Reliability Updating in Linear Opinion Pooling for Multiple Decision Makers.

A thesis submitted to the University of Dublin, Trinity College
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Department of Statistics, Trinity College Dublin

March 2016

Donnacha Bolger

I dedicate this thesis to the memory of my friend Faith Wilder,

5th of September 1985 - 21st of February 2015,

who changed my life and is sorely missed.

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

Donnacha Bolger

Dated: March 9, 2016

Abstract

Accurate information sources are vital prerequisites for good decision making. In this thesis we consider a multiple participant setting, where all decision makers (DMs) have a collection of neighbours with whom they share their beliefs about some common relevant uncertain quantity. When determining which course of action to follow a DM takes into account all the information received from her neighbours. Over time, in light of the returns observed from choices made, DMs update their own beliefs over the uncertain event, and also adjust the degree of consideration that they afford to the opinions of each neighbour based on the level of reliability that the information they provide is ascertained to have. Much of this thesis is concerned with constructing a method that incorporates both of these learning facets in a dynamic fashion. This technique, termed the Plug-in approach, is motivated and derived, and attempts are made to justify its use by consideration of some attractive properties it obeys, in addition to studies conducted using both simulated and real data which compared its performance to some rational alternatives. Generalisations of this method are also provided to a setting where DMs specify their opinions nonparametrically rather than using probability distributions, as well as in a group setting where utilities as well as opinions must be amalgamated. Two subjective approaches are also briefly discussed, before we conclude with numerous suggestions for further research in this field.

Acknowledgements

First and foremost I wish to extend my heartfelt thanks to my supervisor, Dr. Brett Houlding. Throughout my Ph.D. tenure he has been constantly encouraging, and extremely helpful at indicating which ideas deserved further consideration and which should be forgotten about. Brett guided me through my first foray into research, always quick to offer an insightful remark and never too busy for one of my frequent trips across the corridor to his office. He generously funded me to travel to conferences, allowing me to learn, to meet other statisticians, and visit exciting locations (and Kildare). Brett demonstrated that “a Liverpool fan whose company I actually enjoy” is not necessarily a guaranteed oxymoron. I could not have asked for a better supervisor than him.

I next wish to thank the staff of the Statistics Department, who have constantly offered advice when I sought it and been incredibly supportive of my research over the last three years. I’ve greatly enjoyed having a cup of tea with them every day. Special gratitude goes to Simon Wilson, who both supervised my undergraduate final year project (and hence kindled my passion for statistical research) and pointed me towards Brett when I mentioned my contemplation of postgraduate study. Without his encouragement it is doubtful if I would have embarked down the Ph.D. road.

I feel tremendously fortunate regarding the company I have been able to keep during my time in Trinity, in terms of the postdocs and fellow Ph.D. students I’ve worked alongside. I consider Angela, Gernot, Cristina, Joy, Sean, Arthur and Arnab all to be great friends. It was a great pleasure to share an office with Shane O’Meachair, Shuaiwei Zhou and Thinh Doan - all of whom are true Dubliners in different ways, and who make for great company and constant fun. Finally Louis Aslett and Jason Wyse were two wonderful mentor figures, doling out insights and counsel about the world of research and academia both inside and outside of the Lincoln Inn on Friday evenings.

Away from Trinity I want to thank my father Dermot and my brother Diarmuid -

they're my homeboys at home, who have supported me constantly over the last three years. After a tough day at work it was always nice to have an evening of hanging out with them to look forward to. I am not a superstitious person, but a lot of stars aligned extremely conveniently for me both in the lead-up to and during my time as a Ph.D. student, and the unscientific part of me would like to think that my late mother, Bernie, had a non-negligible impact upon this. A lot of the effort that I put into my work was in the hope of doing her memory proud, and I'd like to think that in some small way I achieved that.

Finally, it is utterly inconceivable that this thesis would have come into being without the ever-present love and encouragement of Livi Flynn. Several times over the past few years I felt that the Ph.D. challenge was too great for me to handle - a bottle of Brooklyn Lager and slice of pepperoni pizza with her invariably righted my wrongs. She has more goodness and patience than anyone I have ever met, and is unquestionably my (statistically significantly) better half.

Donnacha Bolger

Trinity College Dublin

March 2016

Contents

List of Tables

List of Figures

Chapter 1

Introduction

When individuals make decisions they generally do so in the face of uncertainty. Decision makers (DMs, who are assumed feminine throughout this thesis) invariably have some unsureness about the underlying process governing their decision environment, *e.g.*, they may not know the probability of a medical operation being successful, or the number of cars that will pass along a motorway during rush hour. Generally the language of probability is that used to express this uncertainty. It seems intuitive that there is a connection between the accuracy of the relevant information that a DM possesses and the satisfaction that she will derive from the outcome resulting from the decision that she makes. In this thesis we conjecture that it is beneficial for a DM to absorb information from as many distinct sources as possible, and to assimilate these into her decision process, incorporating additional knowledge into her task. Having made a decision, a return is witnessed. We are concerned with two types of learning that can subsequently occur. Firstly, a DM can update her own opinion about the inherent decision uncertainty in light of this new evidence that has become available to her. Secondly, she can reassess the respective perceived reliability of her various information sources, who (in the context that we are concerned with) are a collection of fellow non-competing DMs. Much of the work which follows aims at developing a rational methodology that facilitates these two forms of learning and supplies a DM with as accurate an opinion as possible to use in her decision task, as well as providing justifications for the use of this approach in practice. We primarily assume individuals specify their uncertainty via probability distributions, but we also provide an analogous nonparametric method achieving these goals.

Above we have provided a broad outline of the original research that is contained within this thesis. Below we supply detailed comments on the contextual setting that this work is placed within, clearly outline the primary aims of the study conducted, and highlight the research methodology that was adhered to throughout.

1.1 Contextual Setting

Every day important decisions with long-term repercussions are made. The United Nations must reach resolutions on what actions to take concerning global conflicts, governments of countries must determine how capital should be budgeted across their departments, medical organisations must choose which drug trials should be funded and which should not. In scenarios such as those illustrated above it seems unwise for this decision to be made by a single individual, or at least to be made based solely on the judgment held by one. A collection of people will generally possess a broader knowledge span that a single person, and hence more pertinent information can be incorporated into a decision making task by considering the opinions of a multiple individuals. Additionally, on a practical basis, a crucial decision being made by one person alone would lead to a great deal of accountability being placed on the shoulders of that individual, something that can be lessened by a more collective process. Once we concur that decision making via the amalgamation of several points of view holds certain advantages over the alternative, the obvious question concerns the manner in which such a decision should be made. In practical application there are numerous methodologies that are likely to be employed.

Majority rule is an extremely straightforward approach, in which individuals vote for their most favoured choice from a set of possible actions, with the action receiving the greatest number of votes deemed to be best (in some sense) by the collective. This is termed an ordinal decision scheme, with all emphasis explicitly placed on the ranking order of decisions, rather than the degrees of preference inherent within individual rankings and the ranking of the group as a whole. We shall return later in thesis (notably in Section 5.4) to discuss the issues and potential pitfalls entailed within schemes of this nature, but only comment now that they may not provide a complete picture of the true attitudes and outlooks of those involved regarding the viable alternatives considered.

Different individuals will commonly hold diverse opinions based upon their contrasting degrees of knowledge over the uncertainty inherent within the forthcoming decision to be made. Such contrasts may arise due to their potentially disparate backgrounds; for instance their educational or socio-economic circumstances, how long they have been involved with a particular organisation, if they have participated in a process of this nature before, or if they have a vested interest in seeing a particular decision chosen. Methods exist in which the individuals comprising a group will attempt to bridge their intrinsic knowledge gaps by sharing their respective opinions and their personal underlying rationales that led to their formations. Once they have listened to the thoughts of those around them, individuals may potentially alter their own opinions, perhaps conceding that they previously were ill informed about the topic at hand and deferring to the wisdom of colleagues perceived to be wiser. Having done this, the optimal decision is chosen by the group via discussion, which continues until a collective consensus is happened upon. We shall further discuss methodologies of this nature (known as behavioural methods) in more depth in our literature review of Section 2, but for now only note that they too have complications ingrained within them, perhaps most significantly being their susceptibility to biases and their potential lack of rigour.

An obvious technique to be implemented is arguably the most democratic appearing methodology, in which the opinions of all individuals are given an equal consideration. On the surface there are certainly advantages apparent in approaches of this nature. All participants are being treated equally, with no favouritism evident. The concept of the wisdom of crowds is well known, with an averaging of opinions over an uncertain quantity commonly leading to a more accurate estimate than if this was provided by a single individual. Yet, as we shall discuss in our literature review of Section 2, and numerous times in our findings of Section 4, there are clear foibles entailed in this approach. If some individuals possess beliefs that are extremely inaccurate then these have the potential to outweigh the beliefs of their more accurate peers, hence skewing the collective opinion away from the truth. The greatest strength of this approach is its simplicity of implementation, yet there are downsides that may occur as a result of choosing to apply it.

The clearest shortcoming of this equal weighting scheme is that the knowledge of inaccurate individuals are given as much consideration as that of accurate individuals. The easiest way to circumnavigate this would be to listen solely to the individual who is deemed to possess the most reliable point of view, hence ensuring that unsound opinions are discarded. Yet, even discounting our cautioning above about making decisions based on a sole opinion, this leads to another major complication: how can it be assessed who the most accurate individual is? All members of the collective believe the opinion that they hold is accurate; if not they would not hold it, or would change it to one that they felt was a more accurate reflection of the true state of nature. It is extremely unlikely that any individual will willingly have their opinion discarded entirely from the decision making task, especially considering that they may personally believe it to be the most accurate of those proffered.

1.2 Research Aims

Given the above discussion concerning the contextual setting for our research, this thesis can be seen as having three primary aims. First and foremost we aspire to create an original decision making methodology, which takes steps towards solving the complications that are inherent in the schemes discussed above. We mentioned the shortcomings of decisions being made by a single individual, and hence our approach shall be based upon the composition of a collection of (potentially diverse) opinions and judgments from various sources. We commented on the issues inherent with behavioural techniques, and therefore the scheme that we develop shall be strictly mathematical in its formulation. An equal weighting scheme seems intuitive, but suffers from the fact that accurate opinions may be outweighed by inaccurate opinions. However, at the other end of the spectrum, it seems unwise to listen to a single individual deemed to be most accurate, due to the complications in choosing such an individual, as well as the aforementioned problem of basing a decision on a sole opinion. Hence our objective is to create a method that accounts for the opinions of all individuals in an unequal fashion. We aim to weight the opinions of individuals, and to set these weights as proportional to the perceived level of knowledge of participants. Of course, as we shall see in great detail in Section 2.7, this broad concept is not itself a novel idea, but we will expand over the course of this thesis on reasons why our developed technique is indeed truly original, and the advantages it can be seen to hold over existing alternatives.

Once we have developed a methodology of the nature discussed above our second aim is to provide justification for its use in practical contexts. We want to be able to pragmatically advocate for its application in realistic decision scenarios. It is our desire for our approach to be mathematically logical, in the sense that it obeys properties that a rational decision maker would deem important for an internally coherent decision process to adhere to. Hence we intend to investigate the relationship between some attractive statistical and mathematical attributes (for instance the Bayesian paradigm) and our proposed methodology, and to confirm that this methodology does indeed meet these criteria. In addition to justification by mathematical argument we also want to provide data-driven validation, in order to highlight the practical merits of our approach. We mentioned above that equal weighting, listening solely to the individual deemed to be most reliable in a collective, and not taking any other opinions into account at all are three frequently applied practices. We plan to compare the performance of our methodology to each of these alternatives (using a suitably derived comparative metric) to demonstrate the superiority of our technique. In Chapter 4 we see how this is achieved using a mix of simulated data and real world data. By illustrating the merits of our approach against commonly used alternatives we hope to underline the benefit of applying it in practical applications in order to increase the quality of the decision process.

Our final primary aim is to provide some generalisations to our method, in order to increase its flexibility and consequentially broaden the range of scenarios that it is appropriate for usage in. Individuals may express their uncertainty concerned some unknown commodity of interest in a variety of formulations, from a simple point estimate to a fully parameterised probability distribution. Much of the work in this thesis pertains to this latter case (with justification for this choice contained therein) but we investigate in Chapter 6 if the general outline of our principal methodology can be extended to a significantly simplified setting, potentially increasing applicability (albeit at the potential determent of the resulting decision quality). Arguably a strength of the method that we derive is its rigid objectivity. However in Chapter 7 we develop two alternative approaches that adhere to the same generalised ethos of our primary technique but that allow for different degrees of subjectivity to be entailed in their mechanisms. These modifications, different flavours of our central methodology, allow for our decision making scheme to be applied in a variety of contrasting settings rather than a single specific one.

1.3 Research Methodology

We now progress to detail the research methodology that was followed during the course of our research. Statistical decision theory has a rich history, formally dating back at least as far as the 1730s. In order to be able to write knowledgably about this topic, and the various nuanced subsets thereof, a reading of the seminal texts was required. These are discussed in detail in the literature review of Section 2. We began by consideration of the most basic decision making fundamentals (for instance the notions of probability and utility), before expanding upon these in an incremental fashion to explore increasingly deeper matters (such as decision making under imprecise probabilities and the complications inherent within any group decision making scheme). As indicated above, our foremost aim was in the development of a scheme suitable for use in a decision context consisting of multiple individuals with differing degrees of information about the uncertainty at hand. This concept of combining opinions or judgments itself has an abundant amount of publications attributed to it. In Section 2.7 we thoroughly examine the most notable of these, highlighting ways in which they contrast in their execution from what we desire, as well as detailing the contrasts between them, and the various perceived advantages and disadvantages that they can be said to hold over each other. The comprehensive literature review of Section 2 can thus be seen as having a dual function: it provides the interested reader with enough technical details and knowledge that they should be able to comfortably follow the original research which it precedes, as well as providing adequate motivation for the development of our novel decision making methodology.

It is in Section 3 that we derive this decision making methodology, explaining comprehensively how weights (which represent the perceived accuracy of the decision makers they are attached to) are calculated and updated over time, as well as how users may modify their opinions in light of new information witnessed in a statistically coherent fashion. Having produced our approach our desire was then to provide some formal justification for it. We mentioned above that our approach is strictly mathematical, rather than behavioural, and hence an obvious action was to assess its performance in relation to some mathematical principles, and to see if it adhered to these. Crucially, Section 2.6 contains discussion on how it is formally impossible to create a scheme that can adhere to the entirety of a set of attractive criteria. Therefore we simply try to offer a selection of desirable properties that our technique follows, without every making any claims that these are indeed truly exhaustive. Section 3.2 delves into how our approach fits into the Bayesian perspective, an important coherency property for those who wish to adhere to this paradigm while updating weights/opinions over time. Section 3.5 presents four simple and attractive attributes that our method obeys. These pleasing mathematical characteristics provide some initial basic justification for our technique.

As discussed in our aims earlier in this chapter it was our goal to provide a strong justification for the method that we derive. The attractive coherency properties certainly are a step in this direction, but arguably not a full enough one. In order to increase the rigorousness of our validation we strove to use data, and to compare our technique to the aforementioned alternative approaches. An initial question concerned the metric of comparison to be used. We consulted the seminal paper by Gneiting and Raftery (2007) in order to be aware of the broad spectrum of possible metrics that could be considered, eventually choosing one that we felt most appropriate for our setting. Due to the novelty of our approach there were no pre-existing data sets constructed in the commensurable fashion desired. Hence we simulated data to examine a broad range of cases, varying the number of decision makers, the number of decision returns, the confidence and accuracy of predictions, and the statistical distributions used to represent these opinions. We compared our approach to the considered alternatives under the chosen metric. The TU Delft Expert Judgment Data Base is a collection of data sets that have arisen in realistic contexts, and have been used to provide justification for one of the most commonly applied opinion pooling methodologies, the classical method of Cooke (1991). Although the manner in which this data was collected does not fully align with the required specifications of our context of interest we attempted to suitably modify it in order to make it applicable for our approach. Having done this we were able to provide some additional justification for our approach, using real-world data rather than simulated data, in order to strengthen the merits of our technique.

We briefly comment that methods of justification of the ilk discussed here are repeated later in the thesis to deal with the various generalisations of our approach that we develop: we consider a set of attractive properties that our group extension obeys in Chapter 5, look at some axioms that our nonparametric approach in Chapter 6 adheres too in addition to consideration of a brief simulation study, before considering some attractive coherency properties that our primary alternative subjective approach in Chapter 7 obeys.

1.4 Outline of Thesis Chapters

We conclude our introduction with a brief summary of the the material contained in the chapters of this thesis. In addition to this content there are several appendices, highlighted in the main body of writing, that provide samples of code used, calculations omitted, results summarised in the text, and tangential discussion points.

- *Chapter 2 - Literature Review:* The fundamental aspects of statistical decision theory that will be used throughout this thesis are formally introduced. The concept of precise probability is discussed, before the subtler notion of utility theory is treated, as well as the relationship that exists between these two ideas, and some additional comments on risk aversion. Maximisation of expected utility is reviewed, with reference to its axiomisation and a brief numerical example. Imprecise probability is then explored and a short numerical example given, as well as discussion on how decisions can be made in the face of this additional uncertainty, and a brief note on how expectation can be considered as a primitive construct in lieu of probability. Group decision theory is then considered, as well as the issues from Arrow's Impossibility Theorem (Arrow, 1950) and notable attempts to circumnavigate these. Finally, a detailed discussion takes place on how a collection of opinions can be combined into a single opinion. The different philosophies underlying methods of doing this are introduced, and the strengths and weaknesses of various potential approaches supplied. Considerable attention is paid to the classical method of Cooke (1991), and the TU Delft Expert Judgment Data Base used to validate this, as well as the work of Karny & Guy (2004) which we attempt to generalise in the following chapters.

- *Chapter 3 - The Plug-in Approach:* This chapter denotes the commencement of the original research contained within this thesis. The notation used in the remainder of the chapters is formally defined, and some further heuristic justification provided for the concept of linear opinion pooling. The idea of Bayesian updating is introduced, and discussion takes place in relation to three commonly implemented conjugate cases which are used for illustration throughout. A framework is supplied for decision making in the environment of interest, termed the Plug-in (PI) approach. Details are provided regarding how individuals are initially given equal weights, before these are updated in light of their perceived accuracy after returns are witnessed. A short numerical example demonstrates PI weights in a Beta-Binomial conjugate setting. A discussion then takes places relating to the Markovian elements of the PI process and two attractive Bayesian properties it adheres to, as well as its relation to scoring rules. Some coherency properties that the method obeys are examined, a detailed numerical example is provided, and some asymptotic properties of weights and distributions are mentioned. Sample calculations are included regarding the distribution that weights follow in the Normal-Normal conjugate case. The chapter also contains some comments pertaining to the relatively straightforward extension of this approach to a more generalised setting. We briefly compare our method with Bayesian Model Averaging (Hoeting *et al.*, 1999) and discuss some limitations on the application of the PI approach.

- *Chapter 4 - Data-based Justifications:* Leading on from the somewhat informal justifications of the PI approach in the previous chapter, we aim to provide a more rigorous basis for its use. This is attempted using two types of data, simulated and real. Data is simulated for the three distributional cases previously introduced, and the PI approach is compared to a collection of rational alternatives under a particular probability density metric. We also provide the theoretical calculations underlying these simulations, and demonstrate how our simulated proportions asymptotically approximate the true probabilities of the PI approach being superior to alternatives. Next, the performance of the method is assessed using the previously introduced TU Delft Expert Judgment Data Base. Discussion is provided on the contrast between the nature of this data and that naturally arising within the PI approach, before rationalisation is provided for the methods used to bridge these contrasts. The performance of the PI approach is compared for each data set in turn to the performance of alternatives, with the meaning of these results interpreted.

- *Chapter 5 - Group Decision Making:* This short chapter discusses how the PI approach can be applied in a group decision making context. This entails combining utility functions as well as probability distributions, with commentary provided on how this may be done in a manner ensuring commensurability (Boutilier, 2003). The proposed group decision making process is considered in relation to each of the five axioms of Arrow (1950). The chapter concludes with comments on the links between our method and Utilitarianism (Harsanyi, 1955), and the merits of cardinal, rather than ordinal, decision ranking schemes.

- *Chapter 6 - Nonparametric Extension:* Up to this point it has been explicitly assumed that DMs can supply fully parameterised probability distributions to quantify their uncertainty, with the PI approach being reliant upon this premise. Here we discuss how this assumption may be weakened, and consider a far more simplistic method of nonparametric belief specification. The concept of Nonparametric Utility Inference (Houlding & Coolen, 2012) is introduced. Arguments are provided regarding how the rules of this method may be augmented for an opinion updating scenario, where the opinions of a DM over the expectation of the uncertain quantity are represented simply by intervals, referred to as Nonparametric Prevision Intervals (NPPI). In the context that we consider here, expectation (*i.e.*, prevision) naturally arises as the obvious choice of primitive construct. A weighting methodology that explicitly learns over time in multiple ways is suggested, and a detailed heuristic justification is given for its use, invoking the scoring rule discussed in Gneiting & Raftery (2007). A numerical example illustrates how this method may be used in practice. A suitable metric is proposed, before a brief simulation study demonstrates the potential merit of our approach.

- *Chapter 7 - Potential Subjective Alternatives:* The PI approach is said to be fully objective, in the sense that the weights it produces are based solely on the data that is witnessed. In this short chapter two alternative methods that incorporate considerable subjectivity are motivated and derived. The Differing Viewpoints approach explicitly models the utility function of the individual assigning weights, with, for instance, a highly risk prone DM assigning a particular prediction a substantially different weight than the weight assigned by deeply risk averse DM. Some formal statements are provided in relation to the risk aversion metric introduced in Chapter 2, and an appropriate performance metric is detailed. The Kullback-Leibler approach, in which a DM assigns higher weights to those individuals whose beliefs closely mirror her own, is briefly discussed and its strengths and weaknesses examined. The chapter concludes with a short example comparing the performance of the three original methodologies provided in this thesis.

- *Chapter 8 - Summary and Further Research:* This final chapter of this thesis gives a brief summary of the research conducted and the conclusions reached therein. Various potential extensions for further work are offered. An application to a social network setting is suggested, with a notation and linear opinion pooling form provided. A discussion takes place about sequential decision problems, most notably concerning the relation to the polynomial utility class (Houlding *et al.*, 2015), and its possible implementation in the setting from this thesis. A method of decision making with constant learning is suggested using not only imprecise probabilities (as in Chapter 6) but also imprecise utilities (for instance over novel returns), allowing for additional uncertainty for users regarding the specification of relevant quantities. Another idea for complementary research is the introduction of a non-flat hierarchy inherent within a group (*e.g.*, a government), where the weight assigned to individuals incorporates their rank. We also mention learning over time in relation to a series of correlated random quantities, that are either realisations of distinct random variables (with θ potentially being multivariate) or realisations of a single dynamic variable, rather than the static quantity considered in this thesis. Links with the concept of value of information (Howard, 1966) are also alluded to.

Chapter 2

Literature Review

The two pillars upon which statistical decision making is built are probability and utility, both of which shall be given substantial treatment in what follows. Initially we focus on precise probabilities, but we will later turn our attention to imprecise probabilities which shall become relevant in Chapter 6. We discuss maximising expected utility as a decision making criteria, as well as the objections raised, and alternatives proposed, to this approach. Comments are made regarding decision making methodologies in an imprecise setting. We analyse group decision making, issues associated with implementing this fairly (in terms of Arrow, 1950), and some alternative approaches put forward, which will be discussed again in Chapter 5. Lastly we review methods of opinion pooling, referencing the various schools of thought on manners by which this can be done, as well as advantages, disadvantages and justifications for these.

2.1 Precise Probabilities

Probability is a method of quantifying uncertainty. In a non-trivial decision making framework uncertainty will be faced by a DM, who is unsure of the exact consequences that will result from her decisions. Probability describes the unsureness of a DM about the environment she inhabits. In this thesis we consider only DMs with imperfect knowledge, *i.e.*, those who are unsure about some aspect of the true mechanisms controlling their areas of interest, *e.g.*, how a stock price fluctuates, or how many passengers will book a particular flight. Kolmogorov (1950) provides a strict axiomatic framework for probability, with coherent probabilities adhering to his axioms. Here we provide these, where Θ is the space of all possible events, θ is a particular event, and $\mathbb{P}(\theta)$ is the probability of θ occurring.

- **Axiom A1:** If $\theta \in \Theta$ then $\mathbb{P}(\theta) \in \mathbb{R}$ and $\mathbb{P}(\theta) \geq 0$.
- **Axiom A2:** There is a universal event $\Theta^* \subseteq \Theta$ such that $\mathbb{P}(\Theta^*) = 1$.
- **Axiom A3:** For a set of countable mutually exclusive events $\theta_1, \theta_2, \ldots \in \Theta$

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

Further rudimentary properties can be derived from these, *e.g.*, monotonicity and bounds of probability. A set of probabilities failing to meet A1-A3 risks falling prey to a "Dutch book" (*e.g.*, Maher, 1993), which occurs when a DM enters a wager she is doomed to lose irrespective of what outcome occurs. The probability associated with the occurrence of an event θ can be considered the price at which a rational DM would be willing to buy (or equivalently sell) a bet that pays one util (shortly to be defined) if θ occurs, and zero utils if not. When we discuss imprecise probabilities we shall see an intuitive extension of this betting analogy.

2.2 Utility

2.2.1 Expected Value Theory

Early probability studies were centred on gambling, *i.e.*, in determining how likely a player was to win a game and what was a fair stake for them to pay to play. Arguably the foundations of probability theory arise from a series of letters between Blaise Pascal and Pierre de Fermat concerning a question posed to them by Chevalier de Mere regarding a particular game of chance. A fair stake was often considered the expected value of the outcome of playing the game. This method had shortcomings, highlighted by Bernoulli (1738) in the St. Petersburg Paradox. This describes a game in which a player tosses a fair two-sided coin, and wins a pot (doubling with each success) for every consecutive toss resulting in heads. She receives the pot the first time she tosses a tail. The initial pot is \$2. What is a fair stake? Equating this to the expected prize,

E, we find

$$E = \frac{1}{2} \times 2 + \left(\frac{1}{2}\right)^2 \times 2^2 + \left(\frac{1}{2}\right)^3 \times 2^3 + \left(\frac{1}{2}\right)^4 \times 2^4 + \ldots = \sum_{k=1}^{\infty} 1 = \infty$$

This is clearly a ridiculous choice of fair price, indicating that the expected prize from playing is an infinite sum. Hence we see that expected value theory is not always a logical decision making criteria, as here a player will give any finite sum to play. Bernoulli used this argument to motivate a new decision making criteria.

2.2.2 Utility Hypothesis

Bernoulli (1738) wrote 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 said that DMs should specify their own utility functions, u, that describe their personal attitudes over outcomes, risks and gambles. Formally u is a function, $u : \mathcal{R} \to \mathbb{R}$, from the set of possible decision returns $\mathcal{R}$, to the real numbers $\mathbb{R}$. For every possible decision a numerical value (measured in units called utils) can be calculated, with the optimal decision returning the highest value, in a process formally defined in Section 2.3.1. There are many advantages to this method in comparison with expected value theory, *e.g.*, it allows for personalistic interpretation of the merits of outcomes. It will generally yield different numerical values (and hence decisions) for different DMs depending on their utility functions, opinions and (in financial settings) monetary situations. It can be argued as a more useful criteria than expected value theory as it incorporates much more information into the decision making process.

2.2.3 Measures of Risk-Aversion

A DM's utility function, $u(r)$, measures the satisfaction she derives from returns $r \in \mathcal{R}$, and reflects her attitude over gambles, *i.e.*, if she is risk-averse, risk-neutral or risk-prone, and to what degree. Suppose a DM has a fortune of f units, and must determine whether to play a game, raising her fortune to $f + m$ units, or decreasing it to $f - m$ units, with probabilities of 0.5 respectively. Not playing the game has expected utility of $u(f)$, while playing has expected utility of $0.5u(f + m) + 0.5u(f - m)$. She is

- risk averse if $u(f) > 0.5u(f+m) + 0.5u(f-m)$, *i.e.*, she opts not to gamble, and indeed would be prepared to pay to avoid taking the gamble.

- risk prone if $u(f) < 0.5u(f+m) + 0.5u(f-m)$, *i.e.*, she opts to gamble, and indeed would be prepared to pay to take this gamble.

- risk neutral if $u(f) = 0.5u(f+m) + 0.5u(f-m)$, *i.e.*, both decisions are equally favourable, and she would neither pay to take or avoid taking the gamble.

A utility function that is risk-averse over a range is concave over this. Convexity implies risk-proneness and a straight line implies risk-neutrality, *e.g.*, Fig. 2.1. Formal measures exist to determine which classification a function falls under (and to what degree), perhaps most commonly the Arrow-Pratt absolute risk-aversion (ARA) coefficient, from Arrow (1965) and Pratt (1964), given by

$$A(r) = -\frac{u''(r)}{u'(r)} \tag{2.1}$$

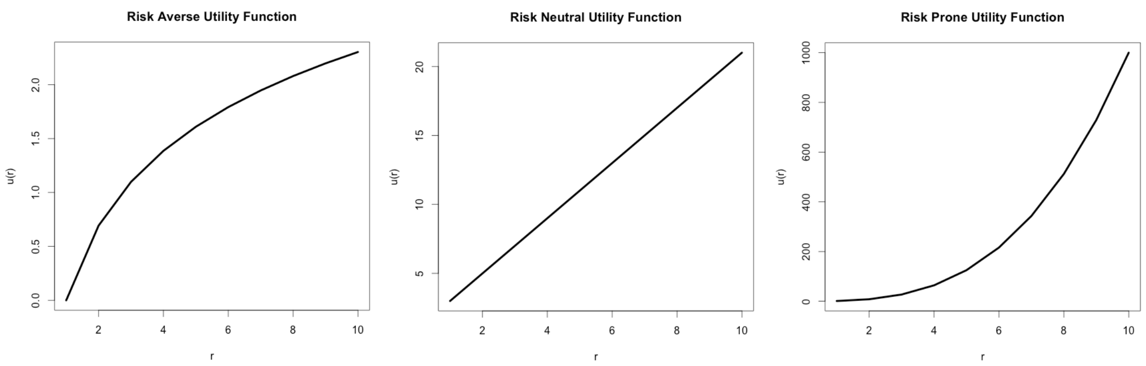

Fig. 2.1: Contrasting utility functions over the range $[1, 10]$. From left to right we have (a) the risk-averse (concave) function $u(r) = \ln(r)$, (b) the risk-neutral (straight) function $u(r) = 2r + 1$ and (c) the risk-prone (convex) function $u(r) = r^3$.

As two examples, $u(r) = \log(r)$ leads to $A(r) = \frac{1}{r}$, positive for all possible values of r (as logarithms take only positive input), implying risk aversion. By contrast, $u(r) = r^3$ yields $A(r) = -\frac{2}{r}$, indicating risk proneness/aversion for positive/negative values of r respectively. Functions of the form $u(r) = 1 - \exp(-\alpha r)$ give $A(r) = \alpha$, termed Constant Absolute Risk Aversion (*e.g.*, Rabin, 2000), with those leading to

$A(r) = \frac{1}{ar+b}$ (for $a, b \in \mathbb{R}$) exhibiting Hyperbolic Absolute Risk Aversion (Merton, 1971). Higher order generalisations of ARA such as absolute prudence and absolute temperance are discussed in Kimball (1990).

2.2.4 Relationship between Utility and Probability

We briefly mention the relationship that exists between utility and probability. Each DM has a subjective utility function $u(r)$. While in practice elicitation of this is difficult, 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 illustrate this.

- **Definition:** A DM's subjective probability for the occurrence of an event is the amount p she is willing to gamble such that she receives 1 util if the event occurs, and 0 utils if not.

- **Definition**: The utility a DM assigns to an outcome r is the value p making her indifferent between

  (a) r for certain and

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

In this utility definition the values are rescaled to the unit interval. We discuss why this is possible in Section 2.3.3. Circularity can be noted between the two definitions. It seems justifiable for French (1994) to refer to utility as "probability's younger twin". Note both concepts are derived from the preference relation ordering, $\succeq$, that we shall shortly define.

2.3 Expected Utility

2.3.1 Decisions

We denote a set of 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 distinct action. The satisfaction derived from a decision is based upon which state of nature occurs in conjunction with it, with DMs uncertain over which state 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 states is true, but a DM is uncertain which one it is prior to choosing to eat there. Suppose there are m potential mutually exclusive and exhaustive states of nature, $\theta_1, \theta_2, \dots, \theta_m \in \Theta$. A DM states her utilities over all possible outcomes (Table 2.1), where $u(d_i, \theta_j)$ denotes the utility resulting from making decision d_i and the occurrence of θ_j.

Table 2.1: Cross-tabulation of utilities for $\mathcal{D}$ and Θ.

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

The expected utility of a decision d_i is the sum over the products of the probability of each possible return occurring and the utility value associated with this, *i.e.*,

$$\mathbb{E}[u(d_i)] = \sum_{j=1}^{m} u(d_i, \theta_j)\mathbb{P}(\theta_j) \tag{2.2}$$

This sum is replaced by an integral if returns and/or probability distributions over returns are continuous. The optimal decision, d^*, maximises Equation (2.2), *i.e.*,

$$d^* = \arg\max_i \mathbb{E}[u(d_i)] \tag{2.3}$$

Lindley (1991) comprehensively treats this approach. In the framework discussed here, it is assumed that DMs can both probabilistically quantify uncertainty over potential returns and specify exact utility values corresponding to these. Realistically this may often not be the case. In the absence of these abilities elicitation methods exist to help discover these unknowns, *e.g.*, the techniques of O'Hagan (1998) are often used for belief elicitation. Regarding utility, a method by which preferences can be ascertained over time is adaptive utility, discussed in, *e.g.*, Cyert and DeGroot (1975), and Houlding & Coolen (2011), and modified to incorporate extreme vagueness in the priors over parameters in Houlding & Coolen (2012). Chajewska *et al.* (2000) also contains information on suitable utility elicitation methods.

2.3.2 Example using Precise Probabilities

A DM must choose whether to play football (d_1), rugby (d_2) or snooker (d_3). There are two possible mutually exclusive states of nature: the weather will be sunny (θ_1) or rainy (θ_2), with Table 2.2 containing her utilities. Suppose she assesses $\mathbb{P}(\theta_1) = 0.7$ implying $\mathbb{P}(\theta_2) = 0.3$. Using Equation (2.2) she finds that $\mathbb{E}[u(d_1)] = 0.73$, $\mathbb{E}[u(d_2)] = 0.48$ and $\mathbb{E}[u(d_3)] = 0.7$, *i.e.*, her optimal decision is d_1, to play football.

Table 2.2: Cross-tabulation of utilities for $\mathcal{D} = \{d_1, d_2, d_3\}$ and $\Theta = \{\theta_1, \theta_2\}$.

	θ_1	θ_2
d_1	1	0.1
d_2	0.6	0.2
d_3	0.7	0.7

2.3.3 The axiomatisation of von Neumann and Morgenstern

The work of von Neumann & Morgenstern (1944) axiomatically justified utility theory. They created four rational axioms. If a DM agreed with these, and was herself rational, then she would make decisions by maximising expected utility. These axioms are:

- *Completeness*: For any $d_1, d_2 \in \mathcal{D}$, a DM can always determine which, if either, she prefers, *i.e.* either $d_1 \succeq d_2$, or $d_2 \succeq d_1$. The relation $\succeq$ implies weak preference, with $d_1 \succeq d_2$ stating d_1 is at least as preferable as d_2. Indifference is indicated by $d_1 \sim d_2$, with $\succ$ denoting strict preference.

- *Transitivity*: If $d_1 \succeq d_2$ and $d_2 \succeq d_3$ then $d_1 \succeq d_3$, for all $d_1, d_2, d_3 \in \mathcal{D}$. This ensures "money-pump" situations cannot arise.

- *Continuity*: If $d_1 \succeq d_2 \succeq d_3$ then there is $p \in [0, 1]$ with $d_2 \sim pd_1 +_g (1-p)d_3$ for all $d_1, d_2, d_3 \in \mathcal{D}$, *i.e.*, there exists a probability p making a DM indifferent between a gamble between the best and worst outcome (with probabilities p and $1 - p$ respectively) and a guaranteed intermediate outcome. Note that the operator $+_g$ is used to denote a gamble between two outcomes rather than standard addition.

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

Von Neumann and Morgenstern showed for a DM agreeing with these that there is a unique (up to positive linear transformation) utility function such that:

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

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

Hence a formal justification was given for maximising expected utility, advocating it as a method for DMs deemed rational in some sense. Expanding on the invariance of utility functions to positive linear transformations, a DM need not be concerned if she specifies her utility function as $u_1(d)$ or $u_2(d) = au_1(d) + b$ for $a, b \in \mathbb{R}$ and $a > 0$, *i.e.*, the decision deemed optimal under u_1 is optimal under u_2, and the converse. This property arises due to the linearity of the expectation operator over its arguments.

2.3.4 Objections and Alternatives

The above method is normative, *i.e.*, providing a formal method by which rational DMs should make decisions. Yet human beings are prone to irrationality and not always behaving in a manner consistent with this. The Allais paradox illustrates the decision making irregularity often exhibited by DMs (Allais, 1953). DMs were questioned on what decisions they would make in two separate hypothetical situations. The first question asked if DMs would rather have \$1,000,000 with certainty (A1), or \$1,000,000 with probability 89%, \$5,000,000 with probability 10% and \$0 with probability 1% (A2). The second question asked if DMs would rather have \$1,000,000 with probability 11% and \$0 with probability 89% (B1) or \$5,000,000 with probability 10% and \$0 with probability 90% (B2). The most common pair of decisions to choose was A1 and B2. This is inconsistent with utility theory, under which choosing A1 implies automatically choosing B1, and choosing A2 implies automatically choosing B2. Irrespective of the utility function given by a DM it is impossible, under the framework of von Neumann and Morgenstern, to choose both A1 and B2. This is presented by Allais as a counterexample to the independence axiom, as the only difference between A1 and

A2, and B1 and B2, is a common increase in the probability of receiving \$0, violating this axiom.

An attempt to resolve this problem was the development of descriptive decision making methods, describing how DMs actually make choices and incorporating human characteristics, perhaps most notably Prospect Theory (Kahneman & Tversky, 1979). This attempted to mirror behaviour exhibited by DMs in realistic scenarios, such as how they feel the pain of a loss more severely than the joy of an equivalent gain, and how they underweight outcomes that are probable compared to those that are guaranteed. Fuzzy logic (Zadeh, 1965) modifies classical set theory, allowing items to belong to several distinct sets at once with varying degrees of membership. Statements need not be strictly true or false, but may have truth values in the $[0, 1]$ interval. There are advantages to this approach but, as we shall shortly see in decision making under imprecise probability, computationally it has significant disadvantages.

2.4 Imprecise Probabilities

Supplying a precise probability is a strong statement, implying that a DM knows enough about the unknown quantity to exactly quantify her uncertainty over it. DMs may have relevant prior experience, or expert advice, to help them do this. What should a DM do if she *a priori* has little relevant information about the topic at hand? Precise probability statements play a vital part in determining optimal decisions, *e.g.*, in Section 2.3.2 using $\mathbb{P}(\theta_1) = 0.7$ led to d_1 being chosen, yet slight augmentation to $\mathbb{P}(\theta_1) = 0.65$ makes d_3 optimal. DMs must be careful in supplying precise probabilities as, if inaccurate, they may lead to negative (low utility) consequences. An alternative is the concept of imprecise probabilities. Rather than a single probability value $\mathbb{P}(\theta)$, a DM gives a lower bound $\underline{\mathbb{P}}(\theta)$, and upper bound $\bar{\mathbb{P}}(\theta)$, that she maintains $\mathbb{P}(\theta)$ lies within. Probability measures uncertainty, with imprecise probability allowing additional vagueness. Analogous to Kolmogorov (1950), there are imprecise probability axioms (Weichselberger, 2000, 2001), B1-B3. It is assumed $\mathbb{P}(\theta)$ obeys A1-A3.

- **Axiom B1:** $0 \leq \underline{\mathbb{P}}(\theta) \leq \bar{\mathbb{P}}(\theta) \leq 1$ for all $\theta \in \Theta$.

- **Axiom B2:** $\underline{\mathbb{P}}(\theta) \leq \mathbb{P}(\theta) \leq \bar{\mathbb{P}}(\theta)$ for all $\theta \in \Theta$.

- **Axiom B3:** $\inf_{\mathbb{P}}\{\mathbb{P}(\theta)\} = \underline{\mathbb{P}}(\theta)$ and $\sup_{\mathbb{P}}\{\mathbb{P}(\theta)\} = \bar{\mathbb{P}}(\theta)$.

From these further properties can be garnered. Denoting by θ^c the complementary event of θ we have $\bar{\mathbb{P}}(\theta^c) = 1 - \underline{\mathbb{P}}(\theta)$ and $\underline{\mathbb{P}}(\theta^c) = 1 - \bar{\mathbb{P}}(\theta)$. Previously we gave a betting price interpretation to precise probabilities. There is an attractive analogy for imprecise probabilities, with the lower bound probability being the smallest price for which a DM is willing to sell a bet (in which she must pay one util if θ occurs and zero utils if not) and the upper bound probability being the largest price for which she is willing to buy a bet (in which she receives one util if θ occurs and zero utils if not). Imprecise probability theory is an expanding topic, with many techniques constructed for tasks that formerly required precise probabilities, *e.g.*, Walley (1991) and Coolen *et al.* (2010).

2.4.1 Decision Making using Imprecise Probabilities

How can DMs determine optimal decisions when opinions about θ are given by imprecise probabilities? Consider Table 2.3, with a DM assessing $\underline{\mathbb{P}}(\theta_1) = 0.4$ and $\bar{\mathbb{P}}(\theta_1) = 0.6$. For $\mathbb{P}(\theta_1) = 0.55$ she deems d_1 optimal, $\mathbb{P}(\theta_1) = 0.45$ means d_2 is optimal and $\mathbb{P}(\theta_1) = 0.5$ gives a tie. Her optimal decision depends upon which probability (in her range) is considered when calculating expected utilities. Often a single decision cannot be declared unanimously optimal under all imprecise configurations.

Table 2.3: Cross-tabulation of utilities for $\mathcal{D} = \{d_1, d_2\}$ and $\Theta = \{\theta_1, \theta_2\}$.

	θ_1	θ_2
d_1	1	0
d_2	0	1

If a single decision cannot be declared optimal it is important to eliminate decisions that are unequivocally not optimal, *i.e.*, inadmissible (*e.g.*, Coolen, 2006). If a DM cannot pick one alternative as maximal then she may choose a decision by a chance mechanism, akin to uniform preference over decisions. It is desirable to omit as many inadmissible decisions as possible before making this choice. We discuss four common decision making methods for imprecise probabilities, two which choose an optimal decision, with two removing inadmissible options. For a particular definition, as discussed in Schervish *et al.* (2003), these methods can be seen as coherent, *i.e.*, avoiding Dutch books. If a decision has maximal expected utility under all belief configurations it is optimal. This will frequently not be the case, but if it is then this decision choice is robust. In a precise setting, Maximin (Wald, 1950) chooses the decision returning the largest minimum expected utility value, *i.e.*, a pessimistic approach. $\Gamma-$Maximin is an imprecise probability extension of this. Maximality (Condorcet, 1785) eliminates inadmissible decisions, *i.e.*, those giving lower values than an alternative for all configurations. E-Admissibility (Levi, 1974) removes all decisions from $\mathcal{D}$ except those which are optimal under at least one belief specification. The set of decisions left after applying E-Admissibility is a subset of that remaining after Maximality.

For the problem in Section 2.3.2, suppose $\underline{\mathbb{P}}(\theta_1) = 0.3$ and $\bar{\mathbb{P}}(\theta_1) = 0.7$, implying $\underline{\mathbb{P}}(\theta_2) = 0.3$ and $\bar{\mathbb{P}}(\theta_2) = 0.7$. Maximising expected utility is inconclusive. The $\Gamma-$Maximum values of d_1, d_2 and d_3 are 0.37, 0.32 and 0.7 respectively, *i.e.*, under this criteria d_3 is optimal. Maximality shows both d_1 and d_3 dominate d_2, so it is eliminated, while under E-Admissibility d_1 is optimal for $\frac{2}{3} < \mathbb{P}(\theta_1) \leq 0.7$ and d_3 is optimal for $0.3 \leq \mathbb{P}(\theta_1) < \frac{2}{3}$. There is no configuration making d_2 optimal, so it is eliminated.

2.4.2 Choosing a Primitive Quantity

Probability is often used in statements of uncertainty. In deriving axioms and theorems the concept of a probability is usually considered "the primitive", *i.e.*, the fundamental quantity upon which all further statements are built. If we are interested in the expectation of a random variable X, and have probabilities for its potential realisations $x_1, \ldots, x_n$ then we define this expectation as

$$\mathbb{E}(X) = \sum_{i=1}^{n} x_i \mathbb{P}(X = x_i) \tag{2.4}$$

This is not the only possible route. As in de Finetti (1974) and Whittle (1992) we can take expectation as the primitive and define other concepts in terms of this. Consider an indicator variable $I_X(x_i)$, taking a value of 1 if $X = x_i$, and 0 if not. The probability of x_i occurring is then the expectation of this indicator variable, *i.e.*,

$$\mathbb{P}(X = x_i) = \mathbb{E}[I_X(x_i)] \tag{2.5}$$

Equation (2.4) defines expectation in terms of probability while Equation (2.5) did the converse. In imprecise probability settings expectation is often considered the primitive (*e.g.*, Walley, 1991). In Chapter 6 this is the framework we adhere to.

2.5 Sequential Problems

The methods above considered a DM making a single decision, yet DMs often need to make decisions for several future epochs simultaneously, *i.e.*, non-myopically. In a myopic setting DMs consider one step into the future, at which time they observe a result, and consider one step into the future again. A non-myopic setting is one in which a DM must decide in May how much money she will need in June, July and August. These problems are solvable using decision trees and the "roll-back" technique (Alghalith, 2012, proposes a "roll-forward" method), discussed via maximising expected utility by Lindley (1991). An issue with this approach is the "curse of dimensionality", *i.e.*, when a large amount of epochs and/or potential decisions are involved trees rapidly become very complicated and computation is slow. Intractability may occur depending on the form of probability distributions and utility functions. A method easing these issues is the polynomial utility class (Houlding *et al.*, 2015), reliant upon the assumptions of polynomial utility functions and Normal probability distributions, creating conjugacy for utility functions analogous to that existing for probability distributions. We return to sequential problems in Chapter 8.

2.6 Group Decision Making

The approaches above dealt with a single DM making a decision, yet group decision making is an important task too. While there is motivation for a normative method to assist groups in making rational choices we shall see, for a certain set of axioms, that this is an unobtainable goal due to the Impossibility Theorem of Arrow (1950). Efforts have been made to circumnavigate this and find an acceptable group decision making method, with strengths and weaknesses of some such approaches discussed below.

2.6.1 Arrow's Impossibility Theorem

Arrow (1950) considered preference ranking among a body of individuals, and if there was a "fair" mapping from the set of n individual rankings $(\succeq_1, \ldots, \succeq_n)$ to a collective ranking ($\succeq$). He contemplated a Social Welfare Function (SWF) operating on a set of individual rankings, which would obey certain basic properties (completeness and transitivity) and lead to a group ranking obeying the same. He put forward five desirable axioms for this SWF to obey. Universality stated that the SWF be defined for every admissible set of individual orderings. Monotonicity declared that if a decision d_1 rose, or did not fall, in the ordering of each DM without any other change in those orderings, and if d_1 was preferred to d_2 before the change to individual orderings, then d_1 is still preferred to d_2. Non-imposition and non-dictatorship respectively ruled that the SWF be neither imposed nor dictatorial. Independence is the most controversial axiom. In a two DM setting, let $\succeq_1$, $\succeq_2$ and $\succeq_1'$, $\succeq_2'$ be two sets of individual orderings. If for both individuals i, and for all $d_1, d_2 \in \mathcal{D}$, $d_1 \succeq_i d_2$ if and only if $d_1 \succeq_i' d_2$, then the choice made is the same whether the individual orderings are $\succeq_1$ and $\succeq_2$, or $\succeq_1'$ and $\succeq_2'$. Monotonicity and non-imposition can be combined to form the Pareto principle axiom (Arrow, 1963). Arrow (1950) showed that for at least two DMs and three distinct decisions, no SWF satisfying the five conditions can be created, *i.e.*, any SWF is irrational in some sense. Independence is the axiom researchers most commonly try to sidestep, as it assumes no preference between any two outcomes is stronger than that between any other two outcomes, *i.e.*, it is ordinal rather than cardinal. Several other interesting impossibility results exist, *e.g.*, 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).

2.6.2 Utilitarianism and the SWF

Utilitarianism is a normative theory, broadly stating that well-being should be maximised and suffering minimised. A leader, P^*, must translate the rankings of n individuals to one collective ranking. From two axioms a cohesive group ranking may be achieved, as in Harsanyi (1955). All DMs assign utilities to each potential option, scaled to $[0, 1]$, giving utility functions $u_1, \ldots, u_n$ that must be translated to a single function, u^*. The axiom of anonymity states that P^* does not know who put forward which ranking, *i.e.*, no bias. The second axiom is the strong Pareto principle, declaring that if each individual is indifferent between two outcomes then so is P^*, *i.e.*, if $u_i(r_1) = u_i(r_2)$ for all $i = 1, \ldots, n$ then $u^*(r_1) = u^*(r_2)$. Similarly if some prefer r_1 to

r_2, and the rest are indifferent, then P^* prefers r_1 to r_2. From these, Harsanyi proved that the only function u^* obeying these axioms is an equally-weighted additive sum of $u_1, \ldots, u_n$. We return to discuss this in Chapter 5. While the approach of Harsanyi was accepted by many, there were detractors, perhaps most notably Buchanan (1954, 1979, 1994a, 1994b) and Buchanan & Tullock (1962), who had reservations about the choice of function, and concerns that the liberal value judgment of individualism was overlooked. Sen (1979, 1990, 1995) discusses the importance of a process not just being mechanically fair, but that its results are fair in a general social sense.

2.6.3 Fully Probabilistic Design

Karny & Guy (2004) consider a setting where each group participant attempts to improve her individual decision quality by sharing her opinions with those around her, optionally augmenting her own opinion in light of what she has learned. Key to their method is the Kullback-Leibler (KL) divergence (Kullback & Leibler, 1951), a measure of the sameness of two probability distributions (which we use in Chapter 7). If f_i and f_j are distributions over Θ then the KL divergence between them is

$$D(f_i||f_j) = \int_{\Theta} f_i(\theta) \log_e \left[\frac{f_i(\theta)}{f_j(\theta)}\right] d\theta \tag{2.6}$$

Another concept of Karny & Guy (2004) is an "ideal distribution". DMs model their beliefs via probabilistic distributions. As well as this modelling distribution Karny & Guy recommend DMs construct an ideal distribution describing their "best-case scenario". This is linked with the KL divergence to determine the optimal decision for a DM using a method called Fully Probabilistic Design (FPD), with the optimal decision minimising the KL value between her modelling and ideal distributions. While a solution does exist it is of a very specific and complex form. Hence, while theoretically possible, it may be highly problematic to use FPD in a non-trivial context.

Karny & Guy (2009) consider cooperation by DMs in a group framework by sharing of probabilistic information. DMs must *"... take all offered information pieces as outputs of noisy information channels and try to estimate parameters of the underlying source"*. Several assumptions are made, crucially that there is no mechanism to measure reliability of information received, nor to compare the differing importance of information from distinct sources. These are issues we seek to resolve in Chapter 3.

Karny & Kracik (2003) describe cases with an inherent group hierarchy, specifically the importance of maximising the amount of information available to DMs. This desire to maximise information flow motivates the research in the remainder of this thesis.

2.6.4 Other Alternatives

Some key group decision making techniques are from Nash (1951), applicable in situations with two or more non-cooperative participants. His Bargaining Theorem showed that working from a set of axioms a fair bargaining point can be reached for a group. Similarly, his Equilibrium Theorem found combinations of strategies that could not be improved upon for all DMs. There are powerful methods, yet there are cases where they are overly simplistic. Nash equilibrium is applicable when all DMs commit to individual rationality, *i.e.*, wanting to maximise their expected utility, assuming all other DMs want to do the same. Thoughts regarding others are purely in terms of the actions that they may take and the consequences of these for themselves. Yet DMs may be willing to sacrifice part of a benefit to themselves to benefit the group as whole. The work of Nash is applicable when DMs contemplate fellow players only in the sense of predicting their actions so as to maximise their own benefit conditional upon these. When fellow DMs are considered in a "good-will" sense, with willingness to make concessions for them, the work of Nash is not applicable. Stirling (2004) proposes a "satisficing" method which searches for a set of decisions deemed "good enough" for all DMs, rather then a single optimal solution. Info-Gap Decision Theory (Ben Haim, 2006, 2007) is a similar satisficing approach that is non-probabilistic and viable in cases of Knightian uncertainty. Rabin (1993) formalises the concept of fairness, deriving axioms and equilibria contrasted to those of Nash (1951).

2.7 Combining Expert Judgments

We now present motivation for the combining of multiple opinions, and provide detailed discussion on the most common methods by which this can be done. Using the terminology of Karny & Guy (2004), we say that the "neighbours" of a DM are the collection of other participants with whom she has a common area of interest. A set of neighbours will frequently have different degrees of knowledge and/or opinions about the uncertain event of interest, and hence may have greatly differing beliefs over what outcome may occur as a result of a decision. Not only this, but they will also 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 (*i.e.*, their opinions have identical means), but one could be almost certain of this, while the other could be deeply unsure (*i.e.*, the associated variance of the latter is substantially higher than that of the former). It seems a logical premise that in order to maximise the decision quality (and minimise the corresponding risk) for individuals within the group, it is imperative to maximise the amount of information that they have access to prior to making their decisions, or to quote Bunn (1975) "*the methodology of combining forecasts is founded upon the axiom of maximal information usage*".

Bates & Granger (1969) can be regarded as a suitably seminal work, in being one of the first to suggest the amalgamation of different forecasts in the hope of gaining more accurate estimation, with a weighted sum of point estimate beliefs being constructed with weights proportional to the absolute error of the predictions of individuals at previous epochs. The performance of this method was tested on several data sets and gave encouraging results, indicating the merits of combining opinions. Newbold & Granger (1974) built upon this, considering several combination rules, amongst them being that of Bates & Granger (1969), as well as the Box-Jenkins, Autoregressive and Holt-Winters models. These were tested on a collection of over eighty financial time series, and although no formal criteria of optimality was demonstrated empirical evidence was strong in favour of the accuracy of combined beliefs in contrast to that of individual beliefs. In a point estimate context consisting of two participants, Bunn (1975) considered the weight assigned to an individual in a linear combination to be proportional to the subjective probability of that individual being the more accurate of the two. Beta-Binomial conjugacy was implemented to make modifications to these reliability measures over time, *i.e.*, this approach contained strong Bayesian elements. This involved consideration of the mean of a Beta prior distribution, with this prior defined over the probability that the first individual was the more reliable of the two. This prior distribution was then updated in light of the data witnessed (*i.e.*, a one if the first individual is more accurate than the second at an epoch, and a zero if not) using a Binomial distribution, with the new weights arising from the mean of the Beta posterior distribution. Performance was contrasted to linear and exponential combining rules, with the Beta method appearing superior in instances were a suitably informative prior could be supplied. Extension to a Dirchlet prior distribution with an associated multinomial updating mechanism is suggested in a setting that consists of more than two individuals.

It stands to reason that a DM who is well informed about the pertinent topic at hand will generally make more astute decisions than a DM who is not. One would expect this more astute decision making to lead to a corresponding increased decision quality, *i.e.*, higher utility for the DM. Within the framework in which we are interested the seemingly rational course of action for all DMs is to share their beliefs with each other, hence incorporating a greater scope of knowledge into their individual decision making tasks. We assume that all DMs are non-competing and therefore have no reason to supply each other with intentionally inaccurate opinions, *i.e.*, there is no motivation for dishonesty amongst participants as the belief shared by a DM with her neighbours is the same one that she herself will use in her own decision making task. Hence, assuming she is rational she will not wish to sabotage her own decision quality, and will therefore supply her neighbours with her true representation of her opinion. All individuals make their own decisions, the consequences, positive or negative, of which will impact solely upon themselves.

Much consideration has been given in the literature to the problem of how to combine the judgments of a set of domain-specific experts into a single opinion to be used by a DM in a decision making problem, for instance by Cooke (1991, 2007) and Clemen & Winkler (1999). Our problem approaches this technique from a somewhat different perspective, with each decision making individual herself being considered an expert (*i.e.*, possessing some form of probability distribution over the uncertainty inherent in the decision task), who combines her opinion with those of other (decision making) experts in order to make her own decision, whose quality she wishes to maximise. Nevertheless, the following discussion, which takes place in the setting of the former context is equally applicable to the latter framework. The aim in both scenarios is to construct a combined belief that is as accurate a representation of some true underlying unknown state of nature as possible, with this output being used by all DMs in our setting and by a single (non-opinion holding) DM in the framework of Cooke (1991).

There are two primary methods by which a collection of opinions can be amalgamated, namely mathematically and behaviourally. The former approach involves some mathematical function, which takes as its inputs the beliefs of individuals and returns as its output a single belief, which is in some sense representative of a collective opinion. The latter is a substantially more heuristic approach in which individuals discuss their opinions together (*i.e.*, the reasoning and rationales underlying these) in some real or virtual setting, in an effort to reach a common consensus. While several behavioural procedures have been developed, perhaps most notably the Delphi method (Dalkey, 1969) and the Nominal Group Technique (Delbecq *et al.*, 1975), these can be seen as suffering from a lack of rigour in their application, and hence may be viewed as somewhat *ad hoc* approaches, with both methods, for example, being potentially susceptible to biases that are inherent within the group structure.

Within mathematical combining there are two dominant methods, namely Bayesian opinion pooling and combining rules. Genest & Zidek (1986) provide a definitive guide to the nuances of both of these, which we highly recommend to the interested reader. Of the two approaches Bayesian pooling is by far the more complex method, with a DM first specifying her own prior distribution over the uncertain quantity, before viewing the opinions of experts as data, which are entered into a likelihood function, and combined with her prior distribution to give a posterior distribution over the parameter of interest. In the case where a DM has no relevant knowledge about the parameter a "flat", or relatively non-informative, prior will be used to reflect this. It will be dominated (*i.e.*, outweighed) by the information contained within the likelihood function in the calculation of the posterior distribution. French (2011) comments favourably on the concept of viewing expert beliefs as data, but concedes that this method has vast problems with implementation in practice, primarily concerning the choice of an appropriate likelihood function. Clemen & Winkler (1999) note that "*at the same time it is compelling, the Bayesian approach is also frustratingly difficult to apply*", and while they propose some suggested forms of likelihood functions to be used in specific problems and/or with certain functional forms of opinions, no generalised method of supplying these has been found, rendering the Bayesian method an extremely problematic one to implement in realistic scenarios.

Lindley & Singpurwalla (1986) illustrate how in a one-period problem concerned with estimating failure rates for a system such a Bayesian approach is technically possible to apply, but it requires participating DMs to specify a vast amount of parameters (pertaining, for instance, to the correlation between neighbours' opinions, as well as their perceived degree of overconfidence), which may be beyond the computational scope of individuals in a realistic application. Jouini & Clemen (1996) cleverly sidestep the difficulties with supplying a suitable form of likelihood function by introducing a copula method that models dependency between the beliefs of individuals. The supplication of a copula function implies that the likelihood function can be rewritten as a product of the individuals' marginal distributions and this copula function, with this function containing the pertinent information about dependencies inherent within opinions of individuals. Nevertheless, as with the content of Lindley & Singpurwalla (1986), this technique still places a large implementation burden on the user in supplying measures of dependence between different individuals, which she may struggle to provide *a priori* in any meaningful fashion.

Combining rules are a more straightforward methodology, in which opinions can be combined by a variety of methods. Perhaps the most straightforward of these is additive linear pooling (*e.g.*, Stone, 1961, Clemen & Winkler, 1999), under which the combined opinion is an additive combination of the opinions of individuals with normalised weights, *i.e.*, those which are non-negative and sum to one. An alternative, discussed for instance in Genest *et al.* (1984) and Heskes (1998) is logarithmic opinion pooling, in which opinions are multiplicatively merged and raised to various powers, prior to normalisation. As discussed in Clemen & Winkler (1999), allowing equal weights for linear and logarithmic pools is equivalent to taking an arithmetic and geometric mean of opinions respectively. Harmonic pooling is another option, discussed in Dawid *et al.* (1995). Cooke (1991) comments on how the linear and logarithmic combining methodologies can be seen as belonging to the same general family of rules, with different parameter choices in each case. Genest & Zidek (1986) provide a detailed discussion on the merits of the various possibilities, with numerous desirable criteria being considered, and strengths and weaknesses of each method assessed. Individuals using logarithmic pooling must be careful with attaching probabilities of zero to events as, if even a single individual does this, a probability of zero will be assigned to this event in the combined distribution irrespective of the magnitude of the probabilities provided for this event by all other participants. Rufo & Perez (2012) note how logarithmic pooling is more likely to lead to unimodal pooling with a lower associated variance than linear pooling, which may lead to multimodal distributions.

The marginalisation property is uniquely obeyed by additive linear opinion pooling, stating that, if a set of individuals provide opinions on some multidimensional unknown parameter, the results yielded from finding the marginal distribution of one element of this from the combined belief will be identical to that found from combining the marginal distributions of all individual opinions over this element. The external Bayesianity property is solely obeyed by logarithmic pooling, giving the same result irrespective of whether the combined belief is updated in light of some new return or if the individual beliefs are updated in light of this and then combined. We shall see in Section 3.2 that the method we propose in Chapter 3 obeys two desirable Bayesian properties that are akin, but not strictly equivalent, to this. Various other potentially attractive properties, for instance the independence preservation property, the strong setwise property and the zero preservation property are discussed in Genest & Zidek (1986). French (1985) comments on the existence of impossibility theorems in a similar vein to that of Arrow (1950), demonstrating that there is no combining rule simultaneously meeting the entirety of a particular set of desirable criteria. In this thesis we choose to use linear opinion pooling for its relative simplicity, its ease of interpretation and in sticking with common practice. We comment that Litchendahl *et al.* (2013) discusses linear averaging of quantiles rather than probabilities, and some motivations for why this area merits further exploration. In what follows we assume that our linear pools are comprised of probability distributions.

