogorov axioms risks falling prey to what is commonly called a Dutch book. This is said to occur when one enters into a wager which they are doomed to lose irrespective of what outcome occurs. In what follows we will assume probabilities that meet the requirements discussed above.

2.2 Utility

2.2.1 Expected Value Theory

Early studies of probabilility were commonly concerned with gambling and games of chance, and were used to determine how likely a player was to win a particular game, and what a fair stake for them to pay to play this game was. Prior to the development of more sophisticated techniques, a fair price was often determined to be the expected value of the outcome of playing the game, *i.e.*, a sum of products of the probability of a given outcome occuring and the return that would be yielded from the occurence of this outcome. However this method has some shortcomings, famously highlighted by Bernoulli (1738) in the St. Petersburg Paradox. This describes a game of a chance in which a player tosses a fair two-sided coin, and wins a pot (which doubles with each success) for every consecutive toss resulting in the uppermost face being a head, which they recieve the first time they toss and observe a tail. The initial pot is \$2. For instance if the player begins with four consecutive heads they would have a current prize of \$16 (*i.e.*, $2 \times 2 \times 2 \times 2$). However if the fifth toss revealed a tail then they would exit the game with a final prize of \$16.

What is the fair price for someone to pay to play this game? Naively at first we consider the method of equating the fair price to the average prize to be attained from the playing of the game, *i.e.*, its expected value E. However when we calculate this we find

$$\begin{aligned} E &= \frac{1}{2} \times 2 + \left(\frac{1}{2}\right)^2 \times 2^2 + \left(\frac{1}{2}\right)^3 \times 2^3 + \left(\frac{1}{2}\right)^4 \times 2^4 + \dots \\ &= \frac{1}{2} \times 2 + \frac{1}{4} \times 4 + \frac{1}{8} \times 8 + \frac{1}{16} \times 16 + \dots \\ &= \sum_{k=1}^{\infty} 1 = \infty. \end{aligned}$$

This divergent sum is clearly a ridiculous choice for the fair price for a player to partake in the game, as it indicates that the expected prize that will be obtained from playing it is an infinite sum. Hence we can see that use of expected value theory is not always a logical criteria with which to make decisions - a DM doing so here would choose to play the game as it claims to yield an infinite amount of money. Bernoulli used this example as a motivation to develop a new (and somewhat deeper) decision making critera, that of utility.

2.2.2 Utility Hypothesis

Bernoulli (1738) commented, with respect to the class of problem outlined above, that

"*..the determination of the value of an item must not be based on the price, but rather on the utility it yields. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount.*"

He developed the idea that individuals should specify their own utility function u, which described their personal attitudes towards potential outcomes, and incorporated their own outlook on risk. Formally a utility function u is a mathematical function, $u : \mathcal{R} \rightarrow \mathbb{R}$, from the set of possible returns $\mathcal{R}$ from a decision, to the set $\mathbb{R}$ of the real numbers. Hence for every possible decision a single numerical value could be calculated. The optimal decision under the method of utility was that which maximises expected utility (which we shall formally define in

Section 2.4.1), *i.e.*, the decision which returned the highest numerical value under the utility function specified.

It is clear that there are many advantages to this method, in comparision with expected value theory. Utility Theory allows for a personalistic interpretation of the merits of various outcomes, and will generally yield different numerical values for different individuals, depending on their beliefs and monetary situation. Hence it can be argued that it is a more useful criteria then expected value theory, as it incorporates much more information into the decision making process.

2.2.3 Measures of Risk-Aversion

A DM has a utility function, $u(r)$, which measures the satisfaction that they derive from some return $r \in \mathcal{R}$. The utility function of a DM should be a reflection of their own feelings towards taking a gamble. It tells whether they are risk-averse, risk-neutral or risk-prone. However it may not be obvious which of these three categories an individual's utility function falls into. Fortunately there are a collection of measures that are used to determine which classification the utility function falls under, and also to what extent it does so. One such commonly used measure, from Arrow (1965) and Pratt (1964) is the Arrow-Pratt absolute risk-aversion (ARA), which involves calculation of the coefficient of absolute risk aversion, given by

$$A(r) = -\frac{u''(r)}{u'(r)}.$$

This method contains the assumption that the utility function $u(r)$ is twice differentiable. We do not discuss this any further, but merely note that utility functions of the form $u(r) = 1 - e^{-\alpha r}$ lead to $A(r) = \alpha$, which is called Constant Absolute Risk Aversion (CARA - see Rabin, 2000), and those which lead to $A(r) = \frac{1}{ar+b}$ (for $a, b \in \mathbb{R}$) are said to exhibit Hyperbolic Absolute Risk Aversion (HARA). Merton (1971) contains several examples of this.

2.3 Relationship between Utility and Probability

We briefly discuss the relationship that exists between the concepts of utility and probability that we alluded to previously. Each DM has her own subjective utility function $u(r)$ over returns $r \in \mathcal{R}$. While in practise elictation of this function can be a difficult task, theoretically it is doable due to the "twinned" relationship of (subjective) probability and utility, discussed in French (1994), whereby it is impossible to define one without the other. The following two formal definitions of the two concepts illustrate this.

- **Definition:** An individual's *subjective probability* for the occurence of an event is the amount p that they are willing to gamble such that they receive a return of 1 util if the event occurs, and 0 if it does not.

- **Definition**: The *utility* that an individual assigns to an outcome r can be thought of as the value p such that they are indifferent between

  (a) r for certain and

  (b) a gamble between the best possible outcome r^* such that $u(r^*)$=1 with probability p, and the worst possible outcome r_* such that $u(r_*)$=0 with probability $1 - p$.

Note that in the definition of utility above, the utility values are rescaled to the $[0,1]$ interval, and we shall discuss why this rescaling is possible in Section 2.4.2. A circularity can be noted between the two definitions, and indeed it seems justifiable for French (1994) to refer to utility as "probability's younger twin". This link between the fundamental definitions of these core decision making concepts is certainly noteworthy, and adds an interesting dimension to the idea of consistency for a DM trying to specify either her probability of an outcome occuring, or the utility that she would derive from its occurence.

2.4 Expected Utility

2.4.1 Decisions

Having introduced the concepts of probability and utility we now consider how these may be combined in order to constitute a decision making procedure. We denote a collection of n potential decisions by $d_1, d_2, \ldots, d_n \in \mathcal{D}$, where $\mathcal{D}$ is the set of admissible decisions, and each d_i is a distinctive potential course of action. Clearly the satisfaction that it is derived from some decision is based upon some state of nature that occurs in conjecture with it, and in a non-trivial case there is uncertainy over which state of nature will obtain. For example, the merits of a trip to a restaurant are dependent upon whether the chef is a good cook or not. One of these two states of nature is true, but the DM is uncertain which is true prior to making the decision to eat at the restaurant. Hence she must make her decision in light of this uncertainty.

We consider the case where there are m potential states of nature, $\theta_1, \theta_2, \ldots, \theta_m \in \Theta$, and a DM is uncertain which state will obtain. Note that Θ is the set of all potential states of nature that may transpire. The merits of a decision are dependent upon the state of nature that occurs in conjunction with it. A DM now must state her utilities over all possible outcomes, as is shown in Table 1 below, where $u(d_i, \theta_j)$ denotes the utility which obtains from making decision d_i and the occurence of state of nature θ_j.

Table 1: Cross-tabulation of $\mathcal{D}$ and Θ,

	θ_1	θ_2	$\ldots$	θ_m
d_1	$u(d_1, \theta_1)$	$u(d_1, \theta_2)$	$\ldots$	$u(d_1, \theta_m)$
d_2	$u(d_2, \theta_1)$	$u(d_2, \theta_2)$	$\ldots$	$u(d_2, \theta_m)$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
d_n	$u(d_n, \theta_1)$	$u(d_n, \theta_2)$	$\ldots$	$u(d_n, \theta_m)$

Now to calculate the expected utility of a decision d_i the DM needs only to sum over the products of the probability of each individual return occuring as a result of that decision, and the utility value associated with the return in conjunction with that decision.

$$\begin{aligned} E[u(d_i)] &= u(d_i, \theta_1)P(\theta_1) + u(d_i, \theta_2)P(\theta_2) + \ldots + u(d_i, \theta_m)P(\theta_m). \\ &= \sum_{j=1}^{m} u(d_i, \theta_j)P(\theta_j). \end{aligned} \tag{1}$$

Note that the sum of Equation (1) will be replaced by an integral in the cases of returns and/or probability distributions over returns being continuous rather then discrete. The optimal decision, d^* is deemed to be the decision d_i which maximises Equation (1). Formally we write

$$d^* = \arg\max_i E[u(d_i)]. \tag{2}$$

Lindley (1991) gives a fully comprehensive treatment of this approach. We make two further comments on the material presented in this section. First, we referred to $\mathcal{D}$ as the set of *admissible* decisions. When is a decision not an element of this set? We say that a decision is inadmissible if it is dominated by all other decisions in the set, *i.e.*, it is not preferred under any state of nature. Secondly, in the framework described above, and in the material that follows, it is assumed that a DM has both the ability to probabilistically quantify her uncertainty over the returns resulting from a decision and also the ability to specify an exact utility value to each potential return. Realistically this may often not be the case. In the absence of the abilities above there exist elictation methods which endevour to help discover these unknowns, see for instance Chajewska *et al.* (2000) for utility elictation and O'Hagan (1998) for belief elication. This is a large and ongoing area of research, which we will not consider any further in this report.

2.4.2 The axiomatisation of von Neumann and Morgenstern

The work of von Neumann & Morgenstern (1944) further developed utility theory when they gave an axiomatic justification for its use. They created four rational axioms, and assuming that a DM concurred with these, and where themselves rational, then they would make their decisions by the method of maximisation of expected utility. The four axioms given by von Neumann and Morgenstern where:

- *Completeness*: For any two possible decisions $d_1, d_2 \in \mathcal{D}$, a DM with well defined preferences can always compare the two, and determine which if either, they prefer, *i.e.* either $d_1 \geq d_2$, or $d_2 \geq d_1$. Note that the relation $\geq$ is used to indicate preference, with $d_1 \geq d_2$ indicating the decision d_1 is at least as preferable as decision d_2. Indifference between the two decisions is indicated by $d_1 \sim d_2$.

- *Transitivity*: If $d_1 \geq d_2$ and $d_2 \geq d_3$ then $d_1 \geq d_3$, for all $d_1, d_2, d_3 \in \mathcal{D}$. This ensures consistency of decision making, and ensures that money-pump situations do not arise.

- *Continuity*: If $d_1 \geq d_2 \geq d_3$ then there exists some $p \in [0,1]$ such that $pd_1 + (1-p)d_3 \sim d_2$ for all $d_1, d_2, d_3 \in \mathcal{D}$, *i.e.*, there exists some probability such that the decision maker is indifferent between a gamble between the best and worst outcome (with probabilities p and $1-p$ respectively) and a guarantee of an intermediate outcome.

- *Independence*: If $d_1 \geq d_2$ and $p \in [0,1]$ then, for any alternative outcome d_3, we must have $pd_1 + (1-p)d_3 \geq pd_2 + (1-p)d_3$ for all $d_1, d_2, d_3 \in \mathcal{D}$. This states that preference is invariant to the introduction of indepedent alternatives.

What von Neumann and Morgernstern showed was that if a DM agreed with these four axioms, then there was a unique (up to positive linear transformation) utility function which satisfied the properties that:

- $u(d_1) \geq u(d_2)$ if and only if $d_1 \geq d_2$ for all $d_1, d_2 \in \mathcal{D}$.

- For all $d_1, d_2 \in \mathcal{D}$, and any $p \in [0,1]$, $u(pd_1 + (1-p)d_2) = pu(d_1) + (1-p)u(d_2)$.

Hence a formal axiomatic justification was provided for the method of maximisation of expected utility, which advocated it as a method to be used by any decision maker who was deemed to be rational (by their agreement with a set of seemingly desirable axioms). To expand on the invariance of utility functions up to positive linear transformations, a DM need not be concerned whether she uses utility function $u_1(d)$ or $u_2(d) = au_1(d) + b$ for $a, b \in \mathbb{R}$ and $a > 0$. Consequentially utility functions are often rescaled to the $[0,1]$ interval, with 0 indicating the utility of the worst possible outcome, and 1 indicating the utility of the best possible outcome. This is commonly done for ease of interpretation.

2.4.3 Sequential Problems

Briefly we note that the decision method considered above was concerned only with a DM having to make a single decision. Often in a realistic setting we have a DM who may need to make decisions for several future epochs in the present. That is to say, this decision-making is not myopic, where a DM only considers one step in the future, at which time they observe some result, and then consider one step in the future again. This is decision-making where a DM must for example decide in 2013 how much money she will need in 2014, 2015 and 2016 in 2013. A method exists by which these problems may be solved, by use of a tree-like diagram called a decision tree, in which the various final outcomes are displayed at the right hand side of the tree, with various potential decision streams represented by different paths to reach these outcomes. The tree grows from left to right, beginning with just a single branch on the left hand side, representing the first decision to be made. A methodology by which these problems may be solved involves the "roll-back" method, in which a DM begins from the right hand side of the tree and works backwards (although recent work in Alghalith (2012) discusses the possibility of the use of a "roll-forward"). We will return to discussion of sequential problems at a later stage. Lindley (1991) gives a treatment of how the method of maximising expected utility may be used to solve these problems. One problem with this decision tree approach is that it may suffer from "curse of dimensionality" issues. This means that when a DM is considering a large amount of decisions epochs, and/or there are a large collection of potential decisions to choose from at each epoch, the tree will rapidly become very complicated. Computation may take a long time, and intractability may occur depending on the distributional form of the probability distributions of the DM and that of their utility function also. A method proposed to decrease these computational problems is that of the polynomial utility class (Houlding, 2008) which uses the assumptions of polynomial utility functions and Normally distributed unknown quantites of interest.

2.5 Objections and Alternatives

The method described above is known as a normative method. This means that it provides a formal method by which a rational DM should choose a decision from a collection of potential alternative options. However a key aspect of this is that by nature human beings are prone to being irrational, and not always behaving in a manner consistent with that described above. While the method of Bernoulli, formally treated by von Neumann and Morgenstern, certainly has a lot of merit, it has some shortcomings in terms of modelling the way in which individuals actually make decisions in real life contexts, as opposed to how they theoretically should.

2.5.1 The Allais Paradox

The Allais paradox is an example of the irregularity of decision making commonly exhibited by individuals, designed by Allais (1953). It involoved questioning individuals on what decisions they would make in two seperate hypothetical monetary situations. The first question, A, asked DMs would they rather have:

- 1: \$1,000,000 with certainty, or
- 2: \$1,000,000 with probability 89%, \$5,000,000 with probability 10% and \$0 with probability 1%.

The second question, B, asked if the DM would rather have:

- 1: \$1,000,000 with probability 11% and \$0 with probability 89%, or
- 2: \$5,000,000 with probability 10% and \$0 with probability 90%.

Of those questioned, it was found that the most common pair of decisions was to choose A1 and B2. However these results were shown to be inconsistent with expected utility theory, under which a choice of A1 implied an automatic choice of B1, and a choice of A2 implied an automatic choice of B2. Irrespective of the utility function selected by an individual, expressing their risk preferences, it is impossible, under the framework developed by von Neumann and Morgenstern to choose both A1 and B2. This is presented by Allais as a counterexample to the most controversial of the four axioms laid down by von Neumann and Morgenstern, that of Indepedence, as the only difference between the choices A1 and A2, and B1 and B2 is a common increase in the probability of receiving \$0. This clearly violates the Indepedence axiom, which states that preference should be invariant to the introduction of independent alternatives.

2.5.2 Prospect Theory

Above we mentioned how the method of maximisation of expected utility was a normative approach which advocated how DMs, deemed rational by their agreement with some fundamental axioms, should make decisions. However as the Allais paradox demonstrates there are clearly some flaws with using this as the sole criterion for decision making, as individuals often violate it. An attempt to resolve this problem was the development of a descriptive method of decision making, which endevoured to describe the way in which individuals did make decisions, and to incorporate genuine human characteristics to the general theory. This is what was done by Kahneman & Tversky (1979) in Prospect Theory. To summarise briefly, each decision (or *prospect*) was first edited into a simplified form using a collection of operations, before being evaluated, a process which involved two arguments, weights and values. We do not discuss the details of this approach here, but it attempted to mirror characteristics exhibited by decision makers in realistic scenarios, such as how individuals feel the pain of a loss more severely then they feel the joy of an equivilant gain, and how individuals underweight outcomes which are probable compared to those which are guaranteed. This approach is arguably one of the most complete alternatives to maximising expected utility.