The next problem that arises concerns the need for an appropriate method of supplying weights to be incorporated in this linear combination, with Genest & Zidek (1986) commenting that this method *"lacks a normative basis for choosing the pooling weights"*. Ideally these weights would be reflective of reliability, with high weights afforded to those individuals who are deemed to be trustworthy (*i.e.*, accurate) and lower weights afforded to those who are deemed to be inaccurate. In Cooke's classical method (Cooke, 1991, 2007) seed variables are used to assess expert accuracy prior to any decision being made, with experts requested to give their opinions regarding a collection of quantities whose exact values are unknown to them, but known to the DM.

These quantities should be similarly themed to that inherent within the DM's decision making task, to ensure good performance (*i.e.*, accuracy) over the seeds is indicative of good performance over the unknown parameter, and the contrary. The opinions of experts are provided in the form of quantiles, *i.e.*, each expert states the values that she believes that 5%, 50% and 95% of realisations for a particular seed variable will respectively be less than. These predictions are then compared with the true value of the seed variables (known to the DM), with weights for the linear additive combination constructed based on the magnitudes of these disparities.

The opinions of experts are ranked on two distinct scoring scales. The first of these is a calibration measure, that rewards individuals whose beliefs contained witnessed values, *i.e.*, the values lay between their lower and upper quantiles. The second is an information metric, that takes the width of an expert's quantile range into account, with narrower (*i.e.*, more confident) opinions garnering a higher score, as they imply more certainty on behalf of the expert in the belief that she holds. This latter metric is akin to consideration of the variance associated with a DM's probability distribution. The aim of the scoring process is to maximise sharpness subject to calibration so, for instance, an expert who supplies the whole real line as her quantile range will receive a good calibration score but will be heavily penalised for the triviality of her belief. The overall weight that is assigned to an expert is a function of her calibration and information scores (for details see Cooke, 1991), with these scores being aggregated across the full collection of seed variables under consideration. Those experts whose scores fail to meet some minimum threshold are awarded a weight of zero. They are excluded from the next stage of the process, in which opinions are supplied over the unknown quantity inherent within the decision task (*i.e.*, that which the DM has no relevant opinion over). Hence in some instances the combined belief may simply be the opinion of a single expert (*i.e.*, she who is considered to be the most reliable) with the opinions of all other experts disregarded. Note that the threshold below which experts are disregarded is constructed relative to the reliability of all experts in a scenario, ensuring there is never a situation where all experts are assigned a weight of zero, *i.e.*, the opinion of at least one expert will always be used by the DM in her decision making task.

DeGroot & Mortera (1991) discuss how weights can be chosen that meet some formal optimality criteria, but this is conducted in a setting somewhat different from that which we are interested in, and is also generally considered unsuitable for the classical framework. All DMs are assumed to have common prior distributions and utility functions, assumptions that are far too strong for our context of interest (and indeed are difficult to conceive of occurring in a realistic scenario). In addition to this, DMs wish to make a common decision, rather than each individual making her own personal decision. Each individual then constructs her posterior distribution based on the data that only she observes, and is unable to share with others in the considered environment, due to "*external forces such as high cost, physical constraints or company confidentiality*". Individuals then share their respective posterior distributions with each other. The optimal weights are determined to be those that minimise the expected value of a particular loss function, which is collectively determined by the group members. As we shall see in the following chapters the above assumptions can be viewed as unreasonable for the problem in which we are interested. We comment that in situations where there is a lack of access to seed variables and other relevant information pertaining to reliability prior to decision making the Laplacian Principle of Indifference (Laplace, 1812) is often applied, giving all individuals initial equal weights in the absence of any significant evidence to favour one individual over another, *i.e.*, given that there is no better evidence to the contrary.

When we discuss the opinions of DMs in Chapter 3 we shall assume that they are expressible in the form of fully parameterised probability distributions, with the amalgamated opinion a linear combination of these, *i.e.*, a mixture model. Rufo *et al.* (2009) and Rufo & Perez (2012) provide discussion in cases where all opinions belong to the "exponential family" (a set of distributions all expressible in a specific form). This assumption is often made for mathematical convenience, as we shall shortly discuss in Section 3.1.2. Two methods are proposed for choosing weights for a linear pool, both involving the aforementioned Kullback-Leibler divergence measure, aiming to choose the weights that minimise the KL value between the combined group posterior distribution and the distributions of individuals. The first scheme proposed suffers from not guaranteeing coherent weights, *i.e.*, potentially providing sets of weights that are not all non-negative and summing to one. The second scheme entails linear programming to ensure these necessary constraints are met. Bayesian elements are inherent within these techniques. They are primarily concerned with amalgamating initial opinions rather than developing the stringent method of learning over time that we are concerned with, and are also heavily based upon the closed form solutions inherent within conjugacy, whereas our approaches, while generally illustrated with conjugate cases for ease of interpretation, can be implemented using numerical techniques in instances where opinions are of intractable forms.

Genest & McConway (1990) provide an excellent summary of a further collection of approaches which can be used to allocate weights in a linear pool. One method involves "veridical" (defined formally as "coinciding with reality") weights, based upon the assumption that the opinion held by exactly one expert is precisely correct, with weights reflecting the relative probabilities of each expert being the correct one. There are clear philosophical issues with this, as it is unlikely that a DM will consider one individual to be entirely correct, with several individuals often contributing distinct complementary information. Another scheme involves weights based on outranking probabilities, *i.e.*, the probability that the next prediction of an individual will outperform those of all others. This method has been shown to be prone to overreaction (in terms of being overly influenced by short-term performance), and by solely considering ordinal rather than cardinal performance ranking it does not fully align with the Bayesian philosophy as it does not consider all the possible information. This outranking scheme may also motivate individuals to dishonestly report their probability distributions in the hope of increasing their weight. A technique that eliminates this motivation involves weights being calculated from proper scoring rules, yet most of these rules (in fact all of them except a logarithmic one) violate the Likelihood Principle. Even this logarithmic rule can potentially return negative weights, which is a clear shortcoming. Bordley & Woolf (1981) supply a method based upon minimising the variance of the combined distribution, but there is no clear rationale for doing so. Finally Barlow *et al.* (1986) built a scheme based on equal veridical weights and logarithmic scoring rules, under which weights are inversely proportional to a distance metric, with high weights awarded to those individuals with opinions similar to those around them. While there are some attractive elements to this approach (its Bayesian associations and the logic of the logarithmic scoring rule) it tends to give conservative weights, and is based upon the assumption of equal veridical weights, which can be seen as extremely restrictive (in addition to the potential philosophical debate against weights of this nature).

Cooke & Goossens (2008) use a vast collection of data sets, named the TU Delft Expert Judgment Data Base, to validate the classical method of Cooke (1991). This database consists of forty-five real world data sets pertaining to a broad spectrum of research fields, for example finance, geography, physics, real estate and health, with each data set having a varying amount of seed variables under consideration, as well as a varying amount of experts who provide their quantile opinions over these. Cooke & Goosen (2008) compare the performance of the classical approach to a collection of alternatives, primarily the simple method of assigning equal weights to all experts irrespective of their perceived reliability, as well as a scheme that only listens to the individual deemed most accurate (*i.e.*, gives her a weight of one) and disregards the opinions of the rest of the experts. The metric considered was a function of the calibration and information scores resulting from different forms of combined belief considered under varying conditions, and some justification was provided for the use of the classical approach. This study also demonstrated how the method could be applied to real data, *i.e.*, that it was implementable in practice in a wide range of realistic scenarios.

Clemen (2008) identifies some potential problems with the validation method that was used in this justification, and cited out-of-sample performance as a metric to be considered in order to remedy this. He demonstrated that, for a particular subset of the data sets used by Cooke & Goossens (2008), it was not conclusive if the classical method outperformed the simple equal weights approach or not, leading him to comment that "*it has been somewhat frustrating to consistently find the simple average performing so well empirically*". Throughout the literature there are examples of well-reasoned and rational performance-based weighting schemes unable to outperform this most basic approach. Clemen (2008) also questions if the classical method incentivises the user to report their genuine predictions, and presents a hypothetical situation in which an individual can potentially manipulate the scheme (and hence the weight that she is allocated) by providing a suitably broad prediction interval. He commented that he would like to see a full scale study carried out, taking into account his aforementioned concerns and preferably using the entire collection of available datasets for justification. Flandoli *et al.* (2011) discuss some conceivable shortcomings of the classical approach, primarily in relation to the forms of cross-validation methods that are used to validate it, recommending that dividing the data into training and test subsets would be a more suitable technique than simple leave-one-out cross-validation, as the latter may potentially be susceptible to biases.

Flandoli *et al.* (2011), proposes an augmented methodology called the Expected Relative Frequency (ERF) model, which involves the construction of triangular distributions from the lower and upper quantiles of a DM as well as her mode. The weight allocated to an expert is then found by integrating this distribution around a neighbourhood (having a size that must be defined by the DM) of the true seed value. Comparisons are made between the performance of the ERF method and the classical method, using a small subset of the data considered in Cooke & Goossens (2008) to conduct this investigation. It could not be clearly indicated which approach gave better results, but it appeared that the ERF model gave a more accurate estimate of some measure of central tendency for a variable, while the classical method gave a better indication of the inherent uncertainty (*i.e.*, the variance associated with this variable). In response to Clemen (2008), Eggstaff *et al.* (2014) conducted a large scale study using the complete TU Delft Expert Judgment Data Base, and successfully showed the merits of the classical approach in opposition to the set of alternatives considered. In addition to this justification they also demonstrated some interesting results concerning the optimal number of seed variables to be used in the process. They showed that this is approximately between five and eight seeds, with performance actually tending to deteriorate after this. They emphasise the importance of choosing the correct seed variables for consideration, because if seeds that are unrelated to the uncertainty inherent in the decision task are used then these may give false indications of the relevant predictive abilities of experts. We shall return to study the TU Delft Expert Judgment Data Base in Chapter 4, where we use it to provide some additional justification for the method discussed in Chapter 3.

We briefly comment that Eggstaff *et al.* (2013) discuss a novel application for the classical approach of Cooke (1991) to assess technical performance in an engineering context. In a setting where no relevant seed variables can be considered (*i.e.*, a unique problem with no appropriate comparisons available) they augment the traditional method to allow for dynamic updating, with previously witnessed values considered as seeds (once a sufficiently large amount of these have been observed, with equal weights used up to this point). This method is applicable in cases where the quantity of interest is expected to vary over time.

Karny & Guy (2004) make reference to combining beliefs in our framework of interest, where each DM herself is considered to be an expert, and can combine her opinions with those of her neighbours in an attempt to increase her own resulting decision quality. In the most simplistic scenario there are two DMs, denoted by P_1 and P_2, who are interested in some pertinent uncertain quantity θ that they have probabilistic beliefs over, given by $f_1(\theta)$ and $f_2(\theta)$ respectively. Karny & Guy (2004) advocate that combined beliefs should be given in the following form:

$$\hat{f}_1(\theta) = \alpha_1 f_1(\theta) + (1-\alpha_1) f_2(\theta) \quad (2.7)$$

$$\hat{f}_2(\theta) = (1-\alpha_2) f_1(\theta) + \alpha_2 f_2(\theta) \quad (2.8)$$

Here α_i is the weight assigned by P_i to her own belief, with $1-\alpha_i$ being the weight she assigns to that of her neighbour, with $0 \leq \alpha_i \leq 1$ for $i=1,2$. Note that if a DM wishes to ignore the opinion of her neighbour entirely she can give herself a dominating weight of one, *i.e.*, giving her neighbour a weight of zero. However, no explanation is provided as to how an individual may choose these weights or, perhaps more importantly, how these weights may be updated over time in light of decision returns witnessed. Karny & Guy (2004) discuss how this method can potentially be extended to a setting with more than two participants with belief sharing conducted in a pairwise iterative manner. However this method is not invariant to the order in which sharing takes place as beliefs shared between neighbours are functions (*i.e.*, weighted sums) of the beliefs of themselves and all the neighbours with whom they have previously shared beliefs. A formal mathematical proof of this statement is included in Appendix B. Much of the work that is presented in Chapter 3 can be seen as an attempt to generalise the work of Karny & Guy (2004) to a setting that consists of n participants, while proposing a method to attempt to solve the problems previously discussed, *i.e.*, of attaching weights to the opinions of each individual that are somehow a reflection of their relative reliability, and that can be augmented over time given new data. Our goal can thus be viewed as providing a non-arbitrary method of assessing the accuracy of neighbours that supplies appropriate weights to be used in a linear opinion pool.

The work that is presented within this thesis is highly applicable to a wide range of realistic problems across a broad spectrum of fields that pertain to risk. The illustrative example contained in Chapter 3 involves a collection of stockbrokers who want to communicate amongst each other to determine the behaviour of a stock price. This is just one of multiple potential implementations of the approach that we derive. A group of computer programmers may wish to pool their beliefs about the average number of bugs occurring per one thousand lines of code, in order to aid their individual decisions on whether to release their software or to continue testing it (*e.g.*, Wilson & McDaid, 2001). Several companies may wish to exchange their opinions on what proportion of a particular demographic (*e.g.*, males under twenty-one) buy a particular product (*e.g.*, computer games, laptops, jeans) to ascertain how large a quantity they should respectively produce. Nuclear power stations may wish to confer between each other as to the perceived risk of a fault occurring to assist in their decision of determining whether additional safety devices need be installed or not (*e.g.*, Starr, 1981). Medical practitioners may seek the opinion of peers as to the probability of a diagnosis being correct given some symptoms witnessed (and similar problems, *e.g.*, Cox, 2012). We include a medical example in Chapter 6 where interest lies in estimating the efficacy of a novel drug treatment. When we consider our results in Chapter 4 for the TU Delft Expert Judgment Data Base we see that our methodology can be implemented to assist in various problems spanning a vast range of research fields. In each case there are clearly risks present that the users wish to avoid, be it a loss of financial wealth, national safety or medical health. Note that in the situations outlined above each individual entity (be it a single person or be it a large multinational company) will have their own personal utility function so even if decisions are made using common beliefs different decisions may well be deemed optimal by different entities.

Chapter 3

The Plug-in Approach

In this chapter we describe a linear opinion pooling method, which we have termed the "Plug-in" (PI) approach. It permits DMs to learn over time, regarding both their own opinion about the uncertain quantity inherent within the decision problem and also their perception of the reliability of the opinions held by their neighbours. We provide discussion on the intuition underlying its calculations, its strong Bayesian associations and its asymptotic behaviour. Some desirable rationality properties that this method obeys are detailed, as well as an illustrated example of the PI approach applied in a financial setting. We derive the distribution that PI weights follow, provide comments on an extension to a more generalised framework, make comparisons with a similar technique and detail some limiations of our methodology.

3.1 The Plug-in Approach

3.1.1 Notation and Basics

We consider a setting consisting of n DMs, labeled $P_1, \ldots, P_n$, with $n \geq 1$. If $n = 1$ then the DM has no neighbours, *i.e.*, she herself is her only source of information, meaning the problem is solved as described in Section 2.3.1. Hence in non-trivial settings we consider the case where $n \geq 2$. There is a common uncertain quantity θ involved in the decision tasks of each DM. In this thesis we assume that θ is static (*i.e.*, it does not change over time dynamically, instead remaining constant and unknown throughout), that all DMs can express their opinions about θ via fully parameterised probability distributions, and that all DMs are willing to listen to the opinions of neighbours, incorporating these into their own decision processes in a non-competing environment. We denote by $f_i(\theta)$ the personal opinion of P_i, and write her combined belief, assimilating the opinions of all her neighbours with her own, as $\hat{f}_i(\theta)$, such that

$$\hat{f}_i(\theta) = \alpha_{i,1} f_1(\theta) + \ldots + \alpha_{i,n} f_n(\theta) \tag{3.1}$$

Much of what follows in this chapter pertains to choosing values of $\alpha_{i,1}, \ldots, \alpha_{i,n}$ that are deemed satisfactory in some sense. Two obvious properties that we constrain these weights to obey are strict positivity (*i.e.*, $\alpha_{i,j} > 0$ for all $i, j = 1, \ldots, n$) and summation to unity (*i.e.*, $\sum_{j=1}^{n} \alpha_{i,j} = 1$ for all $i = 1, \ldots, n$). The first of these conditions implies that the opinions of all DMs have some impact in the combined opinion, with weights potentially becoming arbitrarily small but always exceeding zero. The second property ensures that weights are normalised, leading to their straightforward interpretation. Note that while it is necessary that weight sum to one in order to guarantee that the combined distribution is a valid one, Genest & McConway (1990) provide a brief discussion on how in some specific cases weights may be permitted to be negative while still yielding valid combined distributions. In what follows we shall discuss what we desire weights to represent, and how this may be achieved.

The setting that we consider is one in which DMs enter into a sequence of decision tasks myopically, *i.e.*, they make a decision, see an outcome, and repeat this process. Over time DMs will notice that some neighbours are more accurate information sources than others, *i.e.*, the opinions that they provide seem to more closely mirror the witnessed reality than those of others. The seemingly logical reaction for a DM upon this realisation is to pay more heed to views proffered by neighbours whom she deems reliable, and to somewhat disregard those views of neighbours whom she deems unreliable. Intuitively this should lead to an increased decision quality, as decisions are being made using information sources that are believed, and indeed having shown themselves, to be trustworthy. We wish weights to be reflective of the perceived reliability associated by DMs with neighbours, with $\alpha_{i,j}$ revealing how accurate P_i considers the opinion of P_j to be in comparison with her other neighbours (and herself).

3.1.2 Updating Beliefs

As well as modifying the weights that they associate with the opinions of neighbours, DMs also desire to augment their own opinions in light of new information witnessed from the decisions that they have made. We assume that this updating adheres to the paradigm of Bayes (1763). Each DM considers her initial belief $f_i(\theta)$ to be her prior distribution. A return r is witnessed and an appropriate likelihood function $f(r|\theta)$ is implemented. This prior distribution and likelihood function are then multiplicatively combined to yield her posterior distribution $f_i(\theta|r)$. This expresses her opinions about θ in light of the new information r garnered, written as

$$f_i(\theta|r) = \frac{f(r|\theta)f_i(\theta)}{\int_\Theta f(r|\theta)f_i(\theta)\ d\theta} \tag{3.2}$$

We observe that the likelihood function $f(r|\theta)$ is independent of the decision maker P_i, *i.e.*, we assume it is common for all DMs, who do not control the underlying data generating mechanism. We also assume that all DMs construct prior distributions of the same functional form with (potentially) differently chosen hyperparameters *e.g.*, all DMs have Gamma priors, but their respective shape and scale parameters may differ. This is not a requirement, but we use it in what follows for ease of interpretation, and indeed as it may be a sensible course of action. The (normalising) denominator of Equation (3.2) is independent of θ given the integral. Therefore we frequently rewrite the posterior distribution up to a normalising constant of proportionality as

$$f_i(\theta|r) \propto f(r|\theta)f_i(\theta) \tag{3.3}$$

Note that Bayesian updating is the unique updating strategy that prevents against Dutch Books in a dynamic context, as discussed for instance in Skyrms (1993). Beliefs over $f_i(\theta)$ may be of a form such that conjugacy can be applied, with the product of the likelihood function and the prior distribution leading to a posterior distribution of the same functional form as the prior but with different hyperparameters, *i.e.*, conjugate priors. This is often done for mathematical convenience and elucidation, ensuring tractability and ease of calculations, although it is by no means a necessity, with methods such as Markov Chain Monte Carlo (first accredited to Metropolis *et al.*, 1953) implementable in intractable cases. Probability distributions from the "exponential family" are all conjugate to some likelihood function, and are generally considered suitably flexible to model a wide range of realistic phenomena while simultaneously guaranteeing computational tractability. As the primary research aims of this thesis are in decision theory rather than computational statistics we consider three common conjugate cases for illustration throughout. Beta-Binomial conjugacy is suitable when interest is in the number of successes in a set of (independent and identically distributed) Bernoulli trials, *e.g.*, the number of penalties scored by a footballer in ten attempts. Poisson-Gamma conjugacy is applicable when inference is carried out on the number of occurrences of some event in a fixed period of time, *e.g.*, the number of earthquakes hitting San Francisco over a decade. Finally, Normal-Normal conjugacy is used to learn about the mean of some continuous underlying process, with the associated variance assumed to be known in the simplest case, *e.g.*, the average temperature in Dublin in January.

In the framework that we operate within updating is performed upon individual, rather than combined, beliefs to ensure that DMs can always extract their own personal belief from the linear combination. We discuss Bayesian interpretations of this in Section 3.2. We formally distinguish between $f_i(\theta)$ and $\hat{f}_i(\theta)$; the former reflects the personal probabilistic beliefs of P_i while the latter is a tool that she uses in her decision task, given that she believes herself to not be infallible and is willing to pay heed to the beliefs of others in the hope of increasing her decision quality.

For simplicity we assume that all DMs observe a common return at a particular epoch, *e.g.*, a collection of DMs who are interested in the value that a stock price takes on the last day of May with this value being the same irrespective of how many DMs are involved. Yet there may be cases when this assumption can be viewed as too restrictive, with each individual witnessing their own return as a result of the decision that she makes, *i.e.*, if we have n DMs then there may be n new pieces of information available following an epoch. Commonly all DMs will be interested in the same parameter but observe different noisy realisations of this, *e.g.*, if θ is the proportion of defective goods produced by a wholesaler and different DMs purchase different quantities of goods witnessing different defective proportions respectively. Primarily we consider the simplified case of a common return being witnessed, although in Section 3.8 we demonstrate how straightforward modifications of the common return approach enables implementation in this more generalised setting. This allowance of multiple differing returns further differentiates this method from the setting of a single DM seeking the opinion of a collection of non-decision-making experts (as discussed in Section 2.7), as the decisions made by DMs will impact upon the information that they witness and hence the combined opinions of their neighbours.

As well as a probability distribution over θ we assume that DMs can specify utility functions over potential decision returns r, written as $u_i(r)$ for P_i. In Equation (2.2) we introduced the expected utility associated with a decision, and declared in Equation (2.3) that the optimal decision was that which maximised this. Here we subtly modify this scheme, determining the optimal decision for P_i to be d^* such that

$$\begin{aligned} d^* = \arg\max_{d \in \mathcal{D}} \mathbb{E}[u_i(d)] &= \arg\max_{d \in \mathcal{D}} \Big[\int_\Theta u_i(d,\theta) \hat{f}_i(\theta)\ d\theta \Big] \\ &= \arg\max_{d \in \mathcal{D}} \Big[\int_\Theta u_i(d,\theta) \sum_{j=1}^{n} \alpha_{i,j} f_j(\theta)\ d\theta \Big] \\ &= \arg\max_{d \in \mathcal{D}} \Big[\sum_{j=1}^{n} \alpha_{i,j} \int_\Theta u_i(d,\theta) f_j(\theta)\ d\theta \Big] \qquad (3.4) \end{aligned}$$

For P_i, this is a function of her amalgamated belief $\hat{f}_i(\theta)$ and her utility function $u_i(r)$. DMs explicitly take the beliefs of neighbours into account by using $\hat{f}_i(\theta)$ rather than $f_i(\theta)$ in Equation (3.4). The expected utility P_i assigns to a decision is a linear combination of the expected utilities assigned to it (under her own utility function) by each of her neighbours (and herself), with bigger weights give to those DMs who are deemed reliable. 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 which is meant by the context in which it arises.

3.1.3 Updating Weights

An initial issue to consider is how weights should be specified at the first epoch. In the PI approach it is assumed that *a priori* participants have no knowledge about the accuracy of the opinions of neighbours before they begin making decisions. Hence in light of any better evidence to the contrary DMs initialise weights using the Laplacian Principle of Indifference (Laplace, 1812), assuming themselves and their neighbours to be equally reliable. All individual probability distributions in Equation (3.1) are assigned an equal weight, *i.e.*, at the first decision epoch the combined belief of P_i is

$$\hat{f}_i(\theta) = \frac{1}{n} f_i(\theta) + \ldots + \frac{1}{n} f_n(\theta) \qquad (3.5)$$

Note if a DM does have prior information about the reliability of her neighbours she may allocate initial weights based on this, and proceed in the same manner outlined below. However in what follows we shall focus on the case of equal prior weights. Each DM now determines which decision is optimal for her by combining this with her utility function as in Equation (3.4). This decision is made and some decision outcome observed. The weighting process based on the perceived reliability of DMs now begins. DMs want a method to compare the outcome that occurred to those predicted by neighbours. The return r witnessed is a realisation of the random variable R which follows a distribution according to the true data generating mechanism $f(R = r|\theta)$. We wish to find the probability density each DM placed on r prior to it occurring, *i.e.*, "plugging in" r to their prior predictive distributions. This gives

$$w_i = f_i(R = r) = \int_{\Theta} f(R = r|\theta) f_i(\theta)\ d\theta \tag{3.6}$$

The PI weight of P_i is denoted by w_i. Rational DMs will want to give higher weights to those neighbours with high PI weights (*i.e.*, those who appear reliable) than to those with low weights. For the three conjugate cases previously introduced these PI weights have tractable closed forms given in Equations (3.7), (3.8) and (3.9) for Beta-Binomial (with m the number of trials, and α_i and β_i the respective number of hypothetical successes and failures witnessed by P_i), Poisson-Gamma (with α_i and β_i the respective scale and shape hyperparameters of P_i) and Normal-Normal conjugacy (with σ^2 the know variance, and m_i and s_i^2 the prior mean and variance of P_i) respectively.

$$w_i = \frac{\Gamma(m+1)}{\Gamma(m-r+1)\Gamma(r+1)} \frac{\Gamma(\alpha_i+\beta_i)}{\Gamma(\alpha_i)\Gamma(\beta_i)} \frac{\Gamma(\alpha_i+r)\Gamma(\beta_i+m-r)}{\Gamma(\alpha_i+\beta_i+m)} \tag{3.7}$$

$$w_i = \frac{\alpha_i^{\beta_i}}{\Gamma(r+1)\Gamma(\beta_i)} \frac{\Gamma(r+\beta_i)}{(\alpha_i+1)^{r+\beta_i}} \tag{3.8}$$

$$w_i = \frac{1}{\sqrt{2\pi(\sigma^2+s_i^2)}} \exp\Big(-\frac{(r-m_i)^2}{2(\sigma^2+s_i^2)}\Big) \tag{3.9}$$

We include a Beta-Binomial example in Table 3.1 which gives the prior hyperparameters for four DMs and their corresponding means and standard deviations. We imagine seven successes are seen in ten trials, and use Equation (3.7) to find the PI weights associated with DMs. We see P_4 has the smallest weight, heavily penalised for her inaccurate mean prediction and associated high confidence (*i.e.*, small standard deviation). Note that even though the mean prediction of P_3 is closer than that of P_2 to the Maximum Likelihood Estimate of 0.7, the latter still has a higher weight, arising from her decreased uncertainty, *i.e.*, there is a trade-off between accuracy (in terms of means) and confidence (in terms of standard deviations). Fig. 3.1 presents a graphical interpretation of this information, illustrating the prior predictive distributions of DMs and where PI weights arise from.

Table 3.1: The prior hyperparameters, means and standard deviations of DMs, and their PI and normalised weights after seeing seven successes in ten trials.

P_j	α_j	β_j	$\mathbb{E}_j(\theta)$	$\mathrm{SD}_j(\theta)$	w_j	$\alpha_{i,j}$
P_1	3	3	0.5	0.189	0.120	0.251
P_2	6	4	0.6	0.148	0.172	0.359
P_3	4	2	0.667	0.178	0.160	0.334
P_4	10	20	0.333	0.085	0.027	0.056

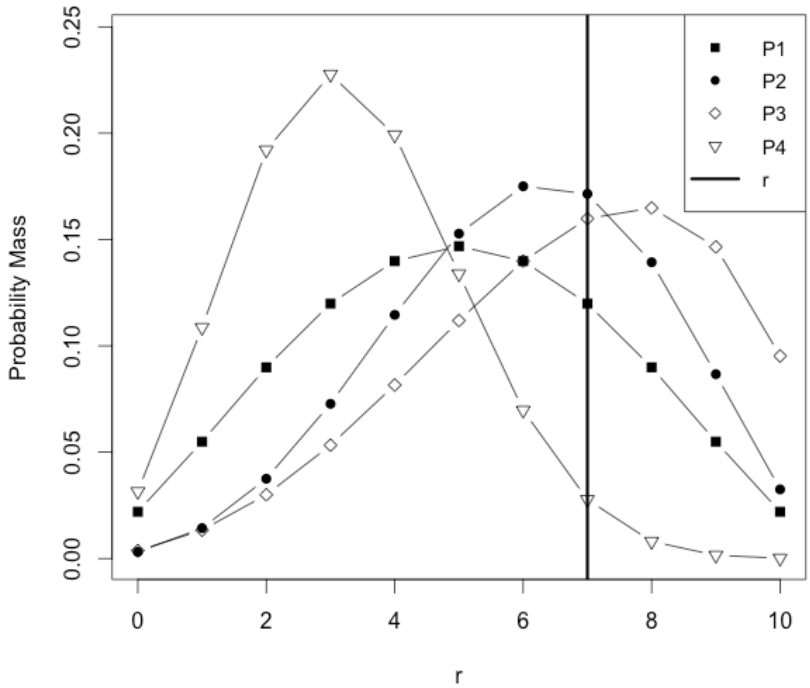

Fig. 3.1: The prior predictive distributions of DMs from their priors in Table 3.1. The data observation is included, with the point of intersection of this vertical line and the prior predictive distribution of a DM yielding her PI weight.

We want to relate the PI weights in Equation (3.6) to the weights in Equation (3.1), *i.e*, if P_i finds w_j then what normalised weight ($\alpha_{i,j}$) should she assign to P_j?

We propose a multiplicative scheme with weights updated from $\frac{1}{n}$ to $\alpha^*_{i,j}$ such that

$$\alpha^*_{i,j} = \frac{\frac{1}{n} w_j}{\sum_{k=1}^{n} \frac{1}{n} w_k} = \frac{w_j}{\sum_{k=1}^{n} w_k} \tag{3.10}$$

Weights are augmented from initially including no measure of the merits of information sources, to accounting for the contrast between their predictions and the revealed noisy realisation of reality. We see this in Table 3.1 with the two most accurate individuals (P_2 and P_3) afforded the larger normalised weights while the seemingly highly inaccurate P_4 is allowed to only make a small contribution to the combined belief at the next epoch. The type of scenario in which this approach is most applicable entails numerous decisions being made and returns witnessed over time, implying that the distributions of DMs and the corresponding associated weights will be modified at each epoch. Loosely speaking, reliability is a flexible concept. DMs who initially appear very accurate may in fact be the opposite, with the first realisation leading to them being considered accurate being a fluke occurrence (*i.e.*, one in the tail probabilities of the data generating mechanism). Conversely DMs with highly accurate opinions may initially receive a low weight if the return witnessed is an outlier. Given this it is clearly important to repeatedly modify the allocated weights over time, augmenting in light of each new piece of evidence that arises. We write a more generalised version of Equation (3.10), where $\alpha^*_{i,j}$ is the new updated weight and $\alpha_{i,j}$ is the previously associated weight with w_j the PI weight for the most recent realisation, as

$$\alpha^*_{i,j} = \frac{w_j \alpha_{i,j}}{\sum_{k=1}^{n} w_k \alpha_{i,k}} \tag{3.11}$$

We can view this scheme as being Markovian, with only the most recent normalised weight ($\alpha_{i,j}$) being used in the calculation of the new normalised weight ($\alpha^*_{i,j}$). Given the weight at an epoch t, the weights at epochs $t+1$ and $t-1$ are conditionally independent. This weighting scheme is dynamically evolving, with elements of the past constantly incorporated into new weights. Issues arise if all DMs have assigned a probability density of zero to the event that occurred, leading to division by zero in Equation (3.11), an invalid operation. Allocating probability densities of zero is strongly advocated against by Cromwell's Rule (Lindley, 1991) as it implies no amount of evidence will persuade the DM to change their mind, *i.e.*, a zero in the prior distribution will always lead to a zero in the posterior distribution, irrespective of the magnitude of contrary data contained in the likelihood function. However this problem will rarely occur in practice, as in most distributions there will always be a non-zero (if in some cases negligibly small) probability density associated with any feasible event occurring. In addition to this, if all individuals did simultaneously assign probabilities of zero to an event that did occur then they all appear equally unreliable, *i.e.*, weights would remain unchanged from the previous epoch.

Gneiting & Raftery (2007) discuss the concept of strictly proper scoring rules. These are methods of assessing the performance of predictions which return the highest possible value (*i.e.*, indicate the greatest merit) when a prediction is equal to the realised value. We consider the PI weight of Equation (3.6) to be our scoring rule. A criticism of this form of scoring rule is that *"it is not sensitive to distance, meaning that no credit is given for assigning probabilities to values near, but not identical to, the one materialising"* (Gneiting & Raftery, 2007). Yet in our setting where individuals constantly update their distributions over θ (and hence by association their predictive distributions) this is not as important an issue, as their distributions will shift accordingly over time in light of the new information witnessed. We include an example of this in Fig. 3.2. Recall that rather than seeing some definitive true parameter value DMs will witness a set of noisy realisations of this, with the seemingly "most accurate" individual potentially changing repeatedly over time given the new evidence available.

3.2 Bayesian Relationship

In Section 2.7 we introduced the concept of fully Bayesian opinion pooling. Under this scheme, an individual P_i could combine her opinion $f_i(\theta)$ with those of her neighbours, $f_j(\theta)$ for $j = 1, \ldots, n$ and $j \neq i$, viewing these opinions as data observations. This led to her posterior distribution over θ being determined as

$$f_i(\theta|\{f_j(\theta)\}_{j\neq i}) = \frac{f(\{f_j(\theta)\}_{j\neq i}|\theta) \times f_i(\theta)}{\int_\Theta [f(\{f_j(\theta)\}_{j\neq i}|\theta) \times f_i(\theta)]\ d\theta} \tag{3.12}$$

It is unclear how the likelihood function (*i.e.*, the first term in the numerator) can be constructed. As previously discussed this difficultly in specifying an appropriate form of likelihood function is the primary reason why updating of this nature is rarely implemented in practice. Hence while it may be desirable to act in a fully Bayesian manner, enabling us to utilise the full benefits associated with this paradigm (*e.g.*, straightforward posterior analysis) it is frequently beyond our computational ability

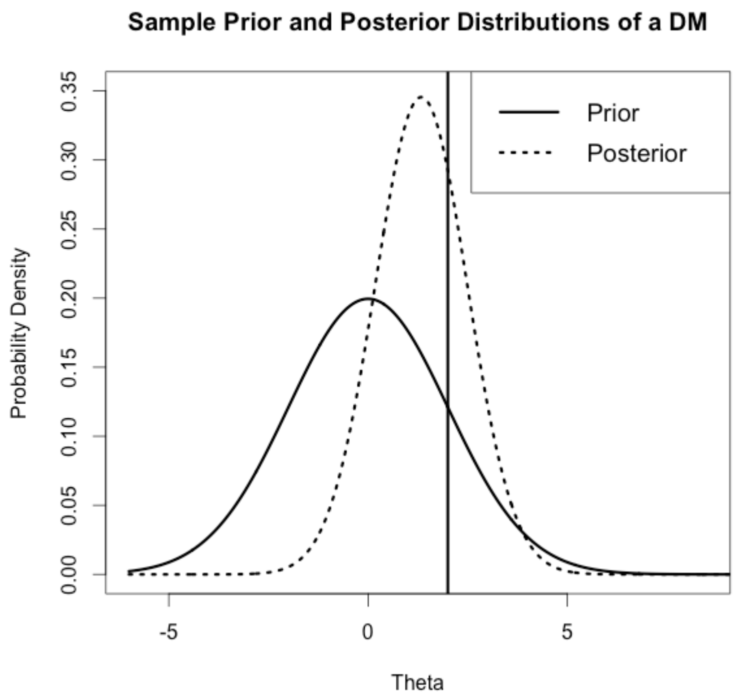

Fig. 3.2: The Normal prior and posterior distributions of a DM, with a vertical line denoting an observation. Opinions shift to the right given new data. Hence while she may initially seem quite inaccurate she appears far more reliable at the next epoch.

to do so. Recall that under the PI approach each DM updates her own probability distribution in a strictly Bayesian manner. In this section we discuss two desirable Bayesian properties that our proposed methodology obeys.

An argument for the functional form of the reweighting scheme in Equation (3.11) is that it ensures updating of a linear combined opinion is conducted in a manner adhering to the Bayesian paradigm, as illustrated for instance by Lee (2012), as part of a discussion on mixture distributions. Suppose the combined belief of P_i is $\hat{f}_i(\theta)$ and that a return r is witnessed with corresponding PI weights of $w_1, \ldots, w_n$ for the n DMs. The posterior distribution $\hat{f}_i(\theta|r)$ can be written as

$$\begin{aligned}\hat{f}_i(\theta|r) &= \frac{f(r|\theta)\hat{f}_i(\theta)}{f(r)} \\ &= \frac{f(r|\theta)\sum_{j=1}^{n}\alpha_{i,j}f_j(\theta)}{\sum_{j=1}^{n}\alpha_{i,j}\int_{\Theta}f(r|\theta)f_j(\theta)\ d\theta} \\ &= \frac{\sum_{j=1}^{n}\alpha_{i,j}f(r|\theta)f_j(\theta)}{\sum_{j=1}^{n}\alpha_{i,j}\int_{\Theta}f(r|\theta)f_j(\theta)\ d\theta}\end{aligned}$$

$$
\begin{aligned}
&= \frac{\sum_{j=1}^{n} \alpha_{i,j} f_j(\theta|r) \int_{\Theta} f(r|\theta) f_j(\theta)\ d\theta}{\sum_{j=1}^{n} \alpha_{i,j} \int_{\Theta} f(r|\theta) f_j(\theta)\ d\theta} \\
&= \frac{\sum_{j=1}^{n} \alpha_{i,j} w_j f_j(\theta|r)}{\sum_{j=1}^{n} \alpha_{i,j} w_j} \\
&\propto w_1 \alpha_{i,1} f_1(\theta|r) + \ldots + w_n \alpha_{i,n} f_n(\theta|r) \qquad (3.13)
\end{aligned}
$$

Hence we see that a strong proponent for the reweighting scheme advocated in Equation (3.11) is that it is unique in providing updating in a manner deemed coherent within a Bayesian framework, and thus can be viewed as a natural choice in some sense assuming one agrees with this paradigm. The combined posterior distribution is a linear combination of individual posterior distributions with the associated weights determined by the PI method being precisely those guaranteeing rational Bayesian updating.

It is also desirable for any coherent updating approach to be consistent with the Likelihood Principle, discussed and illustrated succinctly in Lindley & Phillips (1976). Suppose that over t decision epochs a DM witnesses a collection of returns, $r_1, \ldots, r_t$. We wish to show that the normalised weight she will be assigned after the last return has been observed is invariant to permutations in the order in which these returns are witnessed. This is an appealing property for an approach to obey, arising from the exchangeability that is inherent within Bayesian updating. Let us denote by $\alpha_{i,j}^{(t)}$ the weight assigned by P_i to P_j having seen t returns. All DMs are initially given equal weights at the first epoch, *i.e.*, $\alpha_{i,j}^{(0)} = \frac{1}{n}$ for $j = 1, \ldots, n$ in an environment containing n individuals. We write $w_j^{(t)}$ for the PI weight of P_j having seen the t^{th} return r_t. We note that a DM's normalised weight can be written as a product of her initial weight and all her PI weights up to this point, *i.e.*,

$$
\begin{aligned}
\alpha_{i,j}^{(t)} &\propto w_j^{(t)} \alpha_{i,j}^{(t-1)} \\
&\propto w_j^{(t)} [w_j^{(t-1)} \alpha_{i,j}^{(t-2)}] \\
&\propto w_j^{(t)} [w_j^{(t-1)} [w_j^{(t-2)} \alpha_{i,j}^{(t-3)}]] \\
&\propto \vdots \\
&\propto \prod_{k=1}^{t} w_j^{(k)} \alpha_{i,j}^{(0)} \\
&= \frac{1}{n} \prod_{k=1}^{t} w_j^{(k)} \qquad (3.14)
\end{aligned}
$$

Recall that the PI weight for a DM is the value that her prior predictive distribution takes at the return that was just witnessed. We write R_t for the random variable occurring at epoch t, with r_t being the specific realisation of this that is observed. The data generating mechanism is the same at all epochs by our assumption of the static nature of θ, *i.e.*, $f(R_t = r) = f(R_{t+1} = r)$ for any return r and epoch t. The PI weight at epoch t is written as

$$\begin{aligned} w_j^{(t)} &= f_j(R_t = r_t) \\ &= \int_\Theta f(R_t = r_t|\theta) f_j(\theta|r_1, \ldots, r_{t-1}) \; d\theta \end{aligned} \tag{3.15}$$

Repeatedly substituting Equation (3.15) into Equation (3.14) gives the following:

$$\begin{aligned} \alpha_{i,j}^{(t)} &\propto \frac{1}{n} \prod_{k=1}^{t} w_j^{(k)} \\ &= \frac{1}{n} \prod_{k=1}^{t} \Big[\int_\Theta f(R_k = r_k|\theta) f_j(\theta|r_1, \ldots, r_{k-1}) \; d\theta \Big] \\ &= \frac{1}{n} \prod_{k=1}^{t} f_j(R_k = r_k|r_1, \ldots, r_{k-1}) \\ &= \frac{1}{n} f_j(R_1 = r_1) f_j(R_2 = r_2|R_1 = r_1) \ldots f_j(R_t = r_t|R_1 = r_1, \ldots, R_{t-1} = r_{t-1}) \\ &= \frac{1}{n} f_j(R_1 = r_1, \ldots, R_t = r_t) \end{aligned} \tag{3.16}$$

The last line follows from the law of total probability. We may consider r_0 to be the information that a DM bases her prior opinion upon (which is excluded when necessary for notational convenience). The above identity is invariant to the order in which the values $r_1, \ldots, r_t$ are witnessed, *i.e.*, if they were observed in a permuted order $r_{\sigma(1)}, \ldots r_{\sigma(t)}$ then this would result in

$$\alpha_{i,j}^{(t)} \propto \frac{1}{n} f_j(R_1 = r_{\sigma(1)}, \ldots, R_t = r_{\sigma(t)}) \tag{3.17}$$

Equation (3.17) is equal to Equation (3.16) due to the conditional independence of returns given θ and the exchangeability previously mentioned. It is straightforward to hence infer that the combined posterior distribution, $\hat{f}_i^{(t)}(\theta|\cdot)$ say, is also invariant to the order in which returns are witnessed, as it is a linear combination of weights that have this invariance (as shown above) and Bayesian posterior distributions (which inherently possess the exchangeability property themselves). In conclusion, we have seen that while the PI approach does not equate to a fully Bayesian approach (with opinions of neighbours incorporated into an appropriate likelihood function) its linear technique does adhere to two key Bayesian ideals. The combined posterior distribution is a linear combination of individual updated posterior distributions with associated weights obeying Bayes Theorem, and the approach is invariant to the order in which data is witnessed (*i.e.*, exchangeability).

3.3 Moments of PI Distribution

Interest lies in a probability distribution that is a linear combination (*i.e.*, a mixture) of individual probability distributions. The distribution $f_i(\theta|\cdot)$ of each P_i has some mean and variance, μ_i and σ_i^2, which are functions of its hyperparameters. What is the mean and variance of the combined distribution $\hat{f}_i(\theta|\cdot)$? Generic rules for moments of linear combinations of random variables do not hold in this case, as our concern is in a linear combination of probability distributions over a common random variable θ rather than a linear combination of distributions over distinct random variables, $\theta_1, \ldots, \theta_n$. In general a mixture of distributions that are all of some common distributional form is not guaranteed to be of that distributional form itself, *e.g.*, a mixture of Normal distributions over θ is not ensured to be a Normal distribution. There may be cases when it will be, for instance if one DM is given a weight of one and all others are given weights of zero, or if all individual distributions are identical, but this is not generally precisely the case (McLachlan & Peel, 2000, contains more discussion on this topic). We consider the first moment of $\hat{f}_i(\theta|\cdot)$, *i.e.*, its mean, given as

$$\begin{aligned}\mathbb{E}_{\hat{f}_i}(\theta|\cdot) &= \int_\Theta \theta \hat{f}_i(\theta|\cdot)\ d\theta \\ &= \sum_{j=1}^n \alpha_{i,j} \int_\Theta \theta f_j(\theta|\cdot)\ d\theta \\ &= \sum_{j=1}^n \alpha_{i,j}\mu_j \qquad (3.18)\end{aligned}$$

The mean of the combined distribution is a linear combination of the individual means comprising this. As $\hat{f}_i(\theta|\cdot)$ is a convex combination of distributions its mean is respectively bounded below and above by the minimum and maximum means resulting from these distributions, *i.e.*, $\mathbb{E}_{\hat{f}_i}(\theta) \in [\min_{j=1,\ldots,n} \mu_j, \max_{j=1,\ldots,n} \mu_j]$. This seems an intuitive and important coherency property. Its second moment, its variance, is given by

$$\mathrm{Var}_{\hat{f}_i}(\theta|\cdot) = \Big[\int_\Theta \theta^2 \hat{f}_i(\theta|\cdot)\ d\theta\Big] - [\mathbb{E}_{\hat{f}_i}(\theta|\cdot)]^2$$

$$\Big[\sum_{j=1}^{n}\alpha_{i,j}\int_{\Theta}\theta^2 f_j(\theta|\cdot)\ d\theta\Big] - [\mathbb{E}_{\hat{f}_i}(\theta|\cdot)]^2$$
$$= \sum_{j=1}^{n}\alpha_{i,j}(\sigma_j^2+\mu_j^2) - (\sum_{j=1}^{n}\alpha_{i,j}\mu_j)^2 \tag{3.19}$$

The combined variance is a weighted sum of individual variances plus an additional correction term (zero if the distributions of all DMs have identical means and strictly positive if not) to account for the difference between individual means and the combined mean. If DMs have diverse views over θ then this will lead to a large combined variance even if their individual variances are relatively small. Finally we comment that the probability density that a combined distribution places on a particular value is a linear combination of the density that each individual distribution places upon this value. We illustrate this in Fig. 3.3, in a case where $f_1(\theta) \sim \mathcal{N}(1, 2^2)$ and $f_2(\theta) \sim \mathcal{N}(-2, 3^2)$, with $\alpha_{i,1} = 0.7$ and $\alpha_{i,2} = 0.3$ for $i = 1, 2$. For example as we have $f_1(\theta = 2) = 0.176$ and $f_2(\theta = 2) = 0.055$ this implies $\hat{f}_i(\theta = 2) = 0.7(0.176) + 0.3(0.055) = 0.1397$. Using Equations (3.18) and (3.19) we can see that $\hat{f}_i(\theta)$ is such that $\mathbb{E}_{\hat{f}_i}(\theta) = 0.1$ and $\text{Var}_{\hat{f}_i}(\theta) = 7.39$ with this variance taking into account the respective deviation of the DM's means from the combined one. We observe that this combined distribution is clearly not itself Normally distributed, with a longer left tail resulting from the increased uncertainty of P_2.

3.4 Asymptotic Behaviour

We briefly comment on the behaviour of individual posterior distributions and the weights that these distributions are assigned as the number of returns witnessed grows large. As the amount of data that is observed increases the determination of DM posteriors will be dominated by this, *i.e.*, the prior opinions of individuals will become gradually outweighed by the actual information that they have viewed. Hence in the limit (*i.e.*, as the number of returns witnessed tends towards infinity) all DM posterior distributions will tend towards an identical distribution, with this distribution being an increasingly accurate model of θ by the Law of Large Numbers. Note that this statement is based upon the assumption that a DM does not supply a degenerate prior, *i.e.*, a prior placing a probability mass of one on a particular outcome and zero on all others, as if so this distribution will not augment in light of new data.

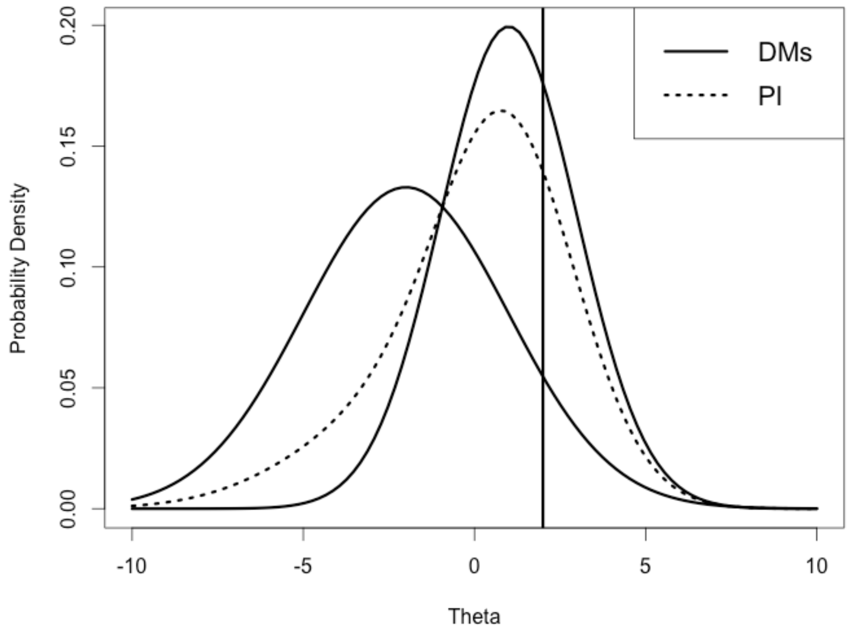

Fig. 3.3: Individual and combined distributions with vertical line denoting $\theta = 2$.

Genest & McConway (1990) comment on weight updating of the form discussed in this chapter (and reference Roberts, 1965, as an early discussant of it), and provide an interesting formal asymptotic result. If we denote by θ_0 the true value of the unknown quantity θ (which of course will be unknown to all DMs throughout the process) then it is demonstrated how in the limit an individuals normalised weight is proportional to the product of their initial weight (which we often assume to be $\frac{1}{n}$) and the density that their prior distribution placed on θ_0, *i.e.*,

$$\lim_{k \to \infty} \alpha_{i,j}^{(k)} \propto \alpha_{i,j}^{(0)} f_j(\theta_0) \tag{3.20}$$

We include an illustration of this phenomenon in Fig. 3.4 with the details concerning prior distributions and the data generating mechanism included in the caption. We comment that in this instance each return was the number of successes in five Bernoulli trials. If this number of trials per epoch was to decrease (*i.e.*, if information was observed at a slower rate) then this convergence would likely occur at a slower rate, while a greater amount of trials per epoch would likely increase the rate of convergence.

Finally we comment that as all DM posterior distributions converge towards a common distribution, an identical combined distribution will be yielded in the limit regardless of weights allocated, *i.e.*, any convex combination of a common distribution will simply yield this common distribution. Nevertheless the above asymptotic result concerning weights highlights how an individual's prior distribution has a lasting impact on their weight, and emphasises the care that they should take in specifying it.

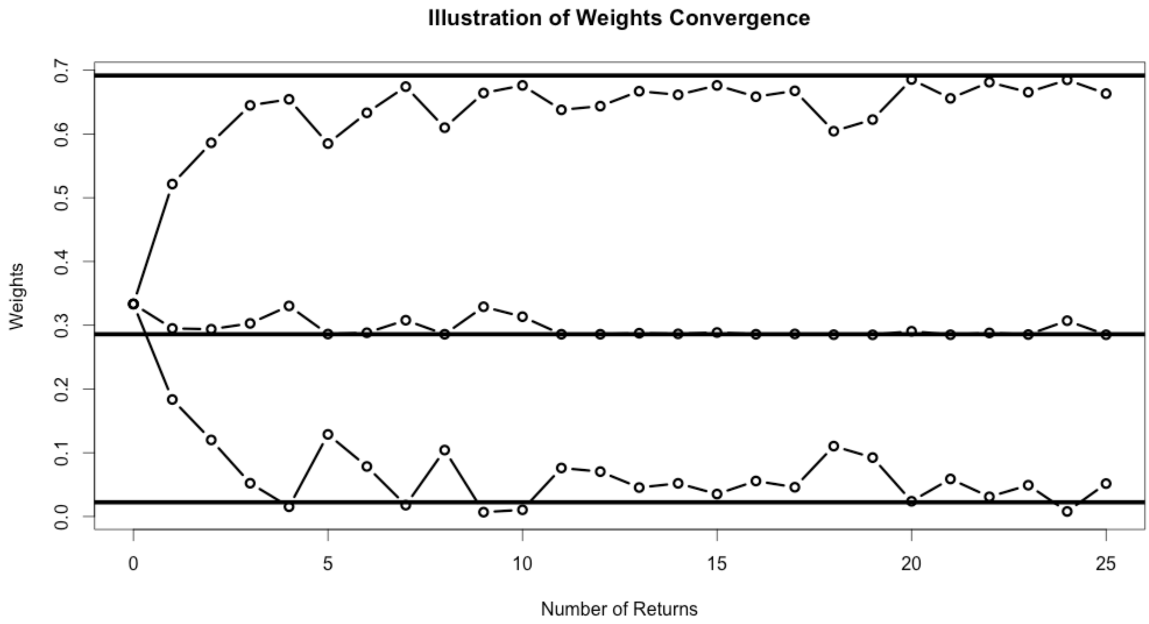

Fig. 3.4: Suppose we have three DMs and a data generating mechanism $R \sim \text{Bin}(5, \theta)$ with $f_1(\theta) \sim \text{Beta}(1, 2)$, $f_2(\theta) \sim \text{Beta}(4, 4)$ and $f_3(\theta) \sim \text{Beta}(3, 10)$. When the true value of θ is 0.6 we can clearly see below that the weights of P_1, P_2 and P_3 respectively converge to their limiting values of 0.29, 0.69 and 0.02, and that this convergence becomes evident quickly.

3.5 Properties and Initial Justifications

As previously discussed the PI approach appears a promising technique as theoretically the associated decision quality should substantially improve over time as the relevant information that decisions are based upon becomes increasingly accurate. It also appears fair to those involved, as individuals are initially assumed to be equally reliable, with increases/decreases in their allocated weight directly proportional to (an objective measure of) how accurate they have previously shown themselves to be. The PI approach takes the previous weight of a DM into account when determining her weight at the next epoch, as Equation (3.11) is a function of the previously allocated weight, *i.e.*, both present (*i.e.*, last realised) accuracy as well as all past performance history are incorporated. We observe that all DMs will have the same combined beliefs $\hat{f}_i(\theta)$

at each individual epoch (assuming they initialise with equal weights) as the weights are determined in a highly 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 $\alpha_{k,j}$. An inductive proof of this is given in Appendix B. When we discuss subjective weighting measures in Chapter 7 we see that this is generally not the case, with different DMs allocating diffuse weights to the same neighbour despite witnessing common information. Below we detail some desirable criteria that the PI approach obeys, demonstrating an underlying rationality and coherency. These are logical properties that one would naturally want a linear opinion pooling methodology of this ilk to adhere to.

- **Property 1**: $w_j \geq 0$ for all $j = 1, \ldots, n$, with $w_j = 0$ if and only if $f_j(R = r) = 0$.
- **Property 2**: If $\alpha_{i,j} < \alpha_{i,k}$ and $w_j < w_k$ then $\alpha^*_{i,j} < \alpha^*_{i,k}$.
- **Property 3**: If $\alpha_{i,j} = \alpha_{i,k}$ and $w_j = w_k$ then $\alpha^*_{i,j} = \alpha^*_{i,k}$.
- **Property 4**: If $\alpha_{i,j} < \alpha_{i,k}$ and $w_j > w_k$ then any of the following may occur depending on differences between initial weights and updated reliability measures:
  - $\alpha^*_{i,j} < \alpha^*_{i,k}$
  - $\alpha^*_{i,j} = \alpha^*_{i,k}$
  - $\alpha^*_{i,j} > \alpha^*_{i,k}$

Proofs are included in Appendix B. Here we briefly interpret the meaning of these properties and discuss why they are desirable criteria for a weighting scheme to possess. Property 1 is relatively trivial, stating that PI weights will always be non-negative and will only equal zero if a DM places no probability density on a (witnessed) outcome occurring. Property 2 states that if P_k is considered more reliable than P_j (based on respective past performances) and then at the next epoch P_k is once again deemed more reliable (*i.e.*, has a larger PI weight) then P_k will still be considered more reliable than P_j. Property 3 is an equality version of Property 2. Lastly Property 4 states that if P_k is considered more reliable than P_j, but that the beliefs of P_j are valued as more accurate at the next epoch (*i.e.*, she has a higher PI weight) then which DM is considered the most reliable after this is determined by the existing discrepancy between their previous normalised weights and the discrepancy between their PI weights.

3.6 Example

Consider a setting consisting of five DMs, $P_1, \dots, P_5$. Each has her own opinion about the true value of θ, which is a latent parameter pertaining to stock performance (and hence the profit or loss resulting from decisions made), and must decide whether to enter into a long forward on the stock (d_1) or not (d_2). Entering into a long forward entails agreeing to buy a stock at a fixed "expiry" time in the future for a fixed "strike" price. If the strike price exceeds the actual value of the stock at the expiry time then the DM has made a loss (as they are buying the stock for more than it is worth), if not then they have made a profit. It is clear that there must be uncertainty over θ, as if not decision making would be trivial, *i.e.*, a DM would always enter the long forward contract if θ is positive, and never do so if θ is negative. In our myopic scenario DMs must decide whether to enter into a long forward or not at a succession of epochs. There is obviously inherent risk, as DMs do not know *a priori* if they will make a profit or a loss in each trade. We assume for convenience that returns are Normally distributed with an unknown mean θ and a known variance of 2, *i.e.*, $R \sim \mathcal{N}(\theta, 2)$. We assume this unknown mean has true value of -2, *i.e.*, on average DMs will make a loss. Participants have prior beliefs over θ which themselves are Normally distributed, *i.e.*, Normal-Normal conjugacy. The prior beliefs of P_i over θ are represented by $f_i(\theta) \sim \mathcal{N}(m_i, s_i^2)$. The decisions that DMs make over whether to enter into 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 influenced by their utility functions and the initial fortune that they have prior to making any decisions. Their opinions about θ, utility functions, and starting fortunes (γ_i) are given in Table 3.2.

In this example we contrast the decisions that would be made by DMs using the PI approach and by the same DMs if they solely heeded their own opinion. We also compare what decisions they would deem optimal using two alternative linear pooling techniques: the Equal Weights (EQ) and Most Reliable (MR) methods. We will discuss these alternatives in depth in Chapter 4, but only comment briefly for now to note that the former involves equal weights being assigned over time to each DM irrespective of accuracy, with the latter entailing a weight of one given to the DM deemed most reliable at an epoch (*i.e.*, with the highest PI weight) and a weight of zero to all other DMs.

Table 3.2: DM information for our financial example.

P_i	$f_i(\theta)$	γ_i	$u_i(\gamma_i + r)$
P_1	$\mathcal{N}(-3, 1)$	\$50	$r + 50$
P_2	$\mathcal{N}(0, 2)$	\$45	$(r + 45)^3$
P_3	$\mathcal{N}(3, 3)$	\$60	$80(r + 60) - 0.5(r + 60)^2$
P_4	$\mathcal{N}(4, 2)$	\$35	$\exp(\frac{r+35}{15})$
P_5	$\mathcal{N}(5, 2)$	\$30	$\log_e(r + 30)$

Note that here the three linear poolings are initialised by the Laplacian Principle of Indifference at the first epoch, so will yield identical distributions (and hence results) in each case. The decisions deemed optimal are given on the left hand side of Table 3.3. We can see that under the PI approach (and hence under the other two linear pooling methods) all DMs opt to enter into the long forward, while only P_1, who is confident that a loss will be made, does not enter into the transaction when listening only to her own opinions.

Table 3.3: Optimal decisions for DMs at the first (LHS) and second (RHS) epoch. Here d^* is the optimal decision for a DM listening only to her own belief, while d^*_{PI}, d^*_{EW} and d^*_{MR} are her optimal decisions using the PI, EQ and MR methods respectively.

P_i	d^*	$d^*_{P.I.}$	$d^*_{E.W.}$	$d^*_{M.R.}$	P_i	d^*	$d^*_{P.I.}$	$d^*_{E.W.}$	$d^*_{M.R.}$
P_1	d_2	d_1	d_1	d_1	P_1	d_2	d_2	d_1	d_2
P_2	d_1	d_1	d_1	d_1	P_2	d_2	d_2	d_1	d_2
P_3	d_1	d_1	d_1	d_1	P_3	d_1	d_2	d_1	d_2
P_4	d_1	d_1	d_1	d_1	P_4	d_1	d_2	d_1	d_2
P_5	d_1	d_1	d_1	d_1	P_5	d_1	d_2	d_1	d_2

Upon entering the long forward all DMs make a loss of \$1.30, *i.e.*, $r = -1.3$, with this value simulated from the true distribution of θ. DMs now update their beliefs in light of this new evidence and combine their augmented opinions. The weights, both PI and normalised, associated with each DM by each of the methods are given in Table 3.4. A graphical interpretation of how these weights arise is included in Fig. 3.5. We can see that the DM deemed most accurate (*i.e.*, having the largest PI weight) was P_2, meaning she is the sole individual whose view is considered in the MR method. She has the highest weight in the PI approach, but the opinions of others are still taken into account with her opinion having a (non-dominant) weight of 0.4763. As we see from the right hand side of Table 3.3 no DM using the PI approach opts to enter into the long forward at the second epoch, while P_3, P_4 and P_5 all would if they listened only to their own beliefs. Using the EQ approach would lead all DMs to enter into the transaction, while the MR method leads, in this case, to the same decisions as the PI approach (although, as we can see from Table 3.4, the PI mean estimate is closer to the true value than that resulting from the MR method, albeit with a larger variance).

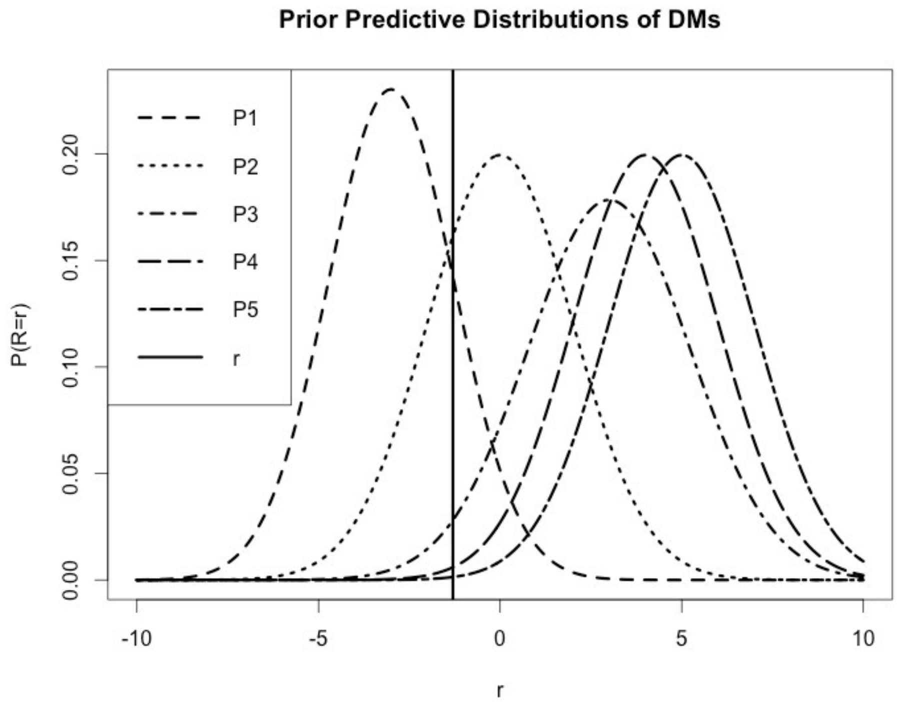

Fig. 3.5: Prior predictive distributions of DMs. The vertical line is the observation $r = -1.3$. The intersection of this and a DM's distribution gives her PI weight.

The decision process has stopped as no DMs enters into a long forward at the second epoch, and hence no new return can be observed, and opinions updated in light of this. We can see the clear advantage of using the PI approach in this example, as individuals who made their decisions without acknowledging the opinions of those around them would continue to enter into long forwards (and quite possibly make losses), while those using the linear pooling technique would have stopped (hence prevented future losses). Fig 3.6 highlights the predictive ability of the various methods, showing that it is the

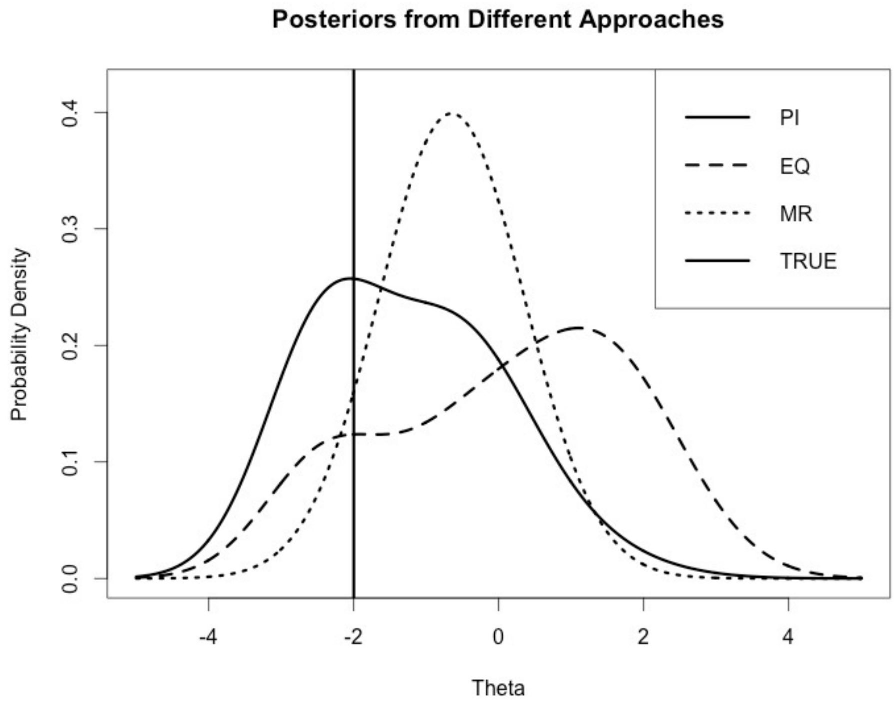

Fig. 3.6: PI, EQ and MR posterior distributions. The vertical line denotes the true θ.

Table 3.4: Updated beliefs of individual DMs and those resulting from the three considered methods, as well as the weights associated in each case, with $\alpha_{i,j}^{PI}$, $\alpha_{i,j}^{EQ}$ and $\alpha_{i,j}^{MR}$ denoting the weights from the PI, EQ and MR weights respectively.

P_j	$\mathbb{E}_j(\theta\|r)$	$\mathrm{Var}_j(\theta\|r)$	w_j	$\alpha_{i,j}^{PI}$	$\alpha_{i,j}^{EQ}$	$\alpha_{i,j}^{MR}$	App	$\mathbb{E}_{\hat{f}_i}(\theta\|r)$	$\mathrm{Var}_{\hat{f}_i}(\theta\|r)$
P_1	-2.43	0.67	0.1423	0.4196	0.2	0	PI	-1.17	2.23
P_2	-0.65	1	0.1615	0.4763	0.2	1	EQ	0.107	3.31
P_3	0.42	1.2	0.0280	0.0826	0.2	0	MR	-0.65	1
P_4	1.35	1	0.0059	0.0174	0.2	0			
P_5	1.85	1	0.0014	0.0041	0.2	0			

PI posterior that places the most density on the true parameter value. Similarly, Fig. 3.7 contrasts the PI posterior distribution with the posteriors of the five individuals, revealing that the PI method leads to more accurate estimation (in terms of posterior density) than the individual posteriors of four of the five DMs. Hence it can be argued that it is in the best interest of DMs to use this combined distribution, as with probability of 0.8 (*i.e.*, with probability over 0.5) it will lead to better estimation than the individual distribution of a randomly chosen DM from within the neighbourhood. We introduce this metric formally in Chapter 4.

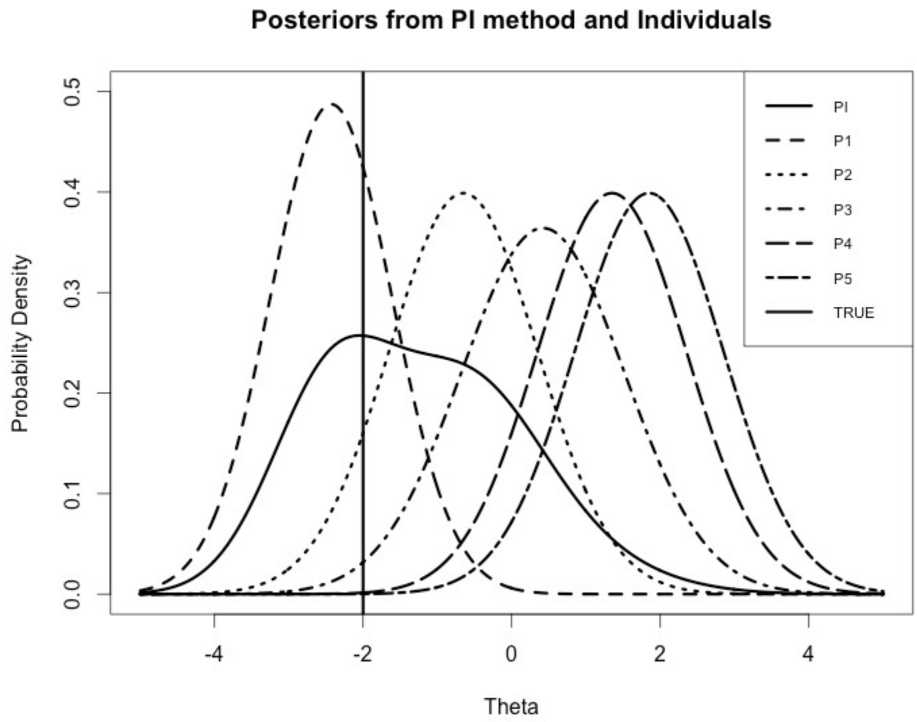

Fig. 3.7: The individual DM posterior distributions and the posterior distribution resulting from the Pl method. The vertical line denotes the true value of $\theta = -2$.

For illustration in Table 3.5 we show the proportion of times that each method was superior (in terms of this posterior density metric) based on 5,000 simulations, using the prior beliefs specified in Table 3.2 and the Normal data generating mechanism. We see, as shall be demonstrated and discussed in depth in the simulation study in Chapter 4, that the PI approach becomes increasingly successful as the number of returns witnessed increases.

Table 3.5: PI, EQ and MR success proportions with the optimal method in bold.

Returns	PI	EQ	MR
2	0.4006	0.1744	**0.4250**
3	**0.4578**	0.1760	0.3662
4	**0.4866**	0.1758	0.3376
5	**0.4938**	0.1670	0.3392

3.7 Distributions of PI Weights

Below we derive the probability distribution that PI weights follow. To ease tractability we consider θ to be a Normal mean with known variance σ^2. In Chapter 8 we discuss potential application of the conjugate utility class (Houlding *et al.*, 2015) to sequential problems of this ilk. The PI weight of P_i, with prior of $f_i(\theta) \sim \mathcal{N}(m_i, s_i^2)$, is

$$w_i = f_i(R = r) = \frac{1}{\sqrt{2\pi(\sigma^2 + s_i^2)}} \exp\Big(- \frac{(r - m_i)^2}{2(\sigma^2 + s_i^2)}\Big) \tag{3.21}$$

This defines w_i in terms of r. By manipulation we can write r in terms of w_i, giving

$$r \;=\; m_i \pm \sqrt{-2(\sigma^2 + s_i^2) \ln \Big[w_i \sqrt{2\pi(\sigma^2 + s_i^2)}\Big]} \tag{3.22}$$

Note w_i is maximised for P_i when the return seen is her prior mean, *i.e.*, when $r = m_i$. This leads to her maximal PI value, $w_{\max_i}$:

$$w_{\max_i} \;=\; \frac{1}{\sqrt{2\pi(\sigma^2 + s_i^2)}} \exp\Big(- \frac{(m_i - m_i)^2}{2\sigma^2}\Big) = \frac{1}{\sqrt{2\pi(\sigma^2 + s_i^2)}} \tag{3.23}$$

The minimum, $w_{\min_i}$, of w_i approaches zero in the limit as the witnessed return deviates further and further from her mean, *i.e.*, $w_i \to 0$ as $|r - m_i| \to \infty$. We produce plots of r against w_i in Fig. 3.8. There are two values of r corresponding to each value of w_i, as realisations both above and below m_i will return identical PI weights. The true distribution for returns is Normally distributed, *i.e.*,

$$\mathbb{P}(R = r|\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\Big(- \frac{(r - \theta)^2}{2\sigma^2}\Big) \tag{3.24}$$

Given this, our definition of r in terms of w_i in Equation (3.22), and mutually exclusiveness, the probability of any particular value of w_i occurring is

$$\begin{aligned}
\mathbb{P}(W_i = w_i|\theta) &= \mathbb{P}\Big(r = m_i \pm \sqrt{-2(\sigma^2 + s_i^2) \ln \Big[w_i \sqrt{2\pi(\sigma^2 + s_i^2)}\Big]}|\theta\Big) \\
&= \mathbb{P}\Big(r = m_i + \sqrt{-2(\sigma^2 + s_i^2) \ln \Big[w_i \sqrt{2\pi(\sigma^2 + s_i^2)}\Big]}|\theta\Big) \\
&+ \mathbb{P}\Big(r = m_i - \sqrt{-2(\sigma^2 + s_i^2) \ln \Big[w_i \sqrt{2\pi(\sigma^2 + s_i^2)}\Big]}|\theta\Big) \\
&= \frac{1}{\sqrt{2\pi\sigma^2}} \Big[\exp\Big(- \frac{(m_i + \sqrt{-2(\sigma^2 + s_i^2) \ln \Big[w_i \sqrt{2\pi(\sigma^2 + s_i^2)}\Big]} - \theta)^2}{2\sigma^2}\Big) \\
&+ \exp\Big(- \frac{(m_i - \sqrt{-2(\sigma^2 + s_i^2) \ln \Big[w_i \sqrt{2\pi(\sigma^2 + s_i^2)}\Big]} - \theta)^2}{2\sigma^2}\Big)\Big]
\end{aligned} \tag{3.25}$$

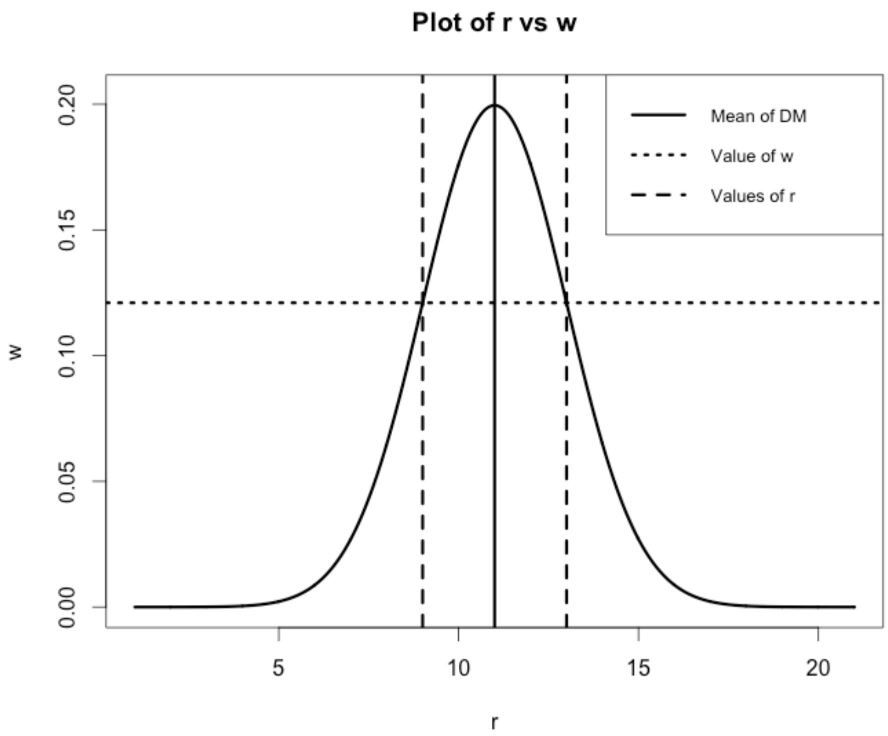

Fig. 3.8: Plot of r *vs.* w_i when $\sigma^2 = 2$, $m_i = 11$ and $s_i^2 = 2$. The dotted horizontal line denotes a w_i with intersecting vertical dashed lines showing the two values of r leading to this w_i. The vertical unbroken line is the r giving $w_{\max_i}$, *i.e.*, $r = m_i$.

We plot this distribution in Fig. 3.9. We must show Equation (3.25) is a valid density function, *i.e.*, non-negativity for all possible w_i and integration to one over its support. The first property is obeyed but the second is not. The missing component is the Jacobian, required due to our change of variables, *i.e.*, finding the probability of w_i by considering the probability of a corresponding r. We calculate $\frac{dr}{dw_i}$, giving us

$$\frac{dr}{dw_i} = \pm \frac{\sqrt{\sigma^2 + s_i^2}}{w_i} \times \frac{1}{\sqrt{-2 \ln \left[w_i \sqrt{2\pi(\sigma^2 + s_i^2)} \right]}} \tag{3.26}$$

The absolute value of the Jacobian is considered so we ignore the $\pm$ in Equation (3.26). As the relationship between r and w_i is not bijective we must be careful with our transformation. This absolute value ensures both additive terms in Equation (3.25) have identical Jacobians. Hence the complete distribution of the PI weights of P_i is the product of the absolute Jacobian and Equation (3.25), *i.e.*,

$$\mathbb{P}(W_i = w_i|\theta) = \frac{\sqrt{\sigma^2 + s_i^2}}{w_i} \times \frac{1}{\sqrt{-2\ln\left[w_i\sqrt{2\pi(\sigma^2+s_i^2)}\right]}} \times \left(\frac{1}{\sqrt{2\pi\sigma^2}}\left[\exp\left(-\frac{(m_i + \sqrt{-2(\sigma^2+s_i^2)\ln\left[w_i\sqrt{2\pi(\sigma^2+s_i^2)}\right]} - \theta)^2}{2\sigma^2}\right) + \exp\left(-\frac{(m_i - \sqrt{-2(\sigma^2+s_i^2)\ln\left[w_i\sqrt{2\pi(\sigma^2+s_i^2)}\right]} - \theta)^2}{2\sigma^2}\right)\right]\right) \quad (3.27)$$

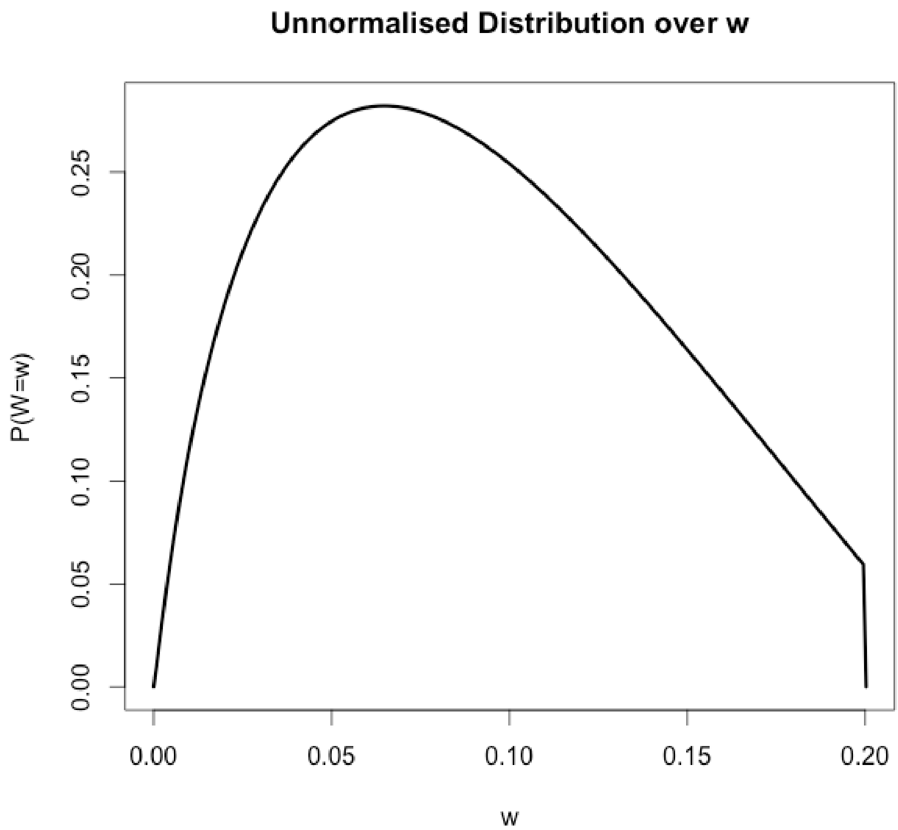

Fig. 3.9: Unnormalised plot of $\mathbb{P}(W_i = w_i|\theta)$ for the parameterisation in Fig. 3.8.

This normalised distribution (integrating to one) is given in Fig. 3.10. The spike as w_i approaches $w_{\max_i}$ arises due to the numerical instability of Equation (3.27) at the point $w_i = w_{\max_i}$. We can calculate expected values and variances of PI weights using the following equations, requiring numerical methods to solve:

$$\mathbb{E}[W_i|\theta] = \int_0^{w_{\max_i}} w_i \mathbb{P}(W_i = w_i|\theta)\, dw_i \quad (3.28)$$

$$\mathrm{Var}[W_i|\theta] = \left[\int_0^{w_{\max_i}} w_i^2 \mathbb{P}(W_i = w_i|\theta)\, dw_i\right] - (\mathbb{E}[W_i|\theta])^2 \tag{3.29}$$

Solving Equations (3.28) and (3.29) give $\mathbb{E}[W_i|\theta] = 0.0769$ and $\mathrm{Var}[W_i|\theta] = 0.003$ for the parameterisation discussed above. We return in Chapter 8 to discuss how this material can be used to attempt to solve sequential problems with multiple DMs.

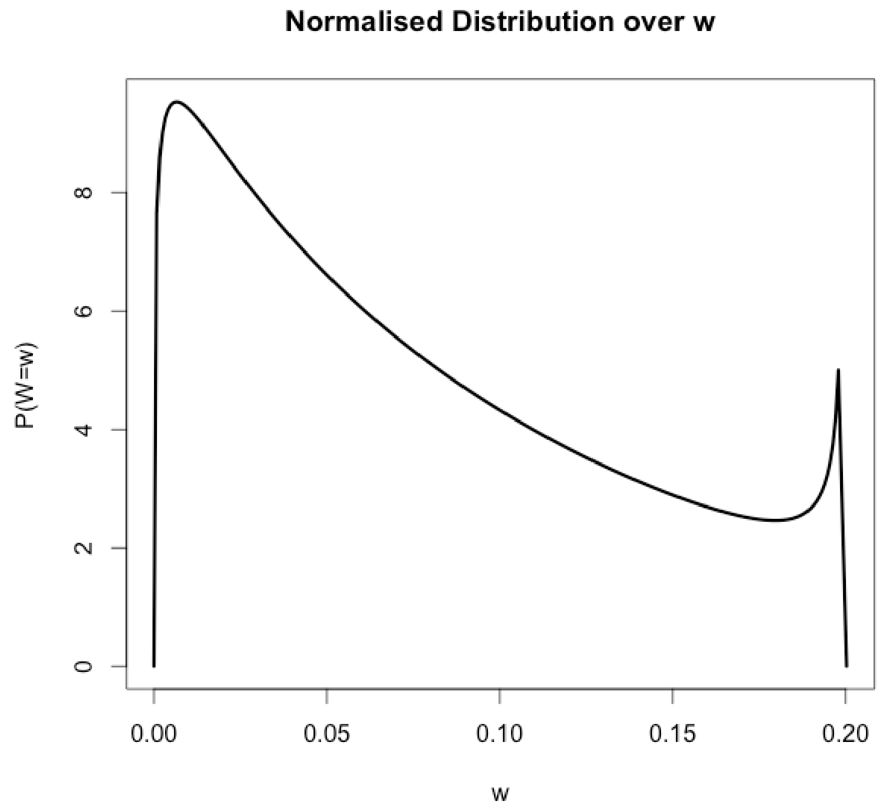

Fig. 3.10: Normalised plot of $\mathbb{P}(W_i = w_i|\theta)$ for the parameterisation in Fig. 3.8.

3.8 Multiple Differing Simultaneous Returns

Above it was assumed that DMs witness a common decision return at each epoch. In this section we extend the previously derived methodology to a framework in which DMs are liable to observe different outcomes from decisions that they simultaneously make. As a motivating example consider n shopkeepers, all of whom buy (potentially varying amounts of) a particular good from the same manufacturing company. All have interest in the same quantity θ which is the probability that an item produced by the manufacturer is defective. Suppose $P_1, \ldots, P_n$ respectively purchase $m_1, \ldots, m_n$ goods, and discover that $k_1, \ldots, k_n$ of these are defective. Each DM has seen a different return but all of these pertain to θ. Here we discuss a method by which this information can be coherently combined.

We introduce vector notation for the decisions made and returns witnessed at the i^{th} epoch, writing $\underline{d_i} = \{d_{i,1}, \ldots, d_{i,n}\}$ and $\underline{r_i} = \{r_{i,1}, \ldots, r_{i,n}\}$ respectively. Here $d_{i,j}$ is the decision made by P_j at epoch i and $r_{i,j}$ is the return witnessed by P_j having made decision $d_{i,j}$. Each DM uses her utility function and the equally weighted combined belief to make an initial decision. We establish the notion of decisions with non-trivial consequences here. Consider a DM who must decide whether to invest in stock A or stock B. Regardless of what decision is made, some information will be witnessed about the uncertainty of interest (*e.g.*, underlying market behaviour) and hence learning will occur about θ, *i.e.*, the DM will update her prior opinion, or perhaps more succinctly, her posterior distribution will differ from her prior distribution. By contrast, consider a DM who must choose whether to invest in stock A or to not. If she chooses the former decision she will witness some information about θ as a consequence, but if she opts for the latter decision then no learning occurs. As her decision was to take no action she observed no new information about θ. In a setting where a common return was witnessed by all DMs the same information was gained regardless of if one DM opted to make a decision with a non-trivial consequence or if they all did. The only instance in which no new data is realised is when no DM makes a decision with a non-trivial consequence, *i.e.*, no updating of prior distributions occurs. In the setting of multiple simultaneous returns the amount of information that is available to DMs is dependent upon how many DMs opt to take gambles. In what follows we provide a method by which this updating can be conducted in complete generality, and then discuss the underlying intuition for our three specific distributional cases.

Once an initial set of decisions $\underline{d_1}$ have been made by DMs, a set of corresponding outcomes $\underline{r_1}$ will be witnessed. As discussed above some of this set will be trivial, and will not be used in the construction of posterior distributions, although in some distributional cases an intuitive interpretation can be provided for these uninformative returns, *e.g.*, in a Binomial setting we can consider observing zero successes in zero trials. Hence the updated belief of P_i is written as

$$f_i(\theta|\underline{r_1}) \quad = \quad f_i(\theta|r_{1,1}, \ldots, r_{1,n})$$

$$
\begin{aligned}
&\propto f(r_{1,1},\ldots,r_{1,n}|\theta)f_i(\theta) \\
&= f(r_{1,1}|\theta)\ldots f(r_{1,n}|\theta)f_i(\theta) \\
&= f_i(\theta)\prod_{j=1}^{n} f(r_{1,j}|\theta) && (3.30) \\
&= f(\underline{r_1}|\theta)f_i(\theta) && (3.31)
\end{aligned}
$$

We see two ways of considering the likelihood function. The set of witnessed returns are conditionally independent given θ so we can consider a product of the likelihood functions for each individual return as in Equation (3.30). Equivalently we may also use a likelihood function for the whole data set at once as in Equation (3.31) with some physical interpretation of this provided shortly. Note that these two approaches will yield identical results as Bayesian updating adheres to the Likelihood Principle.

An alternative to this approach would be for each individual P_i to update her prior distribution solely in light of the information she herself witnesses, yielding $f_i(\theta|r_{1,i})$. Yet this seems unintuitive and self-defeating for the group as a whole. Previously in this thesis we extolled the merits of information pooling and the correlation one would expect to exist between information quality and decision quality. The DMs we consider inhabit a non-competing environment and hence have nothing to lose (in terms of utility) from sharing their new information with all their neighbours, especially when they will receive additional data in return. It is true that if a DM updated solely in light of the return which she herself witnessed this information would propagate into the combined distribution of all DMs as it would be a component of the weighted sum, *i.e.*, each of the elements of $\underline{r_1}$ would be contained in a belief of the form

$$\hat{f}_i(\theta|\underline{r_1}) = \alpha_{i,1}f_i(\theta|r_{1,1}) + \ldots + \alpha_{i,n}f_n(\theta|r_{1,n}) \tag{3.32}$$

Nevertheless this seems counter-productive. In addition, when DMs want to assess the reliability of the beliefs of neighbours it makes sense to use all possible data in order to estimate accuracy as precisely as possible. When P_i is judging the accuracy of P_j she should base this not just on the return she has seen ($r_{1,i}$) as this may be an outlier. Instead, she should consider all available data ($\underline{r_1}$), using as complete a picture as possible of the true state of nature to ascertain a measure of the reliability of neighbours. We propose that the combined posterior distribution of each DM should incorporate all the data witnessed by her neighbours and herself, *i.e.*,

$$\hat{f}_i(\theta|\underline{r_1}) = \alpha_{i,1}f_1(\theta|\underline{r_1}) + \ldots + \alpha_{i,n}f_n(\theta|\underline{r_1}) \tag{3.33}$$

We provide some intuition behind the information contained in the likelihood functions for the three primary distributional cases considered.

- **Binomial Case:** Suppose θ is the success probability for a Bernoulli trial. Each P_i provides initial Beta priors parameterised by α_i and β_i. Following $\underline{d_1}$ each P_i witnesses k_i successes in m_i trials, where it is not necessary that $m_i = m_j$ for all $i \neq j$, *e.g.*, two doctors see three out of four, and seven out of ten, patients cured respectively. When assessing their likelihood function each DM may consider the probability of seeing $k = \sum_{i=1}^{n} k_i$ successes in $m = \sum_{i=1}^{n} m_i$ trials, *e.g.*, the doctors would view their pooled information as ten successes in fourteen trials. Generically we write the updated posterior of P_i as

$$\begin{aligned}
f_i(\theta|k \text{ out of } m \text{ successes}) &= f(k \text{ out of } m \text{ successes}|\theta)f_i(\theta) \\
&= f_i(\theta)\prod_{i=1}^{n} f(k_i \text{ out of } m_i \text{ successes}|\theta) \\
&= \frac{\Gamma(\alpha_i+\beta_i)}{\Gamma(\alpha_i)\Gamma(\beta_i)}\theta^{\alpha_i-1}(1-\theta)^{\beta_i-1}\prod_{i=1}^{n}\binom{m_i}{k_i}\theta^{k_i}(1-\theta)^{m_i-k_i} \\
&\propto \theta^{\alpha_i-1}(1-\theta)^{\beta_i-1}\theta^{\sum_{i=1}^{n}k_i}(1-\theta)^{\sum_{i=1}^{n}m_i-k_i} \\
&= \theta^{\alpha_i+k-1}(1-\theta)^{m-k+\beta_i-1} \qquad (3.34)
\end{aligned}$$

- **Poisson Case:** Interest lies in a rate parameter θ over some fixed unit of time with each P_i having a Gamma prior with scale and shape parameters α_i and β_i respectively. We can adjust for a differing unit of time by scaling θ up or down accordingly. Suppose DMs witness $k_1, \ldots, k_n$ distinct observations in disjoint intervals (if they were not disjoint correlation would need to be taken into account) of length $t_1, \ldots, t_n$ respectively, *e.g*, one DM sees thirty cars in one minute, and another sees fifty cars in two minutes. They may view this pooled information as witnessing $k = \sum_{i=1}^{n} k_i$ observations across a period of time of length $t = \sum_{i=1}^{n} t_i$, *e.g.*, above this would correspond to eighty cars in three minutes. In generality the updated posterior distribution of P_i is

$$\begin{aligned}
f_i(\theta|k \text{ events in time } t) &= f(k \text{ events in time } t|\theta)f_i(\theta) \\
&= f_i(\theta)\prod_{i=1}^{n} f(k_i \text{ events in time } t_i \text{ trials}|\theta) \\
&= \frac{\alpha_i^{\beta_i-1}}{\Gamma(\beta_i)}\theta^{\beta_i-1}e^{-\alpha_i\theta}\prod_{i=1}^{n}\left[\frac{(t_i\theta)^{k_i}}{k_i!}e^{-t_i\theta}\right]
\end{aligned}$$

$$
\begin{aligned}
&\propto \ \theta^{\beta_i-1}e^{-\alpha_i\theta}\theta^{\sum_{i=1}^n k_i}e^{-\theta\sum_{i=1}^n t_i} \\
&= \ \theta^{k+\beta_i-1}e^{-\theta(\alpha_i+t)} \qquad (3.35)
\end{aligned}
$$

- **Normal Case:** Suppose interest is in the mean θ of a Normal distribution. The associated variance σ^2 is known. Each P_i has a Normally distributed prior with mean a_i and variance s_i^2. Each DM witnesses a (potentially different) number m_i of observations, having respective sample means of $r_i = \frac{\sum_{j=1}^{m_i} x_j}{m_i}$. Suppose interest was in the average salary students earn in part time jobs, with one investigator asking ten students and another asking twenty. An overall average can be computed by pooling the data collected by individuals, *i.e.*, considering having witnessed $m = \sum_{i=1}^n m_i$ observations, and writing this grand sample mean as $\bar{x} = \frac{\sum_{i=1}^n r_i m_i}{m}$. The posterior of P_i is given by

$$
\begin{aligned}
f_i(\theta|\text{mean } \bar{x} \text{ of } m \text{ samples}) &= f(\text{mean } \bar{x} \text{ of } m \text{ samples}|\theta) f_i(\theta) \\
&= f_i(\theta)\prod_{i=1}^n f(\text{mean } r_i \text{ of } m_i \text{ samples}|\theta) \\
&= \frac{1}{\sqrt{2\pi s_i^2}}\exp\Big(-\frac{(\theta-a_i)^2}{2s_i^2}\Big)\prod_{i=1}^n\Big[\frac{1}{\sqrt{2\pi\sigma^2}}\exp\Big(-\frac{(r_i-\theta)^2}{2\sigma^2}\Big)\Big] \\
&\propto \ \exp\Big(-\frac{(m\bar{x}-\theta)^2}{2\sigma^2}\Big)\exp\Big(-\frac{(\theta-m_i)^2}{s_i^2}\Big) \\
&\propto \exp\Big[-\frac{1}{2(\frac{1}{s_i^2}+\frac{m}{\sigma^2})^{-1}}\Big(\theta-\frac{\frac{a_i}{s_i^2}+\frac{m\bar{x}}{\sigma^2}}{\frac{1}{s_i^2}+\frac{m}{\sigma^2}}\Big)\Big] \qquad (3.36)
\end{aligned}
$$

Implicit in these instances (and any generic application) is the assumption of homogeneity between observations from different DMs. An area for further development is that in which DMs are interested in quantities $\theta_1, \ldots, \theta_n$ which are not necessarily identical but exhibit substantial correlations and dependencies. For now we see that the posterior distribution of each DM is determined by the information witnessed by the group as a whole, *i.e.*, if P_i was to alter the decision that she makes (*e.g.*, buying shares instead of not buying them) this would change the posterior beliefs for all DMs within the group structure. In the classical linear opinion pooling setting (Cooke, 1991) there is a DM consulting domain-specific experts who provide her with opinions about θ. In this environment there is a single decision made, with the experts providing advice that they do not themselves act upon. Similarities can be seen between this setting and our own in the instance where all DMs witness a common return, with all individuals wanting to assess the accuracy of a set of beliefs in light of a single piece of information, which will be witnessed assuming one of them makes a decision with a non-trivial consequence. However in our extension of this simplified scenario, where all DMs witness their own returns, this is clearly not the case, as the utility functions of DMs will impact upon the decisions that they make, and hence which return will occur and be witnessed by themselves and the group as a whole. This modification to the configuration of our problem provides additional generality, applicability, and originality to our research.

Finally, above we only considered the calculation of posterior distributions of DMs. We turn our attention to the other manner in which learning over time occurs, *i.e.*, weight updating. Weights are found by plugging the witnessed aggregated data into prior predictive distributions of DMs, giving direct analogies of Equations (3.7) – (3.9):

$$w_i = f_i(K = k) = \frac{\Gamma(m+1)\Gamma(\alpha_i+\beta_i)\Gamma(\alpha_i+k)\Gamma(\beta_i+m-k)}{\Gamma(m-k+1)\Gamma(k+1)\Gamma(\alpha_i)\Gamma(\beta_i)\Gamma(\alpha_i+\beta_i+m)} \tag{3.37}$$

$$w_i = f_i(K = k) = \frac{\alpha_i^{\beta_i}}{\Gamma(k+1)\Gamma(\beta_i)} \frac{\Gamma(k+\beta_i)}{(\alpha_i+1)^{k+\beta_i}} \tag{3.38}$$

$$w_i = f_i(\bar{X} = \bar{x}) = \frac{1}{\sqrt{2\pi(\sigma^2+s_i^2)}} \exp\Big(-\frac{(\bar{x}-m_i)^2}{2(\sigma^2+s_i^2)}\Big) \tag{3.39}$$

We see that the PI approach can be extended to a significantly more generalised setting. In what follows we focus on the more simplistic common return scenario, while bearing in mind that such an extension is always achievable.

3.9 Bayesian Model Averaging

We briefly turn our attention to Bayesian Model Averaging (BMA), discussed in detail in Hoeting *et al.* (1999), which we strongly recommend to the interested reader. BMA shares two of the core ideals of our PI approach. Firstly it is a technique for combining a collection of (potentially diverse) probabilistic quantities into a single quantity, and secondly it aims to achieve this while adhering to the Bayesian paradigm. Consider a setting where an individual wishes to model a phenomenon, but is unsure which particular model from a broad class of models (e.g., linear regression models or proportional hazard models) describes this most accurately. For instance, the user may believe the quantity can be modelled via a linear regression, but is unsure which of her available covariates should be included in this. BMA derives an ensemble methodology by which these models may be amalgamated in a linear fashion, with different weights awarded to the different models based on their integrated likelihood/prior predictive distributions. The goal of this process is generally to provide a measure of uncertainty (i.e., a variance) as well as a point estimate (i.e., the mean) for use in a predictive problem, e.g., weather forecasting or financial modelling. Similarities and contrasts are instantly evident here with the PI approach: both are Bayesian techniques for merging several elements into a single element. Yet in the PI approach the elements to be merged are differently parameterised versions of a common probability distribution, while in BMA they are different versions of a common broad statistical family. We comment that BMA is often computational expensive if one wants to include an exhaustive set of all possible models in her ensemble, and that specification of a prior distribution over model types is a challenging task. Solving BMA prior predictive integrals (analogous to our Equation 3.6) is often a very complex task, with exact solutions generally unachievable, and approximations found using the Bayesian Information Criteria approximation (Raftery, 1995).

3.10 Limitations

In this chapter we have provided some simple justifications for the PI approach using mathematical arguments, and in the following chapter we shall attempt additional validation using two types of data. While there is certainly merit in the use of our technique, we conclude this chapter with some brief notes on the limitations associated with it, specifically referencing contexts in which it is not suitable for use.

- It is assumed throughout the process that the uncertain quantity of interest remains constant. The opinion of DMs over this quantity will change over time, in light of new noisy realisations, but in our context we do not consider that the true underlying quantity itself will alter. There are cases when this is very reasonable, for instance if the parameter is the proportion of patients who successfully respond to a new drug trial, or the average temperature in Dublin from year to year. Yet there are certainly cases where it would be useful to be able to explicitly consider the quantity as dynamic, e.g., if it were the price of a piece of real estate in a developing area or the speed at which an athlete in training can run a mile. As it stands the corresponding observations for these phenomena will be incorporated in the Bayesian updating, but they will still be seen as realisations of the initial random variable rather than some evolved one.

- In its current form our model does not explicitly incorporate correlation between potentially similar unknown quantities. As a motivating example suppose interest lies in the time it will take a runner to complete a half marathon, and several realisations of this are available. Intuitively we would imagine a strong correlation between their time for running a half marathon and running a full marathon. Yet, at present, if interest was to shift to the individuals time for a full marathon, the process would need to be begun again from scratch (although obviously DMs could use their knowledge of the half marathon times in the construction of their prior distributions).

- The PI approach is fully parametric, with DMs being required to supply a fully parameterised probability distribution. This may be beyond the scope of some users, even with the available elicitation techniques. In Chapter 6 we develop a nonparametric simplification, but this too is reliant upon the specification of a measure of uncertainty on the behalf of the user, i.e., a class of variance estimate. There may be cases where individuals are interested only in the combining of the point estimates, and do not wish to complicate matters by the consideration of higher order moments. In this situation the PI approach is not suitable for use.

- All individuals have strictly positive weights through the PI scheme, but in the limit these can tend towards zero. DMs may be willing to amalgamate their opinion with those of their neighbours, and to have their own opinion diluted by the individuals around them, especially if they are perceived as being more accurate. Yet there will conceivably be cases where a DM is unwilling to enter into a scheme in which the decision that she herself undertakes, and will bear the consequences of, is one which almost entirely discounts her point of view. Hence users must be made aware before use of the objectivity inherent in the scheme, and must concur with the potential ramifications of this. In Chapter 7 we derive a subjective method, the KL approach, which goes some way towards solving this problem.

Chapter 4

Data-based Justifications

In the previous chapter we introduced the Plug-in approach as a decision making tool for quantifying perceived reliability of information sources, and demonstrated its attractive Bayesian and coherency properties. Yet these advocations lack a certain degree of formality. In this chapter we justify the PI approach with more rigour by considering its performance on sets of data, both simulated and real, providing strong empirical evidence for its use. In addition to discussing and interpreting results we illustrate the theoretical calculations underlying the simulated data study and demonstrate how simulated proportions converge to true probabilities of the PI approach's superiority.

4.1 Alternatives Methods and Metric Choice

In the setting that we are concerned with a DM does not need to listen to her neighbours as the decision that she makes is her own, *i.e.*, she alone feels its consequences. Hence she may wish to make this decision based solely on her own personal beliefs. If she is confident in having a firm understanding of θ she may feel that listening to opinions of (potentially vague/inaccurate) neighbours will only weaken the precision of the distribution that she makes her decision with, potentially diminishing her decision quality. By contrast a DM believing herself to be relatively uninformed about θ may be willing to incorporate views of (potentially far more knowledgeable) neighbours into her decision task, hopeful of increasing her decision quality. Each DM makes a choice: to make decisions using the PI posterior opinion or her own posterior opinion. We attempt to discover which method is superior, by considering which leads to the most accurate belief (*i.e.*, that most closely mirroring θ).

An obvious matter regards the performance metric determining superiority. Initially we considered a metric based on comparing a method's posterior mean with the true θ, with the best method minimising this absolute distance. Yet issues arise with this choice, both in solely considering a distribution's mean (as opposed to, for instance, its median), and also that the corresponding associated uncertainty (*i.e.*, variance) is disregarded. We illustrate this in Fig. 4.1 where the distribution of P_1 has a mean closer to θ but the distribution of P_2 places more probability density on θ. The mean-based metric ignores the overconfidence (small variance) of P_1. Hence when comparing a set of techniques (*i.e.*, the several pooling methods we shortly introduce) we declare the superior method to be that placing the most posterior density on θ. In the individual problem, comparing the PI posterior to those of individuals, we say the PI approach is superior if it places more density on θ than over half the posteriors of DMs, *i.e.*, if a DM is randomly chosen from the group then there is probability over 0.5 that the combined distribution will provide more accurate estimation than her own distribution. DMs do not know *a priori* how reliable they are.

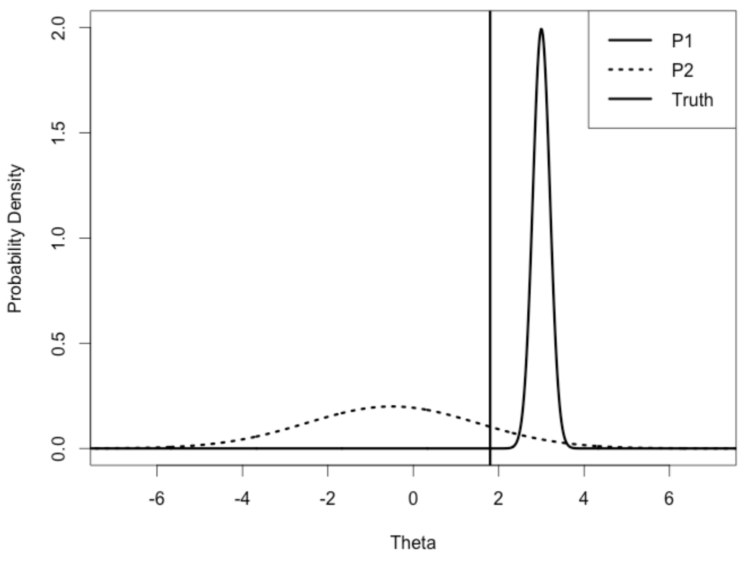

Fig. 4.1: The distributions $f_1(\theta) \sim \mathcal{N}(3, 0.2)$ and $f_2(\theta) \sim \mathcal{N}(-0.5, 2)$, and $\theta = 1.8$.

Suppose a DM opts to listen to her neighbours, *i.e.*, incorporating their opinions into her decision task. We proposed the PI approach as a suitable linear pooling method, but it is not the only option. Several other rational schemes exist. We discussed how at the first epoch all DMs can be given equal weight. This is reasonable and intuitive in the absence of relevant information about their accuracy. A simple weighting scheme is one maintaining these equal weights throughout the process, *i.e.*, weights are unchanged in light of new data. Weaknesses are clear. DMs with concise and accurate opinions are afforded the same merit as those with vague and inaccurate opinions. When more DMs possess opinions of the latter type than the former type this leads to the combined opinion being highly inaccurate and hence to potentially low decision quality for DMs, *i.e.*, the accuracy of knowledgeable DMs is overshadowed. Yet there are situations when this approach may yield good results. Over time opinions of DMs become increasingly accurate (converging to precise identical distributions in the limit) so affording equal weights to a set of reliable DMs will yield a reliable opinion. As more returns are witnessed the probability of the equal weights approach being meritorious increases. The calibre of this approach may grow if DMs predictions are symmetric around θ, *e.g.*, if two DMs underestimate θ by two and three units respectively, and two DMs overestimate θ by two and three units respectively then equal weights leads to prediction centred on θ. This is the "wisdom of crowds" explaining how, as discussed in Section 2.7, the equally weighted combination potentially outperforms more sophisticated schemes. As time progresses predictions of DMs grow closer to θ, *e.g.*, after an epoch we may have two DMs underestimating by one and two units, and two overestimating by one and two units. Hence the symmetry in beliefs and increasing DM accuracy, leads to predictions closer to θ. We investigate this case in Section 4.2.

The PI approach awards high weights to the DMs who appear most accurate in a set of neighbours. An alternative weighting scheme gives a weight of one to the single DM deemed most reliable at the previous epoch (*i.e.* maximised Equation 3.6), and weights of zero to all others, *i.e.*, a significantly simplified version of the PI approach. This may lead to rotation of which DM is deemed most reliable. This approach initialises with the Laplacian Principle of Indifference, with shortcomings evident. Giving weights of one is a dangerous strategy often cautioned against (*e.g.*.Eggstaff *et al.*, 2014). If the outcome witnessed is an unlikely one and the distribution of the DM severely contrasts with θ then subsequent decision quality may be poor, with her inaccurate opinion the only one considered. The PI approach would assign her a large PI weight, but opinions of other DMs would be factored in also, with her past inaccuracy leading to a considerable decrease in her normalised weight. Yet, there are situations where this method could be effective. Suppose all DMs barring one are extremely unreliable, with that DM being accurate. Assuming no unlikely (*i.e.*, tail probability) events occur this DM will be given the dominating weight throughout, ensuring that the combined belief is consistently accurate. Also as time progresses we have discussed how DMs become increasingly accurate. In a setting where all DMs are accurate (*e.g.*, *a priori* or after several epochs), giving a weight of one to a single DM is reasonable as it leads to a similar opinion as a weighted combination (by the asymptotic argument in Section 3.4). The variance of this method will not be impacted upon by the positive correction term in Equation (3.19) so may have a smaller variance than the PI distribution. It is clear that this approach will be most successful if tail events do not occur and a large proportion of DMs possess strongly inaccurate beliefs. Below we assess the merits of the PI distribution against the distributions of individual DMs, and also to the distributions from the Equal Weights (EQ) and Most Reliable (MR) methods.

4.2 Simulation Study

We consider our three conjugate cases: Beta-Binomial (for probabilities), Normal-Normal (for continuous problems) and Gamma-Poisson (for rates), and situations where DMs on average overestimate, underestimate, and have mean predictions centred on, θ. Distributions over prior hyperparameters are in Table 4.1, *e.g.*, for Normal overestimation ($\theta = 0$) prior means are uniformly chosen from $[-2, 8]$ so on average DMs will have a mean of 3. In underestimation cases we set average prior predictions closer to θ, investigating if this impacts performance. The number of DMs involved increases from 2 to 20 and the number of returns witnessed increases from 1 to 12. For each of the 228 (*i.e.*, 19×12) cases we run 5,000 simulations, recording which method is superior each time and aggregating results over the 5,000 simulations. The method with the highest success proportion is deemed superior for that case.

Results for Normal overestimation are in Fig 4.2. After exactly one return the MR approach outperforms the PI and EQ approaches irrespective of the number of DMs involved. After exactly two returns the PI approach gives better estimation if the group contains ten DMs or more, while after exactly three returns it is better if the group

Table 4.1: Data mechanisms and parameters, prior structures and distributions over the simulation of prior parameters, as well as the corresponding average prior means.

	Beta-Binomial	Normal-Normal	Gamma-Poisson
Data Mechanism	$R \sim Binomial(5, \theta)$	$R \sim \mathcal{N}(\theta, 1)$	$R \sim Poisson(\theta)$
θ	0.5	0	5
Priors	$f_i(\theta) \sim Beta(\beta_1, \beta_2)$	$f_i(\theta) \sim \mathcal{N}(m_i, s_i)$	$f_i(\theta) \sim Gamma(\beta_1, \beta_2)$
Overestimation	$\beta_1 \sim U(1, 16)$	$m_i \sim U(-2, 8)$	$\beta_1 \sim U(1, 4)$
	$\beta_2 \sim U(1, 4)$	$s_i \sim U(0, 3)$	$\beta_2 \sim U(1, 32)$
	$\mathbb{E}_{f_i}(\theta) = 0.8$	$\mathbb{E}_{f_i}(\theta) = 3$	$\mathbb{E}_{f_i}(\theta) = 8$
Underestimation	$\beta_1 \sim U(1, 6)$	$m_i \sim U(-6, 2)$	$\beta_1 \sim U(1, 4)$
	$\beta_2 \sim U(1, 14)$	$s_i \sim U(0, 3)$	$\beta_2 \sim U(1, 16)$
	$\mathbb{E}_{f_i}(\theta) = 0.3$	$\mathbb{E}_{f_i}(\theta) = -2$	$\mathbb{E}_{f_i}(\theta) = 4$
Mean-Centred	$\beta_1 \sim U(1, 10)$	$m_i \sim U(-8, 8)$	$\beta_1 \sim U(1, 4)$
	$\beta_2 \sim U(1, 10)$	$s_i \sim U(0, 4)$	$\beta_2 \sim U(1, 20)$
	$\mathbb{E}_{f_i}(\theta) = 0.5$	$\mathbb{E}_{f_i}(\theta) = 0$	$\mathbb{E}_{f_i}(\theta) = 5$

contains seven DMs or more. The more information sources DMs have access to the better PI performance is. As the number of returns grows PI performance becomes increasingly dominant. This is intuitive as the PI method learns over time, gaining increasingly accurate reliability measures in accordance with an increased amount of data. These conclusions are reinforced by the figures in Table 4.2. The proportion of times the PI approach outperforms alternatives steadily increases with the number of DMs. Success proportions for the PI approach for a given number of DMs is higher having seen four returns than three returns. In Fig. 4.3-4.5 we plot the eight other cases (Binomial underestimation, *etc.*), with identical conclusions inferred from these, *i.e.*, PI performance grows with the amount of DMs/returns involved. PI performance is marginally weaker in the underestimation case (in which DMs are *a priori* more accurate) than the overestimation case, but performance actually improves when average predictions are mean-centred. Briefly we comment on the poor EQ performance. In only one case (two DMs and a single return in the Binomial mean-centred scenario) is it superior. Opinions of accurate DMs are sabotaged by those of inaccurate DMs,

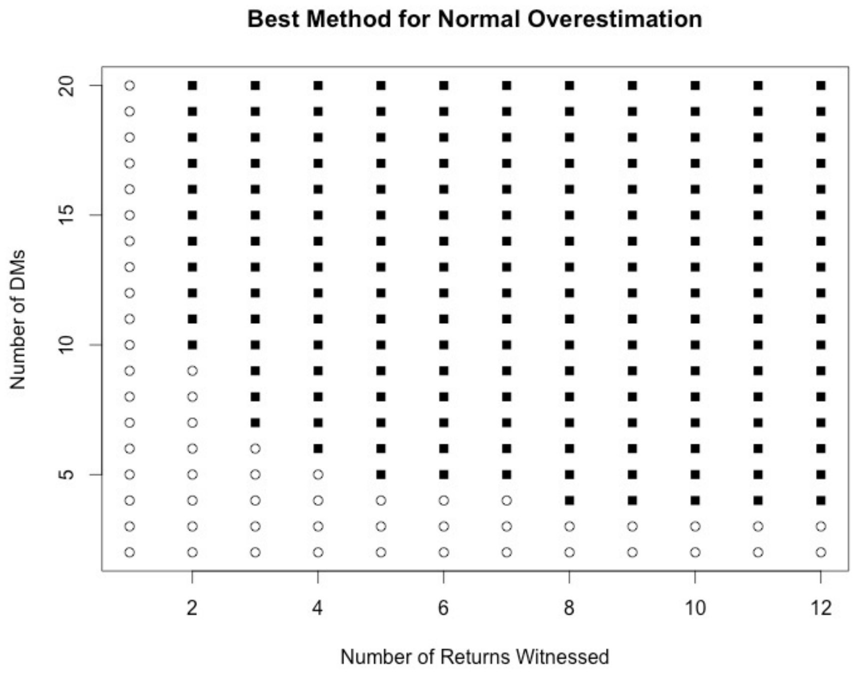

Fig. 4.2: Plot of the method with the highest success proportion in Normal overestimation for a varying amount of DMs/returns. Empty circles imply the MR method is superior, with filled squares/triangles for the PI/EQ methods respectively.

with performance-based weighting giving more accurate estimation. In addition the EQ method often produces higher variances than alternatives, *i.e.*, supplying quite flat (uninformative) distributions.

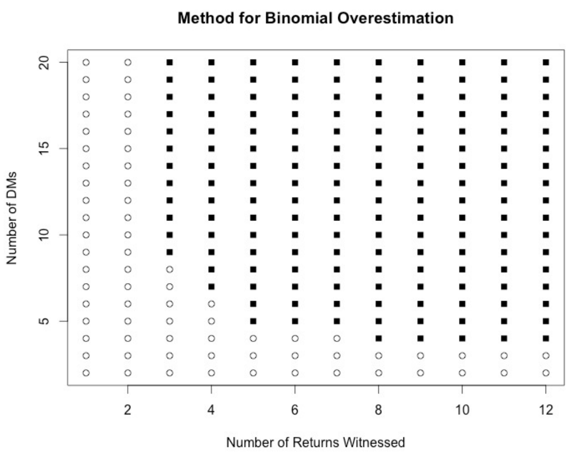

(a)

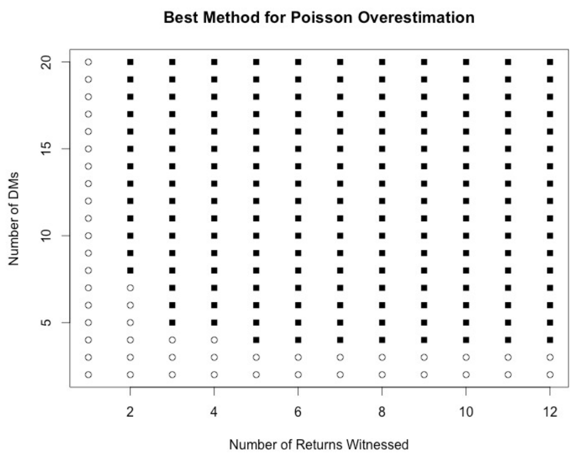

(b)

Fig. 4.3: Optimal methods for Binomial and Poisson overestimation.

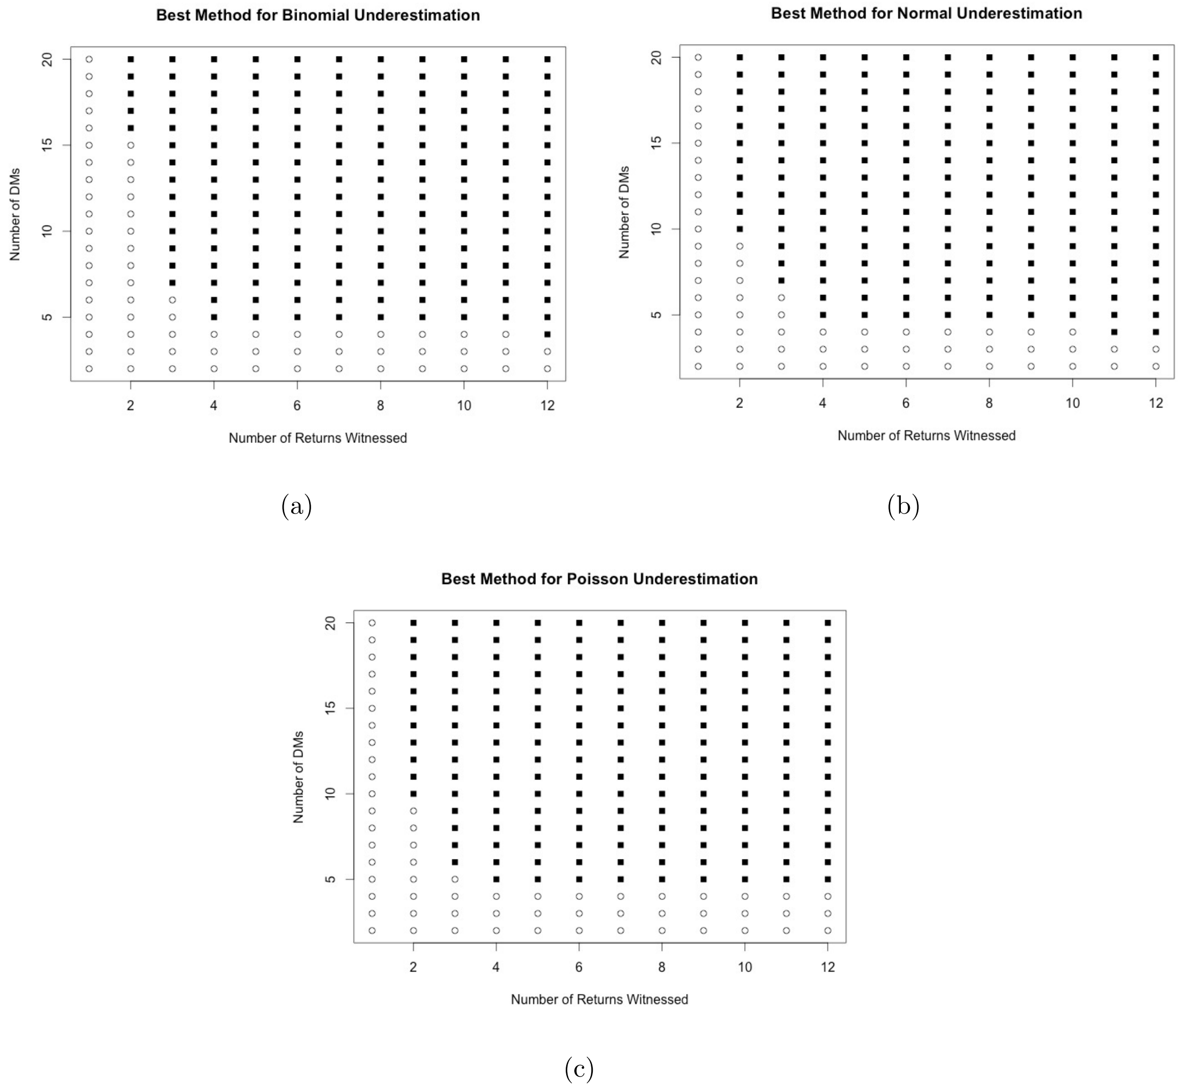

Fig. 4.4: Optimal methods for Binomial, Normal and Poisson underestimation.

We also compare the PI posterior distributions to those posterior distributions of individual DMs. It seems clear that for a reasonable amount of DMs and returns the PI approach is superior to the more simplistic EQ and MR methods, but we have not yet discussed if DMs would be best served simply heeding solely their own opinions and ignoring the additional available information. We conducted a similar study to resolve this using the same initialisations in Table 4.1. For each simulation we considered the density placed on θ by the PI and individual posteriors, declaring the former superior if it placed more density than over half the individual posteriors (ties can occur if there is an even amount of DMs). In each case we aggregated results over 5,000 simulations. In several cases with two DMs individual distributions often tied with the PI distribution,

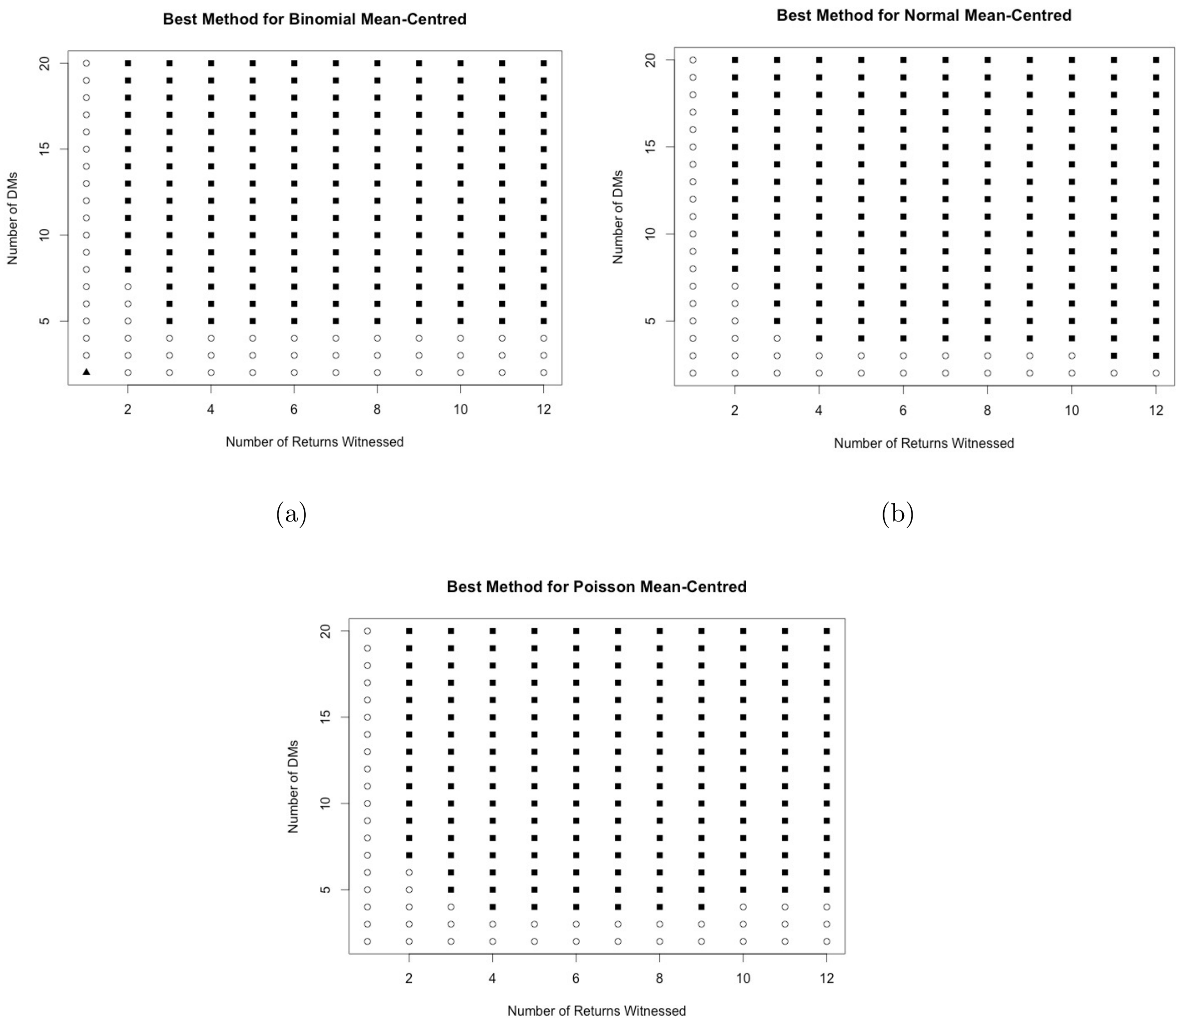

(a) (b) (c)

Fig. 4.5: Optimal methods for Binomial, Normal and Poisson mean-centred beliefs.

i.e., the PI approach gave more accurate estimation for the less accurate DM and less accurate estimation for the more accurate DM. For the case of Binomially mean-centred opinions, individuals own distributions marginally outperformed the weighted combination, although we note that the PI method generally had a more accurate mean but a higher variance. As previously discussed, if most DMs are *a priori* accurate then they will gain little from listening to the opinions of neighbours. However in realistic applications it is unlikely that numerous DMs will have opinions centred on the exact true state of nature. In all other instances the PI approach was superior.