2.5.3 Other Alternatives

Work is ongoing on the search for various alternative ways for which individuals can make decisions. Removing the Independence axiom of von Neumann and Morgenstern is one approach, as well as adding additional axioms, such as that of monotonicity. Two other noteworthy methods are briefly discussed below.

Firstly we mention the method of imprecise probabilities (see for instance Walley (1991)). In this approach, rather then specify an exact probability p for the occurence of some uncertain event, the DM simply specifies a lower and upper bound for this probability, say (p_L, p_U). In the most extreme case where a DM wants to express total ignorance then the interval $(0, 1)$ may be used. Clearly this approach has an advantage in that it allows a DM to express their true opinions in an honest fashion, *i.e.*, they need not be forced into giving some exact probabilistic belief when really they are deeply uncertain about its accuracy. Instead they just need to give an interval, whose width is a measure of their degree of confidence in their predictive ability. However there is also a negative side in that computationally it is more complex, and may lead to indecision in its results, *e.g.*, consideration of a value near the bottom of the interval leads to selection of one decision, while consideration of a value near the top of the interval leads to selection of another.

Finally we mention fuzzy logic, developed by Zadeh (1965). Fuzzy set theory is an extension of the classical crisp theory of sets. In the latter framework there are two mutually exclusive potential states - either an item belongs to a set or it does not. The approach of fuzzy set theory weakens this statement, allowing items to potentially belong to several different sets at once, with weak and strong degrees of membership. Statements need not be strictly true or false, they are allowed to have a truth value, in the $[0, 1]$ interval. Again there are advantages to this approach in that DMs can express their uncertainty rather then give exact values, but computationally it has disadvantages.

2.6 Group Decision Making

The approaches outlined up to this point deal with a single DM making a decision, *i.e.*, they relate to an individual who must make a choice from a collection of alternatives, as opposed to a collective of two or more individuals who must jointly make such a choice. Group decision making is a fundamental part of the world we live in. On a small-scale town councils make decisions which pertain to the welfare of residents, while on a large-scale the UN must reach a resolution on decisions that affect individuals all over the world. Some normative method by which collectives could make decisions would be extremely useful, as it would certainly be advantageous if some approach was put forward which could assist groups to lead them towards a decision which, in some rational sense, was deemed to be optimal. However, as shall be outlined in what follows, under a certain set of axioms, this is an unobtainable goal. While an objective normative procedure exists for the case of an individual solving a (potentially sequential) decision problem, no such methodology has yet been found for the case of a group who face a decision making task. We begin our study of decision making with multiple participants by a look at what is in some sense a crippling blow to it, namely the famous Impossibility Theorem of Arrow (1950). Many efforts have been made to circumnavigate this somewhat damning theorem, in an attempt to find an acceptable methodology for group decision making, which can improve the quality of decisions reached in these cases. Following on from Arrow's Theorem we consider various attempts to provide a solid framework for decision making involving multiple participants, discussing the strengths and weaknesses of these approaches, and the relationships between them.

2.6.1 Arrow's Impossibility Theorem

Arrow (1950) considered the topic of preference ranking among a body of individuals, and whether there was a valid mapping from the collection of individual preference rankings to a collective preference ranking for the body of individuals, which could be considered to be "fair" in some sense. He considered a Social Welfare Function (SWF), which operated on a collection of individual preference rankings, which would obey certain basic properties (namely completeness and transitivity for every DM for every admissible decision), and would lead to a group preference ranking, which would obey the same properties. The four fundamental axioms put forward by Arrow, which he believed were essential to any SWF were:

- *Non-Dictatorship*: There is no individual in the collective such that their individual preference ranking automatically becomes the preference ranking of the group as a whole.

- *Pareto's Principle*: If there are two decisions such that one is prefered to the other by each individual decision maker, then this will be prefered to the other in the collective group preference.

- *Independence of Irrelevant Alternatives*: If some decisions are deleted from the collection of admissible possibilities, but there is no change between the ranking of remaining outcomes in each individuals' preference ranking, then there is no change in the group preference ranking either.

- *Unrestricted Domain*: The collective group preference ranking is defined for any possible collection of individual preference rankings, assuming they all obey the properties of completeness and transitivity.

These axioms all seem to be rational and desirable properties for a logical SWF to obey. However in his famous Impossibility Theorem, Arrow would show that for at least two individuals and at least three distinct decisions to be ranked, no such SWF satistfying all of the above conditions could be created. Hence any SWF must break at least one of the four axioms, and be deemed to be, at least in some sense, irrational. As group decision theory is such a fundamental keystone of human existence there have, of course, been countless efforts to "side-step" the Impossibility Theorem, by creating a function that is coherent, while violating (at least) one of the axioms. The Axiom of Independence is the one that is perhaps the easiest to disregard, as it makes assumptions that no preference between any two outcomes is stronger then any preference between any other two outcomes, that is to say that the ranking is strictly ordinal, with no consideration of the degree of preference involved.

There are several other interesting related impossibility results, which we mention here without delving into their contents, namely May's Theorem (May, 1952), the Liberal Paradox (Sen, 1970), the Gibbard-Satterthwaite Theorem (Gibbard (1973), or Satterthwaite (1975)), and the Duggan-Schwartz Theorem (Duggan & Schwartz, 1992). It is however still the consequences of Arrow's Impossibility Theorem which loom largest.

2.6.2 Utilitarianism and the SWF

Utilitarianism is a normative theory regarding how social choices should be made, namely that, in some sense, happiness or well-being should be maximised, and that suffering should be minimised. We have an individual who is "in charge", and must translate the personal rankings of individuals with respect to a collection of alternatives to one collective ranking of these alternatives. We have seen above that Arrow's Impossibility Theorem causes problems here. Nevertheless, by working from new axioms, a cohesive and legitimate group ranking may be achieved. Harsanyi (1955) created the following framework.

Suppose P^* is the person in change, who must translate the preferences of n individuals to a single preference ranking. Each of the n individuals assigns a utility to each of the potential options, scaled to be the closed $[0, 1]$ interval, *i.e.*, 0 corresponding to the outcome that they consider to be the absolute worst one possible, and 1 corresponding to the outcome that they consider to be the absolute best one possible. Hence we have utility functions $u_1, u_2, \ldots, u_n$ for each of the individuals respectively, which we want to translate to a single utility function, u^*. We have the following axioms:

- *Anonymity*: P^* doesn't know which individual put forward which preference ranking, ensuring no bias in the procedure.

- *Strong Pareto*: If each of the n individuals is indifferent between two outcomes, then so is P^*, *i.e.*, if $u_i(r_1) = u_i(r_2)$ for all $i = 1, 2, \ldots, n$ then $u^*(r_1) = u^*(r_2)$. Similarly if some individuals prefer r_1 to r_2, and the rest are indifferent between the two, then P^* will prefer r_1 to r_2.

With these axioms, Harsanyi demonstrated in his Utilitarian Theorem, that the only possible function, u^*, is one which is an additive weighted sum of the utility functions of the n individuals. Hence we have this as an option for a valid method of group decision making which does not violate any of the terms of Arrow's Impossibility Theorem.

It is worth noting that while the approach of Harsanyi, outlined above, was accepted by many, there were still detractors who felt it didn't merit widespread use. One of the primary figures in opposition to this choice of SWF was Buchanan (1954, 1979, 1994a, 1994b) and Buchanan & Tullock (1962), who had large reservations about the choice of function, as well as the concern that Arrow had overlooked the liberal value judgement of individualism in his theorem's formulation. Sen (1979, 1990, 1995) discusses the importance of a process not just being fair in its mechanism, but also the importance that the results it yields are fair in some social sense.

2.6.3 Fully Probabilistic Design

Karny & Guy (2004) consider a setting where each participant within a group attempts to improve her decision quality by sharing her beliefs with some neighbours (in what is termed the communication phase), and optionally augmenting her own beliefs in light of what she has learned (in what is termed the updating phase). A neighbour is simply another DM with whom one has a common area of interest with. Participants are entitled to be selfish, and to ignore the information shared with them by their neighbours, leaving their beliefs unaltered if they choose to. For example, consider a situation with two individuals, P_1 and P_2, with beliefs $f_1(\theta)$ and $f_2(\theta)$ respectively over some uncertain θ of interest. After the updating phase their distributions will be given by weighted sums,

$$\hat{f}_1(\theta) = \alpha_1 f_1(\theta) + (1 - \alpha_1) f_2(\theta) \tag{3}$$

$$\hat{f}_2(\theta) = (1 - \alpha_2) f_1(\theta) + \alpha_2 f_2(\theta) \tag{4}$$

Note that there are problems with extending this approach to scenarios with three or more participants. Karny & Guy (2004) recommend applying the method used above in an iterated pairwise fashion. However, as shown in Appendix B, the final distributions resulting from this are not invariant to the order in which the sharing is carried out, hence failing what seems to be a desirable and natural property. There are doubtlessly other shortcomings also, including the belief that a DM can supply adequate weights for her confidence in her own belief, and of those of her neighbours, and the lack of any discussion of how learning over time may occur. We will return to discuss potential remedies to these issues in Section 3.

At this juncture it is neccessary to introduce a concept that is perhaps most commonly seen in Information Theory, that of the Kullback-Leibler (KL) Divergence (Kullback & Leibler, 1951). This tool is used as a distance measure, to gauge how far apart two distributions are, and is key here in determining optimal decisions. Formally, consider two probability density functions, f_i and f_j, which both act upon a common support, $\theta \in \Theta$. The the KL divergence is given by

$$D(f_i \| f_j) = \sum\!\!\!\!\!\!\int f_i(\theta) \ln\Big(\frac{f_i(\theta)}{f_j(\theta)}\Big) d\theta. \tag{5}$$

Note that summation is used in the cases of the probability distributions being discrete, with integration used in continuous situations. Some elementary properties of this function are

- Non-negativity: $D(f_i \| f_j) \geq 0$ for all f_i, f_j, *i.e.*, the KL divergence is non-negative for all pairs of distributions.

- Equality: $D(f_i \| f_j) = 0$ if and only if $f_i \equiv f_j$, *i.e.*, $f_i(\theta) = f_j(\theta)$ for all $\theta \in \Theta$. This means that two distributions have a KL divergence of zero only when they are identical at every possible point.

Hence we see that the smaller the proximity between two distributions, the smaller the KL divergence will be, eventually becoming zero if the two distributions are indeed the same. There are some small technical issues with Measure Theory here, in which $D(f_i \| f_j) = 0$ if and only if $f_i \equiv f_j$ almost everywhere (Halmos, 1974), but we ignore these here.

Another important concept that Karny uses is that of an ideal distribution. Individuals model the way in which they believe events will occur via probabilistic distributions, attempting to describe the uncertainty that surrounds their environment. As well as this modelling distribution, Karny recommends individuals to put forward a distribution that in some sense describes a "best-case scenario", or an outcome that will be extremely satisfactory to them. This is their ideal distribution. Karny then ties together this idea, and the KL divergence defined above, to determine the optimal decision for an individual to make, using what he terms Fully Probabilistic Design (FPD). He proposes that the optimal decision is that which minimises the KL distance for the distribution which the DM chooses to model by, and her best-case scenario, *i.e.*, her ideal distribution. We omit the technical details here, but it can be shown that while a solution does exist, it is of a very specific and complicated form. Hence, while it provides a theoretical solution, it seems very possible that it may be highly problematic to endeavour to implement this method of FPD in any non-trivial practical context.

In the situations described above, participants are functioning in an environment devoid of any kind of hierarchy, in which they are all equal, in the sense that they have no superiors or subordinates. However frequently this is not the case, and there is a co-ordinator or facilitator or administrator who oversees the decision procedure, and may force participants to cooperate, or impose their own views upon participants. The most extreme example of this would be a facilitator who acts in an autocratic fashion, ignoring the aspirations and beliefs of all other DMs. Generally one of the key tasks that a facilitator must complete is to determine how much of their own opinion to include in the final distribution - small enough so that they cannot be said to be acting dictatorially, but large enough that they may have a significant effect, as often they may have access to knowledge that the other DMs do not. A prime example of this would be a government, elected by citizens. Karny & Kracik (2003) discusses the importance of minimising the difference between the information available to the citizens, and the surplus information that a government has that the citizens do not. Doing so will help the citizens to make the most informed decision possible, and decrease the need for government involvement (*i.e.*, the government can afford to give itself a lower weight in the weighted sum).

In other work, Karny & Guy (2009) consider further methods of cooperation by participants in a group decision making framework, via the sharing of probabilistic information, in which a DM must *"...take all offered information pieces as outputs of noisy information channels and try to estimate parameters of the underlying source."* A Bayesian ideology is followed, *i.e.*, the DM assigns some prior distribution to the unknown paramater of interest, and a corresponding likelihood function to the data observed (*i.e.*, the information received from neighbours), and performs the usual Bayesian updating to learn about the parameter *a posteriori*. Serveral assumptions are made about participants, perhaps the most interesting being that there is no mechanism to measure the reliability of the information recieved, nor to compare the differing importance of two different pieces of information from two different sources. We will return to discuss this in detail in Section 3, presenting ways that this may be amended.

2.6.4 Equilibria in Game Theory

Some of the foremost techniques in group decision making problems were developed by Nash (1951). These are often applicable in situations with two (or more) non-cooperative participants, which arise frequently in the context of problems within Game Theory. Suppose we have a collection of individuals who play a game together, and as part of the structure of the game they have the ability to bargain with each, to influence the outcome that will occur. How can they determine which set of actions is optimal and fair? Nash (1951) created a collection of six axioms, which he felt were justifiable that a fair bargaining point (*i.e.*, the action which is deemed to be fair taking into account the utilities of all participants) should obey. Given these axioms he proved, in his famed Nash Bargaining Theorem, that there was a unique function that would lead to this fair bargaining point being deemed the optimal decision. A further important contribution was the concept of Nash equilibrium. Here it is assumed that each player knows the equilibrium strategies of the other players, and none of the players involved have anything to gain from making a unilateral change to their own strategy. If a state of play is reached where no individual can benefit from making a change in their strategy whilst all other players leave their respective strategies unchanged, then this collection of choices of strategies, and the payoffs to players that will result from their implementation, comprise a Nash equilibrium.

2.6.5 Objections and Alternatives

Above we have discussed the concept of Nash equilbrium as a game-theoretic solution. While in many cases it is a powerful methodology there are some situations when it may be viewed as overly simplistic, as it does not take into account some potentially key variables, for instance potential altruism amongst participants. Nash equilbrium is applied in cases where all participating individuals are commited to the principle of individual rationality, which is to say that each individual wants to maximise their own expected utility, working under the assumption that all DMs want to do the same. Their thoughts with regard to the other participants of the game are purely in terms of the actions they may take, and what the consequences of these actions will be for themselves. In realistic situations however this may not always be the case, and participants may be willing to sacrifice some part of a benefit to themselves, in order to benefit the group as whole. The difference between the two cases are subtle but crucial. Nash's work is applicable in the case where an individual contemplates his fellow participants only in the sense of predicting their actions so as to maximise her own benefit conditional on their behaviour, while in cases where fellow participants are considered in the sense of some class of "good-will", with a potential willingness to make concessions for them, it is not. Below we briefly present some alternatives to this, which either attempt to incorporate some form of altruism between DMs, or attempt a novel approach to solving the problem at hand.

Stirling (2004) references how the process of finding Nash equilbria is often used in cases when it is not fully appropriate, and assumptions of narrow-minded self-interest are imposed on players when this may be not be the case. He proposes a different approach, called satisficing (a product of satisfying and sacrificing) which searches for a set of decisions which are deemed good enough for participanats, rather then a single optimal solution. Info-Gap Decision Theory (Ben Haim (2006,2007)) is a similar approach to trying to solve a decision problem by satisficing. It is non-probabilistic and considered to be a viable technique in cases where there is extreme (Knightian) uncertainty surrounding some quantity of interest. Lastly we comment on the work of Rabin (1993) attempting to formalise the abstract concept of fairness, which simply indicates that you are nice to those who are nice to you, and motivates you to be vindictive to those who aren't. These characteristics are exhibited more prominently for smaller material gains. From these axioms fairness measures and equilibria are derived, and compared to those equilibria of Nash.

2.7 Combining Expert Judgements

Much consideration has been given to the problem of combining judgements of domain-specific experts into a single belief to be used by a DM in a decision making problem, for instance Cooke (1991, 2007) and Clemen & Winkler (1999) among others. The problem which we will consider in Section 3 approaches this technique from a different viewpoint, with each individual themselves being considered an expert, who must combine her belief with those of other experts, in order to make her own decision. Nevertheless, the following discussion, which takes place in the setting of the former, is equally applicable to the latter.