We illustrate success proportions for Normal overestimation in Table 4.3. We see a decrease in PI success proportions over time, which is to be expected; as DMs witness

Table 4.2: PI, EQ and MR success proportions for Normal Overestimation after three and four returns for a varying amount of DMs. Bold font denotes the optimal method.

	After 3 Returns			After 4 Returns		
DMs	PI	EQ	MR	PI	EQ	MR
2	0.1902	0.1398	**0.6700**	0.2276	0.1470	**0.6254**
3	0.3038	0.1344	**0.5618**	0.3306	0.1298	**0.5396**
4	0.3636	0.1204	**0.5160**	0.3974	0.1344	**0.4682**
5	0.4072	0.1326	**0.4602**	0.4378	0.1212	**0.4410**
6	0.4324	0.1312	**0.4634**	**0.4400**	0.1228	0.4372
7	**0.4388**	0.1228	0.4384	**0.4758**	0.1264	0.3978
8	**0.4656**	0.1214	0.4130	**0.4788**	0.1300	0.3912
9	**0.4824**	0.1246	0.3930	**0.4816**	0.1326	0.3858
10	**0.4838**	0.1094	0.4068	**0.4860**	0.1280	0.3860
11	**0.4798**	0.1290	0.3912	**0.4944**	0.1276	0.3780
12	**0.4858**	0.1288	0.3854	**0.5110**	0.1190	0.3700
13	**0.4922**	0.1322	0.3746	**0.5206**	0.1238	0.3556
14	**0.4960**	0.1200	0.3840	**0.5008**	0.1270	0.3722
15	**0.4986**	0.1284	0.3730	**0.5308**	0.1192	0.3500
16	**0.5020**	0.1186	0.3794	**0.5328**	0.1276	0.3396
17	**0.4972**	0.1292	0.3736	**0.5308**	0.1202	0.3490
18	**0.5122**	0.1194	0.3684	**0.5296**	0.1332	0.3372
19	**0.5004**	0.1316	0.3680	**0.5460**	0.1240	0.3300
20	**0.5162**	0.1186	0.3652	**0.5346**	0.1258	0.3396

more data they themselves become increasingly accurate and less dependent upon the views of their more accurate neighbours. As mentioned above ties are possible when an even number of DMs are involved. Hence we observe slightly lower success proportions in even cases than odd cases, particularly in instances with a small numbers of DMs. In general, success proportions remain relatively constant as the number of DMs increases. We concur that there is doubtlessly merit in DMs using the combined distribution in place of their own, as well as in place of the two alternative pooling methods discussed. The bulk of simulation figures is omitted here for brevity, with several representative illustrations provided in Appendix C. In conclusion, it is clear from the wide range of cases considered (spanning three commonly used distributional assumptions and a vast array of initialisation configurations) that the PI approach is an attractive decision making tool.

Table 4.3: Individual and PI success proportions after three and four returns for a varying amount of DMs. Bold font denotes the optimal method.

	After 3 Returns		After 4 Returns	
DMs	IND	PI	IND	PI
2	**0.5000**	**0.5000**	**0.5000**	**0.5000**
3	0.1728	**0.8272**	0.1876	**0.8124**
4	0.2071	**0.7929**	0.2319	**0.7681**
5	0.1750	**0.8250**	0.1880	**0.8120**
6	0.1931	**0.8069**	0.2085	**0.7915**
7	0.1732	**0.8268**	0.1900	**0.8100**
8	0.1805	**0.8195**	0.1941	**0.8059**
9	0.1774	**0.8226**	0.1844	**0.8156**
10	0.1835	**0.8165**	0.2026	**0.7974**
11	0.1756	**0.8244**	0.2014	**0.7986**
12	0.1712	**0.8288**	0.1894	**0.8106**
13	0.1668	**0.8332**	0.2030	**0.7970**
14	0.1728	**0.8272**	0.1940	**0.8060**
15	0.1628	**0.8372**	0.1988	**0.8012**
16	0.1699	**0.8301**	0.1925	**0.8075**
17	0.1722	**0.8278**	0.1958	**0.8042**
18	0.1678	**0.8322**	0.1927	**0.8073**
19	0.1780	**0.8220**	0.1958	**0.8042**
20	0.1730	**0.8270**	0.1915	**0.8085**

4.3 Theoretical Calculations

4.3.1 True Success Probabilities

Above we demonstrated the merits of the PI approach by simulation, comparing its performance to those of alternatives and recording the proportion of times it was superior to these. As the amount of simulations ran increased, success proportions seemed to converge to a constant value, *i.e.*, the true probability of the PI approach being superior. We demonstrate calculation of this probability and illustrate convergence, focusing on the Beta-Binomial case. We comment on the Poisson-Gamma and Normal-Normal cases, which are given in full in Appendix E. We assume n is odd (*i.e.*, ties cannot occur) but this material is easily modified if not. In the Beta-Binomial case θ is a Bernoulli success probability which each P_i has a Beta prior over:

$$f_i(\theta) = \frac{\Gamma(\alpha_i + \beta_i)}{\Gamma(\alpha_i)\Gamma(\beta_i)} \theta^{\alpha_i - 1}(1-\theta)^{\beta_i - 1} \text{ with } \theta \in [0,1] \tag{4.1}$$

Returns are realisations of a Binomial random variable R. Each r is a number of successes in m independent, identically distributed trials. Each r has probability of

$$f(R = r|\theta) = \binom{m}{r} \theta^r (1-\theta)^{m-r} \text{ with } r = 0, 1, \ldots, m \tag{4.2}$$

Over t epochs there are $(m+1)^t$ possible $t-$tuples of returns witnessable. The unnormalised weight given to P_i after t returns $(u_{i,t})$ is a product of her PI weights over these t returns and her initial equal weight, as in Equation (3.14). Denoting r_k as the k^{th} epoch return and $w_{i,k}$ for the PI weight of P_i over r_k we have:

$$\begin{aligned} u_{i,t} &= \frac{1}{n} w_{i,1} \ldots w_{i,t} \\ &= \frac{1}{n} f_i(R_1 = r_1) \ldots f_i(R_t = r_t | R_{t-1} = r_{t-1} \ldots, R_1 = r_1) \end{aligned} \tag{4.3}$$

We write $\alpha_i^{(k)}$ and $\beta_i^{(k)}$ for the updated hyperparameters of P_i after k returns:

$$\alpha_i^{(k)} = \alpha_i + \sum_{j=1}^{k} r_j \tag{4.4}$$

$$\beta_i^{(k)} = \beta_i + km - \sum_{j=1}^{k} r_j \tag{4.5}$$

Using the convention that $\alpha_i^{(0)} = \alpha_i$ and $\beta_i^{(0)} = \beta_i$, we rewrite Equation (4.3) as

$$u_{i,t} = \frac{1}{n} \prod_{k=1}^{t} \binom{m}{r_k} \frac{(\alpha_i^{(k-1)} + \beta_i^{(k-1)} - 1)!(\alpha_i^{(k-1)} + r_k - 1)!(\beta_i^{(k-1)} + m - r_k - 1)!}{(\alpha_i^{(k-1)} - 1)!(\beta_i^{(k-1)} - 1)!(\alpha_i^{(k-1)} + \beta_i^{(k-1)} + m - 1)!} \tag{4.6}$$

By independence, the probability of return set $\{r_1, \ldots, r_t\}$ is a product of Binomials:

$$f(R_1 = r_1, \ldots, R_t = r_t|\theta) = \prod_{k=1}^{t} \binom{m}{r_k} \theta^{r_k}(1-\theta)^{m-r_k} \tag{4.7}$$

Any value of Equation (4.6) occurs with probability in Equation (4.7). The normalised weight of P_i after t returns (here written as $\gamma_{i,t}$ to avoid notational confusion) is

$$\begin{aligned}
\gamma_{i,t} &= \frac{u_{i,t}}{\sum_{j=1}^{n} u_{j,t}} \\
&= \frac{\frac{1}{n}\prod_{k=1}^{t}\binom{m}{r_k}\frac{(\alpha_i^{(k-1)}+\beta_i^{(k-1)}-1)!(\alpha_i^{(k-1)}+r_k-1)!(\beta_i^{(k-1)}+m-r_k-1)!}{(\alpha_i^{(k-1)}-1)!(\beta_i^{(k-1)}-1)!(\alpha_i^{(k-1)}+\beta_i^{(k-1)}+m-1)!}}{\sum_{j=1}^{n}\frac{1}{n}\prod_{k=1}^{t}\binom{m}{r_k}\frac{(\alpha_j^{(k-1)}+\beta_j^{(k-1)}-1)!(\alpha_j^{(k-1)}+r_k-1)!(\beta_j^{(k-1)}+m-r_k-1)!}{(\alpha_j^{(k-1)}-1)!(\beta_j^{(k-1)}-1)!(\alpha_j^{(k-1)}+\beta_j^{(k-1)}+m-1)!}} \\
&= \frac{\prod_{k=1}^{t}\binom{m}{r_k}\frac{(\alpha_i^{(k-1)}+\beta_i^{(k-1)}-1)!(\alpha_i^{(k-1)}+r_k-1)!(\beta_i^{(k-1)}+m-r_k-1)!}{(\alpha_i^{(k-1)}-1)!(\beta_i^{(k-1)}-1)!(\alpha_i^{(k-1)}+\beta_i^{(k-1)}+m-1)!}}{\sum_{j=1}^{n}\prod_{k=1}^{t}\binom{m}{r_k}\frac{(\alpha_j^{(k-1)}+\beta_j^{(k-1)}-1)!(\alpha_j^{(k-1)}+r_k-1)!(\beta_j^{(k-1)}+m-r_k-1)!}{(\alpha_j^{(k-1)}-1)!(\beta_j^{(k-1)}-1)!(\alpha_j^{(k-1)}+\beta_j^{(k-1)}+m-1)!}}
\end{aligned} \tag{4.8}$$

Weights in Equation (4.8) are merged with distributions in Equation (4.1) with new hyperparameters from Equations (4.4) and (4.5) giving a PI posterior after t returns of

$$\begin{aligned}
\hat{f}_t^{PI}(\theta|r_1, \ldots, r_t) &= \sum_{z=1}^{n} \gamma_{z,t} f_z(\theta|r_1, \ldots, r_t) \\
&= \sum_{z=1}^{n} \Bigg[\frac{\prod_{k=1}^{t}\binom{m}{r_k}\frac{(\alpha_z^{(k-1)}+\beta_z^{(k-1)}-1)!(\alpha_z^{(k-1)}+r_k-1)!(\beta_z^{(k-1)}+m-r_k-1)!}{(\alpha_z^{(k-1)}-1)!(\beta_z^{(k-1)}-1)!(\alpha_z^{(k-1)}+\beta_z^{(k-1)}+m-1)!}}{\sum_{j=1}^{n}\prod_{k=1}^{t}\binom{m}{r_k}\frac{(\alpha_j^{(k-1)}+\beta_j^{(k-1)}-1)!(\alpha_j^{(k-1)}+r_k-1)!(\beta_j^{(k-1)}+m-r_k-1)!}{(\alpha_j^{(k-1)}-1)!(\beta_j^{(k-1)}-1)!(\alpha_j^{(k-1)}+\beta_j^{(k-1)}+m-1)!}} \times \\
&\qquad \frac{\Gamma(\alpha_z^{(t)}+\beta_z^{(t)})}{\Gamma(\alpha_z^{(t)})\Gamma(\beta_z^{(t)})} \theta^{\alpha_z^{(t)}-1}(1-\theta)^{\beta_z^{(t)}-1} \Bigg]
\end{aligned} \tag{4.9}$$

Recall that over t epochs there are $(m+1)^t$ possible return streams. Denote by $\boldsymbol{r}_x$ the x^{th} of these. Consider an indicator variable $I_{j,x}$ giving a 1 if the PI posterior places more density on θ than that of P_j given $\boldsymbol{r}_x$, and 0 if not, *i.e.*,

$$I_{j,x} = \begin{cases} 1 & \text{if } \hat{f}_t^{PI}(\theta|\boldsymbol{r}_x) > f_j(\theta|\boldsymbol{r}_x); \\ 0 & \text{if } \hat{f}_t^{PI}(\theta|\boldsymbol{r}_x) < f_j(\theta|\boldsymbol{r}_x). \end{cases}$$

This forms a matrix of zeros and ones, as in Table 4.4. Interest is in its column sums. If for returns $\boldsymbol{r}_x$ the column sum exceeds $\frac{n}{2}$ then, for $\boldsymbol{r}_x$, the PI approach gives better estimation than over half of DMs, *i.e.*, it is superior. The x^{th} column is S_x:

$$S_x = \sum_{k=1}^{n} I_{k,x} \tag{4.10}$$

Table 4.4: Cross-tabulation of DMs and return streams, with $I_{j,x}$ for each cell $\{j, x\}$.

	$\boldsymbol{r}_1$	$\boldsymbol{r}_2$	$\dots$	$\boldsymbol{r}_{(m+1)^t}$
P_1	$I_{1,1}$	$I_{1,2}$	$\dots$	$I_{1,(m+1)^t}$
P_2	$I_{2,1}$	$I_{2,2}$	$\dots$	$I_{2,(m+1)^t}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
P_n	$I_{n,1}$	$I_{n,2}$	$\dots$	$I_{n,(m+1)^t}$

We introduce an indicator, I_x, returning a 1 if this sum exceeds $\frac{n}{2}$, and a 0 if not:

$$I_x = \begin{cases} 1 & \text{if } S_x > \frac{n}{2}; \\ 0 & \text{if } S_x < \frac{n}{2}. \end{cases}$$

Finally we consider the quantities $\{I_x\}_{x=1,\dots,(m+1)^t}$ and the corresponding probabilities for $\boldsymbol{r}_x$, as in Equation (4.7). The probability that the PI approach is superior to DM distributions is the cross-product of these vectors:

$$\mathbb{P}(\text{PI is superior}|\theta) = \sum_{x=1}^{(m+1)^t} I_x f(\boldsymbol{R} = \boldsymbol{r}_x|\theta) \tag{4.11}$$

We demonstrate that proportions convergence to this probability. Consider a case where $R \sim \text{Bin}(2, \theta)$, for θ=0.7. Suppose there are five DMs with respective Beta$(1, 3)$, Beta$(7, 2)$, Beta$(2, 2)$, Beta$(4, 3)$ and Beta$(2, 1)$ priors. After four returns the true success probability is 0.687. We simulated 5,000 process iterations, recording success proportions at each stage as in Fig. 4.6. These converge to 0.687.

We perform similar calculations in the group problem. We defined the PI distribution after t returns in Equation (4.9) and now turn to the EQ and MR distributions. The former is straightforward with weights independent of returns, *i.e.*,

$$\begin{aligned} \hat{f}_t^{EQ}(\theta|r_1, \dots, r_t) &= \sum_{z=1}^{n} \frac{1}{n} f_z(\theta|r_1, \dots, r_t) \\ &= \frac{1}{n} \sum_{z=1}^{n} \Big[\frac{\Gamma(\alpha_z^{(t)} + \beta_z^{(t)})}{\Gamma(\alpha_z^{(t)})\Gamma(\beta_z^{(t)})} \theta^{\alpha_z^{(t)}-1}(1-\theta)^{\beta_z^{(t)}-1} \Big] \end{aligned} \tag{4.12}$$

The MR posterior depends upon which DM was deemed most reliable (*i.e.*, returned the highest PI weight) at the t^{th} epoch. We consider, $I_{j,x}^{MR}$, giving a 1 if P_j has the biggest PI weight having seen $\boldsymbol{r}_x$, and 0 if not, where $\boldsymbol{r}_x[t]$ is the t^{th} element of $\boldsymbol{r}_x$:

$$I_{j,x}^{MR} = \begin{cases} 1 & \text{if } f_j(R_t = \boldsymbol{r}_x[t]|\cdot) = \max_{k \in \{1,\dots,n\}} f_k(R_t = \boldsymbol{r}_x[t]|\cdot); \\ 0 & \text{if } f_j(R_t = \boldsymbol{r}_x[t]|\cdot) \neq \max_{k \in \{1,\dots,n\}} f_k(R_t = \boldsymbol{r}_x[t]|\cdot). \end{cases}$$

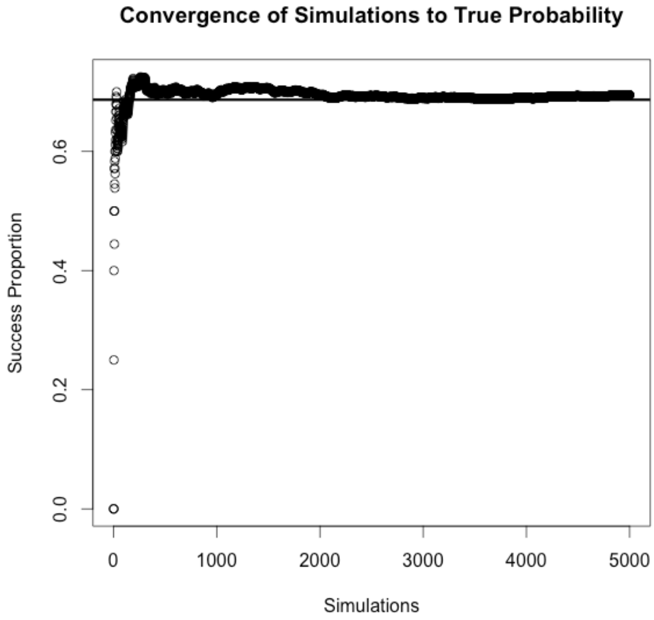

Fig. 4.6: Simulated proportions (the horizontal line denotes the true probability).

The MR posterior after $\boldsymbol{r}_x$ is

$$\begin{aligned}\hat{f}_t^{MR}(\theta|\boldsymbol{r}_x) &= \sum_{z=1}^{n} I_{z,x}^{MR} f_z(\theta|\boldsymbol{r}_x) \\ &= \sum_{z=1}^{n} I_{z,x}^{MR} \times \frac{\Gamma(\alpha_z^{(t)}+\beta_z^{(t)})}{\Gamma(\alpha_z^{(t)})\Gamma(\beta_z^{(t)})} \theta^{\alpha_z^{(t)}-1}(1-\theta)^{\beta_z^{(t)}-1} \qquad (4.13)\end{aligned}$$

Consider three indicators, I_x^{PI}, I_x^{EQ} and I_x^{MR} respectively giving a 1 if the method is superior (in terms of posterior density) given $\boldsymbol{r}_x$ and a zero if not:

$$I_x^{PI} = \begin{cases} 1 & \text{if } \hat{f}_t^{PI}(\theta|\boldsymbol{r}_x) > \max\{\hat{f}_t^{EQ}(\theta|\boldsymbol{r}_x), \hat{f}_t^{MR}(\theta|\boldsymbol{r}_x)\}; \\ 0 & \text{if not.} \end{cases}$$

$$I_x^{EQ} = \begin{cases} 1 & \text{if } \hat{f}_t^{EQ}(\theta|\boldsymbol{r}_x) > \max\{\hat{f}_t^{PI}(\theta|\boldsymbol{r}_x), \hat{f}_t^{MR}(\theta|\boldsymbol{r}_x)\}; \\ 0 & \text{if not.} \end{cases}$$

$$I_x^{MR} = \begin{cases} 1 & \text{if } \hat{f}_t^{MR}(\theta|\boldsymbol{r}_x) > \max\{\hat{f}_t^{PI}(\theta|\boldsymbol{r}_x), \hat{f}_t^{EQ}(\theta|\boldsymbol{r}_x)\}; \\ 0 & \text{if not.} \end{cases}$$

Given these, we define the probabilities of techniques being optimal as

$$\mathbb{P}(\text{PI is superior}|\theta) = \sum_{x=1}^{(m+1)^t} I_x^{PI} f(\boldsymbol{R}=\boldsymbol{r}_x|\theta) \qquad (4.14)$$

$$\mathbb{P}(\text{EQ is superior}|\theta) = \sum_{x=1}^{(m+1)^t} I_x^{EQ} f(\boldsymbol{R}=\boldsymbol{r}_x|\theta) \qquad (4.15)$$

$$\mathbb{P}(\text{MR is superior}|\theta) \quad = \quad \sum_{x=1}^{(m+1)^t} I_x^{MR} f(\boldsymbol{R} = \boldsymbol{r}_x|\theta) \tag{4.16}$$

We demonstrate convergence to these probabilities (using the previous example) in Fig. 4.7. Full calculation details for Poisson-Gamma and Normal-Normal cases are in Appendix E. These follow from the above material with minor modifications. In the Binomial case we can produce a (finite) exhaustive list of all possible values of r, *i.e.*, there are $m+1$. The same cannot be said in the Poisson case which has a (countably) infinite set of possibilities. Equations (4.11) and (4.14) – (4.16) must be finite sums to ensure computability. Hence in the Poisson case we must choose a (finite) upper bound value that returns have negligibly small probability of exceeding. In the continuous Normal distribution, which has an (uncountably) infinite set of returns, we perform a similar process by discretising the range (in a manner that is not too coarse) and choosing suitable lower and upper bounds.

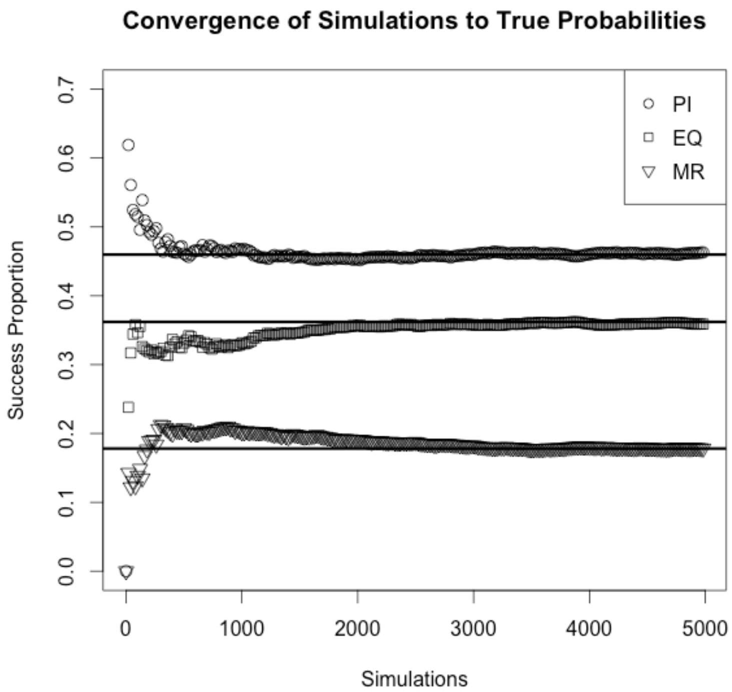

Fig. 4.7: Proportions from simulations (horizontal line denotes the true probabilities).

The examples here are low-dimensional as there are issues related to "curse of dimensionality". If there are N potential returns per epoch then there are N^t potential return streams over t epochs. This term grows quickly, often making computation in a reasonable time impossible. If there are 50 returns considered across five epochs there are over 3×10^8 combinations to compute. This is extremely computationally slow with this number growing unmanageable as epochs increase. Yet we have seen how accurately simulations mirror probabilities. Hence it is adequate to talk in terms of proportions rather than probabilities, as the former are very accurate estimates of the latter with severely decreased computation time. In non-conjugate cases it is impossible to give closed form PI weights meaning simulation is required, *i.e.*, probabilities cannot be calculated, only approximated by proportions. Nevertheless, this material is a nice derivation which confirms the correctness of our simulation method.

4.3.2 Unconditional Probabilities

Above we determined the true probability of the PI approach being superior to a set of alternatives. Implicit within this material is the assumption that the true value of θ is known. For our justification purposes this is perfectly reasonable, but in practical applications DMs will be unaware of the true value of the latent parameter θ prior to (and during) the process, modelling their uncertainty via probability distributions. The question we seek to answer here is the following: if a DM has her prior opinion over θ and receives opinions from neighbours over θ then what is the probability she *a priori* associates with the PI approach giving more accurate estimation than her own distribution? This is akin to "integrating out" θ in Equation (4.11), *i.e.*, we are interested in $\mathbb{P}(\text{PI is superior})$ rather than $\mathbb{P}(\text{PI is superior}|\theta)$.

Here we consider the Beta-Binomial case where $(m+1)^t$ possible return sets can be witnessed over t epochs. We consider each case in turn, and determine if the resulting PI posterior would be more accurate than the DM's own distribution, given her prior opinion over θ. As θ is unknown to the DM she calculates the probability of PI superiority for a range of values of θ, *e.g.*, $\theta \in \{0, 0.01, \ldots, 0.99, 1\}$ in turn. If the PI method has a probability over 0.5 of yielding more accurate estimation than a DM's own distribution then she will be inclined to use it. We produce plots of the corresponding probabilities of the PI method being superior for the range of values of θ in Fig. 4.8 for three DMs with respective $\text{Beta}(4, 2)$, $\text{Beta}(2, 2)$ and $\text{Beta}(1, 3)$ priors, with $\theta \sim \text{Bin}(2, \theta)$, and two returns to be witnessed. P_1 will *a priori* consider the probability of the PI approach being superior to exceed 0.5, given $f_1(\theta)$, if $\theta \leq 0.62$, while for P_2 this holds if $\theta \leq 0.18$ and $\theta \geq 0.78$, and for P_3 if $\theta \geq 0.32$. DMs consider that they will benefit from listening to opinions of others if the true θ transpires to be suitably different from their predictions. Respective prior means of DMs are $\mathbb{E}_{f_1}(\theta) = 0.67$, $\mathbb{E}_{f_2}(\theta) = 0.5$ and $\mathbb{E}_{f_3}(\theta) = 0.25$. At values of θ close to their mean estimates each DM considers her own opinion superior to the PI approach, *i.e.*, that they will not benefit from listening to neighbours' opinions as their own opinion is already sufficiently accurate. DMs do not know *a priori* if their own opinion is accurate or not, and therefore are taking a gamble in some sense: using the PI approach is akin to admitting a possibility that her own belief is incorrect, and hence she is willing to listen to opinions of neighbours, which may improve her accuracy if her own distribution is inaccurate.

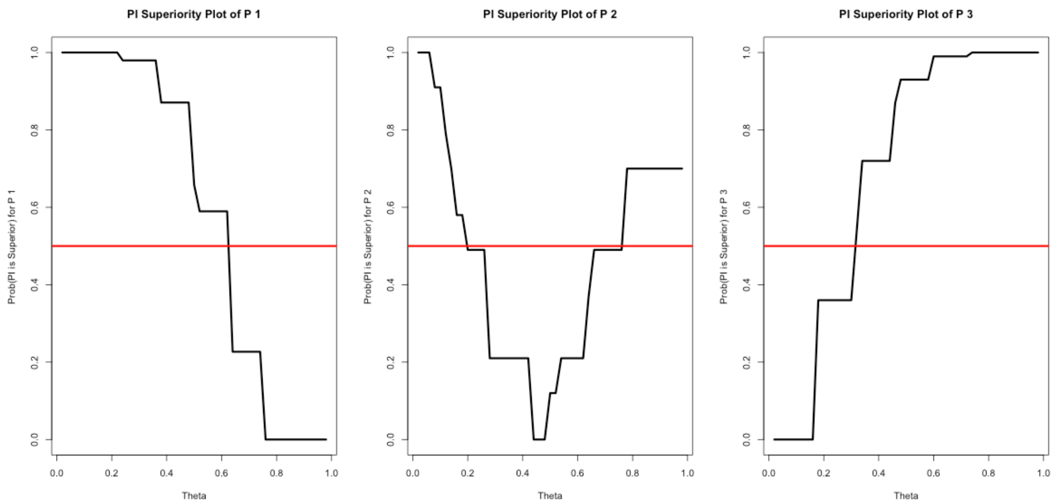

Fig. 4.8: Prior probabilities DMs place on the PI approach yielding superior estimation to their own distributions for varying θ. Horizontal lines denotes a probability of 0.5.

We declared in the individual problem that the PI approach was superior if, for over half the DMs, it led to better estimation over θ. Here we consider, for each value of θ in turn, if the PI approach is considered to have a probability of over 0.5 of giving more accurate estimation for over half of DMs. We find if $0.2 \leq \theta \leq 0.3$ or $0.64 \leq \theta \leq 0.76$ then the PI approach will not be considered best for the set of DMs, while outside this range it will be. We conduct similar analysis in the group problem. For each θ contemplated a DM assesses if *a priori* the PI, EQ or MR approach will yield more accurate estimation for her. We illustrate this is in Fig. 4.9, showing the success probabilities assigned by P_2 to techniques for the θ values considered under her own opinions about θ. The EQ method is deemed optimal for $\theta \leq 0.72$, with PI optimal for $\theta \geq 0.82$, and MR briefly optimal for $0.72 < \theta < 0.82$.

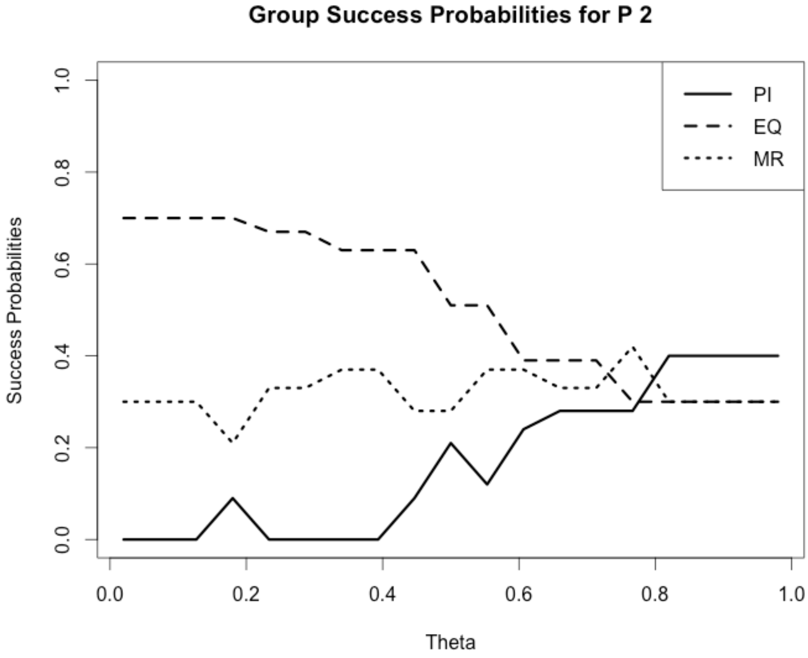

Fig. 4.9: Success probabilities for the PI, EQ and MR approaches for P_2.

Note the distinction between the material here and in the previous subsection. The latter is used for justification, where the value of θ is assumed known in order to calculate the probability of the superiority of the PI method. In the former, DMs wish to know before witnessing returns which values of θ lead to the PI distribution being superior. The resulting probabilities should not impact upon their decision to use the PI approach. A DM naturally believes her opinion over θ is correct: if not she would augment it to one that she felt represented θ more accurately (see the temporal coherence argument in, *e.g.*, Goldstein, 2001). McConway & Genest (1990) provide arguments (and corresponding proofs) similar to this, discussing how a DM should always expect her weight, and those of DMs with opinions similar to her own, to increase at the next epoch, as she believes the data will validate her opinion (although this is based on the assumption of dialectical equilibrium, a somewhat behavourial element we don't directly consider). Yet DMs using the PI approach acknowledge that their own opinion may be incorrect and that they themselves are not infallible, and hence are willing to take into account the opinions of neighbours. They will learn over time about who is reliable and who is not (themselves included), with this additional information used in their weighting scheme. This work serves as a footnote to our theoretical calculations in Section 4.3.1, giving results that are not conditional upon θ.

Ideally we would formulate a theorem that provides details of the necessary conditions required for the PI approach to be superior to the considered alternatives. However from the earliest examinations of methods of opinion combining (e.g., the aforementioned study of Newbold & Granger, 1974) derivation of such a formal theorem has appeared extremely problematic, with comparison to existing methodologies under a particular metric the commonly used technique for assessing the merits of a particular method. Even the classical method of Cooke (1991) is primarily validated using suitable data studies, in addition to some desirable properties that it obeys. Our justification for the PI approach is the same, using our data studies and the properties laid out in Sections 3.2 and 3.4. French (1985) discusses how impossibility theorems exist to show that no combination rule can concurrently obey all of a set of attractive criteria, meaning that while it is doubtlessly good that our method adheres to these criteria, there are certainly other criteria which could be deemed desirable that it does not adhere to.

As previously alluded to, there are a wide range of metrics which can be used to determine superiority, such as the distance between a distributions mean and the true value, which would be most appropriate in a case where only a point estimate was required by the user. However, given our interest in higher order distributional moments, our probability density metric, which explicitly considers both the mean and the variance of the mixture distribution, seems a suitable choice for assessing the quality of our method. As in Section 4.3.2 a DM can try to determine the merits of the PI approach using our theoretical calculations, but this is reliant upon her own opinion about the uncertain quantity of interest, and hence cannot be relied upon: if her understanding of the uncertain quantity is poor then her estimation of the merits of the PI approach are likely to be inaccurate. We can only make the following formal statements, which concern the comparison of the PI approach to the two rational EQ and MR alternatives, under the probability density metric:

- The performance of the PI approach will improve as the number of DMs in the group grows.

- The performance of the PI approach will improve as the number of witnessed returns grows.

- In the limit, as the number of witnessed returns tends to infinity, the performance of all three approaches will become identical (assuming no individual supplies a degenerate prior distribution).

4.4 TU Delft Expert Judgment Data Base

Having validated the PI method with simulated data we attempt to do the same using real data, specifically the TU Delft Expert Judgment Data Base (introduced in Section 2.7). In the framework that we consider all opinion-holding individuals are DMs, using the belief that they hold, as well as those of neighbours, to make decisions. All individuals are completely involved in the decision process, *i.e.*, not just supplying information to be compiled and collated but in sharing the inherent risk. By contrast in the classical method opinions over uncertainty are offered by a set of experts, which are synthesised for use by a single DM who herself is not an expert about the relevant uncertainty, *i.e.*, holds no well-informed opinion over θ. In our PI context each DM makes a decision, while in the classical method only one decision is made, and it is not by an original opinion-holder. Yet the goals of both methods are concurrent, seeking to form a weighted sum of maximum accuracy for a decision task.

We discuss the form of opinions in the TU Delft Expert Judgment Data Base. Each data set consists of n experts giving opinions on m similarly themed seed variables. In Table 4.5 we detail the available forty two data sets explored. In the PI framework we often assume before making a decision (and witnessing a return) that all DMs are equally reliable. Only once relevant returns are seen can reliability statements be made. By contrast in the classical method experts have their accuracy assessed on seed variables (whose true values are known to the DM) with opinions compared to true values and weights based on these disparities. Hence prior to decision making reliability information is available. While seeds are related to the unknown element of the decision problem they are not this element itself, *i.e.*, accuracy on seeds may not correspond to accuracy for the unknown quantity, and the converse. In the PI approach equal weights are only changed given returns that are noisy realisations of θ.

Opinions of DMs in the PI approach must be fully parameterised probability distributions. Under the classical method opinions are quantile values $\{q_{(0.05)}, q_{(0.5)}, q_{(0.95)}\}$

Table 4.5: Details of the data sets in TU Delft Expert Judgment Data Base.

Name	Field	n	m	Name	Field	n	m
A_SEED	Zoology	5	8	LADDERS	Falls	7	10
ACNEXPTS	Chemistry	7	10	MONT1	Volcanos	11	8
ACTEP	Air Traffic	6	10	MVOSEEDS	Volcanos	77	5
AOTDAILY	Trading	5	34	NH3EXPTS	Ammonia	3	10
AOTRISK1	Risk	5	11	OPRISKBANK	Risks	10	16
CARMAG	Health	6	10	PHAC	Health	14	14
DAMS	Dams	11	11	PILOTS	Pilots	31	10
DIKRING	Failure	17	47	PM25	Physics	6	12
DSM1	Safety	10	8	REALESTR	Real Estate	4	31
DSM2	Safety	8	10	RETURN1	Real Estate	5	15
ESTEC1	Space Exp.	4	13	RIVERCHNL	Rivers	6	8
ESTEC2	Space Exp.	7	26	SARS	Health	9	10
ESTEC3	Space Exp.	6	12	SO3EXPT	Physics	4	9
EUNRCRWD	Geography	7	14	TEIDE_MAY	Volcanos	17	10
EUNRCDD	Geography	8	14	THRMBLD	Physics	6	48
EUNRCDIS	Dispersion	8	23	TNO_DISP1	Dispersion	7	36
EUNRCEAR	Health	7	16	TNO_DEPOS1	Geography	4	21
FCEP	Air Safety	5	8	TUD_DISPER1	Dispersion	11	36
GLINVSPC	Species	9	13	VESUVIO	Volcanos	14	10
GROND5	Transport	7	10	VOLCRISK	Risk	45	10
INFOSEC	Security	13	10	WATERPOL	Pollution	11	9

which an expert believes that the realisation has a probability of 0.05, 0.5 and 0.95 of lying below respectively. As discussed in Section 2.7 weights are functions of calibration (accuracy) and information (width) scores, with experts potentially given a weight of zero if their score fails to meet a minimal threshold. In the PI approach all DMs have a strictly positive weight at all epochs. DMs are never totally eliminated from the process due to its learning aspect, *i.e.*, an inaccurate DM becomes more reliable as data is witnessed. DMs who appear initially inaccurate may actually be reliable, with the returns witnessed thus far being tail (*i.e.*, unlikely) events. These reasons provide motivation for constant positivity of weights. The dichotomy between the methods is rationalised by their subtly differing objectives. The PI approach is for use in a dynamic decision setting where new data repeatedly becomes available and old opinions (both over θ and DM accuracy) are augmented. Under the classical method all learning occurs in one step, with weights based on a single set of realisations. A summary of these contrasts is in Table 4.6.

Table 4.6: Differences between the Plug-in and classical methods.

	Plug-in	Classical
Decisions made?	n (per epoch)	1
DMs hold opinions?	Yes	No
Form of opinions?	Probability Distributions	Quantiles
Info. prior to 1^{st} decision?	None (Assume equal reliability)	Yes (Scores from seeds)
Weights of zero?	Never	Sometimes

We illustrate the data using data set A_SEED which consists of five experts and eight seeds. In Table 4.7 we show predictions of the first expert, and true values (having severely differing orders of magnitudes, which we comment on shortly). The central 90% interval for this expert contains true values for only the first three seeds.

Table 4.7: The quantiles of the first expert over the eight seeds in the A_SEED data.

Seed	$q_{(0.05)}$	$q_{(0.5)}$	$q_{(0.95)}$	True	Seed	$q_{(0.05)}$	$q_{(0.5)}$	$q_{(0.95)}$	True
1	0.002	0.019	0.036	0.027	5	70	76	82	92
2	449	2690	4930	3460	6	49	58	66	75
3	673	4040	7400	5090	7	13	20	27	46
4	0.009	0.012	0.015	0.006	8	1.7	3.1	5.8	21

4.4.1 Fitting Distributions to Quantiles

An issue with modifying the data to the PI context is the contrasting belief specification method. A cornerstone of the PI method is the PI weights, *i.e.*, the values of prior predictive distributions for realisations. Consider Normal-Normal conjugacy. Interest is in the mean θ of a Normally distributed process, $R \sim \mathcal{N}(\theta, \sigma^2)$, with σ^2 known.

The prior opinion of P_i is Normal, *i.e.*, $f_i(\theta) \sim \mathcal{N}(m_i, s_i^2)$. As previously demonstrated this gives a prior predictive distribution of $f_i(R = r) \sim \mathcal{N}(m_i, s_i^2 + \sigma^2)$. How can we modify the data, in terms of quantile opinions and returns seen, to suit our purposes? There are numerous reasons for our Normality assumption: its prevalence in realistic scenarios, computational simplicity, that it is straightforward to translate quantiles to distributions under this assumption, and, critically, that many quantiles appear approximately (and often exactly) symmetric around their central quantile, *e.g.*, Table 4.7. Denoting $q_i(x)$ as the x^{th} quantile of P_i we write her mean as

$$m_i \quad = \quad q_i(0.5) \tag{4.17}$$

We considered various forms for the standard deviation, most notably an arithmetic average of the (rescaled) distances from the mean to the lower/upper quantiles. However we found the best representation of an opinion, counteracting occasional extreme underestimation and overestimation, was the minima of these distances, given in Equation (4.18). The 1.645 arises from standard Normal theory, *i.e.*, if the quantity is truly Normally distributed then the distance from its mean to its lower/upper quantile is $1.645 s_i$. This combination of m_i and s_i led generally to close fits to the data.

$$s_i \quad = \quad \min\left(\frac{q_i(0.95) - q_i(0.5)}{1.645}, \frac{q_i(0.5) - q_i(0.05)}{1.645}\right) \tag{4.18}$$

For each seed variable, we consider its true value as that realised by DMs, *i.e.*, the value plugged into their prior predictive distribution. We monitor how well fitted Normal distributions match quantile predictions. For each DM and seed in turn we can graph the Normal cumulative distribution function of the DM over the seed and add their quantiles. If these points lie on the fitted function then this indicates a good fit. We also see if the true value is in a DMs central 90% probabilistic range. An example of this in Fig. 4.10, shows quantiles, fitted distributions and true parameter values for the first four experts over the second A_SEED seed. The Normality assumption seems very acceptable for the first, second and fourth experts, and slightly less so for the third expert. The quantile ranges of all experts contain the true parameter value, although that of the third expert is far wider (*i.e.*, less informative) than those of her peers.

Finally we discuss the variance σ^2 (assumed known) of realised values. If we had a set of realisations for each seed then we could approximate the true variance as an average of their deviation around their sample mean. However we only have a single

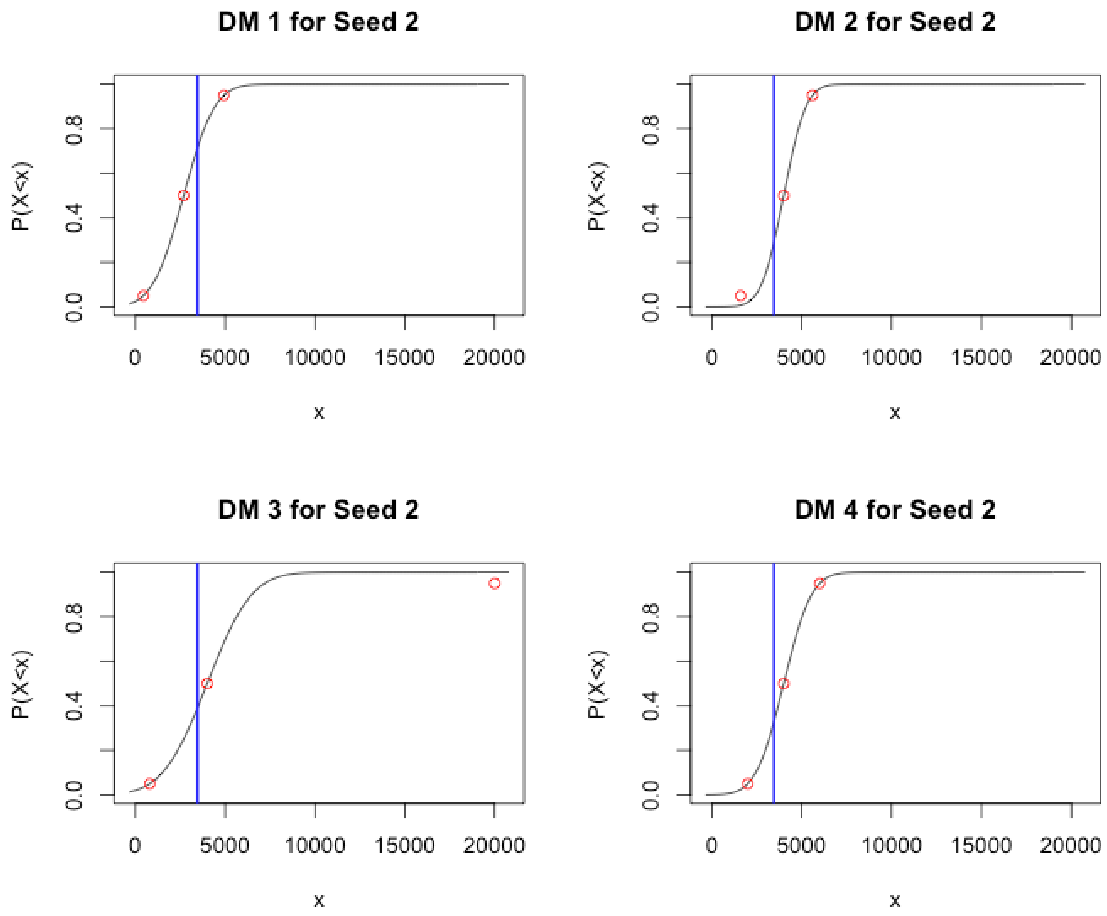

Fig. 4.10: Fitted cumulative distributions of experts using Equations (4.17) and (4.18). Unfilled circles are quantiles and vertical lines represent true parameter values.

value, *i.e.*, the realisation itself. Traditional estimation is not possible. We propose a simple alternative, assuming σ^2 is a fixed fraction of the true/realised value r, *i.e.*,

$$\sigma^2 = \frac{r}{k} \text{ where } k \in \mathcal{R}^+ \tag{4.19}$$

This assumption may appear to let variance be chosen randomly but we provide notes on this choice here. We commented above on the disparity in orders of magnitude of seeds within data sets. Equation (4.19) ensures that the variance of a seed is proportional to its order of magnitude, *e.g.*, if $k = 10$ then a true value of 10 has associated variance of $\sigma^2 = 1$ while a true value of 1000 has $\sigma^2 = 100$. This is more sensible than a variance parameter constant across all seeds independent of magnitudes. Below we conduct analysis for several values of k ensuring robustness of results, *i.e.*, conclusions are not solely based on a specific arbitrary parameterisation choice but hold for numerous possibilities. We shortly see that conclusions are reasonably invariant to k.

4.4.2 Problem Type and Metric Choice

Previously we discussed two contexts where the PI approach can be applied: individual and group problems. In the former DMs decide whether to listen solely to their own opinion or to a PI combination of this and those of her neighbours. The group problem supposes linear pooling must be carried out, *e.g.*, in a full group decision making context or one in which a DM is committed to heeding other opinions but is unsure how to do this. Below we compare the PI distribution to those of DMs (the individual problem) as well as the EQ and MR alternatives (the group problem). These are the same alternatives considered by Cooke & Goossens (2008) and Eggstaff *et al.* (2014), although in the classical method these are used in a static, rather than dynamic, fashion. We consider the metric previously discussed, where densities placed by competing posterior distributions are compared, and the method maximising this density is deemed superior.

The experts in the classical process do not directly learn about unknown quantities over time, *i.e.*, having given quantiles for the first seed, and seen its true value, the expert has not ascertained new information about the second seed. This differs from our context where DMs learn at each epoch, with every realisation being a noisy version of θ. The only learning between seeds in the classical method is that experts whose quantiles have not contained seeds may contemplate widening their ranges for the next seed, *i.e.*, considering that they may be overconfident in their assessments and are providing overly narrow predictions. To quote Cooke & Goossens (2008): *"If after a few realisations the expert was to see that all realisations fall outside his 90% central confidence intervals, he might conclude that these intervals were too narrow and broaden them on subsequent assessments. This means that for this expert, the uncertainty distributions are not independent, as he learns from the realisations. Expert learning is not a goal of an expert judgment study ... rather the decision maker wants experts who do not need to learn from the elicitation."*

This further highlights the contrasting goals of the PI and classical methods. Permuting the original seed ordering is acceptable as there is no clear Markov dependency between consecutive assessments. For a particular permutation we calculate weights for the first seed, observe this value, then update weights for the second seed, and so on, repeating this for various permutations. This contributes an additional robustness to our results as they represent numerous configurations of seed variables rather than a single one. Seed variables should all relate to the uncertain variable in the forthcoming decision task, so they are not fully independent, but their ordering is irrelevant to the classical method with weighting scores aggregated over all seeds. Learning over uncertainty does not occur in the strict Bayesian manner of the PI approach and seeds should be somehow correlated but may not always be. Hence an expert who previously appeared very reliable may suddenly appear highly inaccurate due to a seed being only weakly related to her field of expertise, or differing degrees of seed difficulty. As only minimal learning occurs between seeds (experts possibly widening quantiles to compensate for overconfidence) we permute seeds and still validly measure corresponding results. The dependence between consecutive assessments quoted by Cooke & Goossens (2008) above refers to an experts assessment of her own accuracy rather than learning over seeds.

The PI goal is learning over time so our assessments are conducted on final seeds, *i.e.*, when maximal information has been witnessed and DMs have a good understanding of the reliability of neighbours. If there are m seeds then there are $m(m-1)$ permutations to be considered to take into account every possibility. We saw in Section 3.2 that the PI approach is invariant to the order in which a set of returns is witnessed. Hence identical weights are given at the final seed regardless of if it is preceded by $x_{\sigma_1(1)}, \ldots, x_{\sigma_1(m-1)}$ or $x_{\sigma_2(1)}, \ldots, x_{\sigma_2(m-1)}$ with σ_1 and σ_2 distinct permutations of a common set. Hence for the PI approach we need only consider each seed in turn as the final one. The same holds for the EQ case, with past reliability not factored into weights. However the MR method is dependent upon which DM is deemed most reliable for the $(m-1)^{th}$ seed (*i.e.*, who maximises the PI weight) which is itself independent of previous realisations in this context (as learning occurs here only over weights rather than seeds). Hence we must consider every possible permutation of the final two seeds, *i.e.*, $m(m-1)$ permutations. For each permutation we record which method gives the best estimation and aggregate results over all permutations. The technique with the greatest proportion is superior for that data set. A similar study (needing only m permutations) is conducted contrasting individual distributions to those from the PI approach, with the latter superior if it gives more accurate estimation at the last seed than more than half of DM distributions.

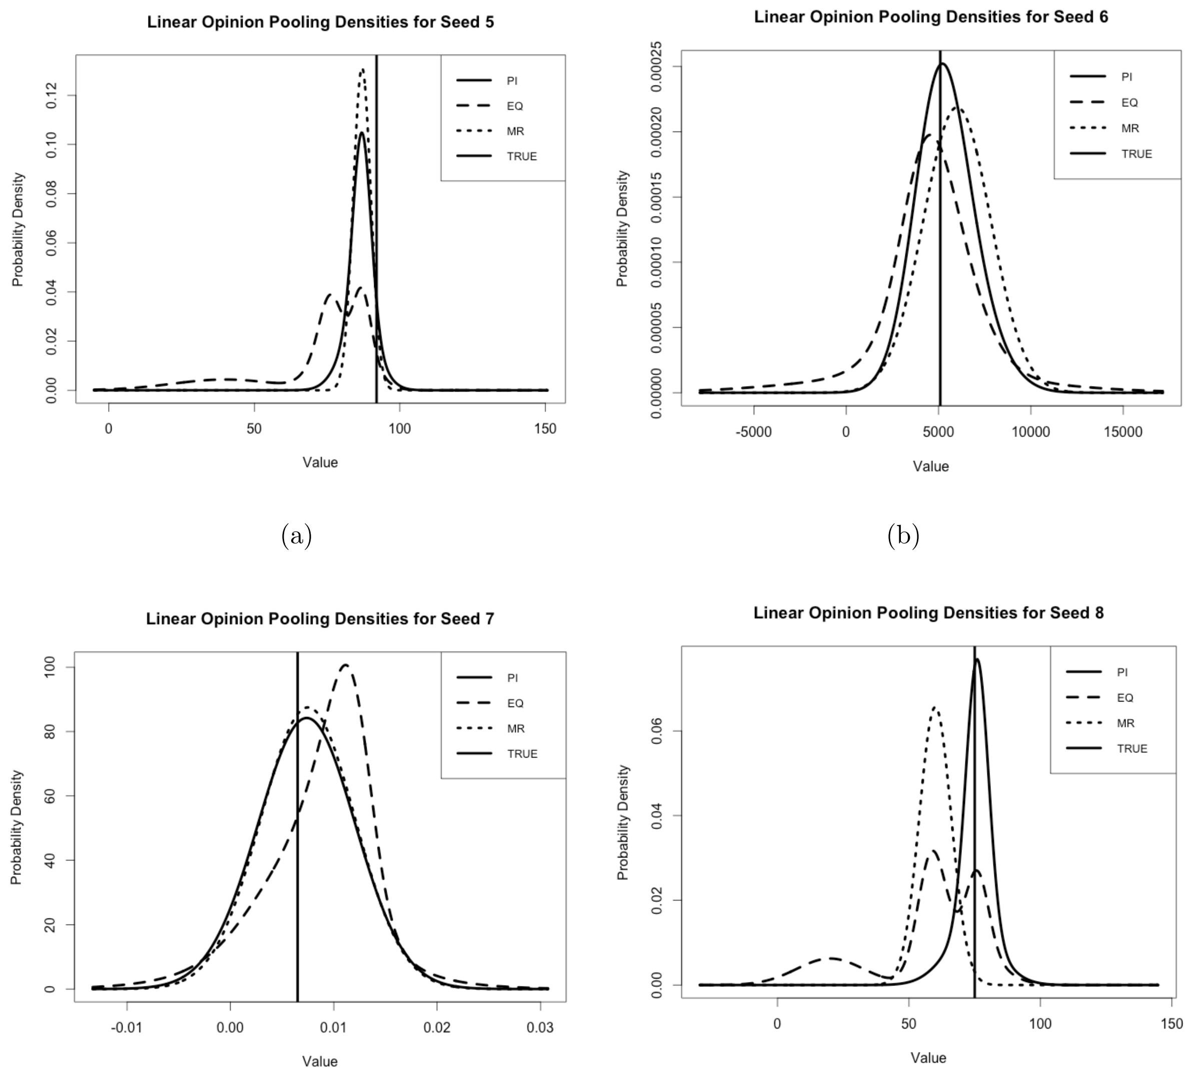

(a) (b) (c) (d)

Fig. 4.11: Distributions of the PI, EQ and MR methods for four A_SEED seed variables. The true parameter value is included (vertical line) in each case.

Fig. 4.11 shows distributions from group methods for four A_SEED seeds. The PI method gives clear better estimation for Seeds 6 and 8 and slightly outperforms the MR method for Seed 5. It is marginally outperformed by the MR method for Seed 7. Table 4.8 shows the (relatively constant) weights allocated to DMs over time, with P_4 and P_5 dominating. For Seeds 5, 6 and 8 listening to a weighted combination of these two DMs gives better estimation than listening to only one of them (even if the DM listened to is the more reliable). We produce similar plots for the individual problem in Fig. 4.12. For all seeds there is at most one DM whose posterior is more accurate

Table 4.8: DM weights across four seeds for a configuration of the A_SEED data rounded to three decimal places (*i.e.*, P_1, P_2 and P_3 have strictly positive weights).

P_i	$\alpha_{i,5}$	$\alpha_{i,6}$	$\alpha_{i,7}$	$\alpha_{i,8}$
P_1	0.000	0.000	0.001	0.000
P_2	0.000	0.000	0.000	0.000
P_3	0.000	0.000	0.000	0.000
P_4	0.405	0.311	0.306	0.166
P_5	0.595	0.689	0.693	0.834

than the PI posterior, *i.e.*, it is in the best interest of DMs to use it.

4.4.3 Results

An issue previously mentioned was the variance scaling parameter in Equation (4.19). Below we let $k = 50, k = 25$ and $k = 5$ to ensure robustness of results. We begin with the individual problem, with results in Table 4.9. We see results are mostly invariant to changes to k. This is highlighted in the aggregated figures in Table 4.11 with the proportion of data sets for which the PI approach is superior constant across values of k. The PI approach is superior to the alternative for the bulk of data sets and hence it is reasonable to declare it meritorious in the individual problem.

We include group problem success proportions in Table 4.10. The PI approach is not dominant in several cases, with the EQ method performing strongly. Aggregated results are in Table 4.12. Results seem relatively invariant to changes in k. The EQ method is successful across the greatest number of data sets. Yet, of the two performance-base schemes the PI method is strongest, outperforming the MR approach for all k considered. We see it is advantageous to use a subtler performance-based scheme. The EQ success is not surprising given our earlier discussion about the nature of seeds. For the data we consider seeds should be correlated and related to the unknown decision variable but this may not always be so, with some seeds potentially being irrelevant and/or expertise on one seed not implying it on another. In cases of this ilk the EQ approach performs strongly using "wisdom of crowds" logic. It seems intuitive that in a setting such as that which the PI approach is intended for, it will produce stronger results against the EQ method, as the PI approach is a technique based upon the

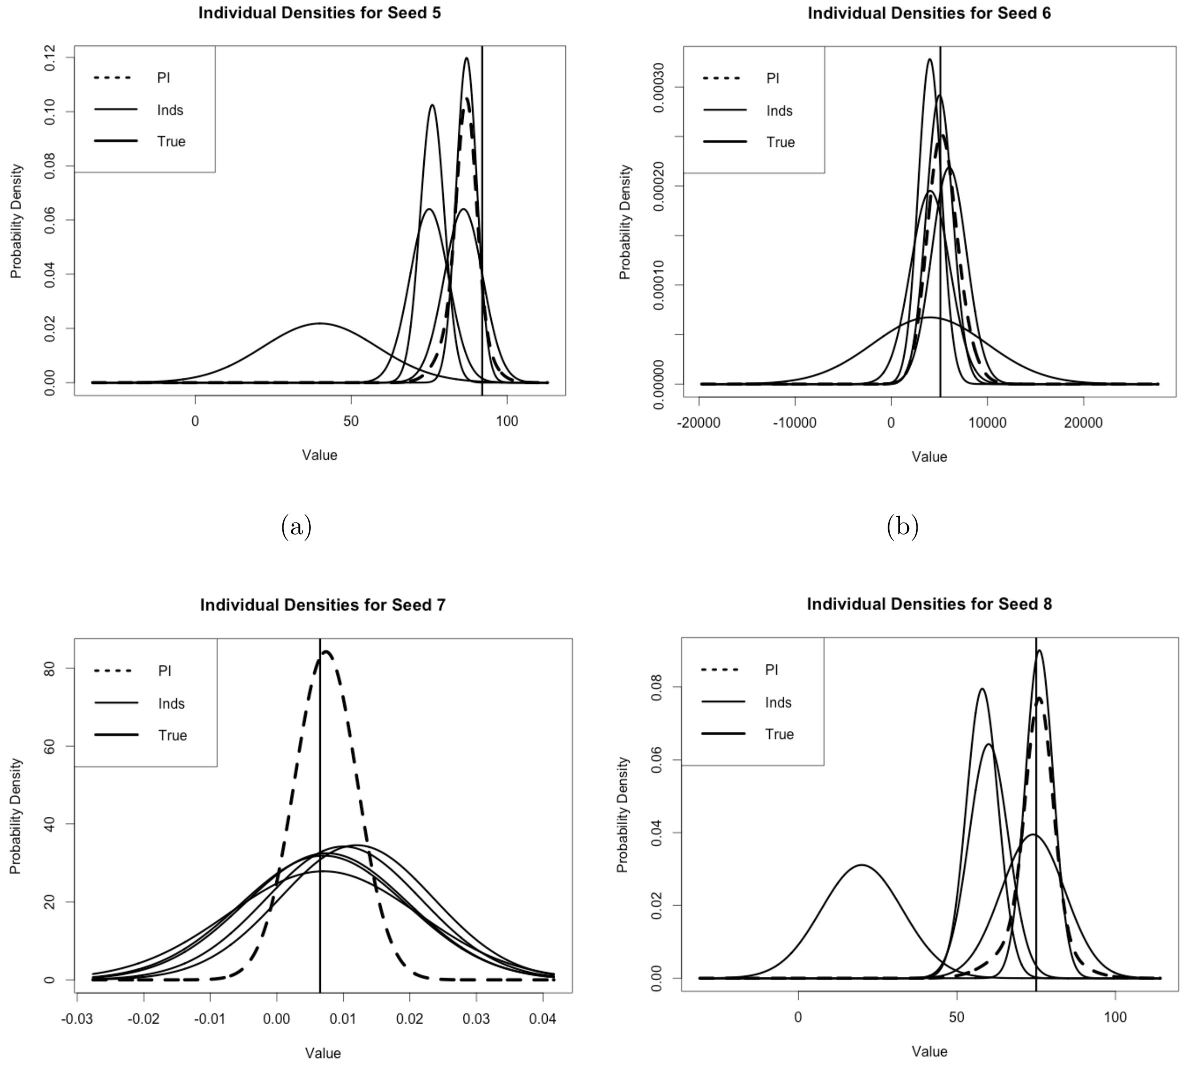

(a) (b) (c) (d)

Fig. 4.12: Distributions of DMs (unbroken curves) and the PI approach (broken curve) for four A_SEED seed variables with true parameter values (vertical lines) included.

principle of learning from information, of which very little (pertaining to seed values) is available here.

The EQ method is dominant for a large amount of datasets in the collection under examination. We have provided discussion above concerning why this may be the case, namely that correlation between seeds may be weak, implying knowledge of one seed does not imply knowledge of the next, and hence the PI method of learning may be bettered by a simple wisdom of crowds averaging. We briefly consider here if the data sets that the EQ approach is superior for share any type of common characteristic.

Table 4.9: Individual problem success proportion. Optimal methods are in bold.

	$k = 5$		$k = 25$		$k = 50$	
Data Set	IND.	P.I.	IND.	P.I.	IND.	P.I.
A_SEED	0.1250	**0.8750**	0.1250	**0.8750**	0.1250	**0.8750**
ACNEXPTS	**0.5000**	**0.5000**	**0.5000**	**0.5000**	**0.5000**	**0.5000**
ACTEP	0.2000	**0.8000**	0.2000	**0.8000**	0.2000	**0.8000**
AOTDAILY	0.4705	**0.5295**	0.4705	**0.5295**	0.4705	**0.5295**
AOTRISK	0.2727	**0.7272**	0.2727	**0.7272**	0.2727	**0.7272**
CARMA	0.2000	**0.8000**	0.2000	**0.8000**	0.2000	**0.8000**
DAMS	0.3636	**0.6363**	0.3636	**0.6363**	0.3636	**0.6363**
DIKRING	0.2340	**0.7659**	0.2304	**0.7659**	0.2340	**0.7659**
DSM1	0.1250	**0.8750**	0.1250	**0.8750**	0.1250	**0.8750**
DSM2	0.4000	**0.6000**	0.4000	**0.6000**	0.4000	**0.6000**
ESTEC1	0.3076	**0.6924**	0.3076	**0.6923**	0.3846	**0.6153**
ESTEC2	0.1153	**0.8846**	0.1153	**0.8846**	0.1153	**0.8846**
ESTEC3	0.2500	**0.7500**	0.2916	**0.7083**	0.2916	**0.7083**
EUNCRWD	0.0714	**0.9285**	0.1428	**0.8571**	0.1428	**0.8571**
EUNRCDD	**0.5714**	0.4285	**0.5000**	**0.5000**	**0.5000**	**0.5000**
EUNRCDIS	0.1739	**0.8260**	0.1956	**0.8043**	0.1521	**0.8478**
EUNCEAR	0.2666	**0.7334**	0.2666	**0.7333**	0.2667	**0.7333**
FCEP	**0.6250**	0.3750	**0.6250**	0.3750	**0.6250**	0.3750
GLINVSPC	0.3846	**0.6153**	0.3846	**0.6153**	0.3846	**0.6153**
GROND5	0.3000	**0.7000**	0.3000	**0.7000**	0.3000	**0.7000**
INFOSEC	0.4000	**0.6000**	0.4000	**0.6000**	0.4000	**0.6000**
LADDERS	0.1000	**0.9000**	0.1000	**0.9000**	0.1000	**0.9000**
MONT1	0.2500	**0.7500**	0.2500	**0.7500**	0.2500	**0.7500**
MVOSEEDS	0.0000	**1.000**	0.0000	**1.0000**	0.0000	**1.0000**
NH3EXPTS	**0.5000**	**0.5000**	**0.5000**	**0.5000**	**0.5000**	**0.5000**
OPRISK	0.2500	**0.7500**	0.2500	**0.7500**	0.2500	**0.7500**
PHAC	0.3571	**0.6428**	0.2857	**0.7142**	0.2857	**0.7142**
PILOTS	0.4000	**0.6000**	0.3000	**0.7000**	0.3000	**0.7000**
PM25	0.4167	**0.5833**	0.3750	**0.6250**	0.4583	**0.5416**
REALESTR	0.2580	**0.7149**	0.2580	**0.7419**	0.3225	**0.6777**
RETURN1	0.2000	**0.8000**	0.2667	**0.7333**	0.4000	**0.6000**
RIVERCHNL	0.2500	**0.7500**	0.2500	**0.7500**	0.2500	**0.7500**
SARS	**0.6000**	0.4000	**0.6000**	0.4000	**0.6000**	0.4000
SO3EXPTS	0.3888	**0.6111**	0.3888	**0.6111**	0.3888	**0.6111**
TEIDE_MAY	0.1111	**0.8889**	0.2000	**0.8000**	0.2000	**0.8000**
THRMBLD	0.3333	**0.6667**	0.2708	**0.7292**	0.3023	**0.6979**
TNO_DISP1	0.3055	**0.6945**	0.3055	**0.6944**	0.3055	**0.6945**
TUD_DEPOS	0.2619	**0.7381**	0.2619	**0.7380**	0.2619	**0.7381**
TUD_DISP	0.1944	**0.8055**	0.1944	**0.8055**	0.1944	**0.8056**
VESUVIO	0.4500	**0.5500**	**0.5000**	**0.5000**	**0.5000**	**0.5000**
VOLCRISK	0.2000	**0.8000**	0.2000	**0.8000**	0.2000	**0.8000**
WATERPOL	0.1111	**0.8889**	0.2222	**0.7778**	0.2222	**0.7778**

Table 4.10: Success proportions for the group problem. Optimal methods are in bold.

	$k = 5$			$k = 25$			$k = 50$		
Data Set	P.I.	EQ.	M.R.	P.I.	EQ.	M.R.	P.I.	EQ.	M.R.
A_SEED	**0.5714**	0.0000	0.4285	**0.5000**	0.0714	0.4285	**0.5000**	0.0714	0.4285
ACNEXPTS	0.2667	**0.5555**	0.1778	0.2667	**0.5556**	0.1778	0.2667	**0.5556**	0.1778
ACTEP	0.2778	**0.3778**	0.3333	0.1555	**0.4778**	0.3667	0.1555	**0.4778**	0.3667
AOTDAILY	0.2932	0.2772	**0.4295**	0.3226	0.2745	**0.4028**	0.3092	0.2905	**0.4001**
AOTRISK1	**0.4000**	0.2363	0.3636	0.3545	0.2636	**0.3818**	0.3545	0.2636	**0.3818**
CARMA	**0.5333**	0.2667	0.2000	**0.5333**	0.2667	0.2000	**0.5333**	0.2667	0.2000
DAMS	0.4454	**0.5000**	0.0545	0.4454	**0.5000**	0.0545	0.4454	**0.5091**	0.0454
DIKRING	**0.3931**	0.3098	0.2969	**0.4028**	0.3075	0.2895	**0.4195**	0.2978	0.2826
DSM1	**0.5714**	0.1785	0.2500	**0.6964**	0.0714	0.2321	**0.5357**	0.0892	0.3750
DSM2	0.1777	**0.4333**	0.3888	0.1667	**0.4333**	0.4000	0.1000	**0.5000**	0.4000
ESTEC1	0.2500	0.3525	**0.3974**	0.2948	0.3076	**0.3974**	0.2500	**0.3782**	0.3717
ESTEC2	0.2769	**0.4507**	0.2723	0.2800	**0.4507**	0.2669	0.2800	**0.4507**	0.2692
ESTEC3	0.3560	0.2727	**0.3712**	**0.3787**	0.2772	0.3484	0.31833	0.3333	**0.3484**
EUNCRWD	**0.3901**	0.2967	0.3131	**0.3462**	0.3076	**0.3462**	**0.3462**	0.3076	**0.3462**
EUNRCDD	0.0329	**0.5989**	0.3681	0.0989	**0.5549**	0.3462	0.1098	**0.5549**	0.3352
EUNRCDIS	0.2766	**0.4308**	0.2924	0.2608	**0.4527**	0.2865	0.3162	**0.4624**	0.2213
EUNCEAR	0.2333	0.3809	**0.3857**	0.2333	0.3809	**0.3857**	0.2284	0.3809	**0.3905**
FCEP	0.3214	0.2857	**0.3928**	0.1071	**0.4464**	**0.4464**	0.1428	**0.4464**	0.4107
GLINVSPC	0.2756	**0.4743**	0.2500	0.2756	**0.4744**	0.2500	0.2756	**0.4743**	0.2500
GROND5	0.2222	**0.4111**	0.3666	0.2555	**0.3888**	0.3555	0.2555	**0.3888**	0.3556
INFOSEC	0.0556	**0.6777**	0.2667	0.2111	**0.5333**	0.2556	0.2111	**0.5333**	0.2555
LADDERS	**0.5555**	0.0888	0.3555	**0.5888**	0.0888	0.3222	**0.6333**	0.0889	0.2777
MONT1	**0.4464**	0.1964	0.3572	**0.4464**	0.2143	0.3393	**0.4464**	0.2142	0.3392
MVOSEEDS	**0.5500**	0.3000	0.1500	**0.6000**	0.2000	0.2000	**0.6500**	0.2000	0.1500
NH3EXPTS	0.1333	**0.6000**	0.2667	0.3333	**0.4667**	0.2000	0.3667	**0.4667**	0.1667
OPRISK	0.3500	**0.4166**	0.2333	0.3538	**0.4208**	0.2208	0.3583	**0.4208**	0.2208
PHAC	0.2637	**0.5329**	0.2032	0.2580	**0.5220**	0.2200	0.2580	**0.5220**	0.2200
PILOTS	0.0778	**0.6667**	0.2555	0.2444	**0.5333**	0.2222	0.2333	**0.5111**	0.2556
PM25	0.1742	0.3787	**0.4469**	0.0984	**0.4772**	0.4242	0.0227	**0.5530**	0.3232
REALESTR	0.0935	**0.6344**	0.2720	0.1505	**0.5612**	0.2881	0.2419	**0.4667**	0.2914
RETURN1	0.3333	0.2619	**0.4047**	0.2809	0.3381	**0.3809**	0.2809	0.3381	**0.3809**
RIVERCHNL	**0.5000**	0.2142	0.2857	**0.4642**	0.2142	0.3214	**0.4642**	0.2142	0.3214
SARS	0.0888	**0.7222**	0.1889	0.0778	**0.7444**	0.1778	0.1778	**0.6555**	0.1667
SO3EXPTS	0.2638	0.3055	**0.4305**	0.2638	0.3055	**0.4305**	0.2638	0.3055	**0.4305**
TEIDE_MAY	**0.4000**	0.3888	0.2111	**0.4000**	0.3888	0.2111	**0.4000**	0.3888	0.2111
THRMBLD	0.2495	0.2721	**0.4782**	0.2859	0.3297	**0.3848**	0.2651	0.3631	**0.3718**
TNO_DISP1	0.2706	**0.3682**	0.3611	0.2706	**0.3682**	0.3611	0.2881	0.3507	**0.3611**
TNO_DEPOS	0.3928	0.1833	**0.4238**	0.3785	0.2000	**0.4214**	0.3785	0.2000	**0.4214**
TUD_DISP1	**0.4785**	0.2406	0.2793	**0.3785**	0.2421	0.2793	**0.4785**	0.2421	0.2793
VESUVIO	0.2333	**0.5222**	0.2444	0.3444	**0.4777**	0.1778	0.3555	**0.4778**	0.1667
VOLCRISK	0.3444	**0.4444**	0.2111	**0.4444**	0.3777	0.1778	**0.4444**	0.3778	0.1788
WATERPOL	0.2778	**0.5416**	0.1805	0.2778	**0.5277**	0.1944	0.2778	**0.5277**	0.1944

Table 4.11: Gross Individual Results: how many data sets methods are optimal for.

k	P.I.	Ind.
$k = 50$	**38**	4
$k = 25$	**38**	4
$k = 5$	**38**	4

Table 4.12: Gross Group Results: how many data sets method are optimal for.

Metric	P.I.	EQ.	M.R.
$k = 50$	11.5	**21.0**	9.5
$k = 25$	12.5	**20.5**	9.0
$k = 5$	12.0	**20.0**	10.0

Without information concerning the experts involved, or a subject-specific knowledge of the seeds that they were assessed upon, it is difficult to use the content of the datasets to draw inference from. In terms of intuition, it appears likely that if all experts involved are highly knowledgeable than the EQ method will perform strongly. In cases where there are a mix of accurate and inaccurate experts the PI approach is likely to be the superior method, as it will give priority to the knowledgeable experts, rather than listening to all equally.

We considered the data sets which each of the three methods were dominant for, and looked at the median of the number of experts involved for these. The median seemed a more sensible measure of central tendency than the mean for this study as, for instance, the 77 MVOSEEDS experts (by far the largest amount of experts for any of the data sets) would have a distorting impact on our conclusions, as means are liable to heavy influence by outliers. In the k=50 case the respective medians for the number of experts in the PI, EQ and MR dominated data sets were 10.5, 8 and 5, in the k=25 case they were 10, 8 and 5 and the k=5 case they were 8.5, 8.5 and 5. This aligns with our intuition, with the EQ and PI methods, which heed opinions from multiple sources (albeit in various ways), generally performing better for larger datasets. By contrast, the MR dominated datasets have by far the smallest median: the wisdom of crowds, weighted or otherwise, is less likely to apply for these small groups, and listening to a single individual may hence be the best course of action. Note that an analysis of the number of seeds involved for the various datasets was also conducted but produced no significant results.

4.4.4 Conclusions

Previously we justified the PI approach by showing it obeyed desirable Bayesian and coherency properties and illustrating its superiority on simulated data. Having done this a natural progression was to use real world validation data. Efforts were made to circumnavigate the differences between the classical and PI methods with probability distributions constructed from quantiles, and PI weights amended in light of the most recent realisation in a Markovian manner. Multiple permutations of original seed orderings have been considered, permissible due to the weak dependence of consecutive seeds and predictions, and that learning about one seed does not directly teach a DM about the next. These permutations ensure robustness of our analysis.

We considered individual and group problems. In the former, overwhelming evidence indicated that it was in the best interest of DMs to use the PI distribution over their own. For the group problem we considered the PI, EQ and MR methods. The PI method was dominant over the MR method but outperformed by the simple EQ method. The strong EQ performance can be partially explained by the fact that consecutive variables may be only weakly correlated, with an accurate prediction for one not necessarily implying accurate prediction on the next. Hence a simple arithmetic averaging may give better results than a performance-based approach due to "wisdom of crowds". As previously alluded to, a facet of research of this nature is that even the most well-reasoned approaches can fail to match this most straightforward combination rule. When considering the analysis contained within this section we must bear in mind the contrasts between the TU Delft data and data which would naturally occur in the PI approach. Our aim was to provide a stronger validation than simply one based upon simulated data, but in using the TU Delft data we have had to bridge a large amount of differences, and may perhaps be trying to "fit square pegs in round holes" to some degree - we are manipulating the raw data from quantiles to Normal distributions, permuting their order, and assuming accuracy on one seed pertains to accuracy on the next. Nevertheless, we do observe interesting results, which provide some additional validation for our PI approach. A closing comment is that there does not appear to be a single type of data set that the PI approach is dominant for, *e.g.*, sets with few experts and lots of variables, lots of experts and few variables, financial data sets, *etc.* We see that the PI approach can be of value in a wide range of potential settings and contexts, *i.e.*, is truly applicable in practice.

Chapter 5

Group Decision Making

In previous chapters we considered a set of n decision makers each with their own decision task, *i.e.*, n decisions were made per epoch. We now turn to group decision making as introduced in Section 2.6. In this environment a single decision is made per epoch that is reflective of the opinions and aspirations of all the DMs forming the cohort. On a small scale, a group of friends may need to decide which film they will collectively see in the cinema, while on a grander scale the United Nations may need to reach a resolution on whether to impose economic sanctions upon a country or not. Regardless of the magnitude of the problem there is a motivation for a mathematical method implementable to determine which decision is optimal for a group as a whole. Below we provide such a technique by combining both probabilities and utilities, and relate this to the terms of Arrow's Impossibility Theorem (Arrow, 1950).

5.1 Group Expected Utility

Previously we discussed how DMs make decisions by maximising expected utility, *i.e.*, calculating the expected utility associated with each possible alternative and determining the optimal decision as that with the largest expected utility, as in Equations (2.2)-(2.3). For an individual P_i this was a function of her probability distribution $f_i(\theta)$ and her utility function $u_i(r)$. In this section we present methods by which probability distributions $f_1(\theta), \ldots, f_n(\theta)$ can be combined to give a single probability distribution $\hat{f}(\theta)$, and utility functions $u_1(r), \ldots, u_n(r)$ can be combined to give a single utility function $u^*(r)$. These are respectively representative of the group consensus over θ

and the group preference over returns. Having done this the expected utility that the group assigns to a decision d_i is

$$\mathbb{E}[u^*(d_i)] = \int_\Theta u^*(d_i, \theta)\hat{f}(\theta)\ d\theta \tag{5.1}$$

The optimal group decision d^* is that maximising this, *i.e.*,

$$d^* = \arg\max_i \mathbb{E}[u^*(d_i)] \tag{5.2}$$

5.1.1 Combining Probabilities

We derived the PI approach for individual decision tasks but it is also suitable in a group setting. We slightly modify the combined belief of Equation (3.1), now writing

$$\hat{f}(\theta) = \alpha_1 f_1(\theta) + \ldots + \alpha_n f_n(\theta) \tag{5.3}$$

We originally denoted by $\alpha_{i,j}$ the weight P_i assigns to P_j. Here we replaced $\alpha_{i,j}$ by α_j, dropping the first subscript to reflect that this weight isn't assigned by a particular DM but by the group as a whole. Hence, in an extension of Equation (3.11), the normalised weight α_i^* afforded to P_i, given her previous normalised weight was α_i, and w_i is her most recent PI weight, is

$$\alpha_i^* = \frac{w_i \alpha_i}{\sum_{k=1}^{n} w_k \alpha_k} \tag{5.4}$$

Individual opinions are updated in the Bayesian manner from Section 3.1.2. The stringent objectivity of the PI approach makes it suitable for use in a group context. All DMs are *a priori* assumed to be equally reliable, with fluctuations from this directly proportional to the accuracy of the information that they provide. Hence no subjectivity enters the analysis, making it appropriate in a group process in which the sole aim should be to increase the decision quality of the group as whole with no concerns for the egos of the DMs involved. In Section 4.2 a simulation study assessed the merits of the PI approach. In a group decision context DMs do not have the choice to listen solely to their own belief and hence opinions must be combined in some fashion to reach a consensus. Given our discussions in Sections 2.7 and 4.2 we believe that linear opinion pooling is an adequate method of doing this. The EQ and MR methods are arguably the two most obvious alternatives to the PI approach. As we have seen in Section 4.2 the PI approach performed admirably against these (especially for a large number of

DMs and/or returns), making it suitable for use in this problem. All previous advocations of the PI approach are applicable in the group decision making context, with its objectivity an even bigger strength in this case.

5.1.2 Combining Utilities

Merging utility functions is a less intuitive task than amalgamating probability distributions. The utility function of a DM takes into account her attitude towards risks and gambles. It is unclear how this could be represented in a combined utility function. The combined probability distribution was a weighted sum with weights reflecting perceived reliability of information sources. If we construct the combined utility function as an additive linear combination of DMs utility functions then equal weight should be given to the utility function of each DM at all epochs to ensure fairness of the process. This indicates that preferences of all DMs are uniformly important, *i.e.*, the method is not biased towards one DM over another. Naively we could construct $u(r)$ as a simple arithmetic average of DMs utility functions, *i.e.*,

$$u(r) = \frac{u_1(r) + \ldots + u_n(r)}{n} \tag{5.5}$$

Yet this motivates DMs to exaggerate their utility functions to make the combined function more similar to their true preferences. Suppose $n = 2$ and P_2 knows $u_1(r) = 2r$, believing her own function to be $u_2(r) = 2r^2$. Thus, if she claims her utility function is $u_2(r) = 4r^2 - 2r$ then this ensures $u(r) = 2r^2$, *i.e.*, the combined function exactly modelling her personal preferences. In reality a DM may not know the exact utility functions of her neighbours when choosing her own, but it nevertheless presents a motivation for potential dishonesty in preference representation. Also, in general, utility functions that are more extreme will tend to dominate the combined utility function. We counteract these issues by rescaling all utility functions to $[0, 1]$, a valid act as utility is invariant to positive linear transformation. Hence each DM now assigns a utility of 1 to their most preferred outcome and a utility of 0 to their least preferred outcome, with this rescaled function denoted as $u_i^*(r)$ for P_i. We assume that there are at least two distinct outcomes, one that is at least as good as any other outcome (denoted r^*), and one that is at least as bad as any other outcome (denoted r_*) with $r^* \succ r_*$. If returns lie on the real line then this involves truncation to ensure finite values, *i.e.*, choosing the best and worst possible values as those that there are negligibly small probabilities of exceeding and being less than respectively. In a setting where initial fortunes are arguments of utility functions of DMs (*e.g.*, a financial context) these functions must be rescaled for all DMs such that they assign a utility of 0 to the worst possible fortune achievable by a group member, and a utility of 1 to the best possible fortune achievable by a group member. If each individual P_i has a fortune of γ_i and a utility function $u_i(\gamma_i + r)$ then rescaling must be conducted ensuring $u_i^*(\min\{\gamma_1, \ldots, \gamma_n\} + r_*) = 0$ and $u_i(\max\{\gamma_1, \ldots, \gamma_n\} + r^*) = 1$ for all $i = 1, \ldots, n$, as illustrated in Section 5.1.3. Generically, the combined utility function, $u^*(r)$, can be written

$$u^*(r) = \frac{u_1^*(r) + \ldots + u_n^*(r)}{n} \tag{5.6}$$

This approach guarantees commensurability (Boutilier, 2003, Houlding & Coolen, 2011), *i.e.*, it is possible to meaningfully compare the respective utility functions of two distinct DMs. Rescaling in this manner ensures absolute comparisons rather than relative comparisons, *i.e.*, we consider the utility that multiple DMs assign a fortune of \$20 rather than the utility that multiple DMs assign a loss of \$20 from their respective starting positions. A useful measure for classifying risk attitudes is ARA, defined in Equation (1.1). For the combined utility function in Equation (5.6) the ARA is given in Equation (5.7). The group ARA is not a simple linear function of the ARA coefficients of the DMs, but a more complex construct.

$$\begin{aligned}
ARA_*(r) &= -\frac{u''^*(r)}{u'^*(r)} \\
&= -\frac{\frac{d^2}{dr^2}[u^*(r)]}{\frac{d}{dr}[u^*(r)]} \\
&= -\frac{\frac{d^2}{dr^2}[\frac{1}{n}u_1^*(r) + \ldots + \frac{1}{n}u_n^*(r)]}{\frac{d}{dr}[\frac{1}{n}u_1^*(r) + \ldots + \frac{1}{n}u_n^*(r)]} \\
&= -\frac{\frac{1}{n}[u_1''^*(r) + \ldots + u_n''^*(r)]}{\frac{1}{n}[u_1'^*(r) + \ldots + u_n'^*(r)]} \\
&= -\sum_{i=1}^{n} \frac{u_i''^*(r)}{u_1'^*(r) + \ldots + u_n'^*(r)}
\end{aligned} \tag{5.7}$$

This approach is similar to Utilitarianism (Harsanyi, 1955) as we discuss later. This method can be seen as reasonable, giving equal weight to all DM preferences to ensure fairness in the process. The utility function of each DM is augmented over time in light of changes in positions (in terms of utils) after decision returns are received.

5.1.3 Example

We illustrate combining utilities for three DMs (detailed in Table 5.1), supposing that the best possible outcome in whatever gamble occurs is a gain of \$20 with the worst being a loss of \$20. As discussed above utilities are rescaled so each DM assigns a utility of 0 to a fortune of $\min\{\$100, \$40, \$80\} - \$20 = \$20$ and a utility of 1 to a fortune of $\max\{\$100, \$40, \$80\} + \$20 = \$120$, *i.e.*, we ensure that utilities are anchored above and below respectively by the best and worst possible state achievable by any group member. Utility functions are rescaled by solving simultaneous equations, *i.e.*, for P_2 we find a and b such that $a(20)^2 - b = 0$ and $a(120)^2 + b = 1$. Rescaled functions are given in Table 5.1 alongside their ARA. We see P_1 is risk-neutral, P_2 is risk-prone (for $r > -40$) and P_3 is risk-averse (for $r < -80$, with values less than this not defined).

Table 5.1: Utility functions of DMs, as well as initial fortunes and ARA coefficients.

P_i	γ_i	$u_i(\gamma_i + r)$	$u_i^*(\gamma_i + r)$	ARA_i
P_1	\$100	$r + 100$	$0.01(r + 100) - 0.2$	0
P_2	\$40	$(r + 40)^2$	$\frac{1}{14000}(r + 40)^2 - \frac{1}{35}$	$-\frac{1}{r+40}$
P_3	\$80	$\log_e(r + 80)$	$0.558 \log_e(r + 80) - 1.67$	$\frac{1}{r+80}$

From Equation (5.6) we find that the combined group utility function is

$$
\begin{aligned}
u^*(r) &= \frac{1}{3}u_1^*(\gamma_1 + r) + \frac{1}{3}u_2^*(\gamma_2 + r) + \frac{1}{3}u_3^*(\gamma_3 + r) \\
&= \frac{1}{3}\Big[0.01(r + 100) - 0.2\Big] + \frac{1}{3}\Big[\frac{1}{14000}(r + 40)^2 - \frac{1}{35}\Big] + \frac{1}{3}\Big[0.558 \log_e(r + 80) - 1.67\Big] \\
&= \frac{1}{300}(r + 100) + \frac{1}{42000}(r + 40)^2 + \frac{93}{500}\log_e(r + 80) - \frac{199}{315}
\end{aligned}
$$

5.1.4 Linear Identity

We defined in Equations (5.1) – (5.2) how a group decision can be made using $\hat{f}(\theta)$ and $u^*(r)$. Using linearity of expectation an attractive and intuitive identity arises:

$$
\begin{aligned}
\mathbb{E}[u^*(d_i)] &= \int_\Theta u^*(d_i, \theta)\hat{f}(\theta)\ d\theta \\
&= \int_\Theta \Big[\frac{1}{n}u_1^*(d_i, \theta) + \ldots + \frac{1}{n}u_n^*(d_i, \theta)\Big]\hat{f}(\theta)\ d\theta \\
&= \int_\Theta \Big[\frac{1}{n}u_1^*(d_i, \theta)\hat{f}(\theta) + \ldots + \frac{1}{n}u_n^*(d_i, \theta)\hat{f}(\theta)\Big]\ d\theta
\end{aligned}
$$

$$
\begin{aligned}
&= \int_\Theta \frac{1}{n} u_1^*(d_i,\theta)\hat{f}(\theta)\ d\theta + \ldots + \int_\Theta \frac{1}{n} u_n^*(d_i,\theta)\hat{f}(\theta)\ d\theta \\
&= \frac{1}{n}\int_\Theta u_1^*(d_i,\theta)\hat{f}(\theta)\ \ d\theta + \ldots + \frac{1}{n}\int_\Theta u_n^*(d_i,\theta)\hat{f}(\theta)\ d\theta \\
&= \frac{1}{n}\mathbb{E}[u_1^*(d_i)] + \ldots + \frac{1}{n}\mathbb{E}[u_n^*(d_i)] \qquad (5.8)
\end{aligned}
$$

The expected utility that the group assigns to a decision is an equally weighted sum of the expected utilities that DMs assign to this. This desirable property serves as partial justification for our method of combining utility functions in Equation (5.6). This identity holds regardless of the manner in which beliefs are combined so long as each DM uses a common distribution in their expected utility assessment, *i.e.*, that resulting from the PI approach, behavioural aggregation, or if each DM had an identical prior opinion. Note therefore that while in what follows we shall assume $\hat{f}(\theta)$ was found via the PI approach this is not the only possibility leading to Equation (5.8) holding. In an analogy of Equation (3.4) we write $\mathbb{E}[u_i^*(d_j)]$ as

$$\mathbb{E}[u_i^*(d_j)] = \int_\Theta u_i^*(d_j,\theta)\hat{f}(\theta)\ d\theta \qquad (5.9)$$

Hence the decision that P_i would make in an individual context is a function of her own utility function and a combined belief, as she is willing to incorporate the opinions of neighbours into her decision task (as rationalised in Section 3.1). Note that while the utility function used in Equation (5.9) is potentially a function of the initial fortunes of neighbours (as discussed in our commensurability argument) this does not impact upon the decision deemed optimal by P_i due to the aforementioned invariance of utility to positive linear transformation.

5.2 Arrow's Impossibility Theorem

Above we created a method by which DMs can combine opinions and utilities to determine an optimal decision for the group as a whole. In Section 2.6.1 we introduced the content of Arrow (1950) which considered the development of a Social Welfare Function (SWF) that would translate individual preference rankings to a group ranking. Five desirable axioms were outlined. Universality stated that a group ranking could be determined from any set of individual rankings. Monotonicity implied that if a decision rose in the preference rankings of an individual, with all other individual rankings remaining unchanged, then this decision would not decrease in the group ranking. Independence of irrelevant alternatives meant that removal of one decision from the ranking scheme would leave preference rankings between all other decisions unchanged. Non-imposition declared that for any conceivable group ranking there was a corresponding set of individual rankings that would lead to this. Finally non-dictatorship decreed that there could be no "dictator" inherent within the group, defined as follows: *"A social welfare function is said to be dictatorial if there exists an individual i such that for all x and y, $x \succ_i y$ implies $x \succ y$ regardless of the orderings of all individuals other than i"*, *i.e.*, that there must be no single individual whose preference ranking automatically becomes the group ranking. Arrow (1950) showed in his Impossibility Theorem that for at least two DMs and three decisions there was no SWF obeying all five axioms.

The symbols $\succ$ and $\succ_i$ denote strict group preference and the preference of individual i respectively, with $\succeq$ and $\succeq_i$, and $\sim_i$ and $\sim$, being respectively analogous for weak preference and indifference. Arrow's SWF deals with ordinal, rather than cardinal, rankings over decisions, *i.e.*, only the order in which decisions are ranked is of importance with the differing degrees of preference inherent within this disregarded. By contrast in our method from Section 5.1 expected utility values play a large part in determining which decision is optimal for the group. For an individual P_i, we relate our terminology to that of Arrow (1950), saying $d_1 \succ_i d_2$ if and only if $\mathbb{E}[u_i(d_1)] > \mathbb{E}[u_i(d_2)]$, $d_1 \succeq_i d_2$ if and only if $\mathbb{E}[u_i(d_1)] \geq \mathbb{E}[u_i(d_2)]$ and $d_1 \sim_i d_2$ if and only if $\mathbb{E}[u_i(d_1)] = \mathbb{E}[u_i(d_2)]$. Hence in our framework we may equivalently talk in terms of which decision is preferred and which decision has the highest expected utility. Arrow (1950) requires preference rankings of DMs to incorporate completeness and transitivity, *i.e.*, DMs must be able to provide a distinct ranking over each pair of decisions (strict, weak or equal preference) and this ranking must itself be coherent. Our procedure obeys both requirements. Regarding completeness, a DM can calculate expected utility for any possible decision using Equation (5.9) and hence can compare the two expected utility values to determine which, if either, she prefers. Transitivity is clear, as if $d_1 \succ_i d_2$ and $d_2 \succ_i d_3$ then $\mathbb{E}[u_i(d_1)] > \mathbb{E}[u_i(d_2)] > \mathbb{E}[u_i(d_3)]$, *i.e.*, $\mathbb{E}[u_i(d_1)] > \mathbb{E}[u_i(d_3)]$, so $d_1 \succ_i d_3$ as required.

5.3 Comparison and Consideration of Axioms

Clearly there are differences between the method in Section 5.1 and that of Arrow (1950), most fundamentally in the manner in which an individual determines her preference ranking. In Arrow (1950) this is done in a solitary fashion with no consideration given to the beliefs or preferences of any of the other DMs in the group. Once each DM has constructed her own opinion these are then amalgamated via the SWF, which returns a group preference ranking. By contrast in our method prior to constructing her own preference ranking each DM shares her opinion over θ with her neighbours in the hope of gaining an increased understanding of it. The motivation for this has been previously documented, *i.e.*, in the hope of increasing decision calibre. Having created their individual preference rankings using their own utility functions and the common shared belief $\hat{f}(\theta)$, these preferences are then synthesised into a single preference ranking using Equation (5.8). The expected utility assigned by the group to a decision is an equally weighted sum of the expected utilities designated by DMs to that decision. Once this has been calculated for all possible decisions these are then ranked based upon which is the largest. In the framework that we have developed there is access to more information than in the minimal setting of Arrow (1950) in which there is no mention of cardinal values. We assume that it is possible for all DMs in our setting to calculate these values and hence derive their individual (and indeed group) preference ranking using these. Finally we comment that Arrow (1950) is concerned with developing a full group preference ranking over all possible admissible decisions, *i.e.*, being able to construct a full ordinal ranking over all possibilities. In the framework that we are concerned with the primary interest is in finding which decision is deemed optimal for the group as a whole, *i.e.*, we are interested in which decision is best for the collective, and less so in which decision is ranked, for instance, fourth or fifth best. Given a full group preference ranking it is trivial to ascertain which decision is optimal for the group, *i.e.*, it is simply that preferred to all alternatives.

Due to the fundamental variations between our method and that of Arrow (1950) it is clear that the techniques in Section 5.1 are not directly pertinent to Arrow's Impossibility Theorem, and the converse. Nevertheless the axioms of Arrow (1950) are desirable coherency properties for a group decision process to obey regardless of how it is defined. We consider our method against each axiom in turn and discover which of these it obeys. In what follows we consider a setting of n DMs $P_1, \ldots, P_n$ who must make one of m possible decisions $d_1, \ldots, d_m$ with $n \geq 2$ and $m \geq 3$.

5.3.1 Universality

Each P_i calculates her expected utility for the m decisions using Equation (5.9), yielding $\mathbb{E}[u_i(d_1)], \ldots, \mathbb{E}[u_i(d_m)]$. Each DM compares these values using operators $\geq, >$ and $=$, and makes an equivalent preference statement using preference operators $\succeq_i, \succ_i$ and $\sim_i$, by the relationship discussed above. The expected utility assigned by the group as a whole to a decision is found by Equation (5.8), giving numeric value to each decision which can be ranked in a manner identical to that discussed above. Hence for each set of individual preference rankings our method generates a group ranking which reveals which decision is optimal for the group as a whole. Our method obeys universality.

5.3.2 Monotonicity

Suppose P_i augments the expected utility she affords to a decision d_j, increasing it so it rises up her individual preference ordering overtaking (without loss of generality) one decision which it now has a higher expected utility than. In doing so it changes from being the $(a+1)^{th}$ ranked decision for P_i to being the a^{th} ranked. No further changes are made to her preference ranking or to those of any other DMs in the group. Clearly in Equation (5.8) an increase in $\mathbb{E}[u_i^*(d_j)]$ corresponds to an increase in $\mathbb{E}[u^*(d_j)]$. As no change has been made to the group expected utility of any other decisions (as no other changes have been made to individual expected utilities) the group ranking of d_j will either remain the same or increase depending on the other expected utility values, but it never decreases. Our approach obeys monotonicity.

5.3.3 Independence of Irrelevant Alternatives

Suppose without loss of generality that the group preference ranking is determined as $d_1 \succ \ldots \succ d_{j-1} \succ d_j \succ d_{j+1} \succ \ldots \succ d_m$. By the identities discussed above this implies $\mathbb{E}[u^*(d_1)] > \ldots > \mathbb{E}[u^*(d_{j-1})] > \mathbb{E}[u^*(d_j)] > \mathbb{E}[u^*(d_{j+1})] > \ldots > \mathbb{E}[u^*(d_m)]$, with each group expected utility found via Equation (5.8). There are two distinct cases now, depending on which decision d_j is removed from the set of possible decisions:

- **Case 1:** Suppose removing d_j does not impact upon r^* or r_*, *i.e.*, the best and worst possible outcomes amongst the decisions. Utility functions of DMs will remain unchanged, and hence the expected utilities assigned by each DM to the remaining $m-1$ decisions will be unchanged also (due to the independence of irrelevant alternatives axiom by von Neumann and Morgenstern, 1944), ensuring that the group expected utilities for the remaining decisions remain the same. This leads to the original group preference ordering, but with d_j omitted and all other preferences remaining unchanged, *i.e.*, $d_1 \succ \ldots \succ d_{j-1} \succ d_{j+1} \succ \ldots d_m$.

- **Case 2:** Suppose removing d_j impacts upon r^* or r_*. Utility functions of DMs will be rescaled (in a setting where initial fortunes are considered) to account for the new best/worst case scenarios achievable. In this instance the ordinal ranking of a DM over the remaining decisions will remain unchanged, but the cardinal difference in utility she associates with these will change. Hence it is possible (although by no means guaranteed) that the group preference ranking may be different from that ultimately occurring from Case 1.

Hence our method does not obey independence of irrelevant alternatives (at least in problems where initial fortunes are incorporated into utility functions), as there will always be (at most two) decisions whose removal can potentially lead to changes in the group ranking.

5.3.4 Non-imposition

It seems straightforward that when using our method every possible combination of group preference ranking is constructible from some combination of individual preference rankings; in fact there are multiple combinations of individual preference rankings (an infinite amount when we consider cardinal utility) leading to these group rankings. A trivial illustration with two DMs and three decisions showing that no social ranking can be imposed is the case where the group preference ranking is identical to the common preference rankings of all DMs, *e.g.*, $d_1 \succ d_2 \succ d_3$ will be constructed if $d_1 \succ_i d_2 \succ_i d_3$ for $i = 1, 2$ and so on. Our method obeys non-imposition.

5.3.5 Non-dictatorship

In the parlance of Arrow (1950) a DM was deemed a dictator if, having witnessed the preferences of the other DMs in her group, she could augment her original preferences in such a fashion that the group preference would mirror her original preferences. Firstly we consider the following example which demonstrates that instances do exist in which our approach may fall prey to a dictator. Suppose that we have two DMs and three possible decisions. The true preferences of P_1 and P_2 are given by $d_1 \succ_1 d_2 \succ_1 d_3$ and $d_3 \succ_2 d_2 \succ_2 d_1$ respectively. The corresponding expected utility values of DMs and the group are given in Table 5.2 with P^* denoting the group as a whole. From Equation (5.8) we see $d_3 \succ d_1 \succ d_2$, which is not the (full) preference ranking of either DM.

Table 5.2: Original expected utilities of the DMs and the group.

	P_1	P_2	P^*
$\mathbb{E}[u_i^*(d_1)]$	0.9	0.1	0.5
$\mathbb{E}[u_i^*(d_2)]$	0.3	0.5	0.4
$\mathbb{E}[u_i^*(d_3)]$	0.2	1	0.6

Now P_1 can be a dictator if there exist $x, y, z \in [0, 1]$ such that

$$\frac{x + 0.1}{2} > \frac{y + 0.5}{2} > \frac{z + 1}{2} \longrightarrow x + 0.1 > y + 0.5 > z + 1$$

This is satisfied, for instance, by $x = 1, y = 0.55$ and $z = 0$ leading to the cardinal preferences in Table 5.3, *i.e.*, $d_1 \succ d_2 \succ d_3$, the true preference of P_1. Hence she is capable by Arrow's definition of being a dictator.

Table 5.3: The augmented expected utilities of P_1 in light of those of P_2, and the new expected utilities of the group.

	P_1	P_2	P^*
$\mathbb{E}[u_i^*(d_1)]$	1	0.1	0.55
$\mathbb{E}[u_i^*(d_2)]$	0.55	0.5	0.525
$\mathbb{E}[u_i^*(d_3)]$	0	1	0.5

Similarly P_2 can be a dictator if there exist $a, b, c \in [0, 1]$ such that

$$\frac{a + 0.2}{2} > \frac{b + 0.3}{2} > \frac{c + 0.9}{2} \longrightarrow a + 0.2 > b + 0.3 > c + 0.9$$

This is satisfied, for instance, by $a = 1, b = 0.7$ and $c = 0$ leading to the cardinal preferences in Table 5.4, *i.e.*, $d_3 \succ d_2 \succ d_1$, the true preference of P_2. Hence she too has the potential to be a dictator. It is not possible for both DMs to simultaneously be dictators. Firstly this would clearly be a contradiction. Secondly in order to behave as a dictator (as defined above) a DM must be cognisant of the preference rankings of all of her neighbours (*i.e.*, for them to be fixed) that she will then alter her own preferences in light of. This is not the case if two (or more) DMs are dictators, as they would need to both be aware of each other's preferences that would constantly be being updated, *i.e.*, a circular argument.

Table 5.4: The augmented expected utilities of P_2 in light of those of P_1, and the new expected utilities of the group.

	P_1	P_2	P^*
$\mathbb{E}[u_i^*(d_1)]$	0.9	0	0.45
$\mathbb{E}[u_i^*(d_2)]$	0.3	0.7	0.5
$\mathbb{E}[u_i^*(d_3)]$	0.2	1	0.6

We have shown that cases do exist where our method is susceptible to a dictator. Now consider the following example, having three DMs and four potential decisions. The respective preferences are $d_1 \succ_1 d_2 \succ_1 d_3 \succ_1 d_4$, $d_3 \succ_2 d_2 \succ_2 d_4 \succ_2 d_1$ and $d_4 \succ_3 d_2 \succ_3 d_1 \succ_3 d_3$. The corresponding expected utility values are in Table 5.5 leading to a group preference ranking of $d_2 \succ d_1 \sim d_3 \sim d_4$.

For P_1 to be a dictator there must exist $x, y, z, t \in [0, 1]$ such that

$$\frac{x + 0.3}{3} > \frac{y + 1.8}{3} > \frac{z + 1.1}{3} > \frac{t + 1.2}{3} \longrightarrow x + 0.3 > y + 1.8 > z + 1.1 > t + 1.2$$

There is no set of values $\{x, y, z, t\}$ satisfying this inequality due to the magnitude of the expected utility for d_2. Irrespective of what augmentations P_1 makes to her own ranking she cannot prevent d_2 from being the decision that is deemed optimal for the group. This is contrary to the decision (d_1) that she herself deems personally

Table 5.5: Original expected utilities of the DMs and the group.

	P_1	P_2	P_3	P^*
$\mathbb{E}[u_i^*(d_1)]$	1	0.1	0.2	0.325
$\mathbb{E}[u_i^*(d_2)]$	0.9	0.9	0.9	0.675
$\mathbb{E}[u_i^*(d_3)]$	0.2	1	0.1	0.325
$\mathbb{E}[u_i^*(d_4)]$	0.1	0.2	1	0.325

preferable. Hence P_1 cannot be a dictator, nor, from the symmetry inherent within the problem, can P_2 or P_3. Due to the strong cardinal utility values placed on d_2 there is no change that any DM can make to her own preference ranking to prevent it being optimal for the group. Therefore in this instance there is nobody capable of being a dictator. We see that there is not always guaranteed to be a DM who can be a dictator. Our method obeys non-dictatorship.

5.4 Discussion

In summary, the method in Section 5.1 is guaranteed to obey universality, monotonicity and non-imposition. The condition of independence of irrelevant alternatives holds if the decision removed does not change the best or worst case outcome achievable from the decision process, but cannot be said to be true in complete generality. If a group consists of three or more DMs there may be a dictator, but this will often not be the case as discussed more below. Despite the contrasting setting of our approach and that considered in Arrow (1950) it was an eventuality that our technique would not be capable of adhering to all of its terms. Nevertheless it seems an admirable facet of our methodology that it performs reasonably well with regard to the axioms of Arrow, in that it can breach at most two of them, and even these breaches are not guaranteed to occur (assuming there are more than two DMs in the group).

Of the five axioms discussed above perhaps the most interesting result pertains to non-dictatorship. We have shown that it is possible to construct a group decision method that will not always contain a DM who is a dictator, *i.e.*, while there are occasions when there may be a DM who can behave dictatorially this is not an inevitable eventuality. Two highly relevant theorems are the previously mentioned Gibbard-Satterthwaite Theorem (Gibbard, 1973, Satterthwaite, 1975) and the Duggan-Schwartz Theorem (1992), both of which are concerned with manipulability, *i.e.*, the potential for a group decision process to be manipulated by a DM within the group. These results show that there is no process meeting some desirable criteria (similar to those of Arrow) that is not potentially susceptible to manipulability, *i.e.*, the ability of DMs to influence group preference by giving a modified version of their true preferences is prevalent in any reasonable decision scheme. However just because manipulability is technically possible does not imply it can always be exacted in practice. This is evident in our decision making scheme from Section 5.1. We have provided an example where a DM can alter her preferences to make the group preference ranking identical to her own (*i.e.*, manipulability). However this is not always the case as our counter-example illustrates. In our setting the preferences of all DMs are given an equal weight in Equation (5.8). It seems clear that as the number of DMs in the group increases the propensity for the existence of a dictator decreases. In a group consisting of only two DMs both are guaranteed to have the ability to be a dictator (formally proved in Appendix B) but for groups with three or more DMs this will not always be the case. Of course such large groups may fall victim to a dictator depending on the expected utility values assigned by DMs to the decisions involved, but it is not guaranteed. It is always worth investigating if a problem is susceptible to a dictator, but in problems with a large amount of DMs this is decreasingly likely. A simple way to combat dictatorship is to not allow DMs to change their preferences once they have initially been put forward, or equivalently to say that DMs are not permitted to witness the preferences of others until they have finalised their own. We have discussed above why it is in the best interest of DMs to share opinions about θ, yet there is no analogous gain to be found from sharing utility functions prior to the decision task.

Our method derived cardinal rankings in addition to ordinal rankings. Utilitarianism (Harsanyi, 1955) dealt with cardinal utilities also, stating that the optimal decision for a group is that maximising group expected utility. The utility function of the group (SWF in the terms of Arrow) is an equally weighted sum of utility functions of individuals, obeying anonymity and strong Pareto as discussed in Section 2.6.2. For motivational reasons akin to those mentioned in Section 5.1.2, Harsanyi (1955) restricted utilities to the unit interval. He demonstrated that the only choice of function obeying his criteria was an equally weighted sum of individual functions. Our method can be seen as an extension of Utilitarianism that ensures commensurability and incorporates additional probabilistic information. We believe consideration of cardinal utility values, rather than simple ordinal preference rankings, leads to a more intuitive, subtler, decision making process. Consider the following simple motivational example. Five housemates must decide whether to watch Planet of the Apes (PA) or The Shawskank Redemption (SS), with four of the five preferring PA to SS. Under a scheme based on ordinal ranking (*e.g.*, majority rule) the group decision would be to watch PA. However, consider the expected utilities assigned by DMs in Table 5.6.

Table 5.6: The expected utilities of the DMs and the group.

	P_1	P_2	P_3	P_4	P_5	P^*
$\mathbb{E}[u_i^*(\text{PA})]$	0.2	0.22	0.17	0.21	0	0.16
$\mathbb{E}[u_i^*(\text{SS})]$	0.18	0.2	0.16	0.2	1	0.348

The first four DMs prefer PA to SS but there is only a fractional difference between their preference for the former over the latter. The respective satisfactions that they will derive from watching either movie are almost indistinguishable. By contrast P_5 has a strong preference for the latter over the former. Despite the apparent preference for PA over SS (in terms of majority rule) the group expected utility for SS is over twice that for PA. Note it is imperative P_5 does not know the preferences of her neighbours prior to stating her own, or else she may manipulate these to ensure a group preference for SS. This is an exaggerated illustration but the point is evident nonetheless. Ordinal rankings do not necessarily tell the whole story of preference as they do not incorporate the degree of preferences involved. We believe this is a fundamental flaw in ranking schemes/group decision processes of this nature, and claim that it is a desired property of any such scheme/process to maximise collective expected utility. This is a property that the method in Section 5.1 obeys. In conclusion, in this chapter we have created a group decision making method based upon maximising expected utility (which can be seen as a rational course of action in its own right) that performs well in terms of the axioms of Arrow (1950). This seems strong evidence in favour of our technique.

Chapter 6

Nonparametric Extension

In Chapter 3 we introduced a method enabling DMs to combine their opinions with those of their neighbours. This approach permitted DMs to update both their own opinions over time and the reliability measures that they assigned to neighbours. Implicit within this approach was the assumption that DMs can express their opinions over θ via fully parameterised probability distributions. This assumption is common (French, 2011), in part due to the availability of elicitation techniques (*e.g.*, O'Hagan 1998). In Chapter 2 we discussed the concept of imprecise probability which is a formulation permitting greater uncertainty (in terms of belief specification) for users. In this chapter we modify our previously derived technique to a nonparametric framework. We consider a very basic method of belief specification and provide a "distribution-free" analogy to the PI approach.

6.1 Belief Specification

In the setting that we consider a DM represents her predictions via a lower and upper bound over the expected value of θ, *e.g.*, between 0.3 and 0.7 for the event of heads in a coin toss. A DM with this belief expects to see, for instance, between 30 and 70 heads in 100 tosses. No values within this range are explicitly deemed more/less likely than any others. Subtly this does not imply that all values are assumed equally likely. The expected value is believed to lie somewhere in this interval but no further statements of consequence can be made. We consider expectation, rather than probability, as our primitive quantity, in contrast to the standard imprecise framework where DMs specify lower and upper bound probabilities of θ lying in a particular range (*e.g.*, Coolen *et al.*, 2010). This increased simplicity leads to a trade-off in terms of a potential decrease in the quality and interpretability of the corresponding results, *i.e.*, often a single decision cannot be declared unanimously optimal. In what follows we frequently use the synonym of expectation, prevision (de Finetti, 1974), which is apt in our context as the opinion of a DM represents the lower and upper bound expected values she judges for θ prior to witnessing data. We create a framework where DMs can linearly combine their prevision bounds, with weights representing reliability, and bounds/weights updated in light of returns observed, *i.e.*, a nonparametric generalisation of the PI approach.

6.2 Nonparametric Utility Inference

A common assumption is that DMs can assign precise utility values to any possible return. This may be unreasonable especially when a DM may receive a return that is novel to her. If this is so how can she be expected to know *a priori* the satisfaction that she will derive from its occurrence? Motivation exists for updating utilities in light of phenomena experienced, *i.e.*, adaptive utility (Cyert & DeGroot, 1975, Houlding & Coolen 2007, 2011). A specialisation of this is Nonparametric Utility Inference (NPUI, Houlding & Coolen, 2012) which contains elements of Nonparametric Inference (Coolen, 2006, Roelofs *et al.*, 2011) and uncertain utility. NPUI relies upon Hill's $A_{(n)}$ assumption (Hill, 1968, 1988, 1993) which states that if the real line is partitioned into distinct, disjoint subintervals, the probability of θ lying within any of these intervals is equal regardless of their respective widths. This models extreme vagueness.

In NPUI utilities are rescaled to the unit interval. A DM is interested in the utility that she will assign to a novel outcome, and has experienced n exchangeable outcomes that (after observation) she believes are similar to this, *e.g.*, a DM who has been to concerts by n distinct rock bands (that she has enjoyed to varying extents) and wants to predict the satisfaction she would derive from attending a concert by a rock band she has never seen before. These known outcomes have utilities labeled $u_{(1)}, u_{(2)}, \ldots, u_{(n)}$ with $0 < u_{(i)} \leq u_{(j)} < 1$ for $i < j$. Modifying Hill's $A_{(n)}$ to the unit interval, the utility for the previously unrealised outcome obeys the following, for $i = 1, \ldots, n-1$:

$$\mathbb{P}(U_{\text{new}} \in [u_{(i)}, u_{(i+1)}]) \;=\; \mathbb{P}(U_{\text{new}} \in (0, u_{(1)}]) = \mathbb{P}(U_{\text{new}} \in [u_{(n)}, 1)) = \frac{1}{n+1} \qquad (6.1)$$

The probability of the utility of the novel return lying in any of the $n+1$ intervals is equal regardless of their widths. From this, lower and upper bounds for the expectation of the unknown utility can be constructed, given respectively by:

$$\underline{\mathbb{E}}[U_{\text{new}}] = \frac{1}{n+1}\sum_{i=1}^{n} u_{(i)} \tag{6.2}$$

$$\bar{\mathbb{E}}[U_{\text{new}}] = \frac{1}{n+1}\Big(1+\sum_{i=1}^{n} u_{(i)}\Big) \tag{6.3}$$

A consequence of this is that the width of the predictive interval will be

$$\Delta(\mathbb{E}(U_{\text{new}})) = \frac{1}{n+1} \tag{6.4}$$

Learning occurs over time with new observations incorporated into Equations (6.2) and (6.3). As n increases the interval width decreases, *i.e.*, the more experiences a DM has had the more confident she will be in making predictions for utility values of previously unrealised similar outcomes.

6.3 Adapting NPUI to Prevision Bounds

Conceptually our goal and that of NPUI are distinct. However as utilities are restricted to the unit interval there is a strong similarity to our topic of interest as the expectation of θ must also lie in this range (by definition of expectation and the constraints upon θ discussed below). In both cases we wish to provide bounds for an unknown quantity in $[0,1]$. We present a method by which nonparametric prevision intervals (NPPI) may be updated over time given observations witnessed following decisions. We begin with a single DM before generalising to a multiple DM setting. Below we consider the restrictive case where θ lies in the unit interval to ensure similarity with NPUI, discussing a real line extension in Section 6.5. Hence θ can be viewed in two potential lights: as a latent parameter value lying in $[0,1]$ or as a Bernoulli trial success probability. In the latter case the data witnessed will be ones and zeros, corresponding to successes and failures respectively. If $X \sim \text{Ber}(p)$ then $\mathbb{E}[X] = p$ which is attractive given our predilection towards expectations. Optionally we can consider (independent and identically distributed) Bernoulli aggregated data, *i.e.*, output which is a realisation of a Binomially distributed quantity. In this instance the witnessed value used in our updating rules is the ratio of the number of successes to the number of trials.

This is less preferable as it counts multiple observations as a single observation and hence will not be as informative (*i.e.*, gives no information on the number of trials witnessed and hence the relative importance of respective pieces of data). Witnessing four successes in five trials is less telling than witnessing four million successes in five million trials, yet using Binomially aggregated data leads to these observations being treated as equal.

6.3.1 Single DM

Prior to the first epoch a DM supplies an interval $[l, u]$, with $l, u \in [0, 1]$ and $l \leq u$, that she believe $\mathbb{E}(\theta)$ lies within, *i.e.*, her initial lower and upper bound values are

$$\underline{\mathbb{E}}_0(\theta) = l \tag{6.5}$$

$$\bar{\mathbb{E}}_0(\theta) = u \tag{6.6}$$

The width of this interval, Δ_0, is a measure of her uncertainty over θ. Narrow intervals indicate confidence in her prediction while wide ones imply a lack of knowledge. There are two extreme cases. When $l = 0$ and $u = 1$ the DM knows nothing about $\mathbb{E}(\theta)$ and can only trivially state that it lies in $[0, 1]$. However learning occurs over time and this interval will become narrower (more informative) once outcomes are witnessed. By contrast when $l = u$ a DM provides a point estimate. We shortly see that this is unreasonable in our context of interest. The DM makes an initial decision, using her bounds in Equations (6.5) and (6.6), and the methods in Section 2.4.1. How can a DM who has witnessed some outcome(s) update her NPPI given this new information?

In Equation (6.4) interval width was $\frac{1}{n+1}$ where n was the number of outcomes realised. In NPUI if a DM had seen no relevant utility values then her interval had a width of one. If not then her inferences were influenced by the pertinent utility values witnessed. The more values she has observed the slimmer her interval is. Modifying the approach from Equations (6.1)-(6.3) to solve our problem we consider which observations (theoretical or real) would a DM have had to witness to form her prior bounds. Thin intervals are equivalent to having witnessed a lot of data. We can view wide intervals as analogous to uninformative priors in a Bayesian setting, *e.g.*, a DM with a Beta$(1, 1)$ prior considers herself to have seen one success in two hypothetical trials while a DM with a Beta$(100, 100)$ prior has the same mean prediction but it is based on two hundred hypothetical trials. In our nonparametric setting the latter DM would have a far narrower NPPI than the former DM due to her increased confidence (*i.e.*, lower variance). Below, rather than use the symbol n (which we have thus far used to denote the number of DMs) we use its Greek equivalent ν. To find the number of hypothetical observations she would have had to witness to construct her prior NPPI a DM solves

$$\bar{\mathbb{E}}_0(\theta) - \underline{\mathbb{E}}_0(\theta) = \Delta_0 = \frac{1}{\nu + 1} \tag{6.7}$$

This reveals $\nu = \frac{1-\Delta_0}{\Delta_0}$ so given the interval $[l, u]$ she can consider the $\frac{1-\Delta_0}{\Delta_0}$ hypothetical observations she would have had to see to create this (under this inferential model). Issues arise if $\Delta_0 = 0$, *i.e.*, division by zero. In a parametric approach this is equivalent to having seen an infinite amount of data. Hill's $A_{(n)}$ assumption and its variants are post-data assumptions related to finite exchangeability, in contrast to the framework developed by de Finetti (1974) that allows for infinite exchangeability. Our inferential focus is in a setting where only a finite amount of values can be observed. We assume $\Delta_0 > 0$ implying that a DM has some vagueness about $\mathbb{E}(\theta)$. Once a DM finds ν she may progress in a manner analogous to NPUI by augmenting prevision bounds after each new observation and increasing ν by 1 for each new piece of data witnessed.

Updating when $\nu \in \mathbb{N}_0$

If a DM solves Equation (6.7) and finds $\nu \in \mathbb{N}_0$ then her NPPI is consistent with having seen a non-negative integer amount of hypothetical observations. She considers which ν observations would lead to this (under this nonparametric model). If the i^{th} observation is x_i then $\{x_1, \ldots, x_\nu\}$ must satisfy the constraints from Equations (6.2)-(6.3):

$$\underline{\mathbb{E}}_0(\theta) = \frac{1}{\nu + 1} \sum_{i=1}^{\nu} x_i \tag{6.8}$$

$$\bar{\mathbb{E}}_0(\theta) = \frac{1}{\nu + 1} \Big(1 + \sum_{i=1}^{\nu} x_i\Big) \tag{6.9}$$

We find that the solutions are $\{x_1, \ldots, x_\nu\}$ such that

$$\sum_{i=1}^{\nu} x_i = l(\nu + 1) = u(\nu + 1) - 1 \tag{6.10}$$

She may choose any values of $x_1, \ldots, x_\nu$ in $[0, 1]$ meeting this constraint as only their sum is used in updating. Once a new set of returns $\{x_{\nu+1}, \ldots, x_{\nu+m}\}$ (with $m = 1$ for

Binomial aggregated data or θ being latent in $[0,1]$, and $m > 1$ if not) is seen after the decision at the first epoch her NPPI bounds become

$$\underline{\mathbb{E}}_1(\theta) = \frac{1}{\nu+m+1}\sum_{i=1}^{\nu+m} x_i \tag{6.11}$$

$$\bar{\mathbb{E}}_1(\theta) = \frac{1}{\nu+m+1}\Big(1+\sum_{i=1}^{\nu+m} x_i\Big) \tag{6.12}$$

Updating when $\nu \notin \mathbb{N}_0$

Our process loses intuitive appeal if $\nu \notin \mathbb{N}_0$ but we can implement it in a similar, albeit somewhat contrived, manner. The value ν is decomposed into $\nu = \lfloor \nu \rfloor + r$ with $\lfloor \nu \rfloor$ the integer part of ν and r the remainder, *i.e.*, $0 < r < 1$. She may now proceed in a manner analogous to that above, *i.e.*, finding the $\lceil \nu \rceil$ values $x_1, \ldots, x_{\lceil \nu \rceil}$ satisfying

$$\underline{\mathbb{E}}_0(\theta) = \frac{1}{\nu+1}\Big(rx_1+\sum_{i=2}^{\lceil \nu \rceil} x_i\Big) \tag{6.13}$$

$$\bar{\mathbb{E}}_0(\theta) = \frac{1}{\nu+1}\Big(1+rx_1+\sum_{i=2}^{\lceil \nu \rceil} x_i\Big) \tag{6.14}$$

We consider that the DM has seen $\lceil \nu \rceil$ values (*i.e.*, v rounded up to the nearest integer), the first of which is rescaled by a factor of r. Any one of the $\lceil \nu \rceil$ values could be chosen to be rescaled, with x_1 used for convenience. This case is less aesthetic than that when $\nu \in \mathbb{N}_0$, but it is still possible to proceed as before with the next m observations denoted $\{x_{\lceil \nu \rceil+1}, \ldots, x_{\lceil \nu \rceil+m}\}$ and a new interval being calculated of width $\frac{1}{\nu+m+1}$.

If a DM finds $\nu \notin \mathbb{N}_0$ she may wish to re-evaluate her NPPI so that it corresponds with the intuitively preferable concept of witnessing an integer number of hypothetical observations. Manipulating Equation (6.10) shows that if a DM supplies an initial lower bound l then her interval will correspond to one formed having witnessed an integer (k) amount of observations if $u = \frac{1}{k+1} + l$, with $k \in \mathbb{N}_+$, *e.g.*, for an initial lower bound $l = 0.15$ choices of k equaling 1, 2 and 3 gives u values of 0.65, 0.483 and 0.4 respectively. Equivalently if she supplies an initial upper bound of u then her interval will correspond to one formed having witnessed an integer (k) amount of observations if $l = u - \frac{1}{k+1}$, for $k \in \mathbb{N}_+$. We include a set of widths and the corresponding (integer) numbers of hypothetical observations that would have had to be observed to lead to these in Table 6.1. Given this, a DM may wish to change her initial NPPI to ease interpretability.

Table 6.1: Values of Δ_0 and corresponding numbers of hypothetical observations, ν.

Δ_0	1	0.5	0.25	0.2	0.1	0.05	0.02	0.01
ν	0	1	3	4	9	19	49	99

Example

A DM provides $[\underline{\mathbb{E}}_0(\theta), \bar{\mathbb{E}}_0(\theta)] = [0.55, 0.75]$ giving $\Delta_0 = 0.2$, *i.e.*, $\nu = 4$ solves Equation (6.7). Using Equation (6.10) she learns hypothetical observations obey $\sum_{i=1}^{4} x_i = 2.75$, *i.e.*, $x_1 = x_2 = 1, x_3 = 0.75$ and $x_4 = 0$ is one suitable observation set. At the first epoch she sees four successes in five trials. These values are exchangeable so without loss of generality $x_5 = x_6 = x_7 = x_8 = 1$ and $x_9 = 0$. From Equations (6.11)-(6.12) she finds

$$\begin{aligned} \underline{\mathbb{E}}_1(\theta) &= \frac{1}{9+1}\sum_{i=1}^{9} x_i = 0.675 \\ \bar{\mathbb{E}}_1(\theta) &= \frac{1}{9+1}\left(1 + \sum_{i=1}^{9} x_i\right) = 0.775 \end{aligned}$$

This interval has width $\Delta_1 = 0.1$, *i.e.*, it has become narrower (and shifted upwards) in light of the data witnessed. We illustrate this in Fig 6.1.