Perhaps the most obvious initital question is how should beliefs be represented? The language of probability seems the obvious choice by which individuals can express uncertainty about unknown quantities. A simplified alternative proposal is the Bayes Linear method (Goldstein & Wooff, 2007) in which an individual states only the first two moments (the mean, variance and potential covariances) of their belief. Wisse *et al.* (2008) apply this method in an opinion pooling framework. However it is commonly regarded that specification of a full probability distribution over uncertain events is possible (French, 2011). Note that in Section 3.8 we do discuss the applicability of our techniques to the Bayes Linear method, and adaption is possible in two of the three cases. In what follows we assume that an individual can either construct this probability distribution herself, or can have it provided via elicitation methods (*e.g.*, O'Hagan, 1998). Nevertheless the Bayes Linear method could be applied in many of the approaches we will develop, as we shall shortly see.

There are two methods in which opinions can be combined, mathematically and behaviourally. The former involves a mathematical function, which takes as an input the various beliefs of individuals and returns a single belief, which in some sense represents collective opinion. The latter is a more heuristic approach, in which individuals discuss their opinions together, in an effort to reach some class of consensus. While several behavioural procedures have been developed, notably the Delphi method (Dalkey, 1969) and the Nominal Group Technique (Delbecq *et al.*, 1975), these suffer from a lack of rigour in application, and can be seen as somewhat *ad hoc* approaches. Within mathematical combining there are two dominant methods, namely linear and Bayesian pooling. The former simply involves constructing a linear combination of the opinions of individuals, for instance in a arithmetic or geometric manner (both of which are discussed in Clemen and Winkler, 1999). The latter is more complex, with a DM first specifying her own prior distribution over the uncertain quantities of interest, before viewing the opinions of the experts as data, which are entered into a likelihood function, and combined with the prior to yield a posterior distribution. French (2011) comments favourably on the concept of expert beliefs as data, but concedes that this method has vast problems with implementation, not least in the need for the choice of an appropriate likelihood function. Clemen & Winkler (1999) propose some suggested forms for specific problems, but no generalised method has been found. Hence linear pooling is the most common method, and is what we will consider in Sections 3 and 4, specifically arithmetically.

The next problem arises with the choosing of the weights for this linear combination. Ideally these would reflect reliability, with high weights given to those individuals who were deemed trustworthy, and lower weights given to those deemed inaccurate. In Cooke (1991) seed variables are used as a method of assessing expert accuracy before a decision needs to be made, with experts asked to give their opinions on quantities whose exact value is unknown to them, but known to the DM. A discussion of the potential shortcomings of this approach, and some appropriate alternatives, takes place in Flandoli *et. al* (2011). In the absence of such seed variables, the Laplacian Principle of Indifference (Laplace, 1812) is often applied, giving all individuals equal weight initially, in the absence of any significant evidence to favour one over another. Aven & Guikema (2011) debate the interpretation of what this linear pool is representative of, and whose belief, if anyones, it is meant to convey. Our own interpretation will become evident in the following sections.

3 Reliability Updating for a single DM

Thus far we have introduced the key concepts of decision theory, both in the context of a single individual who must make some choice from a collection of alternatives, and in the setting of a group who must decide what action is in some sense optimal for them as a whole. The former case is relatively straightforward, while the latter is confounded by Arrow's Impossibility Theorem. In this section we discuss potential methods by which a single DM may listen to her neighbours in order to improve her decision quality. The advantages and disadvantages of these approachs are discussed, and the results that occur in each case are compared by consideration of a simple illustrative example tackled in turn by each method. In Section 4 we will progress to consideration of a more complicated environment in which a collection of DMs must merge their beliefs in order to reach a single collective decision which they wish to be of a high quality. However in what follows in this chapter we discuss the simplified case where beliefs are merged in order to aid the decision quality of each individual DM, *i.e.*, at each decision epoch each DM makes their own decision (as opposed to a single decision being made determined by the beliefs and utilities of all DMs combined), the consequences of which will affect only herself.

3.1 Motivation

Consider a decision making problem which entails multiple DMs. Each is deemed to be imperfect, which is to say that they are uncertain in relation to at least some of the elements of the environment which they inhabit, *i.e.*, no DM is an all-knowing individual. A collection of neighbours will frequently have different degrees of knowledge about the uncertain event of interest, and hence may have greatly differing beliefs over what outcome may obtain from a particular decision. Not only this, but they will commonly have varying degrees of conviction in the belief that they hold, *e.g.*, two neighbours could both predict that the same outcome will occur, but one could be almost certain of this, while the other could be deeply unsure. It seems a logical concept that in order to maximise the decision quality of individual participants within the group it is imperative to maximise the amount of information which they have access to prior to making their decision. It stands to reason that a DM who is well informed about the topic at hand will generally make more astute decisions then one who is not. Within this (non-competing) framework the rational course of action seems to be to get all DMs to share their beliefs with each other, incorporating a greater scope of knowledge into their individual decision making task.

We consider a framework loosely based on that discussed by Karny & Guy (2004) in Section 2.6.3, where each DM shares their beliefs with their neighbour, expressed via their probability distribution over the uncertain parameter of interest, leading to combined beliefs like those in Equations (3) and (4). Now each DM may take into account the beliefs of each of her neighbours when making her decision. The methods that we propose as approachs to determining the weights giving to each neighbours distribution are different then those discussed by Karny & Guy, in an attempt to resolve the problems discussed previously with weights being chosen in a seemingly arbitrary and factless basis. We also seek to address the issue of changing weights over time. Participants will begin to notice over time that some of their neighbours are more accurate sources of information then others. Clearly the seemingly logical course of action for a DM, upon this realisation, is to pay more heed to those views proffered by neighbours whom she deems to be reliable, and to somewhat disregard those views profferd by neighbours whom she deems to be unreliable. Several subjective methods by which each participant can accurately calculate a numerical measure of the reliability of each of her neighbours are detailed in the following sections, with the results of each method being compared. Intuitively this should lead to an increased decision quality, as decisions are now made using information which is believed to be trustworthy, and indeed has shown itself to be so in the past.

As an illustration of the above we considered the example of a collection of stock-brokers with interest in the value of a stock at some future time, but it is clear that the theory being discussed here is potentially applicable to a wide-range of fields. Meteorologists could merge beliefs over what weather will occur at some future date in order to determine which temperature values they should predict, marketing companies could combine opinions on how successful a particular product will be in order to determine whether it should be released or not (and if so how many units should be produced), and software developers could share their thoughts on the number of bugs existing in their programs to determine if additional testing is required or not, *etc.*. Lastly, to consider an approach with non-human DMs we may think of a collection of traffic-lights throughout a city, that are attempting to optimise traffic flow, using sensors to record information which they share amongst each other to try to determine an optimal configuration.

3.2 Reliability Measures

Consider a setting of n decision makers, denoted by $P_1, \dots P_n$, with $n \geq 2$. There is some uncertain event or quantity of interest, θ, for which each participant P_i has their own probability distribution $f_i(\theta)$. We assume the DMs are non-competing, and have no reason to supply each other with intentionally inaccurate beliefs, *i.e.*, there is no motivation for dishonesty amongst participants. They all make their own decisions which will only impact upon themselves. Each DM is willing to listen to the beliefs of her neighbours, and constructs a weighted arithmetic linear combination of the following form:

$$\hat{f}_i(\theta) = \alpha_{i,1} f_1(\theta) + \dots + \alpha_{i,n} f_n(\theta). \tag{6}$$

Here $\alpha_{i,j}$ is the weight given by P_i to P_j, with $\alpha_{i,j} \geq 0$ for all $i, j = 1, \dots, n$ and $\sum_{j=1}^{n} \alpha_{i,j} = 1$ for all $i = 1, \dots, n$. Below we discuss methods by which these weights may be attained as an accurate reflection of the predictive ability of neighbours, in an attempt to ensure a higher decision calibre for DMs.

The decision that is optimal for P_i is that maximising expected utility, *i.e.*, d^* such that

$$d^* = \arg\max_{d \in \mathcal{D}} E[u_i(d)] = \sum\!\!\!\!\!\!\int u_i(d, \theta) \hat{f}_i(\theta) d\theta. \tag{7}$$

Note that this is a generalisation of Equation (2). For P_i this is a function of combined beliefs, $\hat{f}_i(\theta)$, and her utility function $u_i(r)$. For notational ease we use the same symbol u_i for $u_i(d_i, \theta) \equiv u_i(r)$ and $u_i(d)$, and presume it is evident what is meant by the context it appears in. Also here we assume that a DM is able to specify her utility function over potential outcomes. A method by which these preferences could be ascertained over time is adaptive utility, discussed in, for example, Cyert & DeGroot (1975) and Houlding & Coolen (2011), and modified to incorporate extreme vagueness in the priors over preferences in Houlding & Coolen (2012).

3.3 Updating Distributions

Each participant P_i has their own personal prior distribution over θ at the first decision epoch, denoted $f_i(\theta)$. As a result of the decision that they make they will observe some outcome r. In the framework developed here DMs update their beliefs in a Bayesian manner, using the standard "posterior is proportional to likelihood times prior" methodology, *i.e.*,

$$f_i(\theta|r) \propto f_i(r|\theta) f_i(\theta). \tag{8}$$

The idea of Bayesian conjugacy is a very useful one to ensure tractability of distributions, in other words that the posterior distribution is easily obtainable in closed form. This avoids the need for analytical evaluation methods, such as MCMC. For a choice of likelihood model belonging to the exponential family, there exists a conjugate prior distribution, such that the product of the prior and the likelihood will be proportional to a distribution that is of the same functional form of the prior but with different parameters. Common examples of this include Beta-Binomial conjugacy, Poisson-Gamma conjugacy, and, as we shall use later on in our worked example, Normal-Normal conjugacy. Enforcing of this principle may seem like quite a restrictive tactic, however, the exponential family of distributions is a very broad and flexible one, ensuring it is applicable in a wide range of real-world cases, and is hence suitable for realistic modelling. In the example that we consider in this chapter and in subsequent chapters we will consider beliefs that allow for this conjugacy, purely for simplicity and ease of interpretation. Nevertheless it is important to note that we are capable of solving problems using any non-conjugate distributions by numerical methods, such as the aforementioned MCMC approach.

Participants use Equation (8) to update their beliefs on an epoch by epoch basis, with their posterior distribution at epoch n becoming their prior distribution at epoch $n+1$. Hence their distributions are constantly being augmented over time in light of new information, and will become more accurate as the number of epochs increases.

We now progress to the discussion of an important issue - why should participants perform Bayesian updating on their own personal probability distribution of beliefs, $f_i(\theta)$, and not on $\hat{f}_i(\theta)$, the distribution which is a function of the beliefs of all their neighbours and themselves, which, using Equation (7), they have used to make decisions? In the opinion of the author there are several reasons why the approach of updating solely the personal beliefs of a DM should be advocated above the alternative.

Firstly, it is important that a DM, who makes decisions by use of a combined belief, be able to extract her own personal beliefs from this, that is to say that she is able to segregrate her own opinions from those of her neighbours. Under the methodology currently proposed this is the case, with DMs updating the belief that is solely their own over time, before combining it with that of their neighbours in a weighted sum, with weights proportional to the reliability of each neighbour. While, in this framework, where DMs acquiest to use a diluted version of their own beliefs, by subjecting them to becoming a component of a weighted sum instead of leaving them whole, there is still a distinction between $f_i(\theta)$ and $\hat{f}_i(\theta)$. The former reflects their own personal probabilistic beliefs about the outcome of some uncertain event. The latter is what they use to make a decision, bearing in mind that they are not infallible, and while they believe the opinion about unknown θ that they hold to be the right one (for if not then why would they hold it?), they allow for the fact that they may be wrong, and hence listen to, and incorporate, the beliefs of their neighbours into their decision making process. We operate within an environment where there is a willingness among participants to listen to the views of others, and they are prepared to make a concession, allowing their decision to be influenced by external opinions. This is because of the assumed belief that an increase in information will transcend to an increase in decision quality. Hence to refer to $\hat{f}_i(\theta)$ as the beliefs of P_i is not correct. The beliefs of P_i are given by $f_i(\theta)$, while $\hat{f}_i(\theta)$ is just a mechanism that she uses in the hope of improving her decision quality. It is important that we distinguish between the beliefs of a DM, and what she uses to make her decisions with. In our framework there is a subtle difference between these two constructs.

DMs also want true measure of the reliability of their neighbours, and hence want to consider their own personal beliefs, as opposed to their weighted sums which are not necessarily a reflection of what they believe themselves, but a conglomerate of their opinions and those of others. Hence updating over these distributions is essential. Without this, one DM may seem to be highly reliable, when in fact that individual is a deeply inaccurate source, who by chance has a weighted sum that conveys accurate beliefs. In order for DMs to learn about the true underlying accuracy of their neighbours it is important that the distinction between $f_i(\theta)$ and $\hat{f}_i(\theta)$ be maintained, and hence that updating be applied to the former, and not the latter. The latter will only be updated when the DM combines the collection of newly updated individual distributions which she receives from her neighbours, and constructs a new weighted sum based on these. Updating occurs in $\hat{f}_i(\theta)$ through the updates to $f_1(\theta), \ldots, f_n(\theta)$, and via the weighting procedures which we shall discuss in the following section. We will return later to mention how this leads to our approach being not fully Bayesian in the traditional sense, but perhaps could be dubbed as "semi-Bayesian".

3.4 Updating Weights

We see that having made decisions and observed returns r, participants have updated their own distributions using Equation (8) above. Once again each DM shares her beliefs with her neighbours, now proffering their posterior distribution over θ rather than their initial prior. Note that for for notational ease and compactness we simplify and write $f_i(\theta|H)$ to denote the posterior distribution for P_i, where H denotes the history of returns observed from decisions made up to this point. Participant P_i has now received probability distributions from each of her neighbours - but what weight should she afford to each? As we mentioned previously in our discussion on combining expert opinion in Section 2.7, equal weights are often assumed at the first epoch, in lack of any more substantial evidence to the contrary. However, having now observed some outcome they will have received an indication of the relative accuracy of various neighbours, and will want to give more weight to those deemed reliable than to those deemed to be unreliable. Three methodologies are considered below for suitably updating. The first of these is based on the Kullback-Leibler divergence (which we defined in Equation (5)), and differs from the other two approaches in that it does not begin with an assumption of equal weights. While it has some deficiencies and limitations it is still a novel and indeed intuitive approach, and worthy of discussion and potential future development. Secondly we discuss a "plug-in" approach, which while in some sense natural, suffers from some important shortcomings. Finally we construct what we term the "differing viewpoints" approach which resolves many of the problems which the previous two techniques suffer from. We will discuss each of these in turn, and then compare the results that the contrasting methodologies yield via consideration of a common example.

3.5 Kullback-Leibler Approach

DMs will often be biased towards their own belief, and hence may not wish to apply the Principle of Indifference, and give it the same weight as the opinions of neighbours in their weighted sum. They may frequently be inclined to give a high weight to those neighbours who have beliefs similar to their own. We recall the Kullback-Liebler (KL) divergence (Kullback & Leibler, 1951), which serves as a measure of the similarity between two probability distributions. Small values of the KL divergence indicate that the beliefs of two individuals are similar. Hence a DM will want to give high weights to neighbours whose distribution returns a low KL score with her own, and low weights to those resulting in a high score. As $D(f_i||f_i) = 0$ a participant must choose an arbitrary value of $\alpha_{i,i}$ to assign to her own distribution in the weighted sum. Note that this allows a DM to fully discount the beliefs of her neighbours, by allocating herself a weight of $\alpha_{i,i} = 1$. Conversely if she has absolutely no faith in her own beliefs she can set her weight equal to zero. We propose that an individual P_i calculates $D(f_i||f_j)$ for all $j \neq i$, and then inverts these values and rescales so weights sum to $1 - \alpha_{i,i}$. This will ensure that higher weights are accorded to those with beliefs similar to the DM constructing the weighted sum. Note if an individual P_j has an identical distribution to that of P_i then P_i simply halves the weight she had allocated to herself between herself and P_j. Formally the weights are given by

$$\alpha_{i,j} = \frac{\frac{1}{D(f_i||f_j)}}{\sum_{j \neq i} \frac{1}{D(f_i||f_j)}} \times (1 - \alpha_{i,i}). \tag{9}$$

These weights are then used to create $\hat{f}_i(\theta)$, which is used, along with $u_i(r)$ to determine which decision is optimal. A decision is made, some outcome observed, and $f_i(\theta)$ is updated in light of this. The procedure is then repeated, with KL divergence values being calculated for the new beliefs of individuals.

We briefly mention three shortcomings of this approach. Firstly the KL divergence is not guaranteed to be symmetric, which is undesirable. However note that alternative dissimilarity measures which are symmetric could be used, for instance total variation distance, or Euclidian distance. Secondly the choice of the weight $\alpha_{i,i}$ is an arbitrary decision to be made by P_i. However while this may appear to be a problem, it does give added subjectivity to the approach, allowing an individual to ignore the opinions of all neighbours if she wanted to make a decision based solely on her own beliefs or indicate her level of confidence in her own expertise. In this non-competing environment, where the negative consequences of a bad decision only affect the decision maker, this potential stubbornness can be viewed as acceptable. Finally there is a potential problem if a DM has very little faith in her own vague beliefs, and hence gives herself a low weight. Yet beliefs which are similar to hers will receive a high weight, while those which are more confident (*i.e.*, have smaller variance) will, rather unintuitively, receive a low weight.

3.5.1 Step-by Step Methodology for KL approach

We present here a clear step-by-step methodology of the material presented in the previous section, illustrating how this approach should be implemented in practise.

- *Step 1:* Participants $P_1, \ldots, P_n$ specify their beliefs over an uncertain event of interest θ, via their probability distributions $f_1(\theta), \ldots, f_n(\theta)$.
- *Step 2:* Each P_i decides the weight $\alpha_{i,i}$ that they wish to allocate to their own distribution $f_i(\theta)$ in the weighted sum $\hat{f}_i(\theta)$ at the first epoch.
- *Step 3:* Each P_i calculates $D(f_i||f_j)$ for all $j \neq i$. Now they invert each $D(f_i||f_j)$ to find $\frac{1}{D(f_i||f_j)}$, and rescale these so that they sum to $1 - \alpha_{i,i}$. These rescaled values now become the weights $\alpha_{i,j}$ to be used in the weighted sum $\hat{f}_i(\theta)$.
- *Step 4:* The method of maximisation of expected utility is used to determine which decision is optimal for a DM. This calculation involves $\hat{f}_i(\theta)$ and each participant's subjective utility function $u_i(r)$. This is done by use of Equation (7).
- *Step 5:* The outcome of the decisions made are observed, and some results are returned.

- *Step 6:* All participants use Bayes' Theorem, as given in Equation (8) to update their own personal beliefs over the unknown quantity θ.

- *Step 7:* Each P_i may now change the weight they allocate to their own distribution $\alpha_{i,i}$ if they wish to.

- *Step 8*: All participants now share their updated distributions from Step 6 with each other, and Steps 3-7 are repeated for as many epochs as is necessary, or until no new observations can be witnessed, and hence no further updating can occur.

3.6 Plug-in Approach

Unlike the KL approach discussed above, the Plug-in approach initialises with the Principle of Indifference. At the first epoch all DMs give the beliefs of their neighbours, as well as their own beliefs, an equal weight, *i.e.*, $\alpha_{i,j} = \frac{1}{n}$ for all $i, j = 1, \dots, n$. It is only after the first decision is made, and some outcome is observed, that the weighting based on the perceived reliability of individuals begins. We denote by θ^* the value that is witnessed.

For each individual $f_i(\theta^*)$ is calculated. This will return a high value if the belief distribution supplied by P_i placed a high probability of the occurence of θ^*, and a low value if not. Hence it is a precise measure of the predictive ability of an individual. DMs will want to give high weights to neighbours P_i who have high values of $f_i(\theta^*)$. We write $b_i = \frac{1}{f_i(\theta)}$. This is small for reliable DMs, and large for inaccurate ones. The weight allocated to P_i is now changed from $\frac{1}{n}$ to $\frac{1}{n+b_i}$, with these values rescaled to ensure they sum to one. Generally, when $\alpha^*_{i,j}$ denotes the updated weight and $\alpha_{i,j}$ is the previously associated weight, we write

$$\alpha^*_{i,j} = \frac{\frac{1}{(\alpha_{i,j})^{-1}+b_i}}{\sum_{i=1}^{n} \frac{1}{(\alpha_{i,j})^{-1}+b_i}} \tag{10}$$

Note that problems may arise if an individual gives a probability of zero to the event that occured. Allocating probabilities of zero is advocated against by Cromwell's Rule (Lindley, 1991), as it implies no amount of evidence will persuade the individual to change their mind. In the case of an individual breaching this rule, and the event occuring, we may consider omitting them from the analysis.

This method appears promising, as theoretically the decision quality should substantially improve over time, as the relevant information which the decision is based upon becomes increasingly accurate. It also appears fair to the DMs involved, as they are initially assumed to be equally reliable, and the increase/decrease in the weight allocated to them is directly proportional to how accurate they have shown themselves to be. Note that all individuals will have the same combined beliefs $\hat{f}_i(\theta)$ at each individual epoch, as the weights are determined in an objective manner, in an identical fashion for each DM, *i.e.*, at any particular epoch $\alpha_{i,j}$ is guaranteed to be the same as $\alpha_{k,j}$, for all i, k. This is proved formally in Appendix C.

3.6.1 Step-by Step Methodology for Plug-in approach

We present here a clear step-by-step methodology of the material presented in the previous section, illustrating how this approach should be implemented in practise.

- *Step 1:* Participants $P_1, \ldots, P_n$ specify their beliefs over an uncertain event of interest θ, via their probability distributions $f_1(\theta), \ldots, f_n(\theta)$.

- *Step 2:* Each DM creates a weighted sum of beliefs $\hat{f}_i(\theta)$, initially giving herself and all neighbours equal weights.

- *Step 3:* The method of maximisation of expected utility is used to determine which decision is optimal for a DM. This calculation involves $\hat{f}_i(\theta)$ and each participant's subjective utility function $u_i(r)$. This is done by use of Equation (7).

- *Step 4:* The outcome of the decisions made are observed, and some results are returned.

- *Step 5:* All participants use Bayes' Theorem, as given in Equation (8) to update their own personal beliefs over the unknown quantity θ.

- *Step 6:* Each participant plugs in the outcome that they have witnessed, θ^* into the distributions of their neighbours, as well as that of themselves, to find $f_1(\theta^*), f_2(\theta^*), \ldots, f_n(\theta^*)$. They now change from the initial weight $\frac{1}{n}$ to $\frac{1}{n+b_i}$, where $b_i = \frac{1}{f_i(\theta^*)}$. These weights are now rescaled to ensure they sum to 1.

- *Step 7:* Participants now share their updated beliefs from Step 5 with each other, and create their own weighted sum of beliefs for the next decision epoch using these, and the weights calculated in Step 6.

- *Step 8:* Steps 3-8 are now repeated for as many epochs as is necessary using Equation (10) in Step 6, or until no new observations can be witnessed, and hence no further updating can occur.

3.7 Differing Viewpoints Approach

Decision making is built upon two fundamental pillars, probability and utility. The plug-in approach and the KL approach, focused solely on the former in the determination of weights. The final method that we consider incorporates the utility functions in an attempt to increase the procedure's subjectivity.

DMs make their first decision in a manner identical to that of the Plug-in approach, observe some return, and then update their own distributions in light of this. The difference between the methods is in the choosing of the weights. Each individual P_i considers the expected utility she would have assigned to d^*, the decision she made at the first epoch, had she solely listened to the beliefs of each neighbour P_j in turn. We write this value as

$$E_{i|j}[u_i(d^*)] = \oint u_i(d^*, \theta) f_j(\theta) d\theta. \tag{11}$$

Next P_i calculates the absolute value of the difference between the utility that actually resulted from the decision d^* and $E_{i|j}[u_i(d^*)]$, written

$$w_{i,j} = |u_i(r) - E_{i|j}[u_i(d^*)]|. \tag{12}$$

This is a measure of the reliability afforded by P_i to P_j, with large values of $w_{i,j}$ indicating P_j is not an accurate source of information to P_i as there was a large discrepency between what actually happened and what P_j predicted would happen. The converse is also true. Briefly we justify this further. The introduction of the utility function into the analysis is an important step as it allows an increased subjectivity to enter the procedure. Different neighbours will respond differently to the distribution $f_j(\theta)$ depending on their own feelings towards risk. A risk-averse DM may think more highly of the conservative views of a neighbour then a very risk-prone DM for instance. A final point to make is that $w_{i,j}$ deals with the order of magnitude between two quantites, but is invariant to which of the two is bigger. For example consider the case where $u_i(r) = 10$, $E_{i|j}[u_i(d^*)] = 5$ and $E_{i|k}[u_i(d^*)] = 15$. This leads to $w_{i,j} = w_{i,k} = 5$. Should both weights be the same even though the beliefs of P_k lead to a higher expected utility then that of P_j? We argue yes. If a DM is trying to determine which decision to make from a collection of alternatives then of course she will be interested in that with the higher expected utility, as she wants to maximise her return. However in this setting $w_{i,j}$ is simply a measure of the reliability of a neighbour, and a prediction of five utils too high is as accurate/inaccurate as a prediction of five utils too low. It is what a DM could expect to get listening only to a particular neighbour, not what she would get.

Finally if we denote by $\alpha^*_{i,j}$ the updated weight given by P_i to P_j, and $\alpha_{i,j}$ the weight that was previously given before the lastest result was observed, then updating is given by

$$\alpha^*_{i,j} = \frac{\frac{1}{(\alpha^{i,j})^{-1} + w_{i,j}}}{\sum_{j=1}^{n} \frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}} \tag{13}$$

The theorem below shows that this approach meets some desirable conditions. The proof and some added interpretation, are included in Appendix D.

Theorem: The Differing Viewpoints approach obeys the following five properties:

- **1**: $w_{i,j} \geq 0$ for all $i, j = 1, \dots, n$.
- **2**: $w_{i,j} = 0$ if and only if $u_i(r) = E_{i|j}[u_i(d^*)]$. This is impossible in a continuous setting.
- **3**: If $\alpha_{i,j} < \alpha_{i,k}$ and $w_{i,j} > w_{i,k}$ then $\alpha^*_{i,j} < \alpha^*_{i,k}$ and $\alpha^*_{i,k} - \alpha^*_{i,j} > \alpha_{i,k} - \alpha_{i,j}$.
- **4**: If $\alpha_{i,j} = \alpha_{i,k}$ and $w_{i,j} = w_{i,k}$ then $\alpha^*_{i,j} = \alpha^*_{i,k}$.
- **5**: If $\alpha_{i,j} < \alpha_{i,k}$ and $w_{i,j} < w_{i,k}$ then any of the following may occur, depending on the initial differences between the weights, and the updated difference between the reliability measures:
  - $\alpha^*_{i,j} < \alpha^*_{i,k}$, in which case $\alpha^*_{i,k} - \alpha^*_{i,j} < \alpha_{i,k} - \alpha_{i,j}$.
  - $\alpha^*_{i,j} = \alpha^*_{i,k}$
  - $\alpha^*_{i,j} > \alpha^*_{i,k}$

3.7.1 Step-by Step Methodology for Differing Viewpoints approach

We present here a clear step-by-step methodology of the material presented in the previous section, illustrating how this approach should be implemented in practise.

- *Step 1:* Participants $P_1, \ldots, P_n$ specify their beliefs over an uncertain event of interest θ, via their probability distributions $f_1(\theta), \ldots, f_n(\theta)$.
- *Step 2:* Each DM creates a weighted sum of beliefs $\hat{f}_i(\theta)$, initially giving herself and all neighbours equal weights.
- *Step 3:* The method of maximisation of expected utility is used to determine which decision is optimal for a DM. This calculation involves $\hat{f}_i(\theta)$ and each participant's subjective utility function $u_i(r)$. This is done by use of Equation (7).
- *Step 4:* The outcome of the decisions made are observed, and some results are returned.
- *Step 5:* All participants use Bayes' Theorem, as given in Equation (8) to update their own personal beliefs over the unknown quantity θ.
- *Step 6:* Each participant P_i considers the expected value she was have given to the maximal decision from Step 3 under the sole influence of each of her neighbours in turn, using Equation (11).
- *Step 7:* Equations (12) and (13) determine the measure of reliability which P_i associates with P_j, and the weight which she will afford to her at the next decision epoch.
- *Step 8:* Participants now share their updated beliefs from Step 5 with each other, and create their own weighted sum of beliefs for the next decision epoch using these, and the weights calculated in Step 7.
- *Step 9:* Steps 3-8 are now repeated for as many epochs as is necessary, or until no new observations can be witnessed, and hence no further updating can occur.

3.8 Comparisons

In this chapter we have introduced three potential methodologies by which a DM may listen to her neighbours in order to gain more information about some unknown event of interest, and use this information to make a decison. Emphasis was placed on determining which neighbours were accurate sources of information, with DMs wanting to place more emphasis in their decision process on their opinions. Each methodology proposed a different way of doing this, and some strengths and weaknesses of each have been discussed. The KL approach was a realistic description of human behaviour, but did not necessarily guarantee a high decision quality over a short number of decision epochs. An example of this may be a one epoch problem where all the DMs but one have very unreliable views. The plug-in approach was intuitive but lacked the subjectivity of the other two methods, and hence all DMs will have the same $\hat{f}_i(\theta)$ at each individual decision epoch, as we have previously discussed. Lastly the differing viewpoints approach attempted to incorporate the two fundamental concepts of decision making, probability and utility into the process. This weighting scheme was proved to have a collection of logical and desirable properties. Hence it seems to be the most promising of the three approaches in the opinion of the author. In the following section we will compare the performance of the three procedures on a common example, to see if the differing methods yield similar results.

We conclude this section with a brief comment regarding the applicability of Bayes Linear (Goldstein, 1999) belief specifications to the methods we have developed. Clearly they cannot be used in conjecture with the Plug-in approach, as a full distribution is require to enter the observed value into. We also need full probability distributions to calculate KL divergence measures, but if a different non-distributional dissimilarity measure was used then a method analogous to the KL method could be implemented. Lastly it is fully applicable to the Differing-Viewpoints approach.

3.9 Example

Consider a situation with four participants, P_1, P_2, P_3 and P_4. Each individual has her own beliefs about the true value of θ, which is a latent parameter pertaining to stock performance, and hence the profit or loss resulting from the decision made. Each individual must decide whether to enter into a long forward on the stock (d_1) or not (d_2). It is clear that in this setting there must be some uncertainty over the value of θ, as without this decision making would be trivial - a DM should always enter into the long forward contract if θ is positive, and should always not do so if θ is negative. Entering into a long forward involves agreeing to buy a stock at a fixed time (known as the expiry time) in the future for a fixed price, called the strike price. If the strike price exceeds the actual value of the stock at the fixed future time then the DM has made a loss (as they are buying the stock for more then it is worth), if not then they have made a profit. In the myopic scenario considered here, DMs must decide whether to enter into the long forward or not, at a succession of decision epochs.

We assume that beliefs over θ are Normally distributed with unknown mean and known variance. This is done for convenience and elucidation, ensuring tractability of computations via Normal-Normal conjugacy, although of course this is not required in generality and is not restrictive. Let $\theta \sim \mathcal{N}(\mu, 3)$. Participants have prior beliefs over this unknown mean μ, which themselves can be modelled by Normal distributions. The prior beliefs of P_i over μ are represented by $f_i(\mu) \sim \mathcal{N}(m_i, s_i^2)$. The posterior distribution of P_i over μ having observed some relevant return r is then

$$f_i(\mu|r) \sim \mathcal{N}\Big(\frac{\frac{m_i}{s_i} + \frac{r}{\sigma^2}}{\frac{1}{s_i} + \frac{1}{\sigma^2}}, \frac{1}{\frac{1}{s_i} + \frac{1}{\sigma^2}}\Big) \tag{14}$$

The decision participants make over whether to enter the long forward (potentially making a monetary gain but also potentially making a monetary loss) or not (ensuring no monetary gain and no monetary loss) are heavily influenced by their utility functions, expressing their attitudes towards risks and gambles, and the initial fortune which they begin with. Their opinions about θ, utility functions, and starting fortunes are given in Table 2 below.