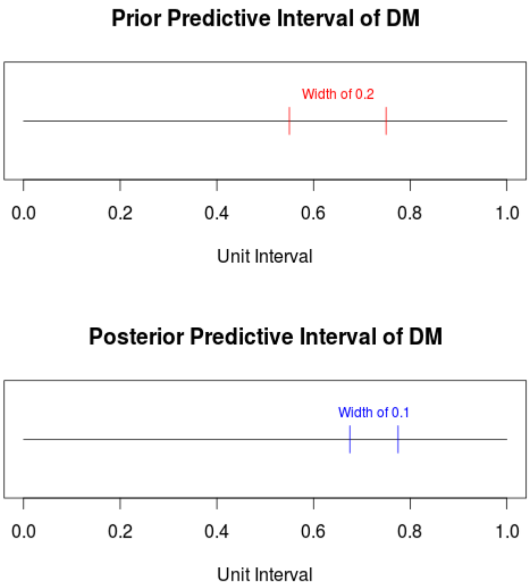

Fig. 6.1: The prior NPPI (top) and posterior NPPI (bottom) in our example.

6.3.2 Multiple DMs

We extend the above theory to a setting of n DMs with prior NPPIs of $[l_1, u_1], \ldots, [l_n, u_n]$ respectively. All are interested in θ. Below we assume $\nu \in \mathbb{N}_0$ for ease. Notation from Section 6.3.1 extends naturally. $\underline{\mathbb{E}}_{t,i}(\theta)$ is the lower NPPI of P_i after the t^{th} epoch, *i.e.*, after t decisions have been made and t sets of data witnessed. This is determined both by the ν_i hypothetical data points $\{x_{i,1}, \ldots, x_{i,\nu_i}\}$ forming her prior, and the subsequent data $\{x_1, \ldots, x_m\}$ seen from t epochs. For Binomially aggregated data $m = t$ as there are as many observations as there are epochs (*i.e.*, one at each) while for Bernoulli data there are more observations than epochs, *i.e.*, $t > m$. We write

$$\underline{\mathbb{E}}_{t,i}(\theta) = \frac{1}{\nu_i + m + 1}\Big(\sum_{j=1}^{\nu_i} x_{i,j} + \sum_{k=1}^{m} x_k\Big) \tag{6.15}$$

Similarly $\bar{\mathbb{E}}_{t,i}(\theta)$ is the upper NPPI bound of P_i after t epochs:

$$\bar{\mathbb{E}}_{t,i}(\theta) = \frac{1}{\nu_i + m + 1}\Big(1 + \sum_{j=1}^{\nu_i} x_{i,j} + \sum_{k=1}^{m} x_k\Big) \tag{6.16}$$

The width of the NPPI of P_i after t epochs is:

$$\Delta_{t,i}(\theta) = \bar{\mathbb{E}}_{t,i}(\theta) - \underline{\mathbb{E}}_{t,i}(\theta) = \frac{1}{\nu_i + m + 1} \tag{6.17}$$

We construct the combined bounds of P_i, having received opinions of neighbours, as weighted sums of the respective lower and upper bounds of her and her neighbours:

$$\hat{\underline{\mathbb{E}}}_{t,i}(\theta) = \sum_{j=1}^{n} \alpha_{i,j} \underline{\mathbb{E}}_{t,j}(\theta) \tag{6.18}$$

$$\hat{\bar{\mathbb{E}}}_{t,j}(\theta) = \sum_{j=1}^{n} \alpha_{i,j} \bar{\mathbb{E}}_{t,i}(\theta) \tag{6.19}$$

Initially all DMs are assumed to be equally reliable, *i.e.*, for $i = 1, \ldots, n$ we have

$$\hat{\underline{\mathbb{E}}}_{0,i}(\theta) = \frac{1}{n}\sum_{j=1}^{n} \underline{\mathbb{E}}_{0,j}(\theta) \tag{6.20}$$

$$\hat{\bar{\mathbb{E}}}_{0,i}(\theta) = \frac{1}{n}\sum_{j=1}^{n} \bar{\mathbb{E}}_{0,j}(\theta) \tag{6.21}$$

Once a return is witnessed and information learned about θ the DMs update equal weights to those reflecting perceived reliability. In Chapter 3 we presented a method by which this was done when DMs supplied opinions via probability distributions. Below we derive a similar method for use in the NPPI case.

6.4 Calculating Weights

It appears that there are three primary elements for assessing predictive ability of a DM. Firstly, how often have their bounds contained corresponding realisations? DMs whose NPPI have contained several realisations appear more accurate than those whose intervals contained few. Secondly, if previously witnessed values were not in the NPPI then how far from the observations have their bounds been? Low values are a sign of more accurate DMs. Thirdly, what was the width of their last NPPI? This is the subtlest of the three measures, requiring contextual interpretation.

In determining the reliability of a neighbour the most obvious fact that a DM will wish to find out is how accurate their last prediction was, as well as how accurate previous predictions have been. We introduce $S_{t,j}$, defined for P_j after t decision epochs as:

$$S_{t,j} \quad = \quad \sum_{m=1}^{t} \mathbb{1}_{x_m \in [\underline{\mathbb{E}}_{m-1,j}(\theta), \bar{\mathbb{E}}_{m-1,j}(\theta)]} \tag{6.22}$$

$S_{t,j}$ is a sum of indicator variables, defined for each m as

$$\mathbb{1}_{x_m \in [\underline{\mathbb{E}}_{m-1,j}(\theta), \bar{\mathbb{E}}_{m-1,j}(\theta)]} = \begin{cases} 1 & \text{if } x_m \in [\underline{\mathbb{E}}_{m-1,j}(\theta), \bar{\mathbb{E}}_{m-1,j}(\theta)]; \\ 0 & \text{if } x_m \notin [\underline{\mathbb{E}}_{m-1,j}(\theta), \bar{\mathbb{E}}_{m-1,j}(\theta)]. \end{cases}$$

Equation (6.22) is the running total of the amount of times that the intervals of P_j have contained corresponding realisations. Bernoulli data is aggregated here, *i.e.*, if $[\underline{\mathbb{E}}_{0,j}(\theta), \bar{\mathbb{E}}_{0,j}(\theta)] = [0.4, 0.6]$ and P_j sees five successes in ten trials then $x_1 = 0.5$ giving $\mathbb{1}_{x_1 \in [\underline{\mathbb{E}}_{0,j}(\theta), \bar{\mathbb{E}}_{0,j}(\theta)]} = 1$. It may seem counter intuitive that $S_{t,j}$ is independent of the amount of trials witnessed (*e.g.*, here $S_{t,j}$=1 regardless of if one success is seen in two trials or five are seen in ten) but the amount of data witnessed is incorporated into the process via the updated NPPI previously derived. The NPPI of a DM will be sharper after five successes in ten trials than one success in two. If a single Bernoulli value is seen per epoch then the NPPI of a DM will rarely contain this. Hence it is the measures of distance and width given below which shall be used in weight updating.

Consider a DM whose interval did not contain the last realisation. The further her NPPI is from this value the less reliable she appears, *e.g.*, if $[\underline{\mathbb{E}}_{0,i}(\theta), \bar{\mathbb{E}}_{0,i}(\theta)] = [0, 0.2]$, $[\underline{\mathbb{E}}_{0,j}(\theta), \bar{\mathbb{E}}_{0,j}(\theta)] = [0.7, 1]$ and a value of 0.3 is witnessed then P_i will appear more reliable than P_j. We introduce a measure of the distance from the NPPI of P_j to the witnessed value at the m^{th} epoch, denoting this distance as $\delta_{j,m}$:

$$\delta_{j,m} = \begin{cases} \min\left\{|x_m - \underline{\mathbb{E}}_{m-1,j}(\theta)|, |x_m - \bar{\mathbb{E}}_{m-1,j}(\theta)|\right\} & \text{if } x_m \notin [\underline{\mathbb{E}}_{m-1,j}(\theta), \bar{\mathbb{E}}_{m-1,j}(\theta)]; \\ 0 & \text{if } x_m \in [\underline{\mathbb{E}}_{m-1,j}(\theta), \bar{\mathbb{E}}_{m-1,j}(\theta)]. \end{cases}$$

This gives the distance from the witnessed value to the closest bound (*i.e.*, either lower or upper) with $\delta_{j,m} = 0$ for a DM whose interval contained the observation. These values are summed over time to give a cumulative measure, δ_j, of the inaccuracy of P_j:

$$\delta_j = \sum_{m=1}^{t} \delta_{j,m} \tag{6.23}$$

The width measure is a more complex construct. Suppose that two DMs have intervals $[\underline{\mathbb{E}}_{0,i}(\theta), \bar{\mathbb{E}}_{0,i}(\theta)] = [0.15, 0.85]$ and $[\underline{\mathbb{E}}_{0,j}(\theta), \bar{\mathbb{E}}_{0,j}(\theta)] = [0.35, 0.6]$ and a value of 0.5 is seen. It seems logical to concur that the latter DM is the more accurate as, while both opinions can be viewed as correct, she has more confidence in her prediction. For accurate DMs narrow intervals are commendable. However, by contrast suppose $[\underline{\mathbb{E}}_{0,i}(\theta), \bar{\mathbb{E}}_{0,i}(\theta)] = [0, 0.1]$ and $[\underline{\mathbb{E}}_{0,j}(\theta), \bar{\mathbb{E}}_{0,j}(\theta)] = [0.7, 1]$, and a value of 0.5 is seen. P_i has the narrower belief but is this something to reward? It implies that she not only holds an inaccurate opinion but is also extremely convinced of its truth. It can be argued in this case that the DM with the wider interval is more reliable, *i.e.*, for inaccurate DMs narrowness should not be heralded as a strictly favourable attribute as it was for accurate DMs. In a Bayesian setting weak prior distributions reflect a lack of knowledge and conviction. When data is witnessed these may often be dominated by the likelihood function, *i.e.*, new data has a large effect *a posteriori*, outweighing the weak prior distribution. Therefore for an inaccurate DM a wide NPPI is not necessarily a negative quality as it implies her opinions are susceptible to change given new information. By contrast consider an extremely inaccurate DM with a very narrow interval. This corresponds to strong conviction in her opinion, *i.e.*, is akin to her basing her prior on a large amount of hypothetical realisations. This belief will be slow to augment towards the truth even given a lot of new information. The inaccurate DM with wide beliefs will learn from her mistakes faster than the inaccurate DM with narrow beliefs. We express this opinion via two axioms, given tabularly in Table 6.2.

- **Axiom 1:** If two DMs have NPPIs containing the witnessed value we penalise the wider interval more.

- **Axiom 2:** If two DMs have NPPIs not containing the witnessed value we penalise the DM with the tighter interval more.

Table 6.2: Merits of different combinations of width and accuracy.

	Accurate Belief	Inaccurate Belief
Narrow Interval	Best Case Scenario	Wost Case Scenario
Wide Interval	2^{nd} Best Case Scenario	3^{rd} Best Case Scenario

Gneiting and Raftery (2007) discuss proper scoring rules for interval predictions. If l and u are the $\frac{\alpha}{2}$ and $1-\frac{\alpha}{2}$ quantiles of an opinion, and x is the data witnessed, then they recommend the negatively oriented rule:

$$S(x;l,u) = (l-u) + \frac{2}{\alpha}(l-x)\mathbb{1}_{(x<l)} + \frac{2}{\alpha}(x-u)\mathbb{1}_{(x>u)} \tag{6.24}$$

In our parlance, for P_j after epoch t this is equivalent to

$$S_j(x;l_j,u_j) = \Delta_{t,j} + \delta_{t,j} \tag{6.25}$$

In contrast to Gneiting and Raftery (2007) we consider a setting in which learning occurs over time. Suppose we witness a return of 1. Equation (6.25) gives the intervals $[0,1]$ and $[0,0.01]$ the same weight. Yet the latter NPPI is radically inaccurate and will be very slow to augment to the truth even if a large amount of realisations are witnessed, while the former although initially trivially vague will quickly converge to the truth. We augment Equation (6.25) to entail the complexity of width:

$$\begin{aligned} S_j(x;l_j,u_j) &= \delta_j + (t-S_{t,j})(1-\Delta_{t,j}) + S_{t,j}\Delta_{t,j} \\ &= \delta_j + t - t\Delta_j + S_{t,j}\Delta_j - S_{t,j} + S_{t,j}\Delta_j \\ &= \delta_j + t - S_{t,j} + \Delta_j(2S_{t,j}-t) \end{aligned} \tag{6.26}$$

The first term penalises past inaccuracy. The second term, $t-S_{t,j}$, is a penalty for DMs who have not contained witnessed values in their NPPI, being minimised if they have been accurate at all t epochs (*i.e.*, $t=S_{t,j}$) and maximised if they have been inaccurate at all epochs (*i.e.*, $S_{t,j}=0$). Consider the final term in this expression, $\Delta_j(2S_{t,j}-t)$. We have $\Delta_j(2S_{t,j}-t)>0$ if $S_{t,j}>\frac{t}{2}$ and $\Delta_j(2S_{t,j}-t)<0$ if $S_{t,j}<\frac{t}{2}$. A DM is penalised for her width if she is mostly accurate (*i.e.*, more than half the time) and rewarded for her width if she is mostly inaccurate (in line with our axioms). If $S_j(x; l_j, u_j)$ is small (*i.e.*, P_j is deemed accurate) then $\Big(S_j(x; l_j, u_j)\Big)^{-1}$ will be large, and the converse. Hence we propose the normalised weight assigned by P_i to P_j is

$$\begin{aligned} \alpha_{i,j} &= \frac{u_{i,j}}{\sum_{k=1}^{n} u_{i,k}} \\ &= \frac{\Big(S_j(x; l_j, u_j)\Big)^{-1}}{\sum_{k=1}^{n} \Big(S_k(x; l_k, u_k)\Big)^{-1}} \end{aligned} \tag{6.27}$$

We initially considered a scheme under which the unnormalised weight assigned to P_j was the positively oriented scoring rule $2t - S_j(x; l_j, u_j)$, with the leading $2t$ chosen to ensure strict positivity of weights. However we found that this approach led to highly inaccurate individuals still receiving high weights as the leading $2t$ had a very large impact; weights seldom tended towards the theoretical lower limit of zero even for DMs with very inaccurate opinions. Returning to our current methodology, we see $\alpha_{i,j}$ is the same for all $i = 1, 2, \ldots, n$ by the objectivity inherent within the process, as was the case for the fully probabilistic PI approach. Under the PI approach the weights at one epoch were directly a function of the weights at the previous epoch, which is not precisely the case here. Nevertheless we can observe that $S_{t,j}$ and δ_j are evolving measures (as cumulative sums across epochs) and $\Delta_{t,j}$ is strictly decreasing, so there is a strong Markovian element to the weights over time, *i.e.*, the future and past weights are independent of each other given the present weight.

6.5 Extension to the Real Line

Above we restricted $\mathbb{E}(\theta)$ to $[0, 1]$. Now we generalise our method so that it can lie anywhere on the real line. We may apply Hill's $A_{(n)}$ assumption to the real line but issues arise with finding analogies to the NPPI updating rules of Equations (6.15)-(6.16). In the second of these rules we added 1 to the sum of the data, yet this only makes sense in a $[0, 1]$ context with any choice of number to add in its place being arbitrary. Straightforward extension is problematic. An alternative, rather than translating our method to the real line, is to translate the real line problem to $[0, 1]$, allowing us to apply the above approach. Translations from the whole real line are not bijective but if we truncate these (as in Coolen, 1996) we may apply our method. We must assume that there is a finite lowest/highest possible parameter value, mapped to 0 and 1 respectively. Intermediate values can now be translated.

Truncated Real Line Example

Suppose $\mathbb{E}(\theta)$ is limited to $[-500, 1000]$ and a DM expects it to fall in $[100, 400]$. We denote by x the original values and by x^* those rescaled to lie in $[0, 1]$. We translate the smallest possible value to $x^* = 0$ and the largest to $x^* = 1$, solving simultaneous equations $a(-500) + b = 0$ and $a(1000) + b = 1$ to give $a = \frac{1}{1500}$ and $b = \frac{1}{3}$, *i.e.*,

$$x^* = \frac{x}{1500} + \frac{1}{3} \text{ for all } x \in [-500, 1000] \tag{6.28}$$

$[\underline{\mathbb{E}}_0(\theta), \bar{\mathbb{E}}_0(\theta)] = [100, 400]$ is mapped to $[\underline{\mathbb{E}}_0(\theta^*), \bar{\mathbb{E}}_0(\theta^*)] = [0.4.0.6]$. This is equivalent to having seen four observations with $\sum_{i=1}^{4} x_i^* = 2$. One possible return set obeying this is $x_1^* = 0.2, x_2^* = 0.4, x_3^* = 0.6$, and $x_4^* = 0.8$. Inverting Equation (6.28) gives

$$x = 1500x^* - 500 \tag{6.29}$$

Hence $x_1^*, x_2^*, x_3^*, x_4^*$ are translated to $x_1 = -200, x_2 = 100, x_3 = 400$ and $x_4 = 700$. Suppose a new observation $x_5 = 350$ is seen, mapped to $x_5^* = 0.567$. Her NPPI becomes

$$\begin{aligned} \underline{\mathbb{E}}_1(\theta^*) = \frac{1}{6}\Big(\sum_{i=1}^{5} x_i^*\Big) &= 0.4278 \\ \bar{\mathbb{E}}_1(\theta^*) = \frac{1}{6}\Big(1 + \sum_{i=1}^{5} x_i^*\Big) &= 0.5945 \end{aligned}$$

Equation (6.29) gives $[\underline{\mathbb{E}}_1(\theta), \bar{\mathbb{E}}_1(\theta)] = [141.7, 391.75]$. Changing opinions are shown in Fig. 6.2. Below we continue in $[0, 1]$, knowing we can always generalise this if needed.

6.6 Example

Consider five doctors, $P_1, \ldots, P_5$. A new drug has come on the market as a potential cure for a disease. It has only been subject to limited testing, *i.e.*, there is uncertainty about its efficacy. If it works then the patient will be cured, if not they will suffer negative side-effects. Each doctor has a set of patients who place different utilities on the possible outcomes, *i.e.*, utilities come from their respective patients (Table 6.3). The potential decisions are d_1 (take the drug, risking side-effects) or d_2 (not taking

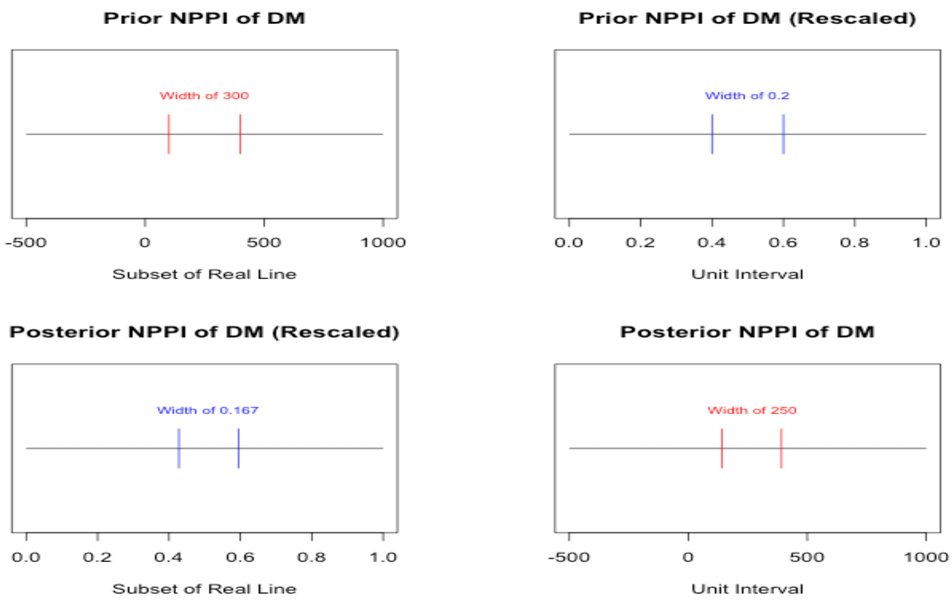

Fig. 6.2: Prior and posterior NPPIs in terms of the real line and unit interval.

the drug). We denote by θ_1 the event of the drug working and by θ_2 the complement. For each patient the utility of d_2 is independent of whether the drug works or not and incorporates their attitude over risks. In Table 6.3 we present the NPPI bounds of doctors over drug efficacy θ, with $\mathbb{P}(\theta_1) = \theta$. Equal weights are given at the first decision epoch resulting in bounds of $[\hat{\underline{\mathbb{E}}}_{0,j}(\theta), \hat{\bar{\mathbb{E}}}_{0,j}(\theta)] = [0.462, 0.822]$ for each DM.

Table 6.3: Initial previsions and utilities, as well as the decisions deemed optimal.

P_j	$[\underline{\mathbb{E}}_{0,j}(\theta), \bar{\mathbb{E}}_{0,j}(\theta)]$	$\Delta_{0,j}$	$u_j(d_1, \theta_1)$	$u_j(d_1, \theta_2)$	$u_j(d_2)$	d^*
P_1	$[0.6, 0.8]$	0.2	1	0	0.85	d_2
P_2	$[0, 1]$	1	1	0	0.3	d_1
P_3	$[0.45, 0.55]$	0.1	1	0	0.4	d_1
P_4	$[0.76, 0.96]$	0.2	1	0	0.2	d_1
P_5	$[0.5, 0.8]$	0.3	1	0	0.35	d_1

Patients must combine their utility over potential outcomes with the corresponding expectations associated with these. They use the combined NPPI of $[0.462, 0.822]$ to do this. Here decisions are monotonic in θ (assuming a DM wishes to use maximising expected utility as her decision making criteria) so she needs only to calculate the expected utility of each decision under the lower and upper bounds. If a decision has the highest expected value in both scenarios then it is optimal for her, as clearly it will be so for all intermediate values. These optimal decisions are in Table 6.3 with only the patient of P_1 opting to not try the drug. Of the four who try the treatment three are cured while one is not and suffers side-effects. Each doctor updates their previsions about the efficacy of the drug, and the reliability of their neighbours, given this Bernoulli data. Out of four trials three successes have been witnessed so the observed sample proportion for the drug working is 0.75, giving the weights in Table 6.4. These seem rational and coherent with the ordinal ranking we would choose. Despite being accurate P_2 is penalised for her wide opinion/vague belief. P_1 and P_5 are deemed the most accurate, with the former achieving a higher weight due to her narrower interval.

Table 6.4: Weights having witnessed three successes out of four trials.

P_j	$S_{1,j}$	$\Delta_{1,j}$	$\delta_{j,1}=\delta_j$	$u_{i,j}$	$\alpha_{i,j}$
P_1	1	0.2	0	5	0.436
P_2	1	1	0	1	0.087
P_3	0	0.1	0.2	0.91	0.079
P_4	0	0.2	0.01	1.24	0.108
P_5	1	0.3	0	3.33	0.290

Doctors update their prevision bounds over θ in light of the new data. Recall, as previously discussed, that they must find the values of the hypothetical observations which would have had to be witnessed to create their initial NPPI, and use Equations (6.15) and (6.16) to update prevision bounds. Table 6.5 shows this information. After only seeing four observations the opinion of P_2 has shifted from triviality to being highly informative, *i.e.*, it has adapted dramatically to the data witnessed. Combining modified opinions with new weights gives a combined prevision of $[\underline{\hat{\mathbb{E}}}_{1,i}(\theta), \hat{\overline{\mathbb{E}}}_{1,i}(\theta)] = [0.648, 0.770]$.

The utilities of the next patients of the doctors are in Table 6.6. All try the treatment. Four are cured while one is not (suffering side-effects). Weights and prevision bounds are updated in light of this, given in Table 6.7. P_2 is deemed the most reliable, having been accurate at both epochs (the only DM to be so). P_3, the only DM to have

Table 6.5: Updated previsions of doctors having witnessed data at first epoch.

P_i	n_i	$\sum_{j=1}^{n_i} x_{i,j}$	$\sum_{j=1}^{4} x_j$	$[\underline{\mathbb{E}}_{1,i}(\theta), \bar{\mathbb{E}}_{1,i}(\theta)]$
P_1	4	3	3	$[0.66, 0.77]$
P_2	0	0	3	$[0.6, 0.8]$
P_3	9	4.5	3	$[0.537, 0.607]$
P_4	4	3.8	3	$[0.755, 0.866]$
P_5	$\frac{7}{3}$	$\frac{7}{3}$ (with $x_1 = 1$)	3	$[0.636, 0.772]$

been inaccurate both times, has the smallest weight. The combined NPPI is updated to $[\hat{\underline{\mathbb{E}}}_{2,i}(\theta), \hat{\bar{\mathbb{E}}}_{2,i}(\theta)] = [0.706, 0.790]$ for all DMs. Fig. 6.3 highlights the change in weights over time. With the possible exception of P_3 it appears that opinions of doctors are starting to converge towards a common value. Fig. 6.4 shows the prevision bounds of individuals over the three epochs. These are decreasing in width over time as each DM witnesses more data. This procedure can be repeated indefinitely.

Table 6.6: Utilities at the second epoch, as well as the decisions deemed optimal.

P_j	$u_j(d_1, \theta_1)$	$u_j(d_1, \theta_2)$	$u_j(d_2)$	d^*
P_1	1	0	0.2	d_1
P_2	1	0	0.3	d_1
P_3	1	0	0.4	d_1
P_4	1	0	0.5	d_1
P_5	1	0	0.45	d_1

This method is applicable in a wide range of instances. A set of retailers can estimate the probability that the goods that they buy are of resalable quality to maximise profit. In system safety DMs may have contrasting thoughts on the probability of a device failing, with utility measured in monetary cost and in terms of safety. Leaving $[0, 1]$, DMs may be concerned with stock market behaviour, in the predicted traffic along a motorway at a particular time (where their aim is minimising congestion), or the number of students applying to a college. There is a vast array of problems which NPPI is equipped to solve, similar to those the PI approach is suitable for but with a

Table 6.7: Weights after four successes out of five trials (*i.e.*, a success probability of 0.8) at the second epoch and the updated NPPI bound of doctors in light of this.

P_j	$S_{2,j}$	$\Delta_{2,j}$	$\delta_{j,2}$	δ_j	$\alpha_{i,j}$	$[\underline{\mathbb{E}}_{2,j}(\theta), \bar{\mathbb{E}}_{2,j}(\theta)]$
P_1	1	0.11	0.03	0.03	0.168	$[0.714, 0.785]$
P_2	2	0.2	0	0	0.433	$[0.7, 0.8]$
P_3	0	0.07	0.193	0.392	0.077	$[0.605, 0.658]$
P_4	1	0.111	0	0.01	0.154	$[0.771, 0.843]$
P_5	1	0.136	0.028	0.028	0.168	$[0.703, 0.784]$

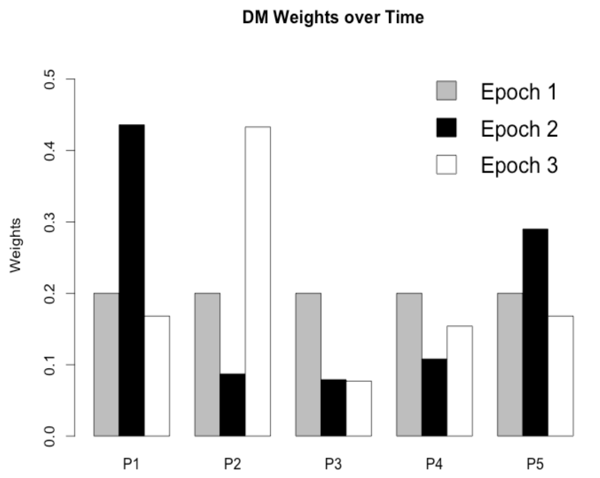

Fig. 6.3: The change in weights awarded to the five DMs over the three epochs.

decreased burden for DMs regarding the manner in which beliefs must be supplied.

6.7 Performance Measure

We wish to assess the merits of various NPPIs. This task is more complex than similar comparisons for the PI approach in Chapter 4. We introduce three rules for determining which NPPI, $[\underline{\mathbb{E}}_{t,i}(\theta), \bar{\mathbb{E}}_{t,i}(\theta)]$ or $[\underline{\mathbb{E}}_{t,j}(\theta), \bar{\mathbb{E}}_{t,j}(\theta)]$, is deemed better.

- **Rule 1:** If both intervals contain $\mathbb{E}(\theta)$ then the better interval is the thinner one, *i.e.*, if $\mathbb{E}(\theta) \in [\underline{\mathbb{E}}_{t,i}(\theta), \bar{\mathbb{E}}_{t,i}(\theta)]$ and $\mathbb{E}(\theta) \in [\underline{\mathbb{E}}_{t,j}(\theta), \bar{\mathbb{E}}_{t,j}(\theta)]$ then we prefer the former if $\Delta_{t,i} < \Delta_{t,j}$.

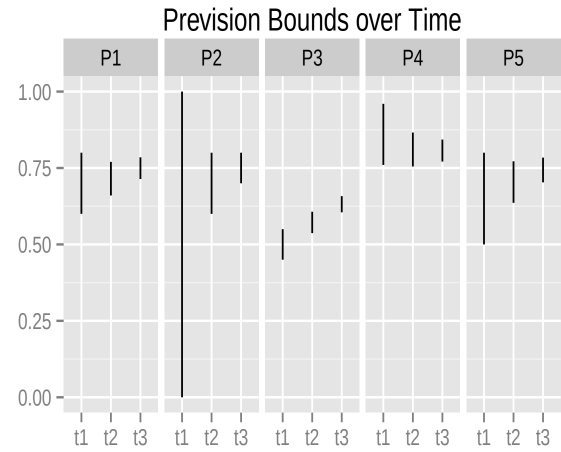

Fig. 6.4: Prevision bounds the five DMs across the three epochs.

- **Rule 2:** If one interval contains $\mathbb{E}(\theta)$ and the other does not then we prefer the former to the latter. When the accurate interval is the thinner one then this is self-evident. When the accurate interval is the wider one we attribute this rule to placing priority on accurate intervals, *i.e.*, a vague and accurate belief is better than a confident and inaccurate belief, as in Table 6.2. Hence if $\mathbb{E}(\theta) \in [\underline{\mathbb{E}}_{t,i}(\theta), \bar{\mathbb{E}}_{t,i}(\theta)]$ and $\mathbb{E}(\theta) \notin [\underline{\mathbb{E}}_{t,j}(\theta), \bar{\mathbb{E}}_{t,j}(\theta)]$ then we prefer the former.

- **Rule 3:** If neither interval contains $\mathbb{E}(\theta)$ then we say the better one is that minimising the sum of its distance to $\mathbb{E}(\theta)$ and one minus its width, *i.e.*, if $\mathbb{E}(\theta) \notin [\underline{\mathbb{E}}_{t,i}(\theta), \bar{\mathbb{E}}_{t,i}(\theta)]$ and $\mathbb{E}(\theta) \notin [\underline{\mathbb{E}}_{t,j}(\theta), \bar{\mathbb{E}}_{t,j}(\theta)]$ then we consider the former more reliable if $d_{t,i} + (1 - \Delta_{t,i}) < d_{t,j} + (1 - \Delta_{t,j})$, where $d_{j,m}$ is

$$d_{j,m} = \min\{|\mathbb{E}(\theta) - \underline{\mathbb{E}}_{m,j}(\theta)|, |\mathbb{E}(\theta) - \bar{\mathbb{E}}_{m,j}(\theta)|\} \tag{6.30}$$

  This is the distance from $\mathbb{E}(\theta)$ to the nearest bound of P_j. This concurs with our earlier discussion about width for inaccurate predictions.

Note that $d_{j,m}$ in Equation (6.30) is distinct from $\delta_{j,m}$ in Equation (6.23), with the latter measuring distance to the last witnessed value rather than the true underlying value $\mathbb{E}(\theta)$. Rule 1 is intuitive. If two NPPIs contain $\mathbb{E}(\theta)$ then the narrower is more desirable, conveying both accuracy and confidence. It is more likely for a single decision to be deemed optimal under a narrow belief as there is a smaller range it must be maximal over. Rule 3 states that for two inaccurate NPPIs we have a preference for wide intervals as previously discussed. The degree of inaccuracy should be incorporated also. Rule 2 is clearly reasonable when the interval containing $\mathbb{E}(\theta)$ is the narrowest. In the case of it being the widest it is harder to concur which is better. Here we place a premium on accuracy, declaring the interval containing $\mathbb{E}(\theta)$ as better. We say that the combined NPPI is optimal if it outperforms more than half the individual DMs, *i.e.*, the metric discussed in Chapter 4. We consider this shortly in our simulation study.

6.8 Links between Individual and Combined NPPIs

Defining an accurate interval as one containing $\mathbb{E}(\theta)$ we have the following theorem.

Theorem 6.1: If the NPPIs of all DMs are accurate then so is the combined NPPI.

Proof: $\mathbb{E}(\theta) \in [\underline{\mathbb{E}}_{t,j}(\theta), \bar{\mathbb{E}}_{t,j}(\theta)]$ for $j = 1, \ldots, n$. As weights are non-negative and sum to one, each amalgamated lower and upper bound is a convex linear combination of DM lower and upper bounds, implying that $\hat{\underline{\mathbb{E}}}_{t,i}(\theta) \in [\min_j \underline{\mathbb{E}}_{t,j}(\theta), \max_j \underline{\mathbb{E}}_{t,j}(\theta)]$ and $\hat{\bar{\mathbb{E}}}_{t,i}(\theta) \in [\min_j \bar{\mathbb{E}}_{t,j}(\theta), \max_j \bar{\mathbb{E}}_{t,j}(\theta)]$. The combined lower bound lies between the smallest and largest lower bound (inclusively) while the combined upper bounds lies between the smallest and largest upper bound (inclusively). As all DMs are accurate this implies $\hat{\underline{\mathbb{E}}}_{t,j}(\theta) \leq \mathbb{E}(\theta)$ and $\hat{\bar{\mathbb{E}}}_{t,j}(\theta) \geq \mathbb{E}(\theta)$, *i.e.*, $\mathbb{E}(\theta) \in [\hat{\underline{\mathbb{E}}}_{t,i}(\theta), \hat{\bar{\mathbb{E}}}_{t,i}(\theta)]$, meaning the combined NPPI is accurate. ■

If it were possible to construct an inaccurate opinion from a set of accurate opinions this would make our method incoherent. We now consider assessing performance of the combined NPPI when all DMs are accurate, starting with an elementary lemma.

Lemma: $\hat{\Delta}_{t,i}(\theta) = \sum_{j=1}^{n} \alpha_{i,j} \Delta_{t,j}(\theta)$ for all $i = 1, \ldots, n$.

Proof: Proof follows from the laws of summations and previous definitions:

$$
\begin{aligned}
\hat{\Delta}_{t,i}(\theta) &= \hat{\underline{\mathbb{E}}}_{t,i}(\theta) - \hat{\bar{\mathbb{E}}}_{t,i}(\theta) \\
&= \sum_{j=1}^{n} \alpha_{i,j} \bar{\mathbb{E}}_{t,j}(\theta) - \sum_{j=1}^{n} \alpha_{i,j} \underline{\mathbb{E}}_{t,j}(\theta) \\
&= \sum_{j=1}^{n} \alpha_{i,j} \Big(\bar{\mathbb{E}}_{t,j}(\theta) - \underline{\mathbb{E}}_{t,j}(\theta) \Big) \\
&= \sum_{j=1}^{n} \alpha_{i,j} \Delta_{t,j}(\theta) \ \blacksquare
\end{aligned}
$$

If the intervals of all DMs are accurate then the combined NPPI is superior if its width is less than that of more than half of individual NPPIs by the rules previously discussed. As $\mathbb{E}(\theta) \in [\underline{\mathbb{E}}_{t,j}(\theta), \bar{\mathbb{E}}_{t,j}(\theta)]$ for all $j = 1, \ldots, n$ we have, by Theorem 6.1, $\mathbb{E}(\theta) \in [\hat{\underline{\mathbb{E}}}_{t,i}(\theta), \hat{\bar{\mathbb{E}}}_{t,i}(\theta)]$. To compare the performance of the combined NPPI against the NPPIs of individuals we use Rule 1, comparing widths, with narrower intervals indicating more reliability. Having seen t sets of data the combined belief is superior if $\hat{\Delta}_{t,i}(\theta) < \Delta_{t,j}(\theta)$, for $j = 1, \ldots, n$ holds for more than $\frac{n}{2}$ of the DMs. Note that if a single DM is inaccurate then the combined NPPI is not ensured to be accurate, but their inaccuracy will be penalised by a low weight and hence their opinion will have only a small impact in the combined NPPI.

6.9 Small Simulation Study

In Chapter 4 we ascertained the merits of the PI approach using simulation, approximating the true probability of it being superior to alternatives in various cases. Here we attempt to do similarly to assess performance of our combined NPPIs. The rules in Section 6.7 are used to compare the relative merits of intervals. We declare the combined NPPI beneficial for a DM if it provides better estimation than the intervals of over half of DMs, *i.e.*, with probability exceeding 0.5 it gives more accurate estimation for a randomly chosen user. Below we run our process from several initialisations and observe how performance varies with adjustments to the number of DMs and/or returns and the prior accuracy of DMs. In each case 5,000 simulations are used to calculate success proportions which approximate true probabilities as the number of iterations grow, converging asymptotically. We conduct testing for the individual and group problems, assessing performance of the combined NPPI against those of DMs, before comparing its merits to those of the equal weights (EQ) scheme.

Combined NPPI *vs.* Individual NPPI

We use four cases to investigate behaviour of our combined NPPI and determine when its performance is strongest. We measure its performance against the intervals of DMs. Results are given in Table 6.8.

- **Case 1a:** We assume $\mathbb{E}(\theta) = 0.4$ with range 0.2, *i.e.*, realisations are uniformly drawn from $[0.2, 0.6]$. DM lower bounds are uniformly chosen from $[0.1, 0.5]$, *i.e.*, on average initial lower bounds of 75% of DMs are less than $\mathbb{E}(\theta)$. Upper bounds are chosen uniformly from $[l + \text{min.width}, l + \min(\text{max.width}, 1 - l)]$. Here l is the lower bound and min.width and max.width are the respective smallest and largest interval width allowed. The term min(max.width,$1 - l$) ensures upper bounds cannot exceed 1. Setting min.width$= 0.05$ and max.width$= 0.5$ the average prior interval width is 0.275. Success proportions denote the proportion of the simulations for which the combined NPPI was superior, *i.e.*, for three DMs and one return in 3,832 of the 5,000 cases (*i.e.*, 0.7664) the combined NPPI was more accurate than the NPPIs of two or more DMs.

- **Case 2a:** The is identical to Case 1a except that we decrease the range of values that realisations take from within 0.2 of $\mathbb{E}(\theta)$ to within 0.1 of it. This should increase the weighted performance as the returns upon which it is based will be more accurate reflections of θ, and therefore should increase the accuracy of the combined NPPI and its success proportion.

- **Case 3a:** This modifies Case 1a so prior accuracy increases. We lower max.width to 0.4 and choose lower bounds uniformly from $[0.25, 0.45]$, meaning intervals now have average width of 0.225. As in Case 1a, on average, 75% of the initial lower bounds of DMs will be less than $\mathbb{E}(\theta)$, but lower bounds will now on average be closer to $\mathbb{E}(\theta)$. Hence DMs are both more accurate and more confident in this accuracy. This should decrease success proportions from Case 1a as opinions of DMs are increasingly accurate prior to sharing with neighbours.

- **Case 4a:** This is identical to Case 3a except that we decrease the range of values which realisations take from within 0.2 of $\mathbb{E}(\theta)$ to within 0.1 of it.

Analysis of Results 1

The primary conclusion that we draw from our results is that combining intervals is worthwhile for DMs. In each case considered the proportion of times that this method is successful exceeds 0.5, *i.e.*, in all instance it was in the best interest of a DM to use it. As the number of DMs increases so do success proportions. This is intuitive. As the amount of DMs involved in the process grows larger the amount of sources of information that a DM has access to increases. We previously conjectured that there

Table 6.8: Success proportion of the combined NPPI for Cases 1a-4a. Highlighted in bold are the instances when our method was optimal.

Case	DMs	Epochs	Prop.	DMs	Epochs	Prop.	DMs	Epochs	Prop.
1a	3	1	**0.7664**	3	3	**0.7814**	3	5	**0.7852**
1a	7	1	**0.8390**	7	3	**0.8946**	7	5	**0.8852**
1a	11	1	**0.8590**	11	3	**0.9204**	11	5	**0.9082**
1a	23	1	**0.8830**	23	3	**0.9316**	23	5	**0.9154**
2a	3	1	**0.8464**	3	3	**0.8516**	3	5	**0.8444**
2a	7	1	**0.9532**	7	3	**0.9622**	7	5	**0.9578**
2a	11	1	**0.9702**	11	3	**0.9876**	11	5	**0.9870**
2a	23	1	**0.9882**	23	3	**0.9950**	23	5	**0.9960**
3a	3	1	**0.6718**	3	3	**0.6828**	3	5	**0.6966**
3a	7	1	**0.7686**	7	3	**0.7934**	7	5	**0.7914**
3a	11	1	**0.8008**	11	3	**0.8282**	11	5	**0.8306**
3a	23	1	**0.8414**	23	3	**0.8542**	23	5	**0.8574**
4a	3	1	**0.7458**	3	3	**0.7532**	3	5	**0.7534**
4a	7	1	**0.8776**	7	3	**0.8782**	7	5	**0.8740**
4a	11	1	**0.9524**	11	3	**0.9244**	11	5	**0.9234**
4a	23	1	**0.9592**	23	3	**0.9572**	23	5	**0.9610**

is a direct correlation between information and accuracy, and between these quantities and decision quality. Success proportions remain relatively constant as the number of epochs increases. We see between Case 1a and Case 2a, and between Case 3a and Case 4a, that there are significant increases in proportions. This relates to the decrease in the range of realisations, with observations falling in $[0.2, 0.6]$ in Cases 1a and 3a and in $[0.3, 0.5]$ in Cases 2a and 4a. In the latter instances learning occurs faster as witnessed data are more genuine reflections of the true underlying phenomenon. Hence success proportions are significantly higher than in the cases with wider ranges. We see increasing prior accuracy (Cases 3a and 4a) leads to decreases in success proportions for the combined opinions. This is expected, as if DMs are themselves more accurate then they are less dependent upon opinions of others.

Combined NPPI *vs.* Equally Weighted NPPI

Here we consider four cases, assessing performance against the EQ method. Success proportions are given in Table 6.9.

- **Case 1b:** We assume $\mathbb{E}(\theta) = 0.4$ and sett the range of possible values as 0.1, *i.e.*, observations lie in $[0.3, 0.5]$. Lower bound priors are uniformly chosen from $[0.1, 0.5]$, *i.e.*, on average, 75% of DMs initially have lower bounds less than $\mathbb{E}(\theta)$. The minimum width is 0.05 and the maximum is 0.5.

- **Case 2b:** Case 1b is modified so lower bounds are now uniformly chosen from $[0.2, 0.6]$. This implies that, on average, half of DMs have initial lower bounds less than $\mathbb{E}(\theta)$, *i.e.*, a higher proportion of DMs overestimate it.

- **Case 3b:** This augments Case 1b, with lower bounds now uniformly chosen from $[0.3, 0.7]$. This implies that, on average, only a quarter of DMs initially have lower bounds less than $\mathbb{E}(\theta)$, *i.e.*, a higher proportion overestimate it.

- **Case 4b:** Case 2b is altered with lower bounds now uniformly chosen from $[0.3, 0.5]$. As in Case 2b this implies that on average half of the DMs have initial lower bounds greater than $\mathbb{E}(\theta)$, but in this instance bounds are closer to $\mathbb{E}(\theta)$, *i.e.*, they are increasingly accurate.

Analysis of Results 2

The dominance of our combining approach is not as unconditionally evident in the group setting as it was for the individual problem, as is clear from Case 1b, where it is outperformed in most instances by the EQ approach. Yet as is obvious from Cases 2b, 3b and 5b, in certain situations it is in the best interest of DMs to merge opinions in our more sophisticated manner. In these instances the prior accuracy of DMs is weakened with on average 50% and 75% of DMs in Cases 2b and 3b respectively overestimating $\mathbb{E}(\theta)$, with success proportions above 0.5 in both scenarios, dramatically so in the latter. Hence we see that if most DMs are *a priori* accurate information sources then a straightforward equally weighted combination will perform strongly, yet as DMs become initially increasingly inaccurate our more subtle method grows increasingly powerful. In Case 4b we made the lower bound predictions of the DMs closer to $\mathbb{E}(\theta)$,

Table 6.9: Success proportion of the combined NPPI for Cases 1b-4b. Highlighted in bold are the instances when our method was optimal.

Cases	DMs	Epochs	Prop.	DMs	Epochs	Prop.	DMs	Epochs	Prop.
1b	3	1	0.4474	3	3	0.3446	3	5	0.3484
1b	7	1	0.4806	7	3	0.1708	7	5	0.1198
1b	11	1	**0.5526**	11	3	0.1274	11	5	0.0542
1b	23	1	**0.6906**	23	3	0.0772	23	5	0.0100
2b	3	1	**0.5570**	3	3	**0.5952**	3	5	**0.6164**
2b	7	1	**0.6932**	7	3	**0.5878**	7	5	**0.5586**
2b	11	1	**0.7548**	11	3	**0.5924**	11	5	**0.5614**
2b	23	1	**0.8560**	23	3	**0.5930**	23	5	**0.5366**
3b	3	1	**0.6936**	3	3	**0.8522**	3	5	**0.8964**
3b	7	1	**0.8522**	7	3	**0.9648**	7	5	**0.9718**
3b	11	1	**0.9158**	11	3	**0.9886**	11	5	**0.9872**
3b	23	1	**0.9780**	23	3	**0.9996**	23	5	**0.9984**
4b	3	1	**0.5544**	3	3	**0.5568**	3	5	**0.5714**
4b	7	1	**0.6670**	7	3	**0.5804**	7	5	**0.5374**
4b	11	1	**0.7452**	11	3	**0.5850**	11	5	**0.5122**
4b	23	1	**0.8302**	23	3	**0.5976**	23	5	**0.5048**

while maintaining the same ratio of accuracy as in Case 2b, leading to decreases in success proportions. Nevertheless they remain above 0.5.

In conclusion, from our brief study it is evident that there is merit in using the linear pooling technique from Sections 6.3-6.4. When comparing its performance to that of the opinions of DMs it appears to be strong, dominant in all cases considered. Results are more ambiguous in comparison with the method of consistent equal weights, which can also be successful. If priors of DMs are all suitably accurate than the EQ approach is the better of the two, while as DMs become increasingly accurate *a priori* the more sophisticated technique becomes stronger. Motivation exists for determining conditions for which each technique is dominant. There is a wide range of variables to consider, *e.g.*, prior accuracy of DMs (expressible in multiple ways), the range which witnessed values can occur in, and the number of DMs/epochs considered. In our fully parameterised setting we instigated a study of this nature but doing so is more complicated in our nonparametric setting, with more parameters to be specified, and with results potentially sensitive to the chosen initialisation. An area for further work, before carrying out a full scale study, is determining a set of cases to consider, covering a suitably broad spectrum of scenarios to provide generalised results. In the above study we only considered the EQ method as a rival to our technique, yet another alternative exists in the form of the Most Reliable approach. We could also compare our combining method to this simplistic performance-based approach.

Chapter 7

Potential Subjective Alternatives

Much of this thesis has considered objective methods of ascertaining the reliability of DM's opinions, with these opinions being either parametrically or nonparametrically expressed. In this chapter we discuss two alternative techniques that add subjectivity to the process. The technique that we shall focus on primarily is called the Differing Viewpoints (DV) approach, incorporating utility as well as probability (*i.e.*, the two core elements of decision making) into its calculations. We compare this method to the PI approach, as well as discussing its relationship with Absolute Risk Aversion (ARA) and a suitable performance metric. The other method, termed the Kullback-Leibler (KL) approach, bases weights upon the degree of similarity between opinions, entailing biases that DMs may feel towards opinions that are akin to their own.

7.1 The Differing Viewpoints Approach

As utility is a fundamental decision making element it seems attractive for a reweighting process to incorporate this. The PI approach was fully objective (assuming equal initial weights are supplied), giving the same results for all DMs, *i.e.*, the weight allocated by P_i to P_j ($\alpha_{i,j}$) is the same as that allocated by P_k to P_j ($\alpha_{k,j}$) and indeed is the same as that P_j affords herself ($\alpha_{j,j}$). Calculation of the PI weight given to P_j was independent of the utility function of the DM assigning this. Yet DMs with different fortunes and attitudes towards risk may react differently to common information, *e.g.*, a risk prone DM may give more weight to an opinion predicting a large profit (with strong associated uncertainty) than to one predicting a small profit (with strong conviction in this belief). A risk averse DM may do the opposite. Arguments of this nature motivate development of a weighting scheme assimilating both probability and the utility function of the user into its calculations.

Initially DMs have no data to base reliability assessments on and combine beliefs as in Equation (3.5). All DMs calculate the decision, d^*, that is deemed optimal for them, *i.e.* that maximises expected utility. Once this decision has been made and a return observed then DMs can start assessing the reliability of their neighbours. Similarly to the PI approach, each DM compares the distributions offered by neighbours to the witnessed outcome, but in the DV approach they do so in a manner taking their own utility function into account as well. Each P_i calculates the expected utility that she would have assigned to d^* had she only listened to the beliefs of P_j, *i.e.*,

$$\mathbb{E}_{i|j}[u_i(d^*)] = \int_{\Theta} u_i(d^*, \theta) f_j(\theta)\ d\theta \tag{7.1}$$

Having done this P_i then computes the absolute value of the difference between this expected utility associated with d^* under the beliefs of P_j and the actual utility that was witnessed following d^* (for all $j = 1, \ldots, n$). We assume that a common return is witnessed by all DMs with comments on an extension to the previously discussed multiple simultaneous returns setting included in Appendix F. This quantity is known as the DV weight, $v_{i,j}$, formally written as

$$v_{i,j} = |u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]| \tag{7.2}$$

If P_j returns a small $v_{i,j}$ then this implies that there is a small discrepancy (in terms of the utils of P_i) between her prediction and the outcome that occurred. Hence she may be considered a reliable information source as her probability distribution seems to accurately model θ in terms of the preferences of P_i. A large $v_{i,j}$ implies the converse. The goal of the DV weight is identical to that of the PI weight, *i.e.*, to compare predictions with observations and to construct a reliability measure based on these disparities. In the PI approach this was based entirely on probability while the DV approach uses both probability and utility. Note accurate individuals return large PI weights in Equation (3.6) but small DV weights in Equation (7.2), and the converse.

The DV method deals only with the order of magnitude between two quantities and is invariant to which is bigger. Suppose for P_i, that $u_i(r) = 10$, $\mathbb{E}_{i|j}[u_i(d^*)] = 5$ and $\mathbb{E}_{i|k}[u_i(d^*)] = 15$, giving $v_{i,j} = v_{i,k} = 5$. It may appear unintuitive that P_i awards the same DV weight to P_j and P_k despite the fact that P_k predicts a higher expected utility, with DMs desiring their utility to be as large as possible. If a DM is trying to determine which decision to choose from a set of alternatives she will be interested in the cardinality of expected utilities as (assuming she is rational) she wishes to maximse her return. However in Equation (7.2) her decision has already been made and cannot be changed. Her sole interest is in ascertaining the reliability of neighbours in the hope of positively influencing the quality of her next decision. Hence 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 to a particular neighbour, not what she would get. Given this, we define the updated normalised weight $\alpha^*_{i,j}$ afforded by P_i to P_j, given she previously allocated her a normalised weight of $\alpha_{i,j}$ and that her most recent DV weight is $v_{i,j}$,as

$$\alpha^*_{i,j} = \frac{\frac{\alpha_{i,j}}{v_{i,j}}}{\sum_{k=1}^{n} \frac{\alpha_{i,k}}{v_{i,k}}} \tag{7.3}$$

Issues arise regarding divisibility by zero if $v_{i,j} = 0$, *i.e.*, if $u_i(r) = \mathbb{E}_{i|j}[u_i(d^*)]$. This can only occur in the limit, when P_j's mean prediction equals the return witnessed and her associated variance tends towards zero. Due to our requirement that all DMs must possess some uncertainty (*i.e.*, positive variance) this limit is never reached in practice. As previously discussed, utility is invariant to positive linear transformation. Hence it is desirable that a utility-based weighting scheme adhere to this invariance. Consider a function $u_i^{'}(r) = au_i(r) + b$, with $a, b \in \mathbb{R}$ and $a > 0$. This gives a DV weight, $v_{i,j}^{'}$, of

$$\begin{aligned} v_{i,j}^{'} &= |u_i^{'}(r) - \mathbb{E}_{i|j}[u_i^{'}(d^*)]| \\ &= |au_i(r) + b - \mathbb{E}_{i|j}[au_i(d^*) + b]| \\ &= |au_i(r) + b - a\mathbb{E}_{i|j}[u_i(d^*)] - b| \\ &= |au_i(r) - a\mathbb{E}_{i|j}[u_i(d^*)]| \\ &= a|u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]| \\ &= av_{i,j}. \end{aligned} \tag{7.4}$$

Inserting this DV weight into Equation (7.3) we see that

$$\alpha^{*\,'}_{i,j} = \frac{\frac{\alpha_{i,j}}{v_{i,j}^{'}}}{\sum_{k=1}^{n} \frac{\alpha_{i,k}}{v_{i,k}^{'}}}$$

$$
\begin{aligned}
&= \frac{\frac{\alpha_{i,j}}{av_{i,j}}}{\sum_{k=1}^{n} \frac{\alpha_{i,k}}{av_{i,k}}} \\
&= \frac{\frac{\alpha_{i,j}}{v_{i,j}}}{\sum_{k=1}^{n} \frac{\alpha_{i,k}}{v_{i,k}}} \\
&= \alpha^*_{i,j} \qquad (7.5)
\end{aligned}
$$

Hence this weighting scheme ensures invariance to affine transformation for a DMs utility function, a desirable and sensible property for a utility-based weighting approach to obey. Implicit in our DV method is the Markovian property and the four criteria (with suitable inversions) discussed in Section 3.5 (proofs are given in Appendix B), making it coherent in some sense. We provide further discussion on the functional form of reweighting scheme given in Equation (7.3). It seems rational that this should be a function of a the previous normalised weight of a DM and her most recent DV weight, *i.e.*, $\alpha^*_{i,j} \propto h(\alpha_{i,j}, v_{i,j})$ for some function $h(\cdot)$. We have seen $h(\alpha_{i,j}, v_{i,j}) = \frac{\alpha_{i,j}}{v_{i,j}}$ ensures invariance. What other choices do? We saw $v_{i,j}$ under $u_i(r)$ transformed to $v'_{i,j} = av_{i,j}$ under $u'_i(r)$ in Equation (7.4). Hence, to obey invariance, h must be such that

$$
\begin{aligned}
\alpha'^*_{i,j} &= \frac{h(\alpha_{i,j}, v'_{i,j})}{\sum_{k=1}^{n} h(\alpha_{i,k}, v'_{i,k})} \\
&= \frac{h(\alpha_{i,j}, av_{i,j})}{\sum_{k=1}^{n} h(\alpha_{i,k}, av_{i,k})} \qquad (7.6) \\
&= \frac{ah(\alpha_{i,j}, v_{i,j})}{a\sum_{k=1}^{n} h(\alpha_{i,k}, v_{i,k})} \qquad (7.7) \\
&= \frac{h(\alpha_{i,j}, v_{i,j})}{\sum_{k=1}^{n} h(\alpha_{i,k}, v_{i,k})} \\
&= \alpha^*_{i,j}
\end{aligned}
$$

We see between Equations (7.6) and (7.7) that it must be possible to write $h(\alpha_{i,j}, av_{i,j})$ as $ah(\alpha_{i,j}, v_{i,j})$ implying $h(\cdot)$ must be multiplicative in its arguments, *i.e.*, must be a product or quotient involving $v_{i,j}$ and $\alpha_{i,j}$. An obvious choice of that h that adheres to this is $h(\alpha_{i,j}, v_{i,j}) = v_{i,j}\alpha_{i,j}$, yet this penalises DMs with low values of $v_{i,j}$ (indicating accuracy), making it incoherent. We choose $h(\alpha_{i,j}, v_{i,j}) = \frac{\alpha_{i,j}}{v_{i,j}}$, as it ensures rational updating, *i.e.*, high weights for those with low values of $v_{i,j}$ and the converse. Various other forms are possible with arguments raised to powers, *e.g.*, $h(\alpha_{i,j}, v_{i,j}) = \frac{\alpha_{i,j}}{v^2_{i,j}}$ yet why this would be done is not obvious. Another property of Equation (7.3) is that, if we denote $\alpha^{(t)}_{i,j}$ by the normalised weight assigned by P_i to P_j after t returns, and by

$v_{i,j}^{(t)}$ the DV weight assigned for the t^{th} return, we can rewrite it as:

$$\begin{aligned}
\alpha_{i,j}^{(t)} &\propto \frac{\alpha_{i,j}^{(t-1)}}{v_{i,j}^{(t)}} \\
&= \frac{\frac{\alpha_{i,j}^{(t-2)}}{v_{i,j}^{(t-1)}}}{v_{i,j}^{(t)}} = \frac{\alpha_{i,j}^{(t-2)}}{v_{i,j}^{(t)} v_{i,j}^{(t-1)}} \\
&= \frac{\frac{\alpha_{i,j}^{(t-3)}}{v_{i,j}^{(t-2)}}}{v_{i,j}^{(t)} v_{i,j}^{(t-1)}} = \frac{\alpha_{i,j}^{(t-3)}}{v_{i,j}^{(t)} v_{i,j}^{(t-1)} v_{i,j}^{(t-2)}} \\
&= \vdots \\
&= \frac{\alpha_{i,j}^{(0)}}{\prod_{k=1}^{t} v_{i,j}^{(k)}} = \frac{\frac{1}{n}}{\prod_{k=1}^{t} v_{i,j}^{(k)}} \\
&= \frac{1}{n \prod_{k=1}^{t} v_{i,j}^{(k)}} \qquad (7.8)
\end{aligned}$$

We see that the unnormalised weight assigned by P_i to P_j is the inverse of the product of her initial (equal) weight and her DV weights to this point. This is akin to the relationship we determined for the PI weights (which guaranteed exchangeability in its setting) in Equation (3.14). Note that a similar exchangeability does not hold for the DV scheme as weights are calculated as a function of a DMs initial fortune which fluctuates over time, *i.e.*, the weight given to a neighbour will differ dependent on if the DM assigning it has a fortune of (for instance) \$50 or \$55.

We conclude this section by commenting briefly that, in Equation (7.1), the DM retrospectively considers the utility that she would have expected to achieve had she solely heeded the opinion of her neighbour P_j. In comparing this, in Equation (7.2), to the utility which was actually realised from the decision that she made, the DM can be seen in some sense as assessing if she regrets not considering this opinion more prominently or the contrary, *i.e.*, if perhaps she may have increased her decision quality by listening more prominently to that individual. As an aside, we comment that the aim of this process can be seen as kindred to that of Regret Theory, derived in Loomes & Sugden (1982), Bell (1982) and Fishburn (1982). This theory was motivated by experimental research, contrasting the regret anticipated and experienced by individuals facing pending uncertainty, and with users assimilating a regret component (generally increasing, continuous and non-negative) to their utility function, with the former subtracted from the latter. Regret Theory has been used to explain phenomena witnessed in auction environments and behavioural finance.

7.1.1 Differing Viewpoints and ARA

Suppose P_i is determining the weight to assign to the opinions of P_j and P_k. A return r has been witnessed and the Normally distributed beliefs of P_j and P_k are symmetric around this with common variance, *i.e.*, the mean prediction of P_j is m units below r and that of P_k is m units above r. Under the PI method these beliefs would receive equal PI weights, *i.e.*, $w_j = w_k$, as a prediction m units below the mean is as accurate as a prediction m units above the mean (given equal variances) by the symmetry of the Normal distribution. The following theorem shows that both distributions may receive different DV weights depending on the nature of $u_i(r)$.

Theorem 7.1: Suppose P_i observes a return r. Two neighbours have the beliefs $f_j(\theta) \sim \mathcal{N}(a, \sigma^2)$ and $f_k(\theta) \sim \mathcal{N}(b, \sigma^2)$, with $|r - a| = |r - b|$ and $a < r < b$. If P_i is

- risk averse over the range of return values she will give a higher DV weight to P_j.
- risk prone over the range of return values she will give a higher DV weight to P_k.
- risk neutral over the range of return values she will give equal DV weights.

Proof: We can find that

- $\mathbb{E}_{i|j}[u_i(d^*)] = \int_{-\infty}^{\infty} \frac{u_i(r)}{\sqrt{(2\pi\sigma^2)}} \exp\Big(- \frac{(r-a)^2}{2\sigma^2}\Big)\, dr$
- $\mathbb{E}_{i|k}[u_i(d^*)] = \int_{-\infty}^{\infty} \frac{u_i(r)}{\sqrt{(2\pi\sigma^2)}} \exp\Big(- \frac{(r-b)^2}{2\sigma^2}\Big)\, dr$

Clearly $\mathbb{E}_{i|k}[u_i(d^*)] > \mathbb{E}_{i|j}[u_i(d^*)]$ as $b > a$.

- If P_i is risk-averse (*i.e.*, $A_i(r) > 0$ for all r considered) then by the nature of her utility function, $|u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]] > |u_i(r) - \mathbb{E}_{i|k}[u_i(d^*)]|$. This implies $v_{i,j} > v_{i,k}$. The DM with the higher prediction is deemed more reliable.

- If P_i is risk-prone (*i.e.*, $A_i(r) < 0$ for all r considered) then by the nature of her utility function, $|u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]| < |u_i(r) - \mathbb{E}_{i|k}[u_i(d^*)]|$. This implies $v_{i,j} < v_{i,k}$. The DM with the lower prediction is deemed more reliable.

- If P_i is risk-neutral (*i.e.*, $A_i(r) = 0$ for all r considered) then by the nature of her utility function, $|u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]| = |u_i(r) - \mathbb{E}_{i|k}[u_i(d^*)]$. This implies $v_{i,j} = v_{i,k}$. Both DMs are deemed equally reliable. ∎

This result may seem counter-intuitive but is rational under proper consideration. By definition, a risk prone DM will have a higher disparity between $u_i(b)$ and $u_i(r)$ than that between $u_i(r)$ and $u_i(a)$. Hence when r is realised the difference between the expected utility promised by P_k and the utility actually occurring will be greater than that between the expected utility promised by P_j and that occurring. Therefore (given the utility preferences of P_i) P_j is more accurate than P_k, as her prediction is closer in utils to the utility that actually occurred than the prediction of P_k. Hence P_j will get a lower DV weight. Similar arguments apply for the risk-averse/risk-neutral cases.