Table 2: DM Information

P_1	$f_i(\mu)$	$u_i(r)$	Initial fortune
P_1	$f_1(\mu) \sim \mathcal{N}(5, 1)$	r	\$50
P_2	$f_1(\mu) \sim \mathcal{N}(2, 3)$	r^3	\$45
P_3	$f_1(\mu) \sim \mathcal{N}(-1, 3)$	$80r - 0.5r^2$	\$60
P_4	$f_1(\mu) \sim \mathcal{N}(-4, 1)$	$e^{\frac{r}{15}}$	\$35

3.9.1 KL example

Participants begin by choosing how much weight to allocate to their own beliefs. After this they use the KL method, as shown in Equation (9) to assign weights to their neighbours, which are included in Table 3. Table 4 shows the decision deemed optimal by each DM using their combined beliefs, and we see that all barring P_4 opt to enter into a long forward. Now assume that those who enter make a loss of \$2, and augment their beliefs in light of this return, as discussed in Equation (8). KL divergence values between the new distributions are calculated, and new weights assigned. Under the new combined beliefs only P_1 enters into a long forward at the next epoch. We see that even with the high levels of subjectivity in this approach (with each DM choosing the weight to assign to their own distribution) learning occurs quickly. Figure 1 gives a graphical interpretation of the changes taking place, focusing on the outlook of P_1 for illustration.

Table 3: Weights at Epoch 1 (LHS) and Epoch 2 (RHS) respectively.

P_i	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,3}$	$\alpha_{i,4}$
P_1	0.6	0.303	0.084	0.013
P_2	0.101	0.5	0.362	0.037
P_3	0.035	0.434	0.4	0.131
P_4	0.024	0.158	0.568	0.25

P_i	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,3}$	$\alpha_{i,4}$
P_1	0.5	0.314	0.149	0.037
P_2	0.052	0.4	0.503	0.045
P_3	0.06	0.12	0.5	0.32
P_4	0.02	0.148	0.432	0.4

Table 4: Optimal Decisions at Epoch 1 (LHS) and Epoch 2 (RHS) respectively.

P_i	d^*
P_1	d_1
P_2	d_1
P_3	d_1
P_4	d_2

P_1	d^*
P_1	d_1
P_2	d_2
P_3	d_2
P_4	d_2

3.9.2 Plug-in example

In the Plug-in approach all participants initially allow each other equal weight, meaning that they all make their decisions using the same combined distribution at the first epoch. They all opt to enter the long forward, as we see in Table 6, and again make a loss of \$2. Weights are then updated by plugging the observed values into the distribution of neighbours to see who is reliable and who is not. Updated distributions are combined with the updated weights (given in Table 5). At the second epoch no DM enters into the long forward. Figure 2 shows the changes to beliefs and weights for P_1, though we note that under this approach these would be the same for all participants.

Table 5: Weights at Epoch 1 (LHS) and Epoch 2 (RHS) respectively.

P_i	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,3}$	$\alpha_{i,4}$
P_1	0.25	0.25	0.25	0.25
P_2	0.25	0.25	0.25	0.25
P_3	0.25	0.25	0.25	0.25
P_4	0.25	0.25	0.25	0.25

P_i	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,3}$	$\alpha_{i,4}$
P_1	0.001	0.101	0.601	0.298
P_2	0.001	0.101	0.601	0.298
P_3	0.001	0.101	0.601	0.298
P_4	0.001	0.101	0.601	0.298

Table 6: Optimal Decisions at Epoch 1 (LHS) and Epoch 2 (RHS) respectively.

P_i	d^*
P_1	d_1
P_2	d_1
P_3	d_1
P_4	d_1

P_1	d^*
P_1	d_2
P_2	d_2
P_3	d_2
P_4	d_2

3.9.3 Differing Viewpoints example

The first phase of this approach is exactly the same as in the Plug-in approach, with the first difference occuring with the changing of the weights at the second epoch. We see in Table 7 that these changes are less extreme then those in the previous example. As Table 8 shows nobody opts to enter into another long forward at the second epoch. Figure 3 gives a graphical representation of changes from the perspective of P_1.

Table 7: Weights at Epoch 1 (LHS) and Epoch 2 (RHS) respectively.

P_i	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,3}$	$\alpha_{i,4}$
P_1	0.25	0.25	0.25	0.25
P_2	0.25	0.25	0.25	0.25
P_3	0.25	0.25	0.25	0.25
P_4	0.25	0.25	0.25	0.25

P_i	$\alpha_{i,1}$	$\alpha_{i,2}$	$\alpha_{i,3}$	$\alpha_{i,4}$
P_1	0.156	0.214	0.343	0.287
P_2	0.068	0.125	0.509	0.298
P_3	0.088	0.152	0.507	0.253
P_4	0.16	0.221	0.35	0.269

Table 8: Optimal Decisions at Epoch 1 (LHS) and Epoch 2 (RHS) respectively.

P_i	d^*
P_1	d_1
P_2	d_1
P_3	d_1
P_4	d_1

P_1	d^*
P_1	d_2
P_2	d_2
P_3	d_2
P_4	d_2

Figure 1: The individual and combined beliefs of individuals at Epoch 1 (a) and 2 (b), and the corresponding changes in weights (c), from the perspective of P_1, using the Kullback-Leibler approach.

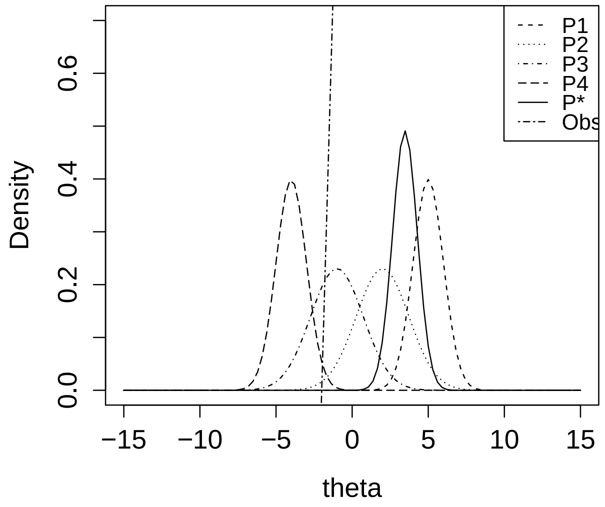

(a)

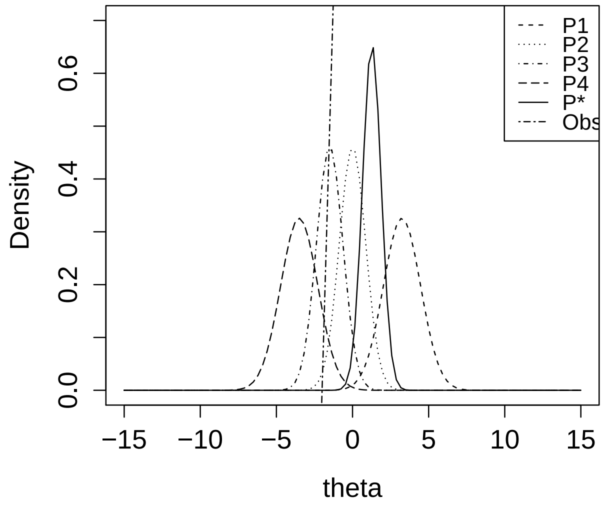

(b)

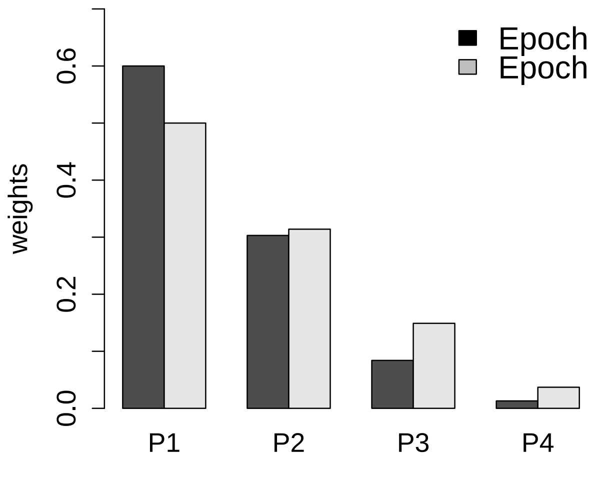

(c)

Figure 2: The individual and combined beliefs of individuals at Epoch 1 (a) and 2 (b), and the corresponding changes in weights (c), from the perspective of P_1, using the Plug-in approach.

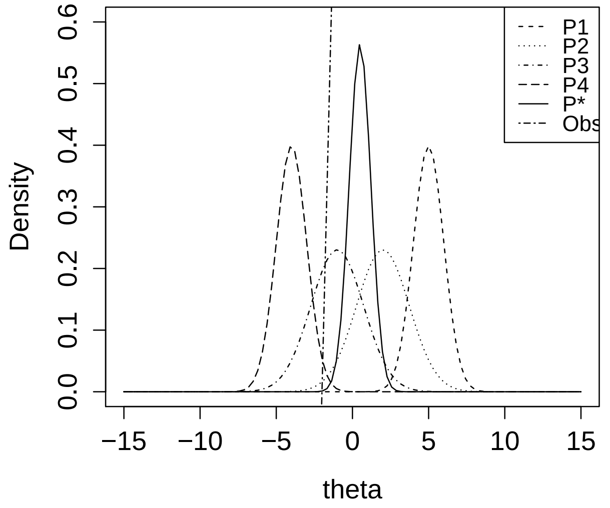

(a)

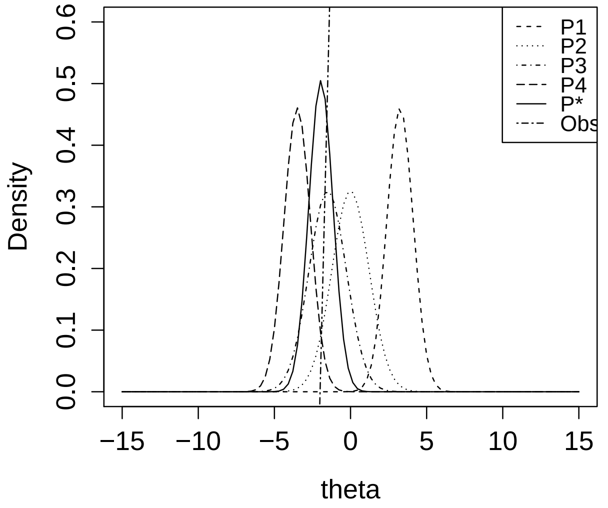

(b)

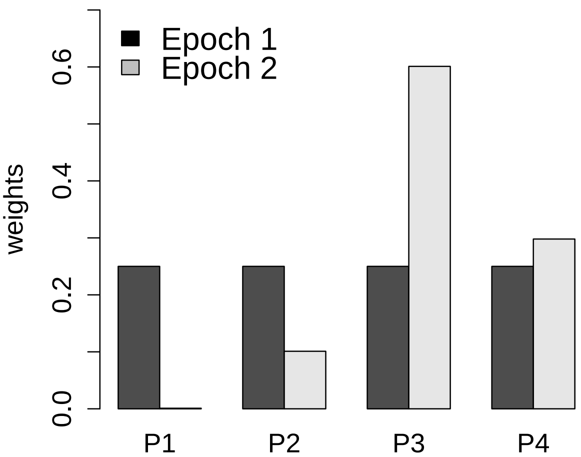

(c)

Figure 3: The individual and combined beliefs of individuals at Epoch 1 (a) and 2 (b), and the corresponding changes in weights (c), from the perspective of P_1, using the Different Viewpoints approach.

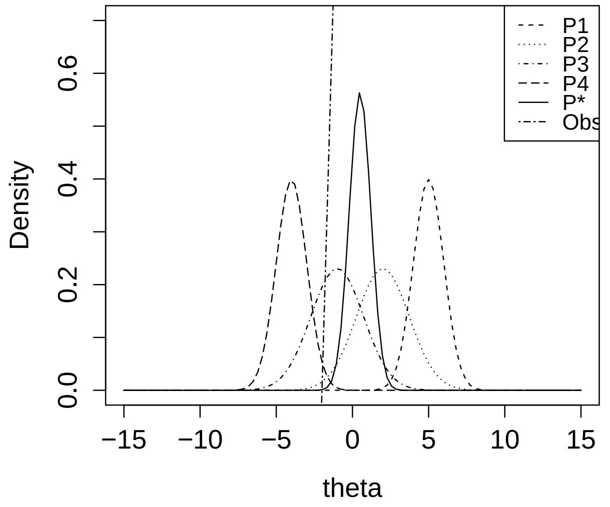

(a)

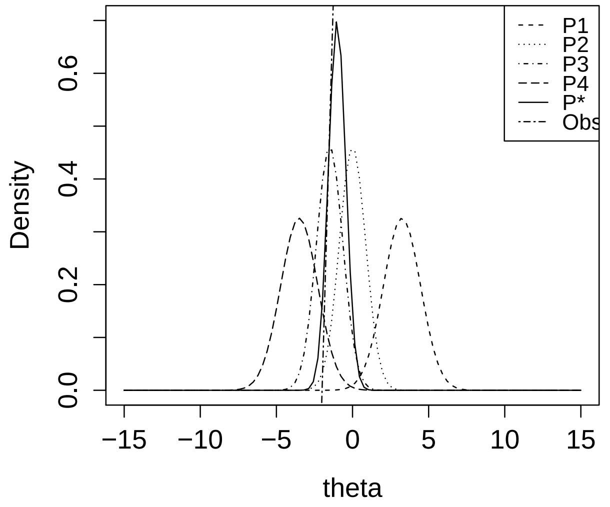

(b)

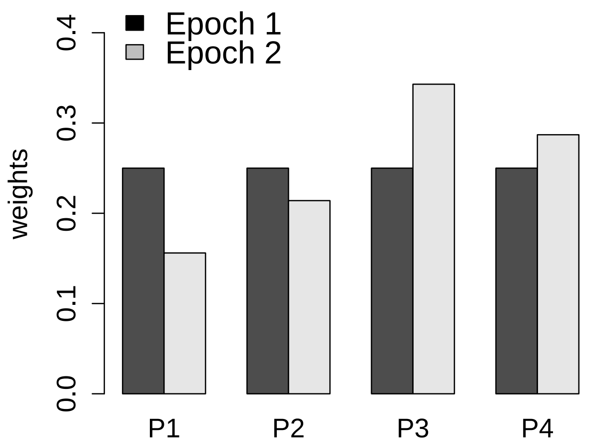

(c)

3.9.4 No Cooperation

Briefly we consider what happens when DMs do not interact with their neighbours, and solely use their own beliefs to make decisions. Table 9 shows the decisions deemed optimal at the first epoch. We see that the two optimistic participants choose to make the decision, while the two pessimistic participants do not. Suppose a loss is made of \$2 by P_1 and P_2, as happens in the previous examples. Now beliefs are updated as we have already seen. Table 10 shows the decision making for the remaining two participants.

Table 9: Decision choices for non-interacting participants at Epoch 1

P_i	$E[d_1]$	$E[d_2]$	d^*
P_1	55.01	50	d_1
P_2	104272	91125	d_1
P_3	2978.75	3000	d_2
P_4	7.92	10.32	d_2

Table 10: Decision choices for non-interacting participants at Epoch 2

P_i	$E[d_1]$	$E[d_2]$	d^*
P_1	51.26	48	d_1
P_2	79720	79507	d_1

Note that we omit P_3 and P_4 at this epoch. As cooperation between participants is no longer occuring it is possible that they will not be informed of the outcome (in this case a loss) made by P_1 and P_2 at the first epoch. We see that, depsite their loss, both of the DMs who opted to enter into the long forward at the first epoch do so again at the second epoch. This is because of the initial strong beliefs that they hold, which will need to be influenced by a lot of evidence to change them. We can see that use of the combined beliefs, as in the three methods provided here, leads to participants learning much quicker abot the uncertain element of interest, and hence making potentially better decisions.

3.9.5 Similarities and Differences

We consider the results in the three different methods applied to the same problem. In both the Plug-in and the Differing Viewpoints approach no DM was willing to enter into the long forward by the second decision epoch. By contrast, in the KL approach there was still one participant, P_1, who was willing to take the gamble and enter into a long forward, despite her loss at the previous epoch. This can be explained somewhat by the ability to specify the participants' confidence in this method, where each DM is able to choose to what degree they heed their own opinion. P_1 choose to give a weighting of 0.6 and 0.5 to her own beliefs at the first and second epoch respectively, and her beliefs over the profit θ were $\theta \sim \mathcal{N}(5, 1)$ initially and $\theta \sim \mathcal{N}(3.25, 0.75)$ having observed a loss. Even if the other participants distributions were indicative of a loss, her belief in a profit (combined with her high degree of confidence in herself) would negate this. This could be seen as either an advantage or a disadvantage - only one outcome has been observed, so it could be argued that P_1 is right to stick to her beliefs, as maybe the loss observed was a fluke result, and she should wait for more concrete evidence before changing her mind. The ability to specify confidence in the KL approach allows for DMs to potentially persist with their beliefs, inspite of evidence to the contrary.