For illustration, suppose $u_i(r) = \exp(\frac{r}{20})$, *i.e.*, P_i is risk prone with $A_i(r) = -\frac{1}{20} < 0$ for all r. A return $r = 5$ is witnessed and $f_j(\theta) \sim \mathcal{N}(4, 1)$ and $f_k(\theta) \sim \mathcal{N}(6, 1)$. If P_i has an initial fortune of 35 then this leads to $v_{i,j}$=0.35 and $v_{i,k} = 0.391$, *i.e.*, the DM with the lower prediction is deemed more reliable. If $u_i(r) = \ln(r)$ (risk prone for $r > 0$ with values less than this not computable), P_i has a fortune of 50, a return of $r = 10$ is seen, and $f_j(\theta) \sim (5, 1)$ and $f_k(\theta) \sim (15, 1)$ then this leads to $v_{i,j} = 0.086$ and $v_{i,k} = 0.081$. The DM with the higher prediction is deemed most reliable.

Theorem 7.1 includes a caveat requiring DMs to be risk-averse/risk-neutral/risk-prone over the return range considered. Yet when the Normal distribution is used the range of plausible values is the real line. A logarithmic utility function is not defined below zero and a utility function $u(r) = r^2$ will change from risk prone to risk averse as returns move from above to below zero. In the examples above it was assumed that the probability of straying into these change-point areas was negligibly small enough to be considered zero. An area for further research is to reformulate our theorem to allow returns lie on the boundary between a risk averse and risk prone utility function.

7.1.2 Different Weights from Different Methods

The weight allocated to a particular DM based on her perceived reliability is the same for all neighbours under the PI approach. Below we illustrate how this is not the case under the DV approach and highlight how weights change as the utility function of a DM changes (as measured by her ARA coefficient). Consider a setting consisting of three DMs each with a Normally distributed belief over the mean of a Normally distributed quantity, *i.e.*, $f_i(\theta) \sim \mathcal{N}(m_i, s_i^2)$ for $i = 1, 2, 3$ where $R \sim \mathcal{N}(\theta, 2)$. The prior beliefs of DMs are in Table 7.1 as well as their utility functions and initial fortunes (γ_i).

These functions are dependent upon the fortunes of DMs, *i.e.*, γ_i is an argument of $u_i(\cdot)$. We comment on our Normality assumption. This is for simplicity with fortunes augmented from γ_i to $\gamma_i + r$ and utilities from $u_i(\gamma_i)$ to $u_i(\gamma_i + r)$. For a setting involving, for instance, the Binomial distribution, additional rules must be specified to relate returns witnessed to the new fortune of a DM.

Table 7.1: Opinions, initial fortunes, utility functions and ARAs of DMs.

P_i	$f_i(\theta)$	γ_i	$u_i(\gamma_i + r)$	ARA_i
P_1	$\mathcal{N}(-2, 2)$	10	$(r + 10)^2$	$-\frac{1}{r+10}$
P_2	$\mathcal{N}(3, 3)$	20	$r + 20$	0
P_3	$\mathcal{N}(5, 2)$	30	$\log_e(r + 30)$	$\frac{1}{r+30}$

Suppose $r = 1.5$ is observed. Initial equal weights were given (under the PI and DV approaches). Once data has been seen weights are updated. Table 7.2 shows weights for the PI method (common to all DMs) and those for each DM under the DV approach.

Table 7.2: The weights allocated by DMs under the PI and DV methods.

	Plug-in	P_1 (Risk Prone)	P_2 (Risk Neutral)	P_3 (Risk Averse)
$\alpha_{i,1}$	0.189	0.297	0.231	0.214
$\alpha_{i,2}$	0.622	0.495	0.538	0.547
$\alpha_{i,3}$	0.189	0.208	0.231	0.239

The return witnessed was halfway between the respective mean predictions of P_1 and P_3 with both having common variances. We see that these DMs are given equal weights under the PI approach and also by risk-neutral P_2 (by Theorem 7.1) using the DV method. Risk-prone P_1 gives a higher weight to the DM with the lower prediction while risk-averse P_3 gives a higher weight to the DM with the higher prediction. Here we see that all DMs agree on who the most reliable individual is (*i.e.*, P_2). Note that DMs have assigned different weights under the DV method to under the PI method. We examine how sensitive weights are to initial fortunes. Table 7.3 shows the weights allocated when fortunes are doubled from those in Table 7.1. Identical ordinal reliability rankings hold but the cardinal values have changed.

Table 7.3: Weights allocated by DMs under the PI and DV methods with γ_i doubled.

	Plug-in	P_1 (Risk Prone)	P_2 (Risk Neutral)	P_3 (Risk Averse)
$\alpha_{i,1}$	0.189	0.265	0.291	0.229
$\alpha_{i,2}$	0.622	0.516	0.418	0.535
$\alpha_{i,3}$	0.189	0.219	0.291	0.236

Also of interest is how weights change dependent on the extremity of a DM's utility function measured by ARA. We consider the situation above, changing the utility function of P_1 each time to make her progressively more risk-prone. We keep her initial fortune fixed at 20 with results given in Table 7.4. We see changes in weights as the exponent of $u_1(\cdot)$ grows, *i.e.*, as she becomes more risk-prone. The DM predicting the higher return is given an increasingly small weight over time as the discrepancy between her prediction (in terms of the utils of P_1) and what occurs grows larger as $u_1(\cdot)$ grows more extreme. We see that the weights allocated by the DV method are heavily influenced by utility functions. It would be an interesting exercise to rewrite the weights assigned by a DM as a function of her ARA, a task that is possible because given the ARA of a DM one can derive her utility function (up to a linear transformation) using ordinary differential equations (*e.g.*, Houlding, 2008).

Table 7.4: Changing $u_1(\cdot)$, $A_1(r)$, and weights assigned by P_1 for $r = 1.5$.

$u_1(20+r)$	$A_1(r)$	$\alpha_{1,1}$	$\alpha_{1,2}$	$\alpha_{1,3}$
$(r+20)^2$	$-\frac{1}{20+r}$	0.265	0.515	0.219
$(r+20)^3$	$-\frac{2}{20+r}$	0.300	0.493	0.207
$(r+20)^4$	$-\frac{3}{20+r}$	0.339	0.468	0.193
$(r+20)^5$	$-\frac{4}{20+r}$	0.378	0.443	0.179
$(r+20)^6$	$-\frac{5}{20+r}$	0.419	0.416	0.165

7.1.3 Differing Viewpoints Metric

A simulation study justified the fully objective probability-based PI approach. Posterior distributions were calculated and contrasted to those from alternative techniques, *i.e.*, the Equal Weights (EQ) and Most Reliable (MR) methods. We considered the density that each distribution placed *a posteriori* on θ and declared the optimal method as that maximising this. The DV method incorporates probability and utility. Hence its justification should be based upon these two concepts. We discuss a suitable comparison metric and illustrate it in practice. A set of DMs make m decisions, seeing m returns, and updating weights m times. Having done this, we consider the expected utility values that are predicted by the various methods for the next epoch and contrast these to the actual utility that would result from the true value of θ being witnessed. Small discrepancies imply predictions are accurate as there is only a small difference between reality and that DM's projection of it. In practice DMs will never know the true value of θ; it is only considered here in order to justify the DV approach. We assume that returns are Normally distributed as discussed above. For our justification methodology we suppose at each epoch DMs make decisions leading to a return being witnessed (*i.e.*, gambles) rather than ones leaving fortunes unchanged. This is done to ensure that we need not determine which decision is optimal for DMs at each epoch, as our goal in this exercise is not to guarantee good decision making but to validate weights/estimates for various linear pooling methods.

Suppose we have three DMs, with interest in the mean θ of a Normal process with $R \sim \mathcal{N}(\theta, 2)$ and $\theta = 5$. Each P_i has a Normal prior over θ, *i.e.*, $f_i(\theta) \sim \mathcal{N}(m_i, s_i^2)$ as well as has her own utility function $u_i(\gamma_i + r)$ and her fortune γ_i given in Table 7.5.

Table 7.5: DM information for example.

P_i	$f_i(\theta)$	γ_i	$u_i(\gamma_i + r)$
P_1	$\mathcal{N}(-5, 4)$	10	$r + 10$
P_2	$\mathcal{N}(6, 3)$	15	$(r + 15)^2$
P_3	$\mathcal{N}(11, 4)$	20	$\exp(\frac{r+20}{20})$

Using the DV method the weight that P_i assigns to P_k will often not be the same as that which P_j allocates to P_k. Likewise all DMs need not agree on which neighbour they consider most reliable (getting a weight of one under the MR method) as this is based upon their own subjective utility functions. Unlike in the PI approach interest is not solely in the means/variances resulting from various techniques and the contrast between these and θ. Instead we are concerned with the expected utilities predicted by DMs and the contrast between these and the true utility occurring. The severity of this discrepancy for P_i is dependent on $u_i(\gamma_i + r)$. Hence to assess the accuracy of distributions from the different methods we take utility into account and compare the expected utility each of these methods would allocate at the next decision epoch, to the utility that would occur from θ being witnessed. This is a strong indicator of the merits of methods as on average realisations will tend towards θ.

The utility contrasts are provided in Table 7.6. We see for P_1 that the DV method leads to an expected utility only 0.438 utils away from the true utility she will experience from θ occurring. Of the three pooling approaches it is the DV method that gives the most accurate estimation for each DM. We see that for P_1 and P_2 that more accurate estimation comes from the DV opinion rather than their own, while P_3 in this instance gets more accurate estimation from her own distribution. Hence here the DV method is declared superior. For the individual problem it gives better estimation than over half of individual distributions. In the group problem it gives better estimation than the other two methods for all DMs.

Table 7.6: Absolute differences between expected utilities predicted by pooling methods (and DMs own distributions) and the true utility occurring from θ being witnessed, respectively denoted by DV_{diff}, EQ_{diff}, MR_{diff} and IND_{diff}.

P_i	DV_{diff}	EQ_{diff}	MR_{diff}	IND_{diff}
P_1	0.438	0.921	0.657	1.942
P_2	35.06	65.10	47.00	47.00
P_3	0.193	0.345	0.251	0.06

A simulation study can be conducted determining how results may vary with the number of DMs in the group and/or returns witnessed. We have seen PI performance growing stronger as more neighbours were involved/returns witnessed. Will the DV approach be the same? How do initial fortunes impact upon weights and hence which method is superior? Perhaps of most interest is the relationship between success proportions and the nature of the utility function of the user in terms of her ARA. Does it provide more accurate estimation for DMs with a particular type of function? An intuition is that the further utility functions deviate from risk neutrality the more accurate estimation will become but this remains to be proved.

7.1.4 Contrasting the DV and PI approaches

We have discussed how the DV and PI methods generally produce different weights. In the example below we calculate weights and resulting posterior distributions using both approaches and compare performance under the probability-based metric used to justify the PI approach and the utility-based metric used to justify the DV approach. Suppose we have three DMs interested in θ with $R \sim \mathcal{N}(\theta, 2)$ and $\theta = 5$, each having a Normally distributed opinion which is given alongside their initial fortune and utility function in Table 7.7. We see that P_2 and P_3 are risk prone, and that P_1 is risk-averse as she has quadratic utility but a negative fortune. We simulate two returns.

Table 7.7: DM information for DV/PI comparison.

P_i	$f_i(\theta)$	γ_i	$u_i(\gamma_i + r)$
P_1	$\mathcal{N}(2,3)$	-15	$(r-15)^2$
P_2	$\mathcal{N}(7,2)$	20	$(r+20)^3$
P_3	$\mathcal{N}(9,2)$	25	$\exp(\frac{r+25}{20})$

Weights given by DMs are in Table 7.8. Under the DV approach risk-prone P_2 and P_3 give similar respective weights to neighbours. The weights allocated by risk-averse P_1 contrast starkly. Interestingly the weights (commonly) given by the PI approach are disparate from those under the DV method for each neighbour, *e.g.*, all DMs consider P_2 the least reliable under the DV approach yet the PI approach considers her the most reliable. Posterior distributions for DMs under the different methods are in Fig. 7.1, with different DV posteriors for each DM as all DMs assign different weights. The PI method places more density on the true value of θ than the DV posteriors of all DMs, *i.e.*, its performance is superior by the probability-based metric.

We contrast the expected utilities that would be predicted by the different posteriors

Table 7.8: Weights allocated by DMs under the PI and DV methods.

	Plug-in	P_1 (DV)	P_2 (DV)	P_3(DV)
$\alpha_{i,1}$	0.242	0.278	0.559	0.509
$\alpha_{i,2}$	0.622	0.271	0.191	0.209
$\alpha_{i,3}$	0.136	0.451	0.250	0.282

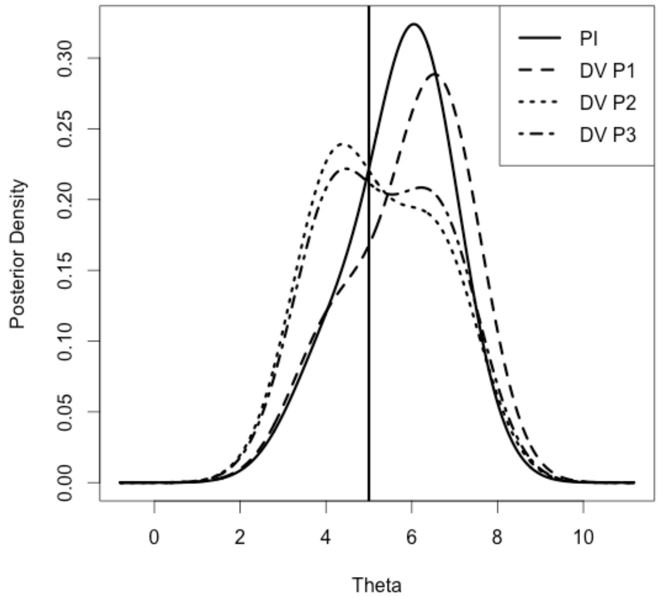

Fig. 7.1: Posteriors of the PI and DV approaches with a vertical line denoting θ.

to the utility that would occur from the true value of θ being realised, *i.e.*, the metric from Section 7.1.3. From Table 7.9, the PI method gives better estimation for P_1 but the DV posteriors are better for P_2 and P_3. Hence as it provides more accurate estimation for over half the DMs the DV method is deemed superior in this instance.

We have seen in this particular example that the PI approach gives more accurate estimation under the probability-based metric while the DV approach gives more accurate estimation under the utility-based metric. Hence which method is deemed superior is a subjective decision based upon the prerogative of the user and if they wish to use their utility function in their weighting calculations or not. Severely contrasting weights are likely to occur dependent on which approach a user opts to implement.

Table 7.9: The difference between the expected utility predicted by the DV/PI approaches and the true utility occurring from θ being witnessed for different DMs.

P_i	DV_{diff}	PI_{diff}
P_1	4.103	3.06
P_2	1128.38	2923.69
P_3	0.157	0.291

7.2 Kullback-Leibler Approach

DMs will often be biased towards their own belief and hence may not wish to assign it the same weight as those of their neighbours in Equation (3.5). They may be inclined to give high weights to neighbours with beliefs akin to their own. We recall that the KL divergence (Kullback & Leibler, 1951) introduced in Section 2.6.3 measured the divergence of two probability distributions, with small values indicating similarity. Hence a DM may want to give high weights to neighbours whose distribution returns a low KL score with her own, and the contrary. As $D(f_i||f_i) = 0$ a participant must choose $\alpha_{i,i}$ to assign to her own distribution. This allows P_i to potentially fully discount the beliefs of neighbours if she allocates $\alpha_{i,i} = 1$. Conversely if she has no confidence in her own opinions she can set $\alpha_{i,i} = 0$. We propose P_i calculates $D(f_i||f_j)$ for each P_j, with $j \neq i$ and then inverts these quantites. Hence if P_j has an opinion similar to P_i then $D(f_i||f_j)$ will be small, and hence its inverse will be large, indicating that P_i considers her to be reliable. The converse also holds. P_i should then rescale these values so weights sum to $1 - \alpha_{i,i}$. Hence higher weights are awarded to those DMs with opinions similar to her own. Formally weights are given by

$$\alpha_{i,j} = \frac{\frac{1}{D(f_i||f_j)}}{\sum_{j \neq i} \frac{1}{D(f_i||f_j)}} \times (1 - \alpha_{i,i}) \text{ for } j \neq i \tag{7.9}$$

Note if P_j has an identical distribution to P_i then P_i may increase her self-weight accordingly, as this appears as evidence favouring her opinion. A decision is made, an outcome observed, and each $f_i(\theta)$ updated in light of this. The procedure is then repeated with KL values calculated for the new opinions of DMs. Note that, by the asymptotic behaviour of distributions discussed in Section 3.4 this will lead to each DM assigning an equal weight to all her neighbours (not including herself) in the limit.

Individuals may also reconsider the weight that they assign themselves, increasing it if they believe themselves to be more accurate than they previously thought, and the contrary. We mention three shortcomings of this approach. Firstly, the KL divergence is not guaranteed to be symmetric (and hence does not meet the formal definition of a distance metric) which could be viewed as undesirable. However in reality the judgment of P_i regarding how similar the opinion of P_j is to her own, may not necessarily be the same as the judgment of P_j regarding how similar the opinion of P_i is to her own. If the user is concerned about asymmetry then alternative symmetric dissimilarity measures could be used, *e.g.*, total variation distance, variational distance or χ^2 distance (Cover & Thomas, 1991). Secondly the choice of $\alpha_{i,i}$ is arbitrary for P_i. While this may appear problematic it gives added subjectivity to the approach, allowing a DM to ignore the opinions of neighbours or indicate her level of confidence in her own expertise. In a non-competing setting, where negative decision consequences only affect the DM, this stubbornness is acceptable. Finally issues arise if a DM has very little faith in her own vague beliefs, and gives herself a low weight. Yet beliefs akin to hers will receive a high weight, while those that are more confident will unintuitively, receive low weights.

We also comment that a DM could potentially initialise her weighting scheme by the KL method and then use the PI method for the subsequent epochs *i.e.*, initially assigning weights in a subjective manner based on the similarity of the opinions of neighbours with her own before proceeding in an objective manner once data has been witnessed.

7.3 Comparing PI, DV and KL approaches

We briefly illustrate the differing results occurring from our three original methods. In reality $\theta = 5$. Suppose $R \sim \mathcal{N}(\theta, 2)$ with all DMs having Normally distributed priors over θ. The PI and DV approaches initialise with the Laplacian Principle of Indifference so, for the sake of comparison, we contrast the resulting posteriors after a single return ($r = 6$) has been witnessed. Hence in Table 7.10 we include respective posterior distributions over θ as well as utility functions, initial fortunes, and the weights that DMs assign their own opinion under the KL method, denoting by $\alpha_{i,i}^{\mathrm{KL}}$ for P_i. P_1 is strongly convinced of the correctness of her opinion while P_3 has little faith in her own opinion.

Table 7.10: DM information for DV/PI/KL comparison.

P_i	$f_i(\theta \vert r=6)$	γ_i	$u_i(\gamma_i + r)$	$\alpha_{i,i}^{\mathrm{KL}}$
P_1	$\mathcal{N}(5, 1.33)$	-15	$(r-15)^2$	0.8
P_2	$\mathcal{N}(6.67, 0.67)$	20	$\log_e(r+20)$	$\frac{1}{3}$
P_3	$\mathcal{N}(7.5, 1)$	10	$r+10$	0.2

The resulting posterior distributions for the various methods are given in Fig. 7.2. We see, for instance that for P_1 it is the KL approach that places the most density on the true θ, *i.e.*, listening heavily to her own opinion transpired to be a good idea as her distribution transpired to be very accurate. In Fig. 7.3 we highlight the difference in the weights allocated by the different DMs under the various techniques. We can observe that in this instance the weights from the PI and DV approach appear to be relatively similar with those allocated by the KL method are strongly contrasting to these. We must recall that when comparing the performances of different approaches it is important to use a metric that takes into account our aspirations, *e.g.*, to maximise posterior density or a utility-based measure.

7.4 Conclusions

In this section we have discussed two subjective methods of ascertaining reliability of information sources. A detailed simulation study demonstrated the merits of the fully objective PI approach. An aim for future research is to derive a similar justification for the methods in this chapter, most notably the DV approach. The KL approach may well yield poor results (in terms of utility) as it allows DMs to completely disregard the opinions of neighbours if they choose to. We may think of this as "unlimited subjectivity". By contrast a DM using the DV approach may choose her utility function but cannot impact further on the weights assigned. This "limited subjectivity" is likely to yield a higher decision quality with objective information witnessed having a more substantive effect on the weights. Finally we note that if P_i does have prior opinions about the accuracy of herself and her neighbours she could subjectively assign these at

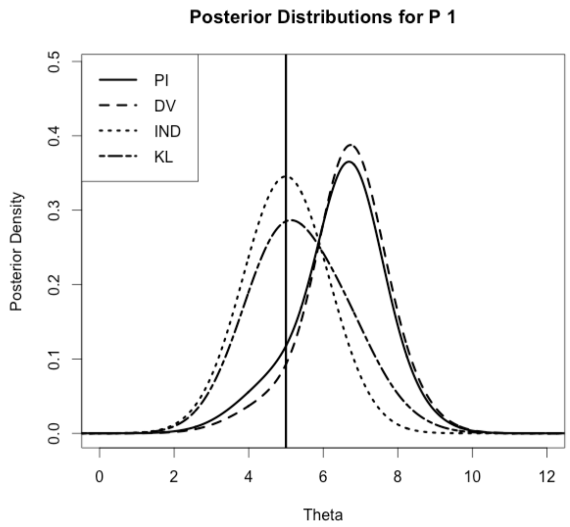

(a)

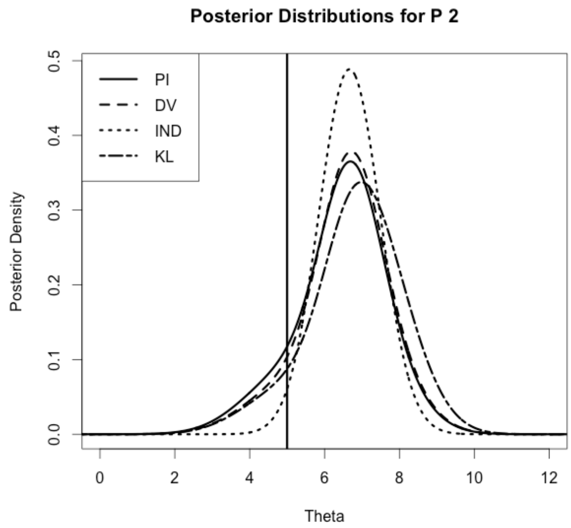

(b)

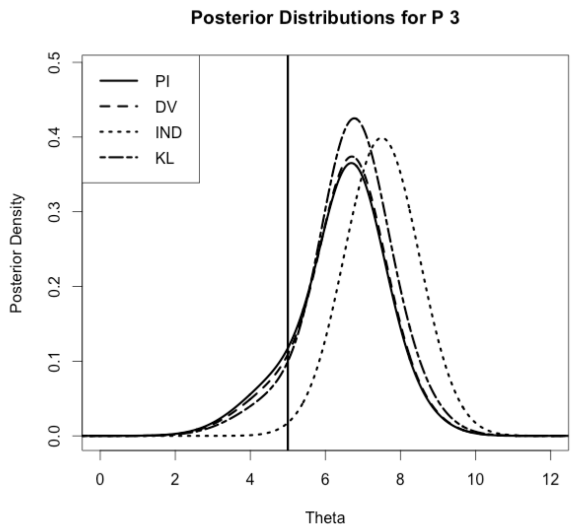

(c)

Fig. 7.2: The posterior distributions for the three DMs respectively, founding using the PI, DV and KL methods. We also include their individual (IND) posterior.

the first epoch (*e.g.*, $\alpha_{i,1} = 0.5, \alpha_{i,2} = 0.3, \alpha_{i,3} = 0.2$), rather than giving equal weights as in Equation (3.5) and proceed using the PI method at all subsequent epochs.

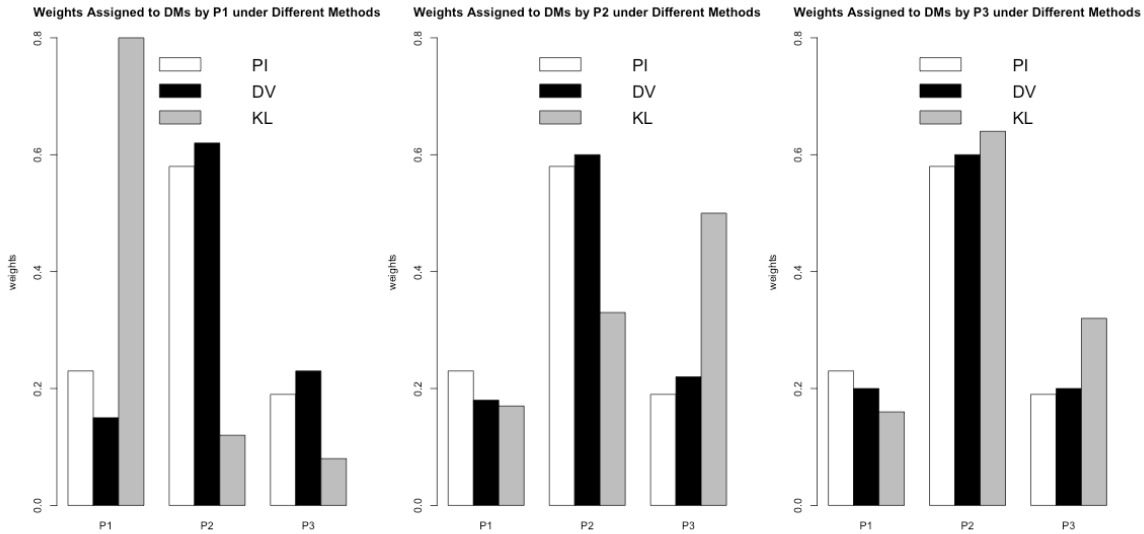

Fig. 7.3: The weights allocated by P_1, P_2 and P_3 respectively (from left to right above) to the opinions of their neighbours under the three different techniques considered, *i.e.*, the PI, DV and KL approaches.

Chapter 8

Summary and Further Research

In this thesis we have developed methods enabling DMs to amalgamate their opinion with (potentially contrasting) opinions from a set of distinct information sources in a fashion that was deemed coherent in some sense. Emphasis was placed on two types of updating: that of the opinions of DMs, and that relating to the perceived accuracy of the neighbours of a DM. In a parametric setting we derived the Plug-in approach. Justifications were provided by attractive Bayesian and coherency properties it obeys, as well as its performance on simulated and real data. We also weakened the need for DMs to supply probability distributions by constructing an analogous framework with uncertainty expressed by simple intervals. We discussed the PI approach in a group decision context, combining it with a utility function that was an equally weighted sum of (suitably rescaled) utility functions of DMs, and noted the relationship between this method and the work of Arrow (1950). Finally we derived two alternative linear opinion pooling methods that both entailed additional subjectivity. Recall the three primary research aims discussed in Chapter 1. These have clearly been achieved over the course of this thesis: we have derived an original decision making methodology that remedied problems inherent within existing techniques, we provided justification for its use in practical contexts using a mix of mathematical arguments and data studies, and we have also provided several generalisations for our approach, increasing its flexibility and hence its applicability.

8.1 Social Networks

An interesting concept would be using the PI approach in a social network framework, *i.e.*, modelling information propagation through a social network. DMs would have neighbours of varying degrees dependent upon their level of connectivity (friends, friends-of-friends, *etc.*) and would have different levels of trust depending on the degree of the neighbour who information was received from. This is an attractive concept, applicable in various realistic settings. Two DMs are neighbours of degree d if they are separated by a minimal series of d dyadic links, and $\mathcal{N}_{i,d}$ is the set of all neighbours of P_i of degree d, *e.g.*, in Fig. 8.1 we have $\mathcal{N}_{3,2} = \{P_1, P_5\}$. This is akin to the standard social network framework discussed in, for instance, Wasserman & Faust (1994). A DM is a neighbour of degree zero with herself. Opinions of DMs are weighted by P_i conditional on the degree to which they are neighbours. Those of degree d have their opinions multiplicatively rescaled by β_i^d with $0 \leq \beta_i \leq 1$, *i.e.*,

$$\hat{f}_i(\theta|\cdot) \quad = \quad \beta_i^0 \alpha_{i,i} f_i(\theta|\cdot) + \beta_i^1 \sum_{j \in \mathcal{N}_{i,1}} \alpha_{i,j} f_j(\theta|\cdot) + \ldots + \beta_i^m \sum_{j \in \mathcal{N}_{i,m}} \alpha_{i,j} f_j(\theta|\cdot) \qquad (8.1)$$

Weights in Equation (8.1) should be normalised, but we leave them in the above form for ease of elucidation. The weights a DM allocates to neighbours are a combination of their "basic weight" (found by the PI method) and their "social weight" (arising from their proximity to the DM). There are two extreme cases: when $\beta_i = 1$ (the DM has no biases based on degree of separation) and when $\beta_i = 0$ (the DM solely heeds her own opinion). This process is subjective with smaller β_i values indicating greater mistrust by P_i for far removed information sources. This is a highly interesting research area with plenty of scope for extension.

8.2 Sequential Problems

In Section 2.5 we touched upon sequential (non-myopic) decision making where DMs must make decisions for multiple future epochs simultaneously without seeing corresponding returns until the end of the process. Methods exist for a single DM involving calculating conditional posterior distributions and "integrating out" uncertainty from right to left using prior information. Fig. 8.2 gives a graphical interpretation of a tree for a problem with two epochs and two potential decisions per epoch. We are

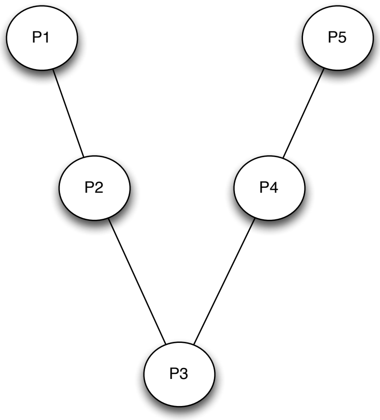

Fig. 8.1: Illustrative social network consisting of five DMs.

interesting in discovering how to solve sequential problems in the setting of this thesis where DMs have numerous information sources, which greatly complicates the standard process. PI weights and the associated normalised weights are conditional upon witnessed returns. Our derivation of the distribution of PI weights in Section 3.7 will be required as without actually observing returns integrals must be solved. We consider the polynomial utility class (Houlding *et al.*, 2015), providing tractable solutions when beliefs are Normally distributed and utility functions are polynomial (or polynomially approximated via Taylor expansions). An aspiration for future work is to solve this problem, increasing the applicability of our PI method. Note that a linear combination of polynomial utility functions is itself polynomial, which will be useful in a group decision making context.

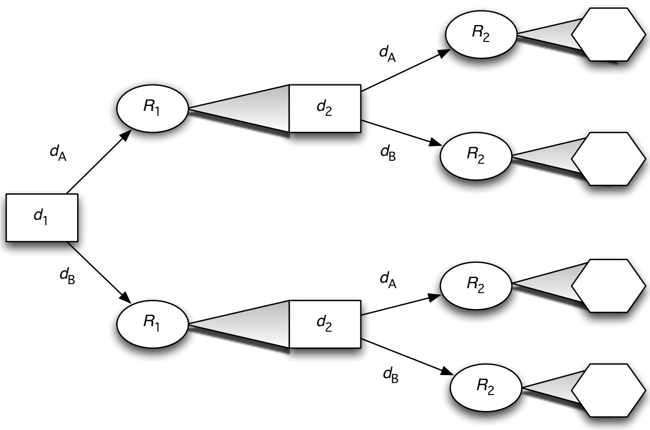

Fig. 8.2: Decision tree for a two-period sequential problem.

8.3 Imprecision of Probabilities and Utilities

We previously discussed how imprecise probabilities allow DMs to more easily express their unsureness in instances where they are not statistically literate, permitting additional uncertainty in decision processes. Another common assumption is that DMs can supply exact utility over possible decision returns, either via discrete values or a continuous function. Yet DMs may often be uncertain of the merits associated with a particular return that they are unfamiliar with, *e.g.*, if they have only eaten octopus once they may lack enough experience to place a precise utility value on the act of doing so again. What if the chef that night had been particularly good/bad? They may instead provide lower and upper utility bounds that they believe their true utility (which will only be derived after multiple experiences) lies within. Note that this subtly differs from NPUI (Houlding & Coolen, 2012) as we consider a DM witnessing noisy realisations of a true utility rather than interpolating from utilities deemed similar. We hope to develop a decision scheme incorporating both imprecise probability (*i.e.*, increased uncertainty over returns from decisions) and imprecise utility (*i.e.*, increased uncertainty about their associated merits). As discussed in Section 2.4.1 this may lead to issues with determining a single decision as optimal, but non-optimal decisions can be eliminated, and if a single decision is deemed maximal under all possible configurations (of both beliefs and utilities) it is extremely robust.

8.4 R Package

The tabulated results and illustrative graphics produced throughout this thesis were created using a collection of original R functions. For the PI approach we created functions calculating distributions and weights under various distributional assumptions, as well as determining if the method was superior to a set of alternatives using both real and simulated data. Functions also calculated the theoretical probabilities of this superiority occurring. In our NPPI setting we have routines taking into account our various updating rules and determining superiority as defined in Section 6.7. Functions for Differing Viewpoints calculations have also been written. An aim is to compile these functions succinctly, ensuring internal consistencies, and creating an R package which individuals who wish to apply the methods described in this thesis can utilise.

8.5 Miscellaneous

We conclude with brief comments on some specific topics within which additional investigation could be conducted.

- In the group decision context of Chapter 5 it was assumed that all DMs had equal standing, *i.e.*, flat hierarchy. Yet this may often not be the case, *e.g.*, Karny & Kracik (2003). There may be different degrees of hierarchy inherent within a group, *i.e.*, a government consisting of a Party leader, ministers, junior ministers, *etc.* All of these may have differing opinions and aspirations, with those of more senior members given more consideration. Could the PI approach be used in this setting with weights assigned in a manner akin to in Equation (8.1)?

- Throughout this thesis we assumed that all DMs are interested in a common θ, being neighbours as this parameter is inherent in all their respective decision tasks. It would increase applicability if our method could be extended to incorporate DMs with correlated, but distinct, uncertainties within their problems, *i.e.*, $\theta_1, \ldots, \theta_n$. For example in a financial context there is an underlying market behavour and different stocks (*i.e.*, different values of θ_i) are correlated to different degrees, *i.e.*, we would imagine that shares in Vodafone and eMobile would exhibit more similar behaviour than shares in Vodafone and Bayern Munich.

- The Differing Viewpoints approach entailed both probability and utility. It is an aim to conduct a full scale simulation study assessing its merits. Doing this was more straightforward in the objective PI approach, but when utility functions must be factored into calculations complexity increases. A logical action would be to find results for a set of DMs with utility functions expressing radically different preferences, ensuring generality of results, and identifying if (for example) the method gives better results for DMs with more extreme utility functions.

- Implicit in this thesis was the assumption that the parameter of interest θ is a scalar, rather than a vector. An area of further investigation would be to explicitly extend the methods presented to those applicable in vector settings, considering the case where the elements of θ could be correlated with each other.

This would require the utility function of a DM to become multi-attributed (*e.g.*, Keeney & Raiffa, 1976).

- Suppose the methods discussed within this thesis were applied in an expert judgment context. In realistic settings a DM would pay experts for providing their opinions. It seems intuitive that the opinion of a highly accurate expert should be considered more valuable than that of an inaccurate expert. This aligns with the concept of value of information (Howard, 1966) which attempts to assign a monetary value to additional information provided to a DM, often contrasting her expected utility before and after receiving this opinion. There appears to be a connection between this and the DV method which, in Equation (7.2), considered a quantity of this ilk. Further investigation could be conducted to examine this relationship in depth and to use it to determine how a DM should divide a fixed sum of money between a collection of domain-specific experts.

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). *Expected Utility Hypothesis and the Allais Paradox*, 27-145. Editors: Allais, M. and Hagen, O. 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. (1963). Uncertainty and the Welfare Economics of Medical Care. *American Economic Review*, 53(5), 941-973.

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

Barlow, R. E., Mensing, R. W. and Smiriga, N. G. (1986). Combination of experts opinions based on decision theory. *Reliability and Quality Control*, 9-19, Basu A. P. (ed.), Amsterdam: North-Holland.

Bates, J. M. & Granger, C. W. J. (1969). The Combination of Forecasts. *Operations Research*, 20(4), 451-468.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, F. R. S. Communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S. *Philosophical Transactions*, 53, 370-418.

Bell, D. E. (1982). Regret In Decision Making Under Uncertainty. *Operations Research*. 30(5), 961-981.

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

Bordley, R. F. and Wolff, R. W. (1981). On the aggregation of individual probability estimates. *Management Science*, 27, 959-964.

Boutilier, C. (2003). On the Foundations of *Expected* Expected Utility. *IJCAI'03 Proceedings of the 18th International Joint Conference on Artificial Intelligence*, 285-290.

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, Valencia, Spain.

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

Bunn, D. W. (1975). A Bayesian Approach to the Linear Combination of Forecasts (1975). *Operational Research Quarterly*, 26(2), 1, 325-329.

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. (2008). Comments on Cooke's Classical Method. *Reliability Engineering and System Safety*, 93, 760-765.

Clemen, R. T. & Winkler, R. L. (1999). Combining Probability Distributions 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.

Cooke, R. & Goossens, L. L. H. J. (2008). TU Delft Expert Judgement Data Base. *Reliability Engineering and System Safety*, 93, 657-674.

Coolen, F. P. A. (1996). Comparing two Populations based on Low Stochastic Structure Assumptions. *Statistics and Probability Letters*, 29, 297-305.

Coolen, F. P. A. (2006). On nonparametric predictive inference and objective

Bayesianism. *Journal of Logic, Language and Information*, 15, 21-47.

Coolen, F. P. A., Troffaes, C.M. & Augustin, T. (2010). Imprecise Probability. *International Encyclopedia of Statistical Science*, Springer.

Condorcet, Marquis de (1785). Essai sur l'Application de l'Analyse a la Probabilite des Decisions Rendues a la Pluralite des Voix. L'Imprimerie Royale.

Cover, T. M. & Thomas, J. A. (1991). Elements of Information Theory. John Wiley & Sons Inc.

Cox, L. A., T. Jr. (2012). Confronting Deep Uncertainties in Risk Analysis. *Risk Analysis*, 32(10), 1607-1629.

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.

Dawid, A. P., DeGroot, M. H., Mortera, J., Cooke, R., French, S., Genest, C., Schervish, M. J., Lindley, D. V., McConway, K. J. & Winkler, R. L. (1995). Coherent Combinations of Experts' Opinions. *Test*, 4(2), 263-313.

de Finetti, B. (1974). Theory of Probability: a Critical Introductory Treatment, Vols. 1 and 2. New York, Wiley.

DeGroot, M. H. & Mortera, J. (1991). Optimal Linear Opinion Pools. *Management Science*, 37(5), 546-558.

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.

Eggstaff, J. W., Mazzuchi, T. A. & Sarkana, S. (2013). The Development of Progress Plans Using a Performance-Based Expert Judgment Model to Assess Technical Performance and Risk. *Systems Engineering*, 17(4), 375-391.

Eggstaff, J. W., Mazzuchi, T. A. & Sarkana, S. (2014). The Effect of the Number of Seed Variables on the Performance of Cooke's Classical Model. *Reliability Engineering and System Safety*, 121, 72-82.

Fishburn, P. C. (1982). The Foundations of Expected Utility. Theory & Decision

Library.

Flandoli, F., Giorgi, E., Aspinall, W. P. & Neri, A. (2011). Comparison 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. (1985). Group Consensus Probability Distributions: A Critical Study. Bayesian Statistics 2, 183-197.

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.

Genest, C. & McConway, K. J. (1990). Allocating the Weights in the Linear Opinion Pool. *Journal of Forecasting*, 9, 53-73.

Genest, C., Weerahandi, S. & Zidek, J. V. (1984). Aggregating Opinions through Logarithmic Pooling. *Theory and Decision*, 17, 61-70.

Genest, C. & Zidek, J. V. (1986). Combining Probability Distributions: A Critique and an Annotated Bibliography. *Statistical Science*, 1(1), 114-135.

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

Gneiting, T. & Raftery, A. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102, 359-378.

Goldstein, M. (2001). Avoiding Forgone Conclusions: Geometric and Foundational Analysis of Paradoxes of Finite Additivity". *Journal of Statistical Planning and Inference*, 94, 73-87.

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

Heskes, T. (1998). Selecting Weighting Factors in Logarithmic Opinion Pools. Advances in Neural Information Processing Systems, The MIT Press, 266-272.

Hill, B. M. (1968). Posterior Distributions of Percentiles: Bayes' Theorem for Sampling from a Population. *Journal of the American Statistical Association*, 63, 677-691.

Hill, B. M. (1988). De Finetti's Theorem, Induction and $A_{(n)}$ or Bayesian Nonparametric Inference. *Bayesian Statistics*, 3, 211-241.

Hill, B. M. (1993). Parametric Models for $A_{(n)}$: Splitting Processes and Mixtures. *Journal of the Royal Statistical Society*, B, 55, 423-433.

Hoeting, J. A., Madigan, D., Raftery, A. & Volinsky, C. T. (1999). Bayesian Model Averaging: A Tutorial. *Statistical Science*, 14(4), 382-417.

Houlding B. & Coolen F. P. A. (2007). Sequential Adaptive Utility Decision Making for System Failure Correction. *Journal of Risk and Reliability* (Proc. IMechE, Part O), 221(4), 285-295.

Houlding, B. (2008). Sequential Decision Problems with Adaptive Utility. PhD Thesis.

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.

Houlding, B., Coolen, F. P. A. & Bolger, D. (2015). A Conjugate Class of Utility Functions for Sequential Decision Problems. *Risk Analysis*, doi: 10.1111/risa.1235.

Howard, R. A. (1966). Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics*, 22-26.

Jouini, M. N. & Clemen, R. T. (1996). Copula Models for Aggregating Expert Opinions. *Operations Research*, 44(3), 444-457.

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. *International Journal of 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.

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

Keeney, R. L. & Raiffa, H. (1976). Decisions with Multiple Objectives: Preference and Value Tradeoffs. Wiley Series in Probability and Mathematical Statistics.

Kimball, M. (1990). Precautionary saving in the Small and in the Large. *Econo-*

metrica, 58, 53- 73.

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

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

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

Lee, P. M. (2012). Bayesian Statistics: An Introduction. Fourth Edition, John Wiley and Sons.

Levi, I. (1974). On Indeterminate Probabilities. *Journal of Philosophy*, 71(13), 391-418.

Lindley, D. V. (1991). Making Decisions (2nd Edition). Wiley & Sons Ltd.

Lindley, D.V. & Phillips, L. D. (1976). Inference for a Bernoulli process (a Bayesian view). *American Statistician*, 30, 112-119.

Lindley, D. V. & Singpurwalla, N. D. (1986). Reliability (and Fault Tree) Analysis using Expert Opinions. *Journal of the American Statistical Society*, 81(393), 87-90.

Litchendahl, K. C. Jr., Grushka-Cockayne, Y. & Winkler, R. L. (2013). Is It Better to Average Probabilities or Quantiles? *Management Science*, 59(7), 1594-1611.

Loomes, G. & Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *Economic Journal*, 92(4), 805-24.

Maher P. (1993). Betting on Theories. New York: Cambridge University Press.

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

McLachlan, G. & Peel. D. (2000). Finite Mixture Models. Wiley.

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

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N. & Teller, A. H. (1953). Equations of State Calculations by Fast Computing Machines. *Journal of Chemical Physics*, 21(6), 1087-1092.

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

Newbold, P. & Granger, C. W. J. (1974). Experience with Forecasting Univariate Time Series and the Combination of Forecasts. *Journal of the Royal Statistical Society, Series A*, 137(2), 131-165.

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.

Raftery, A. E. (1995). Bayesian Model Selection in Social Research. *Sociological Methodology*, 111-195.

Roberts, H. V. (1965). Probabilistic Prediction. *Journal of the American Statistical Association*, 60, 50-62.

Roelofs, V. J., Coolen, F. P. A. & Hart, A. D. M. (2011). Nonparametric Predictive Inference for Exposure Assessment. *Risk Analysis*, 31(2), 218-227.

Rufo, M. J., Martin, J. & Perez, C.J. (2009). Inference on Exponential Families with Mixture of Prior Distributions. *Computational Statistics and Data Analysis*, 53, 3271-3280.

Rufo, J. M. & Perez, C. J. (2012). A Bayesian Approach to Aggregate Experts' Initial Information. *Electronic Journal of Statistics*, 6, 2362-2382.

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, 187-217.

Schervish, M. J., Seidenfeld, T., Kadane, J. B. & Levi, I. (2003). Extensions of Expected Utility and Some Limitations of Pairwise Comparisons. *Proceedings in Informatics*, 18, 496-510.

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 Behavioural 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.

Skyrms, B. (1993). A Mistake in Dynamic Coherency Arguments? *Philosophy of Science*, 60(2), 320-328.

Starr, C. (1981). Risk Criteria for Nuclear Power Plants: A Pragmatic Proposal. *Risk Analysis*, 1(2), 113-120.

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.

Stone, M. (1961). The Opinion Pool. *Annals of Mathematical Statistics*, 32, 1339-1342.

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

Wald, A. (1950). Statistical Decision Functions. John Wiley & Sons.

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

Wasserman, S. & Faust, K. (1994). Social Network Analysis: Methods and Applications. Cambridge, ENG and New York: Cambridge University Press.

Weichselberger, K. (2000). The Theory of Interval Probability as a Unifying Concept for Uncertainty. *International Journal of Approximate Reasoning*, 24, 149-170.

Weichselberger, K. (2001). Elementare Grundbegriffe einer Allgmeineren Wahrscheinlichkeitsrechnung 1. Intervallwahrscheinlichkeit als Umfassendes Konzept. Heidelberg, Physika.

Wilson, S. & McDaid, K. (2001). Deciding how long to test Software. *The Statistician*, 50(2), 117-134.

Whittle, P. (1992). Probability via Expectation. Third Edition, Springer Texts in Statistics.

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

Appendix A

Notation

Below we present the primary notation that is consistently used throughout this thesis.

In a setting with a single DM with a parametric opinion

- $\mathcal{D}$ - the set of admissible decisions.
- $d_1, \ldots, d_n \in \mathcal{D}$ - a collection of n potential admissible decisions.
- $d^* \in \mathcal{D}$ - the decision maximising expected utility.
- $\mathcal{R}$ - the set of potential decision returns.
- $r_1, , \ldots, r_k \in \mathcal{R}$ - a collection of k possible decision returns.
- r^*, r_* - the best and worst outcomes respectively in $\mathcal{R}$.
- γ - the fortune of the DM.
- $u(r)$ or $u(\gamma + r)$ - the utility function of the DM (which may involve her fortune).
- θ - some latent unknown parameter or event 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.

In an information sharing setting where DMs have parametric opinions

- 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 over θ after listening to neighbours.
- γ_i - the fortune of P_i.
- $u_i(r)$ or $u_i(\gamma_i + r)$ - the utility function of P_i (which may involve her fortune).
- w_j - the Plug-in weight given to P_j (by all neighbours).
- $v_{i,j}$ - the Differing Viewpoints weight given by P_i to P_j.
- $\alpha_{i,j}$ - the normalised weight given by P_i to P_j in $\hat{f}_i(\theta)$.

In a group decision setting with parametric opinions

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

In an information sharing setting where DMs have nonparametric opinions

- $\underline{\mathbb{E}}_{t,i}(\theta)$ - the lower bound of the NPPI of P_i after t epochs.
- $\bar{\mathbb{E}}_{t,i}(\theta)$ - the upper bound of the NPPI of P_i after t epochs.
- $\Delta_{t,i}(\theta)$ - the width of the NPPI of P_i after t epochs.
- $\hat{\underline{\mathbb{E}}}_{t,i}(\theta)$ - the lower bound of the combined NPPI after t epochs.
- $\hat{\bar{\mathbb{E}}}_{t,i}(\theta)$ - the upper bound of the combined NPPI after t epochs.
- $\hat{\Delta}_{t,i}(\theta)$ - the width of the combined NPPI after t epochs.

Appendix B

Proofs

This appendix contains proof of several theorems stated in the main text of this thesis.

Proof of Statement in Section 2.6

We prove that a pairwise iterative sharing of distributions, as suggested by Karny & Guy (2004), is not invariant to the order of sharing. For three DMs we consider two different sharing orderings, observing different results in each case.

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.

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

- *Stage 3*: P_2 and P_3 sharing.

$$\hat{f}_2(\theta) = \alpha_2 \hat{f}_2(\theta) + (1-\alpha_2)\hat{f}_3(\theta)$$

$$
\begin{aligned}
&= f_2(\theta)\Big(\alpha_2^2 + (1-\alpha_3)(1-\alpha_1)(1-\alpha_2)\Big) + f_3(\theta)\Big(\alpha_3(1-\alpha_2)\Big) \\
&\quad +f_1(\theta)\Big(\alpha_2(1-\alpha_2) + \alpha_1(1-\alpha_3)(1-\alpha_2)\Big) \qquad \text{(B.2)} \\
\hat{f}_3(\theta) &= \alpha_3\hat{f}_3(\theta) + (1-\alpha_3)\hat{f}_2(\theta) \\
&= f_2(\theta)\Big(\alpha_3(1-\alpha_3)(1-\alpha_1) + \alpha_2(1-\alpha_3)\Big) \\
&\quad +f_1(\theta)\Big(\alpha_1\alpha_3(1-\alpha_3) + (1-\alpha_3)(1-\alpha_2)\Big) + f_3(\theta)\Big(\alpha_3^2\Big) \qquad \text{(B.3)}
\end{aligned}
$$

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.

$$
\begin{aligned}
\hat{f}_1(\theta) &= \alpha_1 f_1(\theta) + (1-\alpha_1)\hat{f}_3(\theta) \\
&= \alpha_1 f_1(\theta) + \alpha_3(1-\alpha_1)f_3(\theta) + (1-\alpha_1)(1-\alpha_3)f_2(\theta) \\
\hat{f}_3(\theta) &= \alpha_3\hat{f}_3(\theta) + (1-\alpha_3)f_1(\theta) \\
&= \alpha_3^2 f_3(\theta) + \alpha_3(1-\alpha_3)f_2(\theta) + (1-\alpha_3)f_1(\theta) \qquad \text{(B.4)}
\end{aligned}
$$

- *Stage 3*: P_2 and P_3 sharing.

$$
\begin{aligned}
\hat{f}_1(\theta) &= \alpha_1\hat{f}_1(\theta) + (1-\alpha_1)\hat{f}_2(\theta) \\
&= f_1(\theta)\Big(\alpha_1^2\Big) + f_2(\theta)\Big(\alpha_1(1-\alpha_1)(1-\alpha_3) + \alpha_2(1-\alpha_1)\Big) \\
&\quad +f_3(\theta)\Big(\alpha_1\alpha_3(1-\alpha_1) + (1-\alpha_1)(1-\alpha_2)\Big) \qquad \text{(B.5)} \\
\hat{f}_2(\theta) &= \alpha_2\hat{f}_2(\theta) + (1-\alpha_2)\hat{f}_1(\theta) \\
&= f_1(\theta)\Big(\alpha_1(1-\alpha_2)\Big) + f_2(\theta)\Big(\alpha_2^2 + (1-\alpha_2)(1-\alpha_1)(1-\alpha_3)\Big) \\
&\quad +f_3(\theta)\Big(\alpha_2(1-\alpha_2) + \alpha_3(1-\alpha_2)(1-\alpha_1)\Big) \qquad \text{(B.6)}
\end{aligned}
$$

Examining coefficients we see that Equation $(B.1)$ is not equal to Equation $(B.5)$, that Equation $(B.2)$ is not equal to Equation $(B.6)$, and that Equation $(B.3)$ is not equal to Equation $(B.4)$, *i.e.*, the beliefs of P_1, P_2 and P_3 respectively differ depending on the order that they receive information from neighbours.

Proof of Statement in Section 3.5

This is an inductive proof that for the PI approach $\hat{f}_i(\theta|\cdot)$ is the same for $i = 1, \ldots, n$ after seeing returns $r_1, \ldots, r_m$. We denote by r_i the i^{th} return witnessed.

- **Proof for** $m = 0$: Having seen zero returns all DMs are assumed equally reliable, *i.e.*, $\alpha_{i,j} = \frac{1}{n}$ for all $i, j = 1, \ldots, n$. The combined belief of P_i will be the same for each DM as they have all received the same distributions $f_1(\theta), \ldots, f_n(\theta)$, *i.e.*,

$$\hat{f}_i(\theta) = \frac{1}{n} f_1(\theta) + \ldots + \frac{1}{n} f_n(\theta) \text{ for all } i = 1, \ldots, n \tag{B.7}$$

- **Assumption for** $m = k$: We assume, having seen k returns, that $\hat{f}_i(\theta|r_1, \ldots, r_k)$ is the same for all DMs, *i.e.*, the combined identical belief for each P_i

$$\hat{f}_i(\theta|r_1, \ldots, r_k) = \alpha_{i,1} f_1(\theta|r_1, \ldots, r_k) + \ldots + \alpha_{i,n} f_n(\theta|r_1, \ldots, r_k) \tag{B.8}$$

- **Proof for** $m = k + 1$. We assume distributions are identical after k returns. Following the decision made at epoch $k + 1$, a return r_{k+1} is seen and the probability attached by each individual's distribution to this outcome is calculated, *i.e.*, $w_1 = f_1(R = r_{k+1}|r_1, \ldots, r_k), \ldots, w_n = f_n(R = r_{k+1}|r_1, \ldots, r_k)$. These values yield the same results regardless of which DM performs the calculations, *i.e.*, objectivity. Equation (3.11) is used to update weights, giving the combined distribution of P_i in Equation $(B.9)$. This will be the same for all $i = 1, \ldots, n$ as the weights are calculated in the same manner for all DMs and all participants have been told the same updated distributions $f_1(\theta|r_1, \ldots, r_k), \ldots, f_n(\theta|r_1, \ldots, r_k)$. Therefore, after $k + 1$ returns all DMs have identical combined beliefs, as required.

$$\hat{f}_i(\theta|r_1, \ldots, r_{k+1}) = \alpha^*_{i,1} f_1(\theta|r_1, \ldots, r_{k+1}) + \ldots + \alpha^*_{i,n} f_n(\theta|r_1, \ldots, r_{k+1}) \tag{B.9}$$

Proof of Properties in Section 3.5

- **Proof of Property 1:** This follows trivially from the non-negativity of prior predictive distributions. In addition w_j will only be zero in the case of P_j placing no probability density on the outcome that was observed occurring. Conversely, if P_j places no probability density on the outcome that was observed occurring then $w_j = 0$.

- **Proof of Property 2:** Given that $\alpha_{i,j} < \alpha_{i,k}$ and $w_j < w_k$ it is evident that $w_j\alpha_{i,j} < w_k\alpha_{i,k}$. Division by the same positive (scaling) constant will not change this inequality, giving us $\alpha^*_{i,j} = \frac{w_j\alpha_{i,j}}{\sum_{l=1}^{n} w_l\alpha_{i,l}} < \frac{w_k\alpha_{i,k}}{\sum_{l=1}^{n} w_l\alpha_{i,l}} = \alpha^*_{i,k}$.

- **Proof of Property 3:** This follows from the proof of Property 2 with inequalities replaced by equalities.

- **Proof of Property 4:** Numerical examples illustrate proof of this proposition. Suppose P_j and P_k are the only DMs in the process, with $\alpha_{i,j} = 0.4$ and $\alpha_{i,k} = 0.6$.

  - If $w_j = 0.1$ and $w_k = 0.05$ we have $\alpha^*_{i,j} = 0.571$ and $\alpha^*_{i,k} = 0.429$, *i.e.*, $\alpha^*_{i,j} > \alpha^*_{i,k}$.
  - If $w_j = 0.3$ and $w_{i,k} = 0.2$ we have $\alpha^*_{i,j} = 0.5$ and $\alpha^*_{i,k} = 0.5$, *i.e.*, $\alpha^*_{i,j} = \alpha^*_{i,k}$.
  - If $w_j = 0.1$ and $w_k = 0.02$ we have $\alpha^*_{i,j} = 0.769$ and $\alpha^*_{i,k} = 0.231$, *i.e.*, $\alpha^*_{i,j} < \alpha^*_{i,k}$.

Proof of Statement in Section 5.4

Suppose that we have two decision makers (P_1 and P_2) who must choose between m distinct courses of action, $d_1, \ldots, d_m$. Their expected utilities are given in Table B.1 with $a_i, b_i \in [0, 1]$ for all $i = 1, \ldots, m$. If $\succ_1$ and $\succ_2$ both have a highest preference for the same decision then there is no need for either DM to be a dictator, *i.e.*, their preference is automatically deemed most favourable in the group ranking. Hence we assume that their highest preferences are for different decisions, or, more formally, that $\arg\max_{d\in\mathfrak{D}} \mathbb{E}[u^*_1(d)] \neq \arg\max_{d\in\mathfrak{D}} \mathbb{E}[u^*_2(d)]$.

Table B.1: Original Expected Utilities of P_1, P_2 and the Group (P^*).

	P_1	P_2	P^*
$\mathbb{E}[u^*_i(d_1)]$	a_1	b_1	$\frac{a_1+b_1}{2}$
$\mathbb{E}[u^*_i(d_2)]$	a_2	b_2	$\frac{a_2+b_2}{2}$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\mathbb{E}[u^*_i(d_m)]$	a_m	b_m	$\frac{a_m+b_m}{2}$

Suppose without loss of generality that $a_1 > a_j$ for all $j = 2, \ldots, n$ (*i.e.*, the most preferred decision for P_1 is d_1) and that $\arg\max_{d\in\mathfrak{D}} \mathbb{E}[u^*_2(d)] \neq d_1$ (*i.e.*, d_1 is not the preferred decision of P_2). We also suppose that $\arg\max_{d\in\mathfrak{D}} \mathbb{E}[u^*(d)] \neq d_1$, *i.e.*, the optimal decision for the group as a whole is not the preferred one (d_1) of P_1. For P_1 to be a dictator she must manipulate $\{a_1, \ldots, a_n\}$ to ensure that $\arg\max_{d\in\mathfrak{D}} \mathbb{E}[u^*(d)] = d_1$. If P_1 sets $\{a_1, \ldots, a_n\}$ so that $a_1 = 1$ and $a_j = 0$ for all $j \neq 1$, *i.e.*, $\mathbb{E}[u_1^*(d_1)] = 1$ and $\mathbb{E}[u_1^*(d_j)] = 0$ for all $j \neq 1$ then this leads to the following:

$$\mathbb{E}[u^*(d_1)] = \frac{1+b_1}{2} \tag{B.10}$$

$$\mathbb{E}[u^*(d_j)] = \frac{0+b_j}{2} \text{ for all } j \neq 1 \tag{B.11}$$

We see that Equation (B.10) is greater than Equation (B.11) if $b_1 > 0$. If $b_1 = 0$ and there is a j such that $b_j = 1$ then this leads to a tie, *i.e.*, the respective preferred decisions of P_1 and P_2 will be equally preferred in the group ranking. Therefore we see that, in a setting with two DMs, P_1 will always be able to adjust her preferences to ensure that her favoured decision is at least as favoured as all others in the group ranking. Hence she can be seen as capable of being a dictator, as, by the symmetry inherent within the problem, can P_2. Hence a scenario with two DMs will always have individuals capable of manipulability.

Proof of Properties for DV approach

Here we formally prove four coherency properties for the DV approach. These are analogous to those proved for the PI approach, introduced in Section 3.5.

- **Property 1:** $v_{i,j} \geq 0$ for all $i, j = 1, \ldots, n$ with $v_{i,j} = 0$ if and only if we have $u_i(r) = \mathbb{E}_{i|j}[u_i(d^*)]$.
- **Proof of Property 1:** Non-negativity holds trivially from the non-negativity of absolute values. We see that if $v_{i,j} = 0$ then this implies $|u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]| = 0$, which only holds if $u_i(r) = \mathbb{E}_{i|j}[u_i(d^*)]$. Conversely if $u_i(r) = \mathbb{E}_{i|j}[u_i(d^*)]$ then $|u_i(r) - \mathbb{E}_{i|j}[u_i(d^*)]| = 0$, *i.e.*, $v_{i,j} = 0$.
- **Property 2:** If $\alpha_{i,j} < \alpha_{i,k}$ and $v_{i,j} > v_{i,k}$ then $\alpha^*_{i,j} < \alpha^*_{i,k}$.
- **Proof of Property 2:** Given that $\alpha_{i,j} < \alpha_{i,k}$ and $w_{i,j} > w_{i,k}$ it is evident that $\frac{\alpha_{i,j}}{v_{i,j}} < \frac{\alpha_{i,k}}{v_{i,k}}$. Division by a common positive scaling value will preserve this inequality, *i.e.*, $\alpha^*_{i,j} = \frac{\frac{\alpha_{i,j}}{v_{i,j}}}{\sum_{l=1}^{n} \frac{\alpha_{i,l}}{v_{i,l}}} < \frac{\frac{\alpha_{i,k}}{v_{i,k}}}{\sum_{l=1}^{n} \frac{\alpha_{i,l}}{v_{i,l}}} = \alpha^*_{i,k}$

- **Property 3:** If $\alpha_{i,j} = \alpha_{i,k}$ and $v_{i,j} = v_{i,k}$ then $\alpha^*_{i,j} = \alpha^*_{i,k}$.
- **Proof of Property 3:** The follows from the proof of Property 2, with inequalities being replaced by equalities.
- **Property 4:** If $\alpha_{i,j} < \alpha_{i,k}$ and $v_{i,j} < v_{i,k}$ then all of the following may occur, depending on differences between initial weights, and updated reliability measures:
  - $\alpha^*_{i,j} < \alpha^*_{i,k}$
  - $\alpha^*_{i,j} = \alpha^*_{i,k}$
  - $\alpha^*_{i,j} > \alpha^*_{i,k}$
- **Proof of Property 4:** Numerical examples illustrate the proof of this proposition. Suppose P_j and P_k are the only two DMs in the process, with $\alpha_{i,j} = 0.4$ and $\alpha_{i,k} = 0.6$. We can let $i = j$ or $i = k$.
  - If $v_{i,j} = 0.25$ and $v_{i,k} = 0.3$ then we have $\alpha^*_{i,j} = 0.444$ and $\alpha^*_{i,k} = 0.556$, *i.e.*, $\alpha^*_{i,j} < \alpha^*_{i,k}$.
  - If $v_{i,j} = 0.4$ and $v_{i,k} = 0.6$ then we have $\alpha^*_{i,j} = 0.5$ and $\alpha^*_{i,k} = 0.5$, *i.e.*, $\alpha^*_{i,j} = \alpha^*_{i,k}$.
  - If $v_{i,j} = 0.05$ and $v_{i,k} = 0.1$ then we have $\alpha^*_{i,j} = 0.57$ and $\alpha^*_{i,k} = 0.43$, *i.e.*, $\alpha^*_{i,j} > \alpha^*_{i,k}$.