The Plug-in approach is the quickest to implement, in that it involves only one weighting calculation at each decision epoch, as opposed to n in the two other methods. However, despite this speed of computation there are some weaknesses evident with this technique. Despite only observing one result the views of P_1 were already given a negligibly small weight, making her contribution insignificant. It seems a harsh measure to eliminate the beliefs of a DM based on one single result, which may be a fluke outcome, especially taking into account that the DM will learn over time in light of the new results observed and become increasingly accurate. Another disadvantage in $\hat{f}_i(\theta)$ being the same for $i = 1, 2, 3, 4$ is that it does not allow any subjectivity to enter into the analysis, in the way that it does in the KL method through the KL divergence measure, and the Differing Viewpoints method through the use of utility in the reweighting procedure. One advantage to be mentioned about the Differing Viewpoints approach is that a brief consideration of the weights that it produces for DMs to assign in their weighted sum show that there is not the same tendency to produce extreme values (*i.e.*, those close to 0 or 1). This is good as it ensures no neighbour of a DM will be entirely discounted early on.

4 Information Sharing Methods in a Group DM setting

Thus far we have discussed methods by which a single DM, who exists in a group environment (*i.e.*, she has some neighbours with whom she has a common interest) makes a decision which is (potentially) based on the information which she receives from those around her. Yet despite the introduction of a group element with respect to the sharing of information, each DM is still making a decision for herself, with no need to heed the utility of her neighbours. If we have n participants then at each epoch n decisions will be made, one by each DM. In this chapter we look to extend the material put forward up to this point to incorporate the kind of group setting discussed in Section 2.6. In this setting the DMs must reach a collective decision between them, which they will all share the common effects of, be it a positive or negative outcome. The methods considered thus far have all involved the maximisation of expected utility, a technique which involves some probability distribution, and some utility function. In what follows we consider methods by which the n participants may combine their probability distributions to produce one probability distribution (we have seen processes like this before, using weighted sums), and also merge their utility functions to produce one utility function. This merged utility function and combined probability distribution will be used to determine which decision is in some sense optimal for the group as a whole. We begin by discussing the somewhat novel topic of combined utility.

4.1 Combining Utilities

The act of merging together the utility functions of a collection of DMs seems an unintuitive task. The utility function of each individual is a personal construct, which is reflective of their feelings towards making gambles and taking risks. Previously when we discussed methods of combining probabilites the method of taking an arithmetic average was used, with weights being adjusted over time depending on the reliability of the participants. However, if an analogous approach was to be taken to constructing a combined utility function it would only be fair that the weights remain equal (and constant) throughout the procedure, *i.e.*, that all DMs utilities are considered as uniformly important in the decison making task. However there is a problem with doing this with the original utility functions, which we shall now illustrate.

Suppose we have a collection of n DMs, whom we label $P_1, P_2, \ldots, P_n$. Initially (and somewhat naively) we may take a simply arithmetic average of their utility functions, $u_1(r), u_2(r), \ldots, u_n(r)$ to form a combined utility function $u(r)$:

$$u(r) = \frac{u_1(r) + u_2(r) + \ldots + u_n(r)}{n} \tag{15}$$

However a big issue with this choice of function is that it gives a motivation for a participant to lie about their original utility function, in order to ensure that the combined utility function is of a form desirable to them. Consider a simplified case with $n = 2$. Suppose for instance that P_2 knows that the utility function of her neighbour is $u_1(r) = 2r$, and believes that her own utility function for returns is $u_2(r) = 2r^2$. Then if she claims that her utility function is actually $u_2(r) = 4r^2 - 2r$ then this will ensure that the shared utility function found by arithmetic averaging is $u(r) = 2r^2$, which is the utility function which exactly models her preferences. Granted this is a very simple example, and in reality one participant may not know the utility function of another when choosing her own, but it shows that there is a motivation to exaggerate your own utility function to ensure that the combined utility function is closer to your actual preferences. Also, in general, utility functions that are more extreme (*i.e.*, very risk-averse or very risk-seeking) will tend to dominate the combined utility function. We propose a way to avoid this problem, and remove the incentive for a DM to embelish her beliefs.

4.1.1 Rescaled Utilities

In Section 2.4.2 we discussed how utility functions are invariant to an positive linear transformation. We will use this crucial property in what follows. Utilities, for ease of interpretation, are often rescaled to the closed $[0, 1]$ interval, with 0 being the utility value assigned to the worst possible outcome, and 1 being the utility value assigned to the best possible outcome.

Here we rescale all utilty functions to this interval, to ensure that the best possible outcome for P_i will give her the same utility as the best possible outcome occuring for P_j will give to P_j. This is the same for worst possible outcomes also. Hence there is no benefit to a DM in exaggerating their utility function, as rescaling will ensure that the highest value they can receive is 1, the same as all of her neighbours. Now here we need to introduce an assumption, to ensure that this method is possible. We assume, in the decision problems considered, that

Assumption 1: There are at least two distinct outcomes, one of which is at least as good as any other outcome (r^*), and one outcome which is at least as bad as any other outcome (r_*).

Note that in a discrete setting it is often easy to see what the best and worst possible outcomes are, for example when a collection of outcomes occur from a Binomial distributon. However consider the problems arising if our potential results are Normally distributed, as in this case we have an infinite amount of potential return (by the continuity of the Normal distribution). Of course as we progress further into the tails of the distribution and further from the mean the probability of actually observing these results becomes negligibly small. Nevertheless, how do we determine which outcome is the best, and which is the worst? By the properities of the Normal distribution there will always be a slightly better/worse outcome by progressing further from the mean, with a slightly smaller probability of occurence. In these cases we need to apply a cut-off point, above and below, for which we assume the probability of a result occuring in these areas are zero. To put this mathematical, if we have some return such that $r \in (-\infty, \infty)$ theoretically, then we assume that for decision making purposes $r \in [a, b]$ and $P(r \in (-\infty, a)) = P(r \in (b, \infty)) = 0$. Obviously this means that we are discounting some possibly plausible event by saying that it is impossible, but by setting the interval $[a, b]$ to be very wide this problem can be minimised.

We also need to make two further assumptions, both of which are quite rational.

Assumption 2: There is a best possible outcome r^*, for which $u_i(r^*) = 1$ for all $i = 1, 2, \ldots, n$.

Assumption 3: There is a a worst possible outcome r_*, for which $u_i(r_*) = 0$ for all $i = 1, 2, \ldots, n$.

That is to say, of all the potential outcomes, there is one outcome r^* which all participants believe to be the best outcome achievable, and similarly there is one outcome r_* which all participants believe to be the worst outcome achievable. Commonly, especially when dealing in a monetary framework, this will be the case, with rational DMs preferring to have more then less. Hence the (rescaled) utility functions of all DMs can be imagined on a graph. All lines will have the same point of origin, as $u_i(r_*) = 0$ for all $i = 1, 2, \ldots, n$, and will end at the same point as $u_i(r^*) = 1$ for all $i = 1, 2, \ldots, n$. However their paths from one point to the other will be potentially very different, expressing the potential differences between their utility functions.

To conclude, the method we propose is that all participants rescale their utility functions to the $[0, 1]$ interval to form $u_1^*(r), u_2^*(r), \ldots, u_n^*(r)$, and then their combined utility function is an arithmetic average of these rescaled values, *i.e.*,

$$u^*(r) = \frac{u_1^*(r) + u_2^*(r) + \ldots + u_n^*(r)}{n} \tag{16}$$

This will ensure that this function, denoted $u^*(r)$ will be in the $[0, 1]$ interval as desired. This guarantees commensurability (Boutilier, 2003).

4.1.2 Example

We demonstrate the applicaton of the above technique in an example containing three DMs. All have different utility functions, and different initial fortunes, which are given in Table 11 below.

Table 11: Utility functions and initial fortunes of DMs

P_i	$u_i(f_i + r)$	f_i
P_1	r	\$100
P_2	$\exp(\frac{r}{20})$	\$80
P_3	r^2	\$75

Suppose that the return observed from a decision is assumed to be in the interval $[-20, 20]$, *i.e.*, that the worst possible result is a loss of \$20, with the best possible outcome being a gain of \$20, so $r_* = -20$ and $r^* = 20$. Calculations for rescaling are shown below. Note that the initial fortune of each DM is included here, so the rescaled utility is a function of the final monetary position of a DM, and not just the deviation in wealth. This is a straightforward exercise, as all that is involved is the solving of linear simultaneous equations.

$$\begin{aligned} au_1(f_1 + r_*) + b &= 0 \rightarrow \ 80a + b = 0 \\ au_1(f_1 + r^*) + b &= 1 \rightarrow \ 120a + b = 1 \end{aligned}$$

$$\begin{aligned} cu_2(f_2 + r_*) + d &= 0 \rightarrow \ 20.085c + d = 0 \\ cu_2(f_2 + r^*) + d &= 1 \rightarrow \ 148.413c + d = 1 \end{aligned}$$

$$\begin{aligned} eu_3(f_3 + r_*) + f &= 0 \rightarrow \ 3025e + f = 0 \\ eu_3(f_3 + r^*) + f &= 1 \rightarrow \ 9025e + f = 1 \end{aligned}$$

We find that $a = 0.025$, $b = -2$, $c = 0.008$, $d = -0.156$, $e = 0.0002$ and $f = -0.504$. Hence the utility functions of DMs can now be rewritten in rescaled form as

$$\begin{aligned} u_1^*(r) &= 0.025u_1(100 + r) - 2 \\ u_2^*(r) &= 0.008u_2(80 + r) - 0.156 \\ u_3^*(r) &= 0.0002u_3(75 + r) - 504 \end{aligned}$$

Finally we can take the arithmetic average of these rescaled utility functions to find $u^*(r)$ given below. We can show that this combined utility function has the desired properties, namely that $u^*(r_*) = 0$ and $u^*(r^*) = 1$.

$$\begin{aligned} u^*(r) &= \frac{1}{3}u_1^*(r) + \frac{1}{3}u_2^*(r) + \frac{1}{3}u_3^*(r) \\ &= \frac{1}{3}0.025u_1(100 + r) - 2 + \frac{1}{3}0.008u_2(80 + r) - 0.156 + \frac{1}{3}0.0002u_3(75 + r) - 0.504 \\ &= 0.0083(100 + r) + 0.0026\exp\left(\frac{r + 80}{20}\right) + 0.00007(r + 75)^2 - 0.887 \end{aligned}$$

Note that this procedure must be repeated at each decision epoch, assuming that the fortunes of at least one of the DMs increases or decreases as a result of the decision made.

4.2 Combining Probabilites

In the previous chapter we considered how a DM could attempt to learn from the environment that she operates within by listening to the information offered by her neighbours, and re-weighting the degree of confidence she had in these sources of information over time in light of the results observed. Here we want to have one single probability distribution representing the beliefs of all DMs, to be used in conjuction with the combined utility function above, in order to determine which decision maximises expected utility. The problem with some of the previously considered methods is that they were personal to each DM. For example, in the KL method each participant could decide to ignore the views of all others if they wanted to. At each epoch, in the KL and Differing Viewpoints approaches, each participant had their own individual weighted sum, which reflected the subjectivity of their own beliefs and utilities. Below we suggest a method that avoid this problem.

4.2.1 The Plug-in Approach revisited

We have disccused the use of the plug-in approach in Section 3.6. In many ways it seems more appropriate for use in a group context then for single DMs. As a technique for a single DM we criticised its lack of subjectivity, but in a group setting this is actually advantageous. Initially all DMs are allowed to have an equal say, with each participant being given an identical weight. However over time, in light of the new information observed, participants who have shown themselves to be unreliable have the weight decreased, while those who have proved themselves to be accurate have theirs increased. This seems to be a fair way to combine probabilities, as the combined distribution that is used will affect all participants, not just one. Hence we propose that the combined distribution of DMs be given by

$$\hat{f}(\theta) = \alpha_1 f_1(\theta) + \alpha_2 f_2(\theta) + \ldots + \alpha_n f_n(\theta) \tag{17}$$

Initially $\alpha_i = \frac{1}{n}$ for all $i = 1, 2, \ldots, n$ and then the weights are updated over time in the plug-in manner previously discussed. The optimal decison is given to be

$$d^* = \arg\max_d \sum\!\!\!\!\!\!\int u^*(d,\theta)\hat{f}(\theta)d\theta \tag{18}$$

4.2.2 Step-by-Step Methodology

We present here a clear step-by-step methodology of the material presented in the previous section, illustrating how this approach should be implemented in practise.

- *Step 1:* Participants $P_1, \ldots, P_n$ specify their beliefs over an uncertain event of interest θ, via their probability distributions $f_1(\theta), \ldots, f_n(\theta)$, and their utility functions $u_1(r), u_2(r), \ldots, u_n(r)$.

- *Step 2:* The combined probability distribution $\hat{f}(\theta)$ is constructed as an arithmetic average of $f_1(\theta), f_2(\theta), \ldots, f_n(\theta)$, *i.e.*, $\alpha_i = \frac{1}{n}$ for all $i = 1, 2, \ldots, n$.

- *Step 3:* Each DM rescales their utility function to the $[0, 1]$ interval to produce $u_i^*(r)$.

- *Step 4:* The combined utility function $u^*(r)$ is now constructed as an arithmetic average of the rescaled utility functions $u_1^*(r), u_2^*(r), \ldots, u_n^*(r)$.

- *Step 5*: The optimal decision d^* is that which maximises expected utility, as given in Equation (18). The calculation of the expected utility of each decision is found by use of $\hat{f}(\theta)$ from Step 2, and $u^*(r)$ from Step 4.

- *Step 6*: Some outcome is observed from the decision made.

- *Step 7:* All participants use Bayes' Theorem, as in Equation (8), to update their own personal beliefs over the unknown quantity θ.

- *Step 8:* The value of θ observed, θ^*, is plugged into each $f_i(\theta)$ to determine the new weights α_i.

- *Step 9*: Rescaled utility functions are now updated to factor in profits or losses made as a result of the decision made.

- *Step 10:* Steps 3-9 are now repeated for as many epochs as is necessary, or until no new observations can be witnessed, and hence no further updating can occur.

4.3 Example

We return to the three DMs discussed in Section 4.1.2. Their utility functions and initial fortunes are as stated in Table 11, and their combined utility function is given by

$$u^*(r) = 0.0083(100+r) + 0.0026\exp\Big(\frac{r+80}{20}\Big) + 0.00007(r+75)^2 - 0.887$$

Suppose that collectively they must decide if they should enter into a long forward on a particular stock or not. Their beliefs over the parameter θ, pertaining to stock performance and hence to the profit or loss that will will result are given in Table 12 below. It is assumed that θ has unknown mean μ, but known variance $\sigma^2 = 2$, in a manner akin to that from Section 3.9. Steps 1,3 and 4 we have seen previously in Section 4.1.2.

Table 12: Utility functions and initial fortunes of DMs

P_i	$f_i(\theta)$
P_1	$\mathcal{N}(3,1)$
P_2	$\mathcal{N}(1,2)$
P_3	$\mathcal{N}(-2,1)$

Combining these we find in Step 2 that

$$\begin{aligned}\hat{f}(\theta) &\sim \frac{1}{3}\mathcal{N}(3,1) + \frac{1}{3}\mathcal{N}(1,2) + \frac{1}{3}\mathcal{N}(-2,1)\\ &\sim \mathcal{N}\Big(\frac{2}{3},\frac{4}{9}\Big)\end{aligned}$$

With $\hat{f}(\theta)$ and $u^*(r)$ we now have a single probability distribution represented the beliefs of DMs, and a single utility function representing their preferences. DMs may now combine these, using the method of maximising expected utility, to determine which course of action to take, either d_1 (to enter the long forward) or d_2 (to not). We see that in Step 5

$$\begin{aligned}E[d_1] &= \int_{-20}^{20} \frac{0.0083(r+100) + 0.0026\exp\Big(\frac{80+r}{20}\Big) + 0.00007(75+r)^2 - 0.8867}{\sqrt{\frac{8}{9}\pi}} \exp\Big(-\frac{(r-\frac{2}{3})^2}{\frac{8}{9}}\Big)dr\\ &= 0.473\\ E[d_2] &= 0.0083(100) + 0.0026\exp\Big(\frac{80}{20}\Big) + 0.00007(75)^3 - 0.887 = 0.406\end{aligned}$$

We see that the optimal decision $d^* = d_1$. Suppose in Step 6 that a loss of -\$1 is observed. In

Step 7 Normal-Normal conjugacy is used to update the distributions of each DM over unknown θ.

$$
\begin{aligned}
f_1(\mu|r=-1) &\sim \mathcal{N}\Big(\frac{\frac{3}{1}+\frac{-1}{2}}{\frac{1}{1}+\frac{1}{2}}, \frac{1}{\frac{1}{1}+\frac{1}{2}}\Big) \sim \mathcal{N}\Big(\frac{5}{3}, \frac{2}{3}\Big) \\
f_2(\mu|r=-1) &\sim \mathcal{N}\Big(\frac{\frac{1}{2}+\frac{-1}{2}}{\frac{1}{2}+\frac{1}{2}}, \frac{1}{\frac{1}{2}+\frac{1}{2}}\Big) \sim \mathcal{N}(0,1) \\
f_3(\mu|r=-1) &\sim \mathcal{N}\Big(\frac{\frac{-2}{1}+\frac{-1}{2}}{\frac{1}{1}+\frac{1}{2}}, \frac{1}{\frac{1}{1}+\frac{1}{2}}\Big) \sim \mathcal{N}\Big(-\frac{5}{3}, \frac{2}{3}\Big)
\end{aligned}
$$

Now $\theta^* = -1$ is substituted into the original distributions of the DMs to gauge their reliability. The following results ensue, in Step 8 shown in Table 13.

Table 13: Updated weights for $\hat{f}(\theta)$

P_i	$f_i(\theta^*)$	b_i	$\frac{1}{n+b_i}$	α_i
P_1	0.0001336	7485	0.0001335	0.001
P_2	0.171	5.848	0.113	0.446
P_3	0.2419	4.134	0.14	0.553

These are now combined with updated probability distributions give a new weighted sum of probability beliefs, $\hat{f}(\theta)$:

$$
\begin{aligned}
\hat{f}(\theta) &\sim 0.001\mathcal{N}\Big(\frac{5}{3}, \frac{2}{3}\Big) + 0.446\mathcal{N}(0,1) + 0.553\mathcal{N}\Big(-\frac{5}{3}, \frac{2}{3}\Big) \\
&\sim \mathcal{N}(-0.92, 0.403)
\end{aligned}
$$

Now all participants need to update their rescaled utility functions in light of the loss suffered in Step 9. The calculations are omitted, and the resulting functions are just given below.

$$
\begin{aligned}
u_1^*(r) &= 0.025(98+r) - 1.95 \\
u_2^*(r) &= 0.0086\exp\Big(\frac{78+r}{20}\Big) \\
u_3^*(r) &= 0.00017(73+r)^2 - 0.481 \\
u^*(r) &= 0.0083(98+r) + 0.0029\exp\Big(\frac{78+r}{20}\Big) + 0.000057(73+r)^2 - 0.862
\end{aligned}
$$

This updated utility function is combined with $\hat{f}(\theta)$ to produce

$$\begin{aligned}
E[d_1] &= \int_{-20}^{20} \frac{0.0083(98+r) + 0.0029\exp\left(\frac{78+r}{20}\right) + 0.000057(73+r)^2 - 0.862}{\sqrt{0.806\pi}} \exp\left(-\frac{(r+0.92)^2}{0.806}\right) dr \\
&= 0.378 \\
E[d_2] &= 0.0083(98) + 0.0029\exp\left(\frac{78}{20}\right) + 0.000057(73)^3 - 0.862 = 0.4
\end{aligned}$$

We see now that the optimal decision for the group is $d^* = d_2$, *i.e.*, they decide not to enter into the long forward contract. The decision making procedure ends here, as no further information can be observed.

4.4 Conclusion

The problem of determining the optimal decision in a group decision making problem is far more complicated then doing so for a single DM. Here we have attempted to reconfigure the parameters of the problem to use the method of maximisation of expected utility used in the single DM setting, by merging probabilities and utilites of a collection of DMs into a single probability distribution and utility function. The preceeding example details how this may be applied in a realistic context. We need to consider which of Arrow's axioms we are breaking, and how problematic this is. This area of study is an active one, with researchers trying to develop methods of dealing with a problem that occurs worldwide in countless ways on a daily basis.

5 Further Research

In this section we will discuss some areas for future work, most of which are based upon expanding the ideas previously presented in Sections 3 and 4. As a starting point it appears that the techniques which we have developed in Section 3 are quite intuitively aesthetic, and from the example considered it appears that they provide a notable increase in decision quality. However in order to provide a more formal justification for the implementation of these methods it is clear that some class of rigourous proof is required, to serve as an adequate advocation for the use of these approaches. One such idea would be to try to create an axiomatic basis. Another, perhaps more straightforward thought would be to investigate rates of convergence. If we were able to show that the combined belief, $\hat{f}_i(\theta)$ converged to the true value of θ faster then in excess of 50% of the beliefs of individuals, $(f_i(\theta))_{i=1,2,\ldots,n}$, at least half of the time, then this would seem a strong statement in favour of the method proposed. A first step in trying to prove a general result of this nature would be run simulations for various numbers of neighbours with various beliefs and to study the emperical evidence which results.

A very interesting idea would be consideration of our information sharing methods in the context of a social network framework. Could our methods be used to describe propagation of information being passed through a social network? Here each individual would have neighbours of varying degrees depending upon their level of connectivity (friends, friends-of-friends, friends-of-friends-of-friends, *etc.*), and would have different levels of trust depending on the degree of the neighbour from which information was received from. This is clearly a complicated idea, but is undoubtedly an attractive concept, and one which could be applied in a wide range of realistic settings. One such example in a machine-learning context, which we have touched upon before, would be a collection of sensors passing information between each other in order to, for instance, control traffic-flow in a city.

An assumption that we have made in Section 3 is that the beliefs held by individuals are independent of each other. This assumption is made for the sake of convenience, but in many realistic settings may not be the case, with considerable correlation existing between the opinions of a collection of domain-specific experts, which could be factored into our analysis. This would certainly add an extra dimension of complexity to the problem, but could potentially lead to a significant increase in the accuracy of our modelling. Further thought will need to be had, and reading to be done, to begin considering how such a covariance structure between experts could be estimated.

In the preceding two sections we have always assumed that all DMs are able to supply a fully probabilistic distribution over some uncertain event of interest. However, especially for those not fluent in the language of probability, this may not always be an achievable goal. Hence we have an interest in trying to translate the previously developed methods to an analogous non-parametric framework, and indeed we have already begun some preliminary research into this. We consider a setting where each individual i supplies their belief about some uncertain parameter in the form of an interval $[l_i, u_i]$, *i.e.*, they place lower and upper bounds on the parameter of interest, creating an interval which they believe the true value lies within. How can individuals' combine beliefs of this nature? How can an individuals own belief, and their belief in the reliability of their neighbours, be updated in light of evidence witnessed? It is clear that the width of an interval has an important role to play in this process, similar to that of the variance supplied in a probabilistic framework. Tentative early efforts have been made to solve these problems, and more attention will certainly be directed towards them in the future.

In several of our techniques, in both Sections 3 and 4, the Laplacian Principle of Indifference is invoked at the first epoch of a decision problem, assuming all individuals to be equally reliable. Is there some way in which prior information could be filtered into this process? For instance in an academic context the initial weights could correspond to the number of papers published by each individual, as this could be seen as a measure of their relative expertise. Such priors could prove to greatly increase decision quality in a problem with a small (potentially single) amount of epochs, in which the learning that results from the outcomes witnessed will not have time to be used in the analysis. Investigation could be done into seeing if such suitable priors could be developed, although intuitively it seems that these may only be constructed on a specific case-by-case basis, as opposed to a generic formulation covering all possible situations.

Traditionally Bayesian statistics involve use of the belief that the posterior distribution is proportional to the product of the prior distribution and the appropriate likelihood function. In the updating that features in our approaches in Sections 3 and 4, each individual i performs updating on their own belief, $f_i(\theta)$, in a Bayesian manner, but the combined belief $\hat{f}_i(\theta)$ is not modified in the standard Bayesian way. This is due to problems touched upon already with the difficultly of specifying a valid and fitting likelihood function (which would require a covariance structure of the type discussed above). Yet it is evident that there is something somewhat Bayesian going on here, in that there is a combined belief, then some evidence (data) is witnessed, and the new combined belief is influenced by this data. Further research needs to be carried out with regards to the almost philisophical interpretation of how Bayes' Theorem is incorporated in our methods, and if we are using some type of "semi-Bayesian" updating.

While Section 3 deals with an individual who is making a decision, the consequences of which will pertain solely to themselves, Section 4 expanded on this to a group setting, in which a single decision must be made, which somehow represents the beliefs and utilities of all participants, the consequences of which will be shared by all of them. Evidently we need to expand on this, but it can serve as a starting point. Interest potentially lies in the consideration of a hierarchy among decision makers, of the style of that discussed previously by Karny & Kracik (2003). Can we extend our existing framework to incorporate a benevolent dictator (*i.e.*, a facilitator), who must combine the beliefs and desires of her subordinates with her own, in a way which is fair, but also assimilates the fact that she may have access to more information then them, and that their desires may be selfishly motivated? This would certainly be a cause worthy of investigation, as complicated cases like this are a regular occurence in real-life.

Finally, all of the methods which we have discussed in Sections 3 and 4 are myopic, in that decisions are made in a singular fashion, *i.e.*, some choice is made, some evidence witnessed, another decision is made, *etc.*. Yet sequential decision problems, as mentioned in Section 2.4.3, occur naturally in many realistic settings. Can any of our ideas be extended to a sequential setting? One idea we have concerns the approach developed in Section 4. The methods discussed in Houlding (2008), can be used to solve sequential decision problems, where a DM has

Normally distributed beliefs and a polynomial utility function, in a quick and accurate manner. In a case where all our DMs had Normally distributed beliefs then their combined belief would also be Normal, and if they had polynomial utility functions (or functions that could be approximated by polynomials via Taylor expansions) then the weighted sum of these would itself be polynomial, making it a potentially suitable problem to be tackled by the method developed by Houlding. Investigation will be carried out to see if this "roll-back" method can be applied here, and if so it can be seen as a valuable extension of the existing technique.

6 Bibliography

Alghalith, M. (2012). Forward Dynamic Utility Functions: A New Model and New Results. *European Journal of Operational Research*, 223(3), 842-845.

Allais, M. (1953). The Foundations of a Positive Theory of Choice involving Risk and a Criticism of the Postulates and Axioms of the American School (translation). In M. Allais and O. Hagen, editors, *Expected Utility Hypothesis and the Allais Paradox*, 27-145. D. Reidel Publishing Company.

Arrow, K. J. (1950). A Difficulty in the Concept of Social Welfare. *Journal of the Political Economy*, 58(4), 328-346.

Arrow, K. J. (1965). The theory of Risk Aversion. *Aspects of the Theory of Risk Bearing*, Markham Publ. Co., Chicago, 90-109.

Aven, T. & Guikema, S. (2011). Whose Uncertainty Assesments (Probability Distributions) does a Risk Assesment Report: The Analyst's or the Experts'? Reliability Engineering and System Safety, 96(10), 1257-1262.

Ben-Haim, Y. (2006). Info-Gap Decision-Theory: Decisions under Severe Uncertainty, Academic Press, London.

Ben-Haim, Y. (2007). Info-Gap Robust-Satisficing and the Probability of Survival. DNB Working Papers, Netherlands Central Bank, Research Department.

Bernoulli, D. (1738). *Comentarii Academiae Scientiarum Imperialis Petropolitanae*. Exposition of a New Theory on the Measurement of Risk (translation, 1954). *Econometrica*, 22, 22-36.

Boutilier, C (2003). On the Foundations of Expected Expected Utility. 18th International Joint Conference on AI, 712-717.

Buchanan, J. M. (1954). Social Choice, Democracy and Free Markets. *Journal of Political Economy*, 62(2), 334-343.

Buchanan, J. M. & Tullock, G. (1962). The Calculus of Consent. Ann Arbor, University of Michigan Press.

Buchanan, J. M. (1979). What should Economists do? Liberty Press, Indianapolis.

Buchanan, J. M. (1994a). Foundational Concerns: A Criticism of Public Choice Theory. Unpublished Manuscript presented at the European Public Choice Meeting, Valenia, Spain.

Buchanan, J. M. (1994b). Dimensionality, Rights and Choices among Relevant Alternatives. Unpublished Manuscript presented at a meeting honouring Peter Bernholz, Basel, Switzerland.

Cantor, G. (1874). Uber eine Eigenschaft des Inbergrifes aller reellen algebraischen Zahlen. *J. Reine Agnew, Math*, 84, 242-258.

Chajewska, U., Koller, D. & Parr, R. (2000). Making Rational Decisions using Adaptive Utility, Elicitation, Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-00), 363-369.

Clemen, R. T. & Winkler, R. L. (1999). Combining Probability Distributons from Experts in Risk Analysis. *Risk Analysis*, 19, 187-203.

Cooke, R. (1991). Experts in Uncertainty. *Opinion and Subjective Probability in Science, Environmental Ethics and Science Policy Series*, Oxford University Press.

Cooke, R. (2007). Expert Judgement Studies. *Reliability Engineering and System Safety*, 93, 655-677.

Cyert, R. M. & DeGroot, M. H. (1975). Adaptive Utility. *Adaptive Economic Models*, 223-246, eds. R.H. Day & T. Groves, Academy Press, New York.

Dalkey, N. C. (1969). The Delphi Method: An Experimental Study of Group Opinions. Report No. RM-5888-PR. The Rand Corporation.

Delbecq, A. L., van de Ven, A. H. & Gustafson, D. H. (1975). Group Techniques for Program Planning. Glenview, IL: Scott Foresman.

Duggan, J. & Schwartz, T. (1992). Strategic Manipulability is Inescapable: Gibbard-Satterthwaite without Resoluteness. Division of the Humanities and Social Sciences, California Institute of Technology, Social Science Working Paper 817.

Flandoli, F., Giorgi, E., Aspinall, W. P. & Neri, A. (2011). Comparision of a New Expert Elicitation Model with the Classical Model, Equal Weights and Single Experts, Using a Cross-Validation Technique. Reliability Engineering and System Safety, 96 (10), 1292-1310.

French, S. (1994). Utility: Probability's Younger Twin? *Aspects of Uncertainty, A Tribute to D.V. Lindley*, Chapter 12, 171-180.

French, S. (2011). Aggregating Expert Judgement. Revista de la Real Academia de Ciencias Exactas, Fisicas y Naturales. Serie A. Matematicas, 105(1), 181-206.

Gibbard, A. (1973). Manipulation of Voting Schemes: A General Result. *Econometrica*, 41(4), 587-601.

Goldstein, M. & Wooff, D. A. (2007). Bayes Linear Statistics: Theory and Methods. Chichester: John Wiley.

Halmos, P. R. (1974). Measure Theory. New York, Springer-Verlag.

Harsanyi, J. (1955). Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisions of Utility. *Journal of Political Economy*, 63, 309-321.

Houlding, B. (2008). Sequential Decision Making with Adaptive Utility. PhD Thesis, Appendix C.

Houlding, B. & Coolen, F. P. A. (2011). Adaptive Utility and Trial Aversion, Journal of Statistical Planning and Inference, 141(2), 734-747.

Houlding, B. & Coolen, F. P. A. (2012). Nonparametric Predictive Utility Inference. European Journal of Operational Research, 221(1), 222-230.

Kahneman, D. & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2), 263-291.

Karny, M. and Guy, T. V. (2004). On Dynamic Decision-Making Scenarios with Multiple Participants. *Multiple Participant Decision Making*, Advanced Knowledge International, 17-28.

Karny, M. & Guy, T. V. (2009). Cooperation via sharing of probabilistic information. *Int. J. Computational Intelligence Studies*, 1(2), 139-162.

Karny, M. & Kracik, J. (2003). A Normative Probabilistic Design of a Fair Government Decision Strategy. *Journal of multi-criteria Decision Analysis*, 12, 111-125.

Kolmogorov, A. N. (1950). Foundations of the Theory of Probability (translation). Chelsea Publishing Company.

Kullback, S. & Leibler, R. (1951). On Information and Suffciency. *Annals of Mathematical Statistics*, 22, 79-87.

Laplace, S. (1812). Theorie analytique des probabilites. Paris, Ve. Courcier.

Lindley. D. V. (1991). Making Decisions (2 ed.). Wiley & Sons Ltd.

May, K. O. (1952). A Set of Independent Necessary and Sufficient Conditions for Simply Majority Decision. *Econometrica*, 20(4), 680-684.

Merton, R. C. (1971). Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*, 3, 373-413.

Nash, J. F. (1951). Non-cooperative Games. *Annals of Mathematics*, 54, 289-295.

O'Hagan, A. (1998). Eliciting Expert Beliefs in Substantial Practical Applications, *Journal of the Royal Statistical Society, Series D*, 47(1), 21-35.

Pratt, J. W. (1964). Risk Aversion in the Small and in the Large. *Econometrica*, 32 122-136.

Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics. *The American Economic Review*, 83(5), 1281-1302.

Rabin, M. (2000). Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica*, 68(5), 1281-1292.

Satterthwaite, M. A. (1975). Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory*, 10, 87-217.

Sen, A. K. (1970). The Impossibility of a Paretian Liberal. *Journal of Political Economy*, 78(1), 152-157.

Sen, A. K. (1979). Personal Utilities and Public Judgements: or What's Wrong with Welfare Economics? *The Economic Journal*, 89, 537-558.

Sen, A. K. (1990). Rational Fools: A Critique of the Behavorial Foundations of Economic Theory. *Beyond Self Interest*, Chapter 2, University of Chicago Press, Chicago.

Sen, A. K. (1995). Rationality and Social Choice. *The American Economic Review*, 85(1), 1-24.

Stirling, W.C. (2004). Satisficing Games for Multiple-Participant Coordinated Decision Making. TED ESF Workshop, CMP'04: Multiple Participants Decision Making, Institute of Information Theory and Automation Academy of Sciences of the Czech Republic .

von Neumann, J. & Morgenstern, O. (1944). Theory of Games and Economic Behaviour. Second Edition, Princeton University Press.

Walley, P. (1991). Statistical Reasoning with Imprecise Probabilities. Chapman and Hall.

Wisse, B., Bedford, T. & Quigley, J. (2008). Expert Judgement Combination using Moment Methods. *Reliability Engineering and System Safety*, 93, 675-686.

Zadeh, L. A. (1965). Fuzzy Sets. *Information and Control*, 8(3), 338-353.

7 Appendices

7.1 Appendix A

Below we present the notation used throughout this report.

Single DM

- $\mathcal{D}$ - the set of admissible decisions.
- $d_1, d_2, \ldots, d_n \in \mathcal{D}$ - a collection of n potential admissible decisions.
- $d^* \in \mathcal{D}$ - the decision which maximises expected utility.
- $\mathcal{R}$ - the set of potential decision returns.
- $r_1, r_2, \ldots, r_n \in \mathcal{R}$ - a collection of n possible decision returns.
- r^*, r_* - the best and worst outcomes respectively in $\mathcal{R}$.
- θ - some unobservable unknown parameter of interest.
- Θ - the set of possible states of nature.
- $\theta_1, \ldots, \theta_m$ - a collection of m possible states of nature.
- $f(\theta)$ - the probability distribution of a DM over θ.

Single DM with information sharing

- P_i - the i^{th} DM.
- $f_i(\theta)$ - the probability distribution of P_i over θ.
- $\hat{f}_i(\theta)$ - the probability distribution of P_i after listening to her neighbours.
- $u_i(r)$ - the utility function of P_i.
- $\alpha_{i,j}$ - the weight given by P_i to P_j

Multiple DM Problem

- $u_i^*(r)$ - the utility function of P_i rescaled to the $[0, 1]$ interval.
- $u^*(r)$ - the combined utility function of the n participants.
- $\hat{f}(\theta)$ - the combined probability distribution of the n participants.
- α_i - the weight given to P_i in $\hat{f}(\theta)$

7.2 Appendix B

This appendix contains the proof of the fact that the pairwise interative sharing of distributions framework, which was suggested by Karny & Guy (2004), is not invariant to the order of sharing.

We consider the case where we have three participants, P_1, P_2 and P_3, who have distributions $f_1(\theta), f_2(\theta)$ and $f_3(\theta)$ respectively over unknown θ. We consider two different sharing orderings here, and observe different results in both cases.

Ordering 1

- *Stage 1*: P_1 and P_2 sharing.

$$\begin{aligned}\hat{f}_1(\theta) &= \alpha_1 f_1(\theta) + (1-\alpha_1) f_2(\theta)\\ \hat{f}_2(\theta) &= \alpha_2 f_2(\theta) + (1-\alpha_2) f_1(\theta)\end{aligned}$$

- *Stage 2*: P_1 and P_3 sharing.

$$\begin{aligned}\hat{f}_1(\theta) &= \alpha_1 \hat{f}_1(\theta) + (1-\alpha_1) f_3(\theta)\\ &= \alpha_1^2 f_1(\theta) + \alpha(1-\alpha_1) f_2(\theta) + (1-\alpha_1) f_3(\theta) \qquad (19)\\ \hat{f}_3(\theta) &= \alpha_3 f_3(\theta) + (1-\alpha_3)\hat{f}_1(\theta)\\ &= \alpha_3 f_3(\theta) + \alpha_1(1-\alpha_3) f_1(\theta) + (1-\alpha_3)(1-\alpha_1) f_2(\theta)\end{aligned}$$

- *Stage 3*: P_2 and P_3 sharing.

$$\begin{aligned}\hat{f}_2(\theta) &= \alpha_2 \hat{f}_2(\theta) + (1-\alpha_2)\hat{f}_3(\theta)\\ &= f_2(\theta)\Big(\alpha_2^2 + (1-\alpha_3)(1-\alpha_1)(1-\alpha_2)\Big) + f_3(\theta)\Big(\alpha_3(1-\alpha_2)\Big)\\ &\quad + f_1(\theta)\Big(\alpha_2(1-\alpha_2) + \alpha_1(1-\alpha_3)(1-\alpha_2)\Big) \qquad (20)\\ \hat{f}_3(\theta) &= \alpha_3 \hat{f}_3(\theta) + (1-\alpha_3)\hat{f}_2(\theta)\\ &= f_2(\theta)\Big(\alpha_3(1-\alpha_3)(1-\alpha_1) + \alpha_2(1-\alpha_3)\Big)\\ &\quad + f_1(\theta)\Big(\alpha_1\alpha_3(1-\alpha_3) + (1-\alpha_3)(1-\alpha_2)\Big) + f_3(\theta)\Big(\alpha_3^2\Big) \qquad (21)\end{aligned}$$

Ordering 2

- *Stage 1*: P_2 and P_3 sharing.

$$\begin{aligned}\hat{f}_2(\theta) &= \alpha_2 f_2(\theta) + (1-\alpha_2) f_3(\theta)\\ \hat{f}_3(\theta) &= \alpha_3 f_3(\theta) + (1-\alpha_3) f_2(\theta)\end{aligned}$$

- *Stage 2*: P_1 and P_3 sharing.

$$\begin{aligned}\hat{f}_1(\theta) &= \alpha_1 f_1(\theta) + (1-\alpha_1)\hat{f}_3(\theta)\\ &= \alpha_1 f_1(\theta) + \alpha_3(1-\alpha_1) f_3(\theta) + (1-\alpha_1)(1-\alpha_3) f_2(\theta)\\ \hat{f}_3(\theta) &= \alpha_3 \hat{f}_3(\theta) + (1-\alpha_3) f_1(\theta)\\ &= \alpha_3^2 f_3(\theta) + \alpha_3(1-\alpha_3) f_2(\theta) + (1-\alpha_3) f_1(\theta) \qquad (22)\end{aligned}$$

- *Stage 3*: P_2 and P_3 sharing.

$$\begin{aligned}\hat{f}_1(\theta) &= \alpha_1 \hat{f}_1(\theta) + (1-\alpha_1)\hat{f}_2(\theta)\\ &= f_1(\theta)\Big(\alpha_1^2\Big) + f_2(\theta)\Big(\alpha_1(1-\alpha_1)(1-\alpha_3) + \alpha_2(1-\alpha_1)\Big)\\ &\quad + f_3(\theta)\Big(\alpha_1\alpha_3(1-\alpha_1) + (1-\alpha_1)(1-\alpha_2)\Big) \qquad (23)\\ \hat{f}_2(\theta) &= \alpha_2 \hat{f}_2(\theta) + (1-\alpha_2)\hat{f}_1(\theta)\\ &= f_1(\theta)\Big(\alpha_1(1-\alpha_2)\Big) + f_2(\theta)\Big(\alpha_2^2 + (1-\alpha_2)(1-\alpha_1)(1-\alpha_3)\Big)\\ &\quad + f_3(\theta)\Big(\alpha_2(1-\alpha_2) + \alpha_3(1-\alpha_2)(1-\alpha_1)\Big) \qquad (24)\end{aligned}$$

It is clear by examination of the coefficents involved that Equation (18) is not equal to Equation (22), that Equation (19) is not equal to Equation (23) and that Equation (20) is not equal to Equation (21), *i.e.*, the beliefs of P_1, P_2 and P_3 respectively differ depending on the order which they receive information from their neighbours.

7.3 Appendix C

This is an inductive proof that, in the Plug-in Approach, $\hat{f}_i(\theta)$ is the same for $i = 1, \ldots, n$ at each individual epoch m.

- **Proof for** $m = 1$: At the first epoch it is assumed, by the Laplacian Principle of Indifference, that each DM is as reliable as any other, and hence $\alpha_{i,j} = \frac{1}{n}$ for all $i, j = 1, \ldots, n$. The combined belief of P_i is given by Equation (25) which will clearly be the same for each DM as they have all been told the same distributions $f_1(\theta), \ldots, f_n(\theta)$.

$$\hat{f}_i(\theta) = \frac{1}{n} f_1(\theta) + \ldots + \frac{1}{n} f_n(\theta). \tag{25}$$

- **Assumption for** $m = k$: At the k^{th} epoch we assume that $\hat{f}_i(\theta)$ is the same for each individual. Now the combined identical belief for each P_i is given by

$$\hat{f}_i(\theta|H) = \alpha_{i,1} f_1(\theta|H) + \ldots + \alpha_{i,n} f_n(\theta|H). \tag{26}$$

- **Proof for** $m = k + 1$. We are assuming that distributions are identical at the k^{th} epoch. Now, following the decision made at epoch k, some observation θ^* is made, and the probability attached by each individual's distribution to this outcome is calculated, *i.e.*, $f_1(\theta^*), \ldots, f_n(\theta^*)$. Note that these values will yield the same results regardless of which DM is performing the calculation, as this is an objective process. Equation (10) is used to update weights, giving the combined distribution of P_i in Equation (26). This will be the same for $i = 1, \ldots, n$ as the weights are all calculated in the same way for all DMs and all participants have been told the same updated distributions $f_1(\theta|H), \ldots, f_n(\theta|H)$. Therefore, at the $k + 1^{th}$ epoch all DMs have identical combined beliefs, as required.

$$\hat{f}_i(\theta|H) = \alpha^*_{i,1} f_1(\theta|H) + \ldots + \alpha^*_{i,n} f_n(\theta|H). \tag{27}$$

7.4 Appendix D

- **Proposition 1**: $w_{i,j} \geq 0$ for all $i = 1, \ldots, n$ and $j = 1, \ldots, n$.

  Proof: $w_{i,j} = |u_i(r) - E[u_{i|j}(d^*)]| \geq 0$. This is trivially true by the non-negativity of absolute values.

- **Proposition 2:** $w_{i,j} = 0$ if and only of $u_i(r) = E[u_{i|j}(d)]$.

  Proof: $w_{i,j} = 0$ implies that $|u_i(r) - E[u - i|j(d^*)]| = 0$, which is only possible when $u_i(r) = E[u_{i|j}(d^*)]$, by elementary properties of absolute values. Conversely if $u_i(r) = E[u_{i|j}(d^*)]$ then this implies $|u_i(r) - E[u_{i|j}(d^*)]| = 0$ and hence $w_{i,j} = 0$.

- **Proposition 3:** If $\alpha_{i,j} < \alpha_{i,k}$ and $w_{i,j} > w_{i,k}$ then $\alpha^*_{i,j} < \alpha^*_{i,k}$ and $\alpha^*_{i,k} - \alpha^*_{i,j} > \alpha_{i,k} - \alpha_{i,j}$. . In non-mathematical language this states that if P_i initially considers P_k to be more reliable then P_j, and then P_k shows herself to more reliable then P_j at the following decision epoch, then P_i will still consider P_k to be more reliable then P_j in her next weighted sum, and the magnitude between the weights will have increased.

  Proof: $\alpha_{i,j} < \alpha_{i,k}$, which implies $(\alpha_{i,j})^{-1} > (\alpha_{i,k})^{-1}$. Given $w_{i,j} > w_{i,k}$ this means $(\alpha_{i,j})^{-1} + w_{i,j} > (\alpha_{i,k})^{-1} + w_{i,k}$, and hence $\frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}} < \frac{1}{(\alpha_{i,k})^{-1} + w_{i,k}}$. Division by the same figure on both sides of the inequality will lead to no changes, giving us $\alpha^*_{i,j} = \frac{\frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}}{\sum_{j=1}^{n} \frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}} < \frac{\frac{1}{(\alpha_{i,k})^{-1} + w_{i,k}}}{\sum_{j=1}^{n} \frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}} = \alpha^*_{i,k}$ as required. That $\alpha^*_{i,k} - \alpha^*_{i,j} > \alpha_{i,k} - \alpha_{i,j}$ is evident from the defintions of $\alpha^*_{i,j}$ and $\alpha^*_{i,k}$ and the fact that $w_{i,j} > w_{i,k}$.

- **Proposition 4:** If $\alpha_{i,j} = \alpha_{i,k}$ and $w_{i,j} = w_{i,k}$ then $\alpha^*_{i,j} = \alpha^*_{i,k}$. This states that if P_i initially considers P_j and P_k to be equally reliable, and then P_j and P_k show themselves to be equally reliable at the following decision epoch, then P_i will still consider P_j and P_k to be equally reliable in her next weighted sum.

  Proof: $\alpha_{i,j} = \alpha_{i,k}$, which implies $(\alpha_{i,j})^{-1} = (\alpha_{i,k})^{-1}$. Given $w_{i,j} = w_{i,k}$ this means $(\alpha_{i,j})^{-1} + w_{i,j} = (\alpha_{i,k})^{-1} + w_{i,k}$, and hence $\frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}} = \frac{1}{(\alpha_{i,k})^{-1} + w_{i,k}}$. Division by the same figure on both sides of the inequality will lead to no changes, giving us $\frac{\frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}}{\sum_{j=1}^{n} \frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}} = \frac{\frac{1}{(\alpha_{i,k})^{-1} + w_{i,k}}}{\sum_{j=1}^{n} \frac{1}{(\alpha_{i,j})^{-1} + w_{i,j}}}$, *i.e.*, $a^*_{i,j} = \alpha^*_{i,k}$ as required.

- **Proposition 5:** If $\alpha_{i,j} < \alpha_{i,k}$ and $w_{i,j} < w_{i,k}$ then any of the following may occur, depending on the initial differences between the weights, and the updated difference between the reliability measures:

  - $\alpha^*_{i,j} < \alpha^*_{i,k}$ in which case $\alpha^*_{i,k} - \alpha^*_{i,j} < \alpha_{i,k} - \alpha_{i,j}$.
  - $\alpha^*_{i,j} = \alpha^*_{i,k}$
  - $\alpha^*_{i,j} > \alpha^*_{i,k}$

  This states that if P_i initially considers P_k to be more reliable P_j, and then P_j shows herself to more reliable then P_k at the following decision epoch, then in the next weighted sum any of three things may occur:

  - P_i will still consider P_k to be more reliable then P_j.
  - P_i will consider P_j and P_k to be equally reliable.
  - P_i will now consider P_j to be more reliable then P_k.

  Proof: Some simple examples illustrate the proof of this proposition. Consider the case where $\alpha_{i,j} = 0.4$ and $\alpha_{i,k} = 0.6$.

  - Suppose $w_{i,j} = 1$ and $w_{i,k} = 1.5$. This leads to $\alpha_{i,j} = 0.475$ and $\alpha_{i,k} = 0.529$, *i.e.*. $\alpha^*_{i,j} < \alpha^*_{i,k}$.
  - Suppose $w_{i,j} = 0.5$ and $w_{i,k} = \frac{4}{3}$. This leads to $\alpha_{i,j} = 0.5$ and $\alpha_{i,k} = 0.5$, *i.e.*. $\alpha^*_{i,j} = \alpha^*_{i,k}$.
  - Suppose $w_{i,j} = 1$ and $w_{i,k} = 2$. This leads to $\alpha_{i,j} = 0.511$ and $\alpha_{i,k} = 0.489$, *i.e.*. $\alpha^*_{i,j} > \alpha^*_{i,k}$.

  These examples show that the three different relations between updated weights are possible in this case.