Appendix C

Simulation Study Results

In Tables C.1-C.9 we include success proportions for the nine cases of the simulation study in Section 4.2. For brevity we only include a subset in each instance, showing results for initialisations with between two and ten DMs having witnessed between two and five returns. The superior technique has its success proportion highlighted in each instance. The behaviour of techniques is clear from the proportions below.

Table C.1: Success proportions: Normal overestimation in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1170	0.1458	**0.7372**	2	3	0.1902	0.1398	**0.6700**
3	2	0.2628	0.1308	**0.6064**	3	3	0.3038	0.1344	**0.5618**
4	2	0.3168	0.1320	**0.5512**	4	3	0.3636	0.1204	**0.5160**
5	2	0.3748	0.1348	**0.4904**	5	3	0.4072	0.1326	**0.4602**
6	2	0.3850	0.1356	**0.4794**	6	3	0.4324	0.1312	**0.4364**
7	2	0.4154	0.1238	**0.4608**	7	3	**0.4388**	0.1228	0.4384
8	2	0.4180	0.1360	**0.4460**	8	3	**0.4656**	0.1214	0.4130
9	2	0.4338	0.1276	**0.4386**	9	3	**0.4824**	0.1246	0.3930
10	2	**0.4394**	0.1312	0.4294	10	3	**0.4838**	0.1094	0.4068
2	4	0.2276	0.1470	**0.6254**	2	5	0.2640	0.1386	**0.5974**
3	4	0.3306	0.1298	**0.5396**	3	5	0.3476	01344	**0.5180**
4	4	0.3974	0.1344	**0.4682**	4	5	0.4158	0.1398	**0.4444**
5	4	0.4378	0.1212	**0.4410**	5	5	**0.4324**	0.1362	0.4314
6	4	**0.4400**	0.1228	0.4372	6	5	**0.4610**	0.1346	0.4044
7	4	**0.4758**	0.1264	0.3978	7	5	**0.4856**	0.1216	0.3928
8	4	**0.4788**	0.1300	0.3912	8	5	**0.4966**	0.1264	0.3770
9	4	**0.4816**	0.1326	0.3858	9	5	**0.5120**	0.1246	0.3634
10	4	**0.4860**	0.1280	0.3860	10	5	**0.5088**	0.1300	0.3612

In Tables C.10-C.18 we document success proportions for the PI approach in the individual problem, highlighting cases where it is superior to individual distributions.

Table C.2: Success proportions: Normal understimation in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1372	0.1992	**0.6636**	2	3	0.1818	0.1760	**0.6422**
3	2	0.2762	0.1680	**0.5558**	3	3	0.3284	0.1636	**0.5088**
4	2	0.3338	0.1754	**0.4908**	4	3	0.3772	0.1608	**0.4620**
5	2	0.3706	0.1624	**0.4670**	5	3	0.4142	0.1600	**0.4258**
6	2	0.3888	0.1582	**0.4530**	6	3	0.4190	0.1590	**0.4220**
7	2	0.3938	0.1646	**0.4416**	7	3	**0.4410**	0.1606	0.3984
8	2	0.4070	0.1672	**0.4258**	8	3	**0.4392**	0.1660	0.3948
9	2	0.4114	0.1654	**0.4232**	9	3	**0.4586**	0.1492	0.3922
10	2	**0.4240**	0.1674	0.4086	10	3	**0.4680**	0.1484	0.3836
2	4	0.2154	0.1708	**0.6138**	2	5	0.2444	0.1612	**0.5864**
3	4	0.3398	0.1608	**0.4994**	3	5	0.3622	0.1582	**0.4826**
4	4	0.3914	0.1574	**0.4512**	4	5	0.3958	0.1646	**0.4396**
5	4	**0.4276**	0.1530	0.4194	5	5	**0.4300**	0.1592	0.4108
6	4	**0.4378**	0.1550	0.4072	6	5	**0.4470**	0.1510	0.4020
7	4	**0.4602**	0.1540	0.3858	7	5	**0.4726**	0.1500	0.3774
8	4	**0.4600**	0.1532	0.3868	8	5	**0.4726**	0.1480	0.3764
9	4	**0.4784**	0.1564	0.3652	9	5	**0.4794**	0.1524	0.3682
10	4	**0.4714**	0.1602	0.3654	10	5	**0.4966**	0.1558	0.3476

Table C.3: Success proportions: Normal mean-centred in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1500	0.1774	**0.6276**	2	3	0.2086	0.1666	**0.6248**
3	2	0.2652	0.1600	**0.5748**	3	3	0.3266	0.1590	**0.5144**
4	2	0.3374	0.1606	**0.5020**	4	3	0.4028	0.1520	**0.4452**
5	2	0.4728	0.1510	**0.4762**	5	3	**0.4410**	0.1396	0.4164
6	2	0.4066	0.1466	**0.4468**	6	3	**0.4620**	0.1422	0.3958
7	2	0.4214	0.1474	**0.4312**	7	3	**0.4822**	0.1390	0.3788
8	2	**0.4302**	0.1450	0.4248	8	3	**0.4858**	0.1346	0.3796
9	2	**0.4344**	0.1354	0.4302	9	3	**0.4846**	0.1346	0.3808
10	2	**0.4502**	0.1334	0.4164	10	3	**0.5004**	0.1352	0.3644
2	4	0.2378	0.1522	**0.6100**	2	5	0.2734	0.1620	**0.5646**
3	4	0.3684	0.1530	**0.4786**	3	5	0.3788	0.1556	**0.4656**
4	4	**0.4284**	0.1530	0.4186	4	4	**0.4402**	0.1470	0.4128
5	4	**0.4752**	0.1434	0.3814	5	5	**0.5058**	0.1336	0.3606
6	4	**0.4972**	0.1376	0.3652	6	5	**0.5130**	0.1424	0.3446
7	4	**0.5124**	0.1370	0.3506	7	5	**0.5442**	0.1308	0.3250
8	4	**0.5094**	0.1400	0.3506	8	5	**0.5402**	0.1442	0.3156
9	4	**0.5402**	0.1262	0.3336	9	5	**0.5522**	0.1290	0.3188
10	4	**0.5460**	0.1292	0.3248	10	5	**0.5602**	0.1398	0.3000

Table C.4: Success proportions: Poisson overestimation in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1260	0.2044	**0.6696**	2	3	0.1588	0.1796	**0.6616**
3	2	0.2648	0.1934	**0.5418**	3	3	0.3216	0.1676	**0.5108**
4	2	0.3302	0.1902	**0.4796**	4	3	0.3760	0.1776	**0.4464**
5	2	0.3726	0.1872	**0.4402**	5	3	**0.4426**	0.1786	0.3988
6	2	0.3808	0.1908	**0.4284**	6	3	**0.4378**	0.1718	0.3904
7	2	0.4000	0.1882	**0.4118**	7	3	**0.4508**	0.1784	0.3708
8	2	**0.4130**	0.1838	0.4032	8	3	**0.4600**	0.1708	0.3692
9	2	**0.4328**	0.1768	0.3904	9	3	**0.4964**	0.1710	0.3596
10	2	**0.4208**	0.1770	0.4022	10	3	**0.4742**	0.1638	0.3620
2	4	0.1964	0.1762	**0.6724**	2	5	0.2096	0.1696	**0.6208**
3	4	0.3466	0.1734	**0.4800**	3	5	0.3078	0.1654	**0.4638**
4	4	0.4026	0.1802	**0.4172**	4	5	**0.4158**	0.1724	0.4118
5	4	**0.4384**	0.1768	0.3848	5	5	**0.4746**	0.1752	0.3502
6	4	**0.4724**	0.1660	0.3616	6	5	**0.4798**	0.1724	0.3478
7	4	**0.4858**	0.1746	0.3234	7	5	**0.5330**	0.1620	0.3050
8	4	**0.5020**	0.1746	0.3234	8	5	**0.5330**	0.1620	0.3050
9	4	**0.5000**	0.1766	0.3234	9	5	**0.5222**	0.1694	0.3084
10	4	**0.5148**	0.1632	0.3220	10	5	**0.5336**	0.1540	0.3104

Table C.5: Success proportions: Poisson underestimation in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1456	0.2224	**0.6320**	2	3	0.2070	0.1996	**0.5934**
3	2	0.2722	0.2040	**0.5238**	3	3	0.3204	0.1892	**0.4904**
4	2	0.3272	0.2132	**0.4596**	4	3	0.3614	0.1908	**0.4478**
5	2	0.3454	0.2016	**0.4530**	5	3	0.3794	0.2070	**0.4136**
6	2	0.3728	0.1952	**0.4320**	6	3	**0.4162**	0.1886	0.3970
7	2	0.3798	0.1950	**0.4252**	7	3	**0.4250**	0.1976	0.3774
8	2	0.3844	0.2024	**0.4132**	8	3	**0.4342**	0.0.1802	0.3856
9	2	0.3894	0.2096	**0.4090**	9	3	**0.4346**	0.1920	0.3704
10	2	**0.3922**	0.2188	0.3890	10	3	**0.4368**	0.1970	0.3662
2	4	0.2372	0.1828	**0.5800**	2	5	0.2392	0.1946	**0.5662**
3	4	0.3290	0.1836	**0.4874**	3	5	0.3552	0.1808	**0.4640**
4	4	0.3780	0.1868	**0.4352**	4	5	0.3836	0.1918	**0.4246**
5	4	**0.4170**	0.1844	0.3986	5	5	**0.4154**	0.1828	0.4018
6	4	**0.4238**	0.1930	0.3832	6	5	**0.4338**	0.1962	0.3700
7	4	**0.4538**	0.1886	0.3576	7	5	**0.4478**	0.1864	0.3658
8	4	**0.4538**	0.1886	0.3576	8	5	**0.4638**	0.1920	0.3442
9	4	**0.4556**	0.2018	0.3426	9	5	**0.4674**	0.1938	0.3388
10	4	**0.4700**	0.1864	0.3436	10	5	**0.4766**	0.1950	0.3284

Table C.6: Success proportions: Poisson mean-centred in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1578	0.2414	**0.6008**	2	3	0.1920	0.2168	**0.5912**
3	2	0.2740	0.2274	**0.4986**	3	3	0.3154	0.2182	**0.4664**
4	2	0.3366	0.2278	**0.4356**	4	3	0.3584	0.2108	**0.4308**
5	2	0.3682	0.2186	**0.4132**	5	3	**0.4048**	0.2066	0.3886
6	2	0.3862	0.2118	**0.4020**	6	3	**0.4280**	0.2050	0.3670
7	2	**0.3928**	0.2286	0.3786	7	3	**0.4456**	0.2078	0.3466
8	2	**0.3880**	0.2250	0.3870	8	3	**0.4496**	0.2038	0.3466
9	2	**0.4140**	0.2196	0.3664	9	3	**0.4490**	0.2120	0.3390
10	2	**0.3998**	0.2192	0.3810	10	3	**0.4652**	0.2038	0.3310
2	4	0.2242	0.2062	**0.5696**	2	5	0.2366	0.2080	**0.5554**
3	4	0.3408	0.2022	**0.4570**	3	5	0.3606	0.1922	**0.4472**
4	4	**0.4080**	0.2042	0.3878	4	5	**0.4054**	0.2014	0.3932
5	4	**0.4160**	0.2080	0.3760	5	5	**0.4430**	0.1938	0.3572
6	4	**0.4600**	0.1932	0.3468	6	5	**0.4744**	0.1924	0.3332
7	4	**0.4680**	0.1984	0.3356	7	5	**0.4794**	0.1930	0.3276
8	4	**0.4762**	0.2032	0.3206	8	5	**0.4816**	0.2034	0.3150
9	4	**0.4852**	0.1962	0.3186	9	5	**0.4890**	0.2146	0.2964
10	4	**0.5000**	0.1938	0.3062	10	5	**0.5004**	0.1992	0.3004

Table C.7: Success proportions: Binomial overestimation in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1760	0.1914	**0.6326**	2	3	0.2354	0.1820	**0.5826**
3	2	0.2594	0.1662	**0.5744**	3	3	0.3112	0.1540	**0.5348**
4	2	0.2850	0.1686	**0.5464**	4	3	0.3500	0.1552	**0.4948**
5	2	0.3230	0.1624	**0.5146**	5	3	0.3768	0.1528	**0.4704**
6	2	0.3150	0.1726	**0.5124**	6	3	0.3926	0.1512	**0.4562**
7	2	0.3288	0.1680	**0.5032**	7	3	0.3908	0.1608	**0.4484**
8	2	0.3328	0.1714	**0.4958**	8	3	0.4136	0.1574	**0.4290**
9	2	0.3326	0.1716	**0.4958**	9	3	**0.4184**	0.1744	0.4072
10	2	0.3320	0.1666	**0.5014**	10	3	**0.4292**	0.1586	0.4122
2	4	0.2656	0.1724	**0.5620**	2	5	0.2802	0.1762	**0.5436**
3	4	0.3406	0.1568	**0.5026**	3	5	0.3536	0.1622	**0.4842**
4	4	0.3872	0.1526	**0.4602**	4	5	0.4008	0.1582	**0.4410**
5	4	0.3990	0.1616	**0.4394**	5	5	**0.4324**	0.1782	0.3894
6	4	0.4118	0.1722	**0.4160**	6	5	**0.4253**	0.1804	0.3942
7	4	**0.4262**	0.1796	0.3942	7	5	**0.4390**	0.1744	0.3866
8	4	**0.4400**	0.1790	0.3810	8	5	**0.4500**	0.1848	0.3652
9	4	**0.4390**	0.1850	0.3760	9	5	**0.4518**	0.1878	0.3604
10	4	**0.4400**	0.1884	0.3716	10	5	**0.4504**	0.3536	0.3536

Table C.8: Success proportions: Binomial underestimation in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1840	0.2622	**0.5338**	2	3	0.2212	0.2496	**0.5292**
3	2	0.2526	0.2628	**0.4846**	3	3	0.3026	0.2406	**0.4568**
4	2	0.2870	0.2452	**0.4678**	4	3	0.3440	0.2364	**0.4196**
5	2	0.3124	0.2510	**0.4366**	5	3	0.3614	0.2390	**0.3990**
6	2	0.3334	0.2540	**0.4126**	6	3	0.3748	0.2372	**0.3880**
7	2	0.3308	0.2524	**0.4168**	7	3	**0.3860**	0.2462	0.3678
8	2	0.3452	0.2548	**0.4000**	8	3	**0.3874**	0.2578	0.3548
9	2	0.3542	0.2542	**0.3916**	9	3	**0.4074**	0.2506	0.3420
10	2	0.3470	0.2528	**0.4002**	10	3	**0.4140**	0.2428	0.3778
2	4	0.2582	0.2396	**0.5022**	2	5	0.2526	0.2576	**0.4898**
3	4	0.3214	0.2420	**0.4366**	3	5	0.3312	0.2406	**0.4282**
4	4	0.3602	0.2452	**0.3946**	4	5	0.3720	0.2420	**0.3860**
5	4	**0.3830**	0.2422	0.3748	5	5	**0.3832**	0.2350	0.3818
6	4	**0.3952**	0.2382	0.3666	6	5	**0.3956**	0.2558	0.3486
7	4	**0.4200**	0.2502	0.3298	7	5	**0.4256**	0.2392	0.3352
8	4	**0.4080**	0.2550	0.3370	8	5	**0.4328**	0.2456	0.3216
9	4	**0.4212**	0.2598	0.3190	9	5	**0.4286**	0.2454	0.3260
10	4	**0.4312**	0.2504	0.3184	10	5	**0.4402**	0.2546	0.3052

Table C.9: Success proportions: Binomial mean-centred in the group problem.

DMs	Returns	PI	EQ	MR	DMs	Returns	PI	EQ	MR
2	2	0.1748	0.3416	**0.4836**	2	3	0.2166	0.3102	**0.4732**
3	2	0.2458	0.3364	**0.4178**	3	3	0.2818	0.3092	**0.4090**
4	2	0.2804	0.3380	**0.3816**	4	3	0.3173	0.3032	**0.3796**
5	2	0.3126	0.3256	**0.3618**	5	3	**0.3482**	0.3180	0.3338
6	2	0.3226	0.3382	**0.3392**	6	3	**0.3648**	0.3070	0.3282
7	2	0.3244	0.3258	**0.3498**	7	3	**0.3724**	0.3062	0.3214
8	2	**0.3456**	0.3182	0.3362	8	3	**0.3824**	0.3130	0.3046
9	2	**0.3554**	0.3164	0.3282	9	3	**0.4062**	0.3090	0.2848
10	2	**0.3770**	0.2928	0.3302	10	3	**0.4002**	0.3238	0.2760
2	4	0.2380	0.2980	**0.4640**	2	5	0.2430	0.2870	**0.4700**
3	4	0.3058	0.2978	**0.3964**	3	5	0.3070	0.3008	**0.3922**
4	4	0.3386	0.2980	**0.3634**	4	5	0.3382	0.3016	**0.3602**
5	4	**0.3788**	0.2884	0.3228	5	5	**0.3822**	0.2968	0.3210
6	4	**0.3798**	0.2946	0.3256	6	5	**0.4058**	0.2976	0.2966
7	4	**0.4158**	0.2926	0.2916	7	5	**0.3994**	0.2898	0.3108
8	4	**0.4270**	0.2820	0.2910	8	5	**0.4290**	0.2856	0.2854
9	4	**0.4430**	0.2932	0.2638	9	5	**0.4220**	0.3094	0.2686
10	4	**0.4440**	0.2892	0.2668	10	5	**0.4348**	0.2918	0.2734

Table C.10: Success proportions: Normal overestimation in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	**0.5000**	2	3	**0.5000**
3	2	**0.8422**	3	3	**0.8272**
4	2	**0.8063**	4	3	**0.7929**
5	2	**0.8538**	5	3	**0.8250**
6	2	**0.8326**	6	3	**0.8069**
7	2	**0.8536**	7	3	**0.8268**
8	2	**0.8486**	8	3	**0.8195**
9	2	**0.8570**	9	3	**0.8226**
10	2	**0.8531**	10	3	**0.8165**
2	4	**0.5000**	2	5	**0.5000**
3	4	**0.8124**	3	5	**0.8068**
4	4	**0.7681**	4	5	**0.7697**
5	4	**0.8120**	5	5	**0.7918**
6	4	**0.7915**	6	5	**0.7794**
7	4	**0.8100**	7	5	**0.7908**
8	4	**0.8059**	8	5	**0.7802**
9	4	**0.8156**	9	5	**0.7842**
10	4	**0.7974**	10	5	**0.7838**

Table C.11: Success proportions: Normal underestimation in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	**0.5000**	2	3	**0.5000**
3	2	**0.7600**	3	3	**0.7546**
4	2	**0.7381**	4	3	**0.7187**
5	2	**0.7730**	5	3	**0.7560**
6	2	**0.7553**	6	3	**0.7408**
7	2	**0.7820**	7	3	**0.7720**
8	2	**0.7683**	8	3	**0.7496**
9	2	**0.7884**	9	3	**0.7684**
10	2	**0.7761**	10	3	**0.7578**
2	4	**0.5000**	2	5	**0.5000**
3	4	**0.7368**	3	5	**0.7290**
4	4	**0.7033**	4	5	**0.6929**
5	4	**0.7294**	5	5	**0.7270**
6	4	**0.7211**	6	5	**0.7082**
7	4	**0.7352**	7	5	**0.7270**
8	4	**0.7357**	8	5	**0.7151**
9	4	**0.7384**	9	5	**0.7406**
10	4	**0.7386**	10	5	**0.7236**

Table C.12: Success proportions: Normal mean-centred in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	**0.5000**	2	3	**0.5000**
3	2	**0.8048**	3	3	**0.7922**
4	2	**0.7782**	4	3	**0.7552**
5	2	**0.8368**	5	3	**0.8130**
6	2	**0.8268**	6	3	**0.8060**
7	2	**0.8608**	7	3	**0.8344**
8	2	**0.8436**	8	3	**0.8227**
9	2	**0.8616**	9	3	**0.8426**
10	2	**0.8566**	10	3	**0.8450**
2	4	**0.5000**	2	5	**0.5000**
3	4	**0.7784**	3	5	**0.7722**
4	4	**0.7434**	4	5	**0.7401**
5	4	**0.7942**	5	5	**0.7754**
6	4	**0.7855**	6	5	**0.7637**
7	4	**0.8055**	7	5	**0.7966**
8	4	**0.8093**	8	5	**0.8008**
9	4	**0.8236**	9	5	**0.8244**
10	4	**0.8216**	10	5	**0.8078**

Table C.13: Success proportions: Poisson overestimation in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	**0.5000**	2	3	**0.5000**
3	2	**0.7622**	3	3	**0.7552**
4	2	**0.7365**	4	3	**0.7268**
5	2	**0.7972**	5	3	**0.7922**
6	2	**0.7814**	6	3	**0.7719**
7	2	**0.8316**	7	3	**0.8150**
8	2	**0.8139**	8	3	**0.8021**
9	2	**0.8470**	9	3	**0.8236**
10	2	**0.8281**	10	3	**0.8191**
2	4	**0.5000**	2	5	**0.5000**
3	4	**0.7550**	3	5	**0.7368**
4	4	**0.7195**	4	5	**0.7068**
5	4	**0.7688**	5	5	**0.7704**
6	4	**0.7642**	6	5	**0.7466**
7	4	**0.7940**	7	5	**0.7778**
8	4	**0.7833**	8	5	**0.7636**
9	4	**0.8110**	9	5	**0.7898**
10	4	**0.8006**	10	5	**0.7780**

Table C.14: Success proportions: Poisson underestimation in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	**0.5000**	2	3	**0.5000**
3	2	**0.7102**	3	3	**0.6992**
4	2	**0.6522**	4	3	**0.6647**
5	2	**0.7122**	5	3	**0.7200**
6	2	**0.7074**	6	3	**0.6954**
7	2	**0.7358**	7	3	**0.7214**
8	2	**0.7200**	8	3	**0.7167**
9	2	**0.7514**	9	3	**0.7360**
10	2	**0.7445**	10	3	**0.7368**
2	4	**0.5000**	2	5	**0.5000**
3	4	**0.6902**	3	5	**0.6940**
4	4	**0.6577**	4	5	**0.6490**
5	4	**0.6974**	5	5	**0.8842**
6	4	**0.6830**	6	5	**0.6699**
7	4	**0.7122**	7	5	**0.7128**
8	4	**0.7000**	8	5	**0.6922**
9	4	**0.7238**	9	5	**0.7110**
10	4	**0.7278**	10	5	**0.6934**

Table C.15: Success proportions: Poisson mean-centred in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	**0.5000**	2	3	**0.5000**
3	2	**0.6704**	3	3	**0.6680**
4	2	**0.6310**	4	3	**0.6312**
5	2	**0.6940**	5	3	**0.6744**
6	2	**0.6734**	6	3	**0.6479**
7	2	**0.7008**	7	3	**0.6848**
8	2	**0.6963**	8	3	**0.6819**
9	2	**0.7114**	9	3	**0.6946**
10	2	**0.6994**	10	3	**0.6849**
2	4	**0.5000**	2	5	**0.5000**
3	4	**0.6460**	3	5	**0.6400**
4	4	**0.6191**	4	5	**0.6219**
5	4	**0.6646**	5	5	**0.6514**
6	4	**0.6447**	6	5	**0.6330**
7	4	**0.6586**	7	5	**0.6596**
8	4	**0.6596**	8	5	**0.6483**
9	4	**0.6836**	9	5	**0.6616**
10	4	**0.6670**	10	5	**0.6597**

Table C.16: Success proportions: Binomial overestimation in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	0.4889	2	3	0.4916
3	2	**0.7066**	3	3	**0.7056**
4	2	**0.6783**	4	3	**0.6795**
5	2	**0.7636**	5	3	**0.7510**
6	2	**0.7342**	6	3	**0.7259**
7	2	**0.7898**	7	3	**0.7676**
8	2	**0.7750**	8	3	**0.7552**
9	2	**0.8024**	9	3	**0.7856**
10	2	**0.7830**	10	3	**0.7715**
2	4	0.4871	2	5	0.4878
3	4	**0.7008**	3	5	**0.6830**
4	4	**0.6668**	4	5	**0.6494**
5	4	**0.7366**	5	5	**0.7164**
6	4	**0.7131**	6	5	**0.7056**
7	4	**0.7540**	7	5	**0.7364**
8	4	**0.7301**	8	5	**0.7213**
9	4	**0.7696**	9	5	**0.7444**
10	4	**0.7520**	10	5	**0.7425**

Table C.17: Success proportions: Binomial undestimation in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	0.4911	2	3	0.4914
3	2	**0.5904**	3	3	**0.5864**
4	2	**0.5623**	4	3	**0.5609**
5	2	**0.5950**	5	3	**0.5880**
6	2	**0.5807**	6	3	**0.5679**
7	2	**0.5968**	7	3	**0.5990**
8	2	**0.5888**	8	3	**0.5910**
9	2	**0.6238**	9	3	**0.6200**
10	2	**0.6072**	10	3	**0.6029**
2	4	0.4916	2	5	0.4941
3	4	**0.5702**	3	5	**0.5608**
4	4	**0.5486**	4	5	**0.5443**
5	4	**0.5930**	5	5	**0.5764**
6	4	**0.5728**	6	5	**0.5547**
7	4	**0.6048**	7	5	**0.5832**
8	4	**0.5917**	8	5	**0.5728**
9	4	**0.6108**	9	5	**0.6168**
10	4	**0.6043**	10	5	**0.5896**

Table C.18: Success proportions: Binomial mean-centred in the individual problem.

DMs	Returns	PI	DMs	Returns	PI
2	2	0.4952	2	3	0.4932
3	2	0.4510	3	3	0.4520
4	2	0.4543	4	3	0.4466
5	2	0.4270	5	3	0.4272
6	2	0.4295	6	3	9.4201
7	2	0.4034	7	3	0.4140
8	2	0.4173	8	3	0.4001
9	2	0.4000	9	3	0.3950
10	2	0.4007	10	3	0.4069
2	4	0.4924	2	5	0.4922
3	4	0.4596	3	5	0.4532
4	4	0.4561	4	5	0.4587
5	4	0.4254	5	5	0.4216
6	4	0.4282	6	5	0.4305
7	4	0.4100	7	5	0.4156
8	4	0.4075	8	5	0.4142
9	4	0.3942	9	5	0.3918
10	4	0.3987	10	5	0.3998

Appendix D

Sample Code

Simulation Study Code

Here we provide an illustration of the code used in our group problem simulation study in Section 4.2. Below is commented code for the function **PlugIn.Group**.

```
N<-Number of DMs
n<-Number of returns
true<-true value of theta
par1<-vector of first parameter of DMs priors
par2<-vector of second parameter of DMs priors
dist<-"binomial","normal" or "poisson"
plot<-"TRUE" or "FALSE" dependent on if graphs are required or not
x<-extra parameter (number of trials in binomial case, known variance in
     normal case, omitted in poisson case).

PlugIn.Group<-function(N,n,true,par1,par2,dist,plot,x)
{
a<-matrix(NA,N,n+1) #matrix for "a" hyperparameter of distributions.
b<-matrix(NA,N,n+1) #matrix for "b" hyperparameter of distributions.

means<-matrix(NA,N,n+1) #mean matrix. (i,j) entry is mean of DM i at epoch j.
vars<-matrix(NA,N,n+1)  #var matrix. (i,j) entry is var of DM i at epoch j.
```

```
w<-matrix(NA,N,n+1) #matrix for PI weights with (i,j) as above.
w[,1]<-rep(0,N) #sets initial values equal to 0.

ualpha<-matrix(NA,N,n+1) #matrix for unnormalised weights. (i,j) as above.
alpha<-matrix(NA,N,n+1) #matrix for normalised weights. (i,j) as above.
alpha[,1]<-rep(1/N,N) #initial equal unnormalised weights.
ualpha[,1]<-rep(1/N,N) #initial equal normalised weights.

relalpha<-matrix(NA,N,n+1) #matrix of most reliable DMs
relalpha[,1]<-rep(1/N,N) #initial weights are equal

if (dist=="binomial"){ #learning in Binomial setting.
a[,1]<-par1 #inserts prior
b[,1]<-par2 #inserts prior

r<-rbinom(n,x,true) #simulates Binomial data

means[,1]<-a[,1]/(a[,1]+b[,1]) #initial means
vars[,1]<-(a[,1]*b[,1])/((a[,1]+b[,1])^2*(a[,1]+b[,1]+1)) #initial vars

for (j in 2:(n+1)){ #updates hyperparameters given data
a[,j]<-a[,j-1]+r[j-1]
b[,j]<-b[,j-1]+(x-r[j-1])}

for (j in 2:(n+1)){ #updates means and variances
means[,j]<-a[,j]/(a[,j]+b[,j])
vars[,j]<-(a[,j]*b[,j])/((a[,j]+b[,j])^2*(a[,j]+b[,j]+1))}

for (j in 2:(n+1)){#PI weights from predictive distribution
w[,j]<-exp((lgamma(x+1)+lgamma(a[,j-1]+b[,j-1])+lgamma(a[,j-1]+r[j-1])+
         lgamma(b[,j-1]+x-r[j-1])-lgamma(x-r[j-1]+1)-lgamma(r[j-1]+1)-
         lgamma(a[,j-1]) -lgamma(b[,j-1])-lgamma(a[,j-1]+b[,j-1]+x)))}}
```

```
if (dist=="normal"){ #learning in Normal setting.
means[,1]<-par1 #inputs prior means
vars[,1]<-par2 #inputs prior variances

r<-rnorm(n,true,sqrt(x)) #simulates Normal data

for (j in 2:(n+1)){#updates posterior means and variances
means[,j]<-((means[,j-1]/vars[,j-1])+(r[j-1]/x))/((1/vars[,j-1])+(1/x))
vars[,j]<-1/((1/vars[,j-1])+(1/x))}

for (j in 2:(n+1)){#creates plug-in weights
w[,j]<-dnorm(r[j-1],means[,j-1],sqrt(x+vars[,j-1]))}
}

if (dist=="poisson"){ #learning in Poisson setting.
a[,1]<-par1 #inserts prior scale parameters.
b[,1]<-par2 #inserts prior shape parameter.

r<-rpois(n,true) #simulates Poisson data.

means[,1]<-b[,1]/a[,1] #calculates prior mean.
vars[,1]<-b[,1]/(a[,1]^2) #calculates prior variance.

for (j in 2:(n+1)){#data witnessed used to update shape and scale
a[,j]<-a[,j-1]+1
b[,j]<-b[,j-1]+r[j-1]}

means[,]<-b[,]/a[,] #updates means
vars[,]<-b[,]/(a[,]^2)#updates variances
```

```
for (j in 2:(n+1)){#creates plug-in weights
w[,j]<-exp(b[,j-1]*log(a[,j-1])+lgamma(r[j-1]+b[,j-1])-lgamma(r[j-1]+1)
         -lgamma(b[,j-1])-(r[j-1]+b[,j-1])*log(a[,j-1]+1))}
}

for (j in 2:(n+1)){ #updates unnormalised and normalised weights.
ualpha[,j]<-alpha[,j-1]*w[,j]
alpha[,j]<-ualpha[,j]/(sum(ualpha[,j]))}

#divides MR weight if multiple DMs have the same view
for (j in 2:(n+1)){for (i in 1:N){ ###1 to MR, 0 to others
ifelse(max(w[,j])==w[i,j],relalpha[i,j]<-1/length(w[,j][w[,j]==max(w[,j])]),
       relalpha[i,j]<-0)}}

#means of the three methods
mean.PI<-sum(alpha[,n+1]*means[,n+1])
mean.EQ<-sum(rep(1/N,N)*means[,n+1])
mean.MR<-sum(relalpha[,n+1]*means[,n+1])

#variances of the three methods
var.PI.comp<-rep(NA,n)
var.PI.comp<-sum(alpha[,n+1]*(means[,n+1]^2+vars[,n+1]))
var.PI<-var.PI.comp-(mean.PI)^2
var.EQ.comp<-rep(NA,n)
var.EQ.comp<-sum(rep(1/N,N)*(means[,n+1]^2+vars[,n+1]))
var.EQ<-var.EQ.comp-(mean.EQ)^2
var.MR<-sum(relalpha[,n+1]*vars[,n+1])

w.methods<-rep(NA,3)

if(dist=="binomial"){
dens<-rep(NA,N)
```

```
for (i in 1:N){
dens[i]<-dbeta(true,a[i,n+1],b[i,n+1])}}

if(dist=="normal"){
dens<-rep(NA,N)
for (i in 1:N){ # density each DM places on truth
dens[i]<-dnorm(true,means[i,n+1],sqrt(vars[i,n+1]))}}

if(dist=="poisson"){
dens<-rep(NA,N)
for (i in 1:N){
dens[i]<-dgamma(true,b[i,n+1],a[i,n+1])}}

w.methods[1]<-sum(alpha[,n+1]*dens) #PI density
w.methods[2]<-sum(rep(1/N,N)*dens) #EQ density
w.methods[3]<-sum(relalpha[,n+1]*dens) #MR density

#which method is best?
w.res<-ifelse(max(w.methods)==w.methods[1],0,
        (ifelse(max(w.methods)==w.methods[2],1,2)))

if(plot=="TRUE"){
if(dist=="normal"){
#range considered
t<-seq(min(true,means[,n+1]-5*vars[,n+1]),max(true,means[,n+1]+5*vars[,n+1]),
   length.out=1000)

#EQ densities
dens.t<-matrix(NA,N,length(t))
for (i in 1:N){
dens.t[i,]<-dnorm(t,means[i,n+1],sqrt(vars[i,n+1]))}
```

```
#PI densities
dens.PI<-rep(NA,length(t))
for (i in 1:length(t)){
dens.PI[i]<-sum(alpha[,n+1]*dens.t[,i])
}

#MR
plot(t,dnorm(t,mean.MR,sqrt(var.MR)),type="l",col="green",lwd=3,
     main="Different Densities",xlab="Theta", ylab="Prob(Theta)",
     ylim=c(0,max(colMeans(dens.t),dens.PI,dnorm(t,mean.MR,sqrt(var.MR)))))

#EQ
lines(t,colMeans(dens.t),col="blue",lwd=3)

#PI
lines(t,dens.PI,col="red",lwd=3)

abline(v=true,lwd=3)
legend("topright",c("PI","EQ","MR","True"),lty=c(1,1,1,1),lwd=c(3,3,3,3),
       col=c("green","blue","red","black"),cex=0.6)
}

if(dist=="binomial"){
#range considered
t<-seq(0,1,0.001)

#EQ densities
dens.t<-matrix(NA,N,length(t))
for (i in 1:N){
dens.t[i,]<-dbeta(t,a[i,n+1],b[i,n+1])}

#PI densities
```

```
dens.PI<-rep(NA,length(t))
for (i in 1:length(t)){
dens.PI[i]<-sum(alpha[,n+1]*dens.t[,i])
}

#MR
plot(t,dbeta(t,sum(relalpha[,n+1]*a[,n+1]),sum(relalpha[,n+1]*b[,n+1])),
      type="l", col="green",lwd=3,main="Different Densities",xlab="Theta",
       ylab="Prob(Theta)", ylim=c(0,max(colMeans(dens.t),dens.PI,
      dbeta(t,sum(relalpha[,n+1]*a[,n+1]),sum(relalpha[,n+1]*b[,n+1])))))

#EQ
lines(t,colMeans(dens.t),col="blue",lwd=3)

#PI
lines(t,dens.PI,col="red",lwd=3)

abline(v=true,lwd=3)
legend("topright",c("PI","EQ","MR","True"),lty=c(1,1,1,1),lwd=c(3,3,3,3),
col=c("red","blue","green","black"),cex=0.6)
}

if(dist=="poisson"){
#range considered
t<-seq(0,max(true,means[,n+1]+5*vars[,n+1]),length.out=1000)

#EQ densities
dens.t<-matrix(NA,N,length(t))
for (i in 1:N){
dens.t[i,]<-dgamma(t,b[i,n+1],a[i,n+1])}

#PI densities
```

```
dens.PI<-rep(NA,length(t))
for (i in 1:length(t)){
dens.PI[i]<-sum(alpha[,n+1]*dens.t[,i])
}

#MR
plot(t,dgamma(t,sum(relalpha[,n+1]*b[,n+1]),sum(relalpha[,n+1]*a[,n+1])),
      type="l",col="green", lwd=3,main="Different Densities",xlab="Theta",
      ylab="Prob(Theta)",ylim=c(0,max(colMeans(dens.t),dens.PI,
      dgamma(t,sum(relalpha[,n+1]*b[,n+1]),sum(relalpha[,n+1]*a[,n+1])))))

#EQ
lines(t,colMeans(dens.t),col="blue",lwd=3)

#PI
lines(t,dens.PI,col="red",lwd=3)

abline(v=true,lwd=3)
legend("topright",c("PI","EQ","MR","True"),lty=c(1,1,1,1),lwd=c(3,3,3,3),
         col=c("green","blue","red","black"),cex=0.6)
}
}

#returns data frame of key attributes
dfa<-data.frame(c(mean.PI,mean.EQ,mean.MR),c(var.PI,var.EQ,var.MR),
        c(true,rep("",2)),c(w.methods),c(w.res,rep("",2)))
rownames(dfa)<-c("PI","EQ","MR")
colnames(dfa)<-c("Means","Variances","True","Density","Winner")
dfa
}
```

Real Data Study Code

We also provide commented code used for our real data study in the group problem in Section 4.4.

```
data<-the name of the data set of interest, e.g., "A_SEED".
k<-the associated variance scaling parameter.

PlugInCooke<-function(data,k)
{
#data sets for expert opinions and real values respectively
data.experts<-read.csv(paste("~/Documents/Statistics/Orginial/Third Year/
                       Cooke/DATA/",data,".EXPERTS.csv",sep=""),header=TRUE)
data.real<-read.csv(paste("~/Documents/Statistics/Orginial/Third Year/
                  Cooke/DATA/",data,".REAL.csv",sep=""),header=TRUE)

#number of experts and seeds respectively
n<-dim(data.experts)[1]/dim(data.real)[1]
m<-dim(data.real)[1]

#matrices for means and sds for DMs distributions, as discussed in thesis.
means<-matrix(NA,n,m)
sdevs<-matrix(NA,n,m)

for (i in 1:m){
means[,i]<-data.experts[c(0:(n-1))*m+i,6]}

for (i in 1:m){
for (j in 1:n){
sdevs[j,i]<-min((data.experts[(c(0:(n-1))*m+i)[j],7]-
                  data.experts[(c(0:(n-1))*m+i)[j],6])/1.645,
                 (data.experts[(c(0:(n-1))*m+i)[j],6]-
                  data.experts[(c(0:(n-1))*m+i)[j],5])/1.645)}}
```

```
#associated variance - function of k
sigsq<-abs(data.real[,3])/k

#PI weights. Initialised at zero. (i,j) entry is weight i^th DM
#associates with j^th seed (in original ordering).
w<-matrix(NA,n,m+1)
for (i in 1:n){w[i,1]<-0}
for (i in 1:n){
for (j in 2:(m+1)){
w[i,j]<-dnorm(data.real[j-1,3],means[i,j-1],sqrt(sigsq[j-1]+sdevs[i,j-1]^2))}}

#vector for storing optimal method for last seed for each permutations
last.seed<-rep(NA,m*(m-1))

#determines the m(m-1) permutations
grid<-expand.grid(c(1:m),c(1:m))
omit<-which(grid[,1]==grid[,2],)
grid<-grid[-omit,]

#measures which method is best for each permutation in turn
for (q in 1:(m*(m-1)))
{
#chooses q^th permutation and permutes relevant quantities
perm<-c(c(1:m)[-as.numeric(grid[q,])],as.numeric(grid[q,]))
means.new<-means[,perm]
sdevs.new<-sdevs[,perm]
data.real.new<-data.real[perm,]
sigsq.new<-sigsq[perm]
w.new<-w[,c(1,(perm+1))]

#matrices of unnormalised and normalised weights, initialised
#as equal. (i,j) entry is weight for i^th DM after seeing j seeds.
```

```
ualpha<-matrix(NA,n,m+1)
alpha<-matrix(NA,n,m+1)
for (i in 1:n){alpha[i,1]<-1/n}
for (i in 1:n){ualpha[i,1]<-1/n}

#matrix for which DM is deemed most reliable. Initialised with
#equal weights. (i,j) entry is 1 if i^th DM maximises PI weight
#after seeing (j-1) returns, 0 if not.
relalpha<-matrix(NA,n,m)
relalpha[,1]<-rep(1/n,n)

#updates weights over time.
for (j in 2:(m+1)){
for (i in 1:n){
for (k in 1:n){
ualpha[k,j]<-alpha[k,j-1]*w.new[k,j]}
alpha[i,j]<-ualpha[i,j]/(sum(ualpha[,j]))}}

#determines which DM is most reliable at each epoch,
#allowing for potential ties.
for (j in 2:m){for (i in 1:n){
ifelse(max(w.new[,j])==w.new[i,j],relalpha[i,j]<-1/
         length(w.new[,j][w.new[,j]==max(w.new[,j])]),relalpha[i,j]<-0)}}

#the density each DM places on the true value of each seed.
dens<-matrix(NA,n,m)
for (i in 1:n){
for (j in 1:m){
dens[i,j]<-dnorm(data.real.new[j,3],means.new[i,j],sdevs.new[i,j])}}

#the density placed by each method in turn on true value.
w.methods<-matrix(NA,3,m)
```

```
for (j in 1:m){
w.methods[1,j]<-sum(alpha[,j]*dens[,j]) #PI density
w.methods[2,j]<-sum(rep(1/n,n)*dens[,j]) #EQ density
w.methods[3,j]<-sum(relalpha[,j]*dens[,j]) #MR density
}

#determines which method is optimal for each seed. By equal
#weights all methods are equal for first seed. Hence i^th entry
is which method is best for seed i+1. 0 denotes PI, 1 denotes
#EQ and 2 denotes MR.
res.methods<-rep(NA,m-1)
for (i in 2:m){
ifelse(max(w.methods[,i])==w.methods[1,i],res.methods[i-1]<-0,
ifelse(max(w.methods[,i])==w.methods[2,i],res.methods[i-1]<-1,
         res.methods[i-1]<-2))}

#notes which method is optimal for final seed for this permutation
last.seed[q]<-res.methods[m-1]
}

#aggregates results over the m(m-1) permutations and
#prints results
last.scores<-c(sum(last.seed==0),sum(last.seed==1),
                       sum(last.seed==2))/(m*(m-1))
last.scores
}
```

Appendix E

Theoretical Calculations

In Section 4.3.1 we derived the theoretical calculations underlying the true probability of the PI approach being dominant over a set of alternatives, with simulated proportions approximating this true probability as the number of simulations increases. We provided calculations for the Beta-Binomial case in the main text of this thesis. Below we supply similar calculations for the Poisson-Gamma and Normal-Normal cases.

Poisson-Gamma conjugacy

Here θ is a rate parameter over a unit of time. Each P_i has a Gamma prior over θ:

$$f_i(\theta) = \frac{\alpha_i^{\beta_i}}{\Gamma(\beta_i)}\theta^{\beta_i - 1}e^{-\alpha_i\theta} \text{ with } \theta \in (0, \infty) \tag{E.1}$$

Returns follow a Poisson distribution, with the probability of a particular r being

$$f(R = r|\theta) = \frac{\theta^r}{r!}e^{-\theta} \text{ with } r = 0, 1, \ldots \tag{E.2}$$

When $R \sim \text{Bin}(m, \theta)$ there was a finite amount of returns witnessable per epoch, yet the Poisson distribution can produce a (countably) infinite number of potential returns. To ensure computability we assume that there is a finite upper bound u with $\mathbb{P}(R > u|\theta) < \epsilon$ where ϵ is small. The hyperparameters of P_i after k returns are

$$\alpha_i^{(k)} = \alpha_i + k \tag{E.3}$$

$$\beta_i^{(k)} = \beta_i + \sum_{j=1}^{k} r_j \tag{E.4}$$

We write $u_{i,t}$ as

$$u_{i,t} = \frac{1}{n}\prod_{k=1}^{t}\frac{\Gamma(r + \beta_i^{(k-1)})}{\Gamma(r_k + 1)\Gamma(\beta_i^{(k-1)})}\frac{\left(\alpha_i^{(k-1)}\right)^{\beta_i^{(k-1)}}}{(\alpha_i^{(k-1)} + 1)^{r_k + \beta_i^{(k-1)}}} \tag{E.5}$$

The probability of return set $\{r_1, \dots, r_t\}$ is a product of Poisson distributions, *i.e.*,

$$f(R_1 = r_1, \dots, R_t = r_t|\theta) = \prod_{k=1}^{t} \frac{\theta^{r_k}}{r_k!} e^{-\theta} \tag{E.6}$$

The normalised weight given to P_i after t returns is

$$\begin{aligned}
\gamma_{i,t} &= \frac{\frac{1}{n}\prod_{k=1}^{t} \frac{\Gamma(r+\beta_i^{(k-1)})}{\Gamma(r_k+1)\Gamma(\beta_i^{(k-1)})} \frac{(\alpha_i^{(k-1)})^{\beta_i^{(k-1)}}}{(\alpha_i^{(k-1)}+1)^{r_k+\beta_i^{(k-1)}}}}{\sum_{j=1}^{n} \frac{1}{n}\prod_{k=1}^{t} \frac{\Gamma(r+\beta_j^{(k-1)})}{\Gamma(r_k+1)\Gamma(\beta_j^{(k-1)})} \frac{(\alpha_j^{(k-1)})^{\beta_j^{(k-1)}}}{(\alpha_j^{(k-1)}+1)^{r_k+\beta_j^{(k-1)}}}} \\
&= \frac{\prod_{k=1}^{t} \frac{\Gamma(r+\beta_i^{(k-1)})}{\Gamma(r_k+1)\Gamma(\beta_i^{(k-1)})} \frac{(\alpha_i^{(k-1)})^{\beta_i^{(k-1)}}}{(\alpha_i^{(k-1)}+1)^{r_k+\beta_i^{(k-1)}}}}{\sum_{j=1}^{n} \prod_{k=1}^{t} \frac{\Gamma(r+\beta_j^{(k-1)})}{\Gamma(r_k+1)\Gamma(\beta_j^{(k-1)})} \frac{(\alpha_j^{(k-1)})^{\beta_j^{(k-1)}}}{(\alpha_j^{(k-1)}+1)^{r_k+\beta_j^{(k-1)}}}}
\end{aligned} \tag{E.7}$$

The weights in Equation $(E.7)$ are combined with the distributions in Equation $(E.1)$, with hyperparameters updated as in Equations $(E.3)$ and $(E.4)$, yielding a PI posterior distribution after t returns of

$$\begin{aligned}
\hat{f}_t^{PI}(\theta|r_1, \dots, r_t) &= \sum_{z=1}^{n} \Big[\frac{\prod_{k=1}^{t} \frac{\Gamma(r+\beta_z^{(k-1)})}{\Gamma(r_k+1)\Gamma(\beta_z^{(k-1)})} \frac{(\alpha_z^{(k-1)})^{\beta_z^{(k-1)}}}{(\alpha_z^{(k-1)}+1)^{r_k+\beta_z^{(k-1)}}}}{\sum_{j=1}^{n} \prod_{k=1}^{t} \frac{\Gamma(r+\beta_j^{(k-1)})}{\Gamma(r_k+1)\Gamma(\beta_j^{(k-1)})} \frac{(\alpha_j^{(k-1)})^{\beta_j^{(k-1)}}}{(\alpha_j^{(k-1)}+1)^{r_k+\beta_j^{(k-1)}}}} \times \\
&\qquad \frac{(\alpha_z^{(t)})^{\beta_z^{(t)}}}{\Gamma(\beta_z^{(t)})} \theta^{\beta_z^{(t)}-1} e^{-\alpha_z^{(t)}\theta} \Big]
\end{aligned} \tag{E.8}$$

We can proceed in a manner analogous to that from the Beta-Binomial case. We demonstrate convergence. Suppose we have $R \sim \text{Pois}(\theta)$ with $\theta = 2$ and three DMs with respective prior distributions of $f_1(\theta) \sim \text{Gamma}(1,3)$, $f_2(\theta) \sim \text{Gamma}(7,2)$, $f_3(\theta) \sim \text{Gamma}(2,2)$ and $\epsilon = 0.001$. After four returns have been witnessed the true success probability is 0.962. We simulated 5,000 process iterations and recorded the success proportion at each stage. We see in Fig. E.1 that convergence is clear.

The EQ posterior is a special case of Equation $(E.8)$, *i.e.*,

$$\begin{aligned}
\hat{f}_t^{EQ}(\theta|r_1, \dots, r_t) &= \sum_{z=1}^{n} \Big[\frac{1}{n} \times \frac{(\alpha_z^{(t)})^{\beta_z^{(t)}}}{\Gamma(\beta_z^{(t)})} \theta^{\beta_z^{(t)}-1} e^{-\alpha_z^{(t)}\theta} \Big] \\
&= \frac{1}{n} \sum_{z=1}^{n} \Big[\frac{(\alpha_z^{(t)})^{\beta_z^{(t)}}}{\Gamma(\beta_z^{(t)})} \theta^{\beta_z^{(t)}-1} e^{-\alpha_z^{(t)}\theta} \Big]
\end{aligned} \tag{E.9}$$

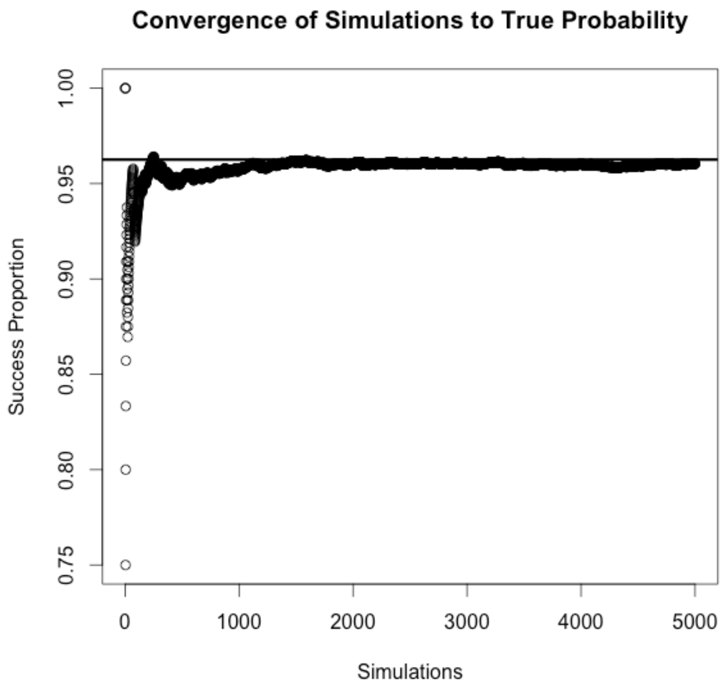

Fig. E.1: Success proportions/the true probability in a Poisson-Gamma case.

The MR posterior requires calculation of the indicators, $I^{MR}_{z,x}$, leading to

$$\hat{f}^{MR}_t(\theta|\underline{r}_x) = \sum_{z=1}^{n} I^{MR}_{z,x} \times \frac{(\alpha_z^{(t)})^{\beta_z^{(t)}}}{\Gamma(\beta_z^{(t)})} \theta^{\beta_z^{(t)}-1} e^{-\alpha_z^{(t)}\theta} \tag{E.10}$$

We proceed as in the Beta-Binomial case to find which method is superior. Convergence (using the example previously considered) is shown in Fig. E.2, with the PI method being superior in this case.

Normal-Normal conjugacy

Here θ is the mean of a Normally distributed process with known variance σ^2. The prior of each P_i over θ is Normally distributed with mean m_i and variance s_i^2, *i.e.*,

$$f_i(\theta) = \frac{1}{\sqrt{2\pi s_i^2}} \exp\Big(- \frac{(\theta - m_i)^2}{2s_i^2} \Big) \tag{E.11}$$

Returns are realisations of a continuous Normal random variable. In the Beta-Binomial case there were $(m+1)$ possible returns per epoch, and in the Poisson-Gamma case we restricted there to be only $(u+1)$ possibilities with u chosen so that the probability of seeing a return exceeding this was negligibly small. However, as the Normal distribution

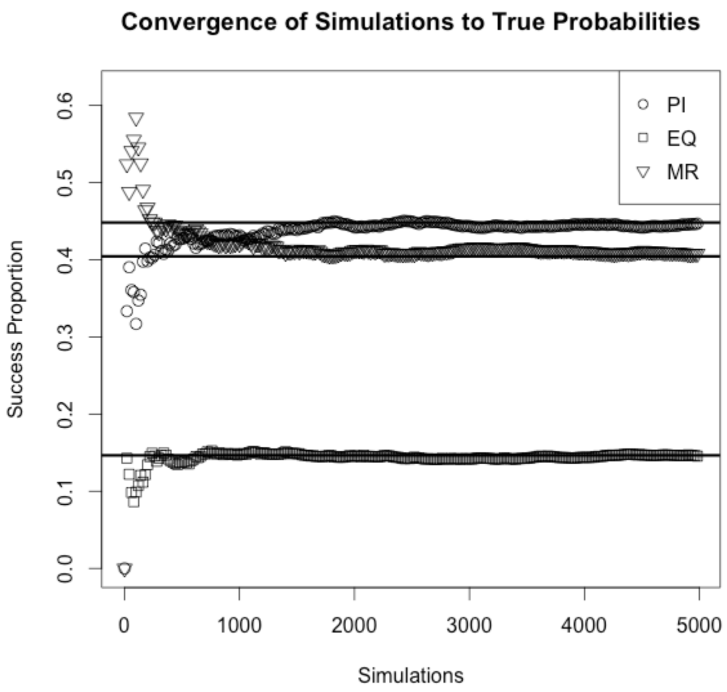

Fig. E.2: Success proportions/the true probability in a Poisson-Gamma case.

is continuous the probability of seeing any particular value is zero (*i.e.*, an uncountably infinite amount of potential returns). We introduce discretisation and suppose that $r \in [a_1, \ldots, a_m]$, *i.e.*, that there are m possibly returns that are equally spaced apart ($a_q - a_{q-1} = d$ for all $q = 2, \ldots, m$ with $a_{q-1} < a_q$ for all $q = 2, \ldots, m$). Clearly the smaller d is/the bigger m is, the less coarse the method is. We have a trade-off between accuracy of results and computational ability. The set of considered returns is centred on θ, *i.e.*, $a_{\frac{m+1}{2}} = \theta$ if m is odd, and $\frac{a_{\frac{m}{2}} + a_{\frac{m}{2}+1}}{2} = \theta$ if not. Values should be chosen so that returns below a_1 or above a_m have negligibly small probability of occurring. We find the probability of seeing any value a_i by integrating over the Normal distribution within a range of $\frac{d}{2}$ either side of a_i, *i.e.*,

$$f(R = a_i|\theta) = \int_{a_i - \frac{d}{2}}^{a_i + \frac{d}{2}} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\Big(- \frac{(a_i - \theta)^2}{2\sigma^2}\Big)\, d\theta \tag{E.12}$$

This integral must be numerically solved as no closed form solution exists. The posterior hyperparameters of P_i, $m_i^{(k)}$ and $s_i^{(k)}$, after k returns are

$$m_i^{(k)} = \frac{\frac{m_i}{s_i^2} + \frac{\sum_{j=1}^{k} r_j}{\sigma^2}}{\frac{1}{s_i^2} + \frac{k}{\sigma^2}} \tag{E.13}$$

$$s_i^{(k)} = \sqrt{\frac{1}{\frac{1}{s_i^2} + \frac{k}{\sigma^2}}} \tag{E.14}$$

We can write $u_{i,t}$ as

$$u_{i,t} = \frac{1}{n}\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2 + (s_i^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_i^{(k-1)})^2}{2(\sigma^2 + (s_i^{(k-1)})^2)}\Big) \tag{E.15}$$

The probability of seeing returns $\{r_1, \ldots, r_t\}$ is a product of Normal distributions:

$$f(R_1 = r_1, \ldots, R_t = r_t|\theta) = \prod_{k=1}^{t} \int_{r_k - \frac{d}{2}}^{r_k + \frac{d}{2}} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\Big(- \frac{(r_k - \theta)^2}{2\sigma^2}\Big) \tag{E.16}$$

The normalised weight afforded to P_i after t sets of returns is

$$\begin{aligned}
\gamma_{i,t} &= \frac{\frac{1}{n}\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2+(s_i^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_i^{(k-1)})^2}{2(\sigma^2+(s_i^{(k-1)})^2)}\Big)}{\sum_{j=1}^{n}\frac{1}{n}\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2+(s_j^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_j^{(k-1)})^2}{2(\sigma^2+(s_j^{(k-1)})^2)}\Big)} \\
&= \frac{\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2+(s_i^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_i^{(k-1)})^2}{2(\sigma^2+(s_i^{(k-1)})^2)}\Big)}{\sum_{j=1}^{n}\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2+(s_j^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_j^{(k-1)})^2}{2(\sigma^2+(s_j^{(k-1)})^2)}\Big)}
\end{aligned} \tag{E.17}$$

The weights in Equation ($E.17$) are combined with the Normal distributions in Equation ($E.11$), with hyperparameters updated as in Equations ($E.13$) and ($E.14$), yielding a PI posterior distribution after t returns of

$$\begin{aligned}
\hat{f}_t^{PI}(\theta) &= \sum_{z=1}^{n}\Big[\frac{\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2+(s_z^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_z^{(k-1)})^2}{2(\sigma^2+(s_z^{(k-1)})^2)}\Big)}{\sum_{j=1}^{n}\prod_{k=1}^{t} \frac{1}{\sqrt{2\pi(\sigma^2+(s_j^{(k-1)})^2)}} \exp\Big(- \frac{(r_k - m_j^{(k-1)})^2}{2(\sigma^2+(s_j^{(k-1)})^2)}\Big)} \times \\
&\qquad \frac{1}{\sqrt{2\pi(s_z^{(t)})^2}} \exp\Big(- \frac{(\theta - m_z^{(t)})^2}{2(s_z^{(t)})^2}\Big)\Big]
\end{aligned} \tag{E.18}$$

We proceed as previously outlined to determine the probability that the PI approach is superior. We demonstrate convergence where $R \sim \mathcal{N}(\theta, 1)$ with $\theta = 0$. Suppose there are three DMs with $f_1(\theta) \sim \mathcal{N}(-1, 1)$, $f_2(\theta) \sim \mathcal{N}(4, 4)$, $f_3(\theta) \sim \mathcal{N}(5, 3)$. After two returns have been observed the true success probability is 0.549. We discretised the return space somewhat coarsely ($m = 50$ leading to $d = 0.122$) and simulated 5,000 process iterations (as in Fig. E.3) with convergence again being evident.

In the group problem the EQ posterior is a special case of Equation ($E.18$), *i.e.*,

$$\hat{f}_t^{EQ}(\theta|r_1, \ldots, r_t) = \sum_{z=1}^{n}\Big[\frac{1}{n} \times \frac{1}{\sqrt{2\pi(s_z^{(t)})^2}} \exp\Big(- \frac{(\theta - m_z^{(t)})^2}{2(s_z^{(t)})^2}\Big)\Big]$$

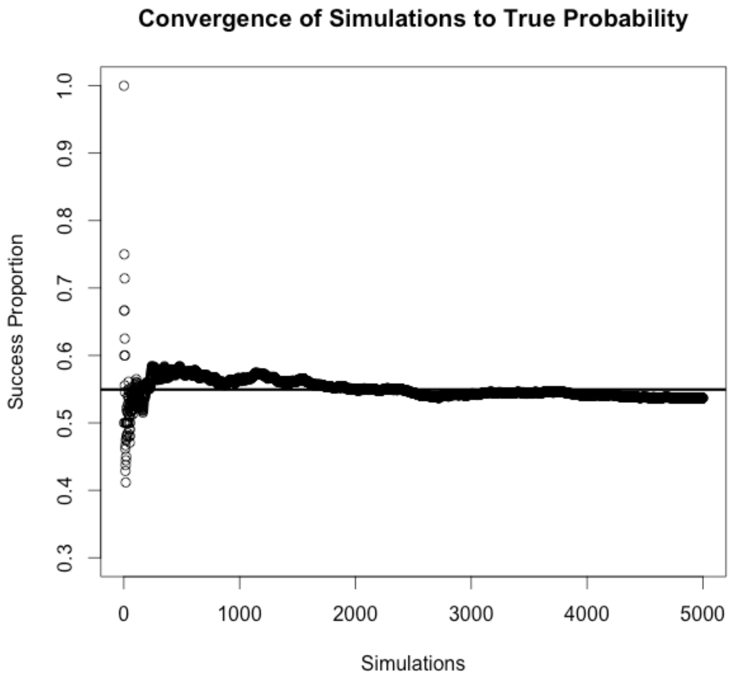

Fig. E.3: Success proportions/the true probability in a Normal-Normal case.

$$= \frac{1}{n}\sum_{z=1}^{n}\Big[\frac{1}{\sqrt{2\pi(s_z^{(t)})^2}}\exp\Big(-\frac{(\theta-m_z^{(t)})^2}{2(s_z^{(t)})^2}\Big)\Big] \tag{E.19}$$

The MR posterior requires calculation of the indicators, $I_{z,x}^{MR}$, leading to

$$\hat{f}_t^{MR}(\theta|\underline{r}_x) = \sum_{z=1}^{n} I_{z,x}^{MR} \times \frac{1}{\sqrt{2\pi(s_z^{(t)})^2}}\exp\Big(-\frac{(\theta-m_z^{(t)})^2}{2(s_z^{(t)})^2}\Big) \tag{E.20}$$

We proceed as before to find which method is superior. Convergence is illustrated in Fig. E.4. The MR method is optimal in this case.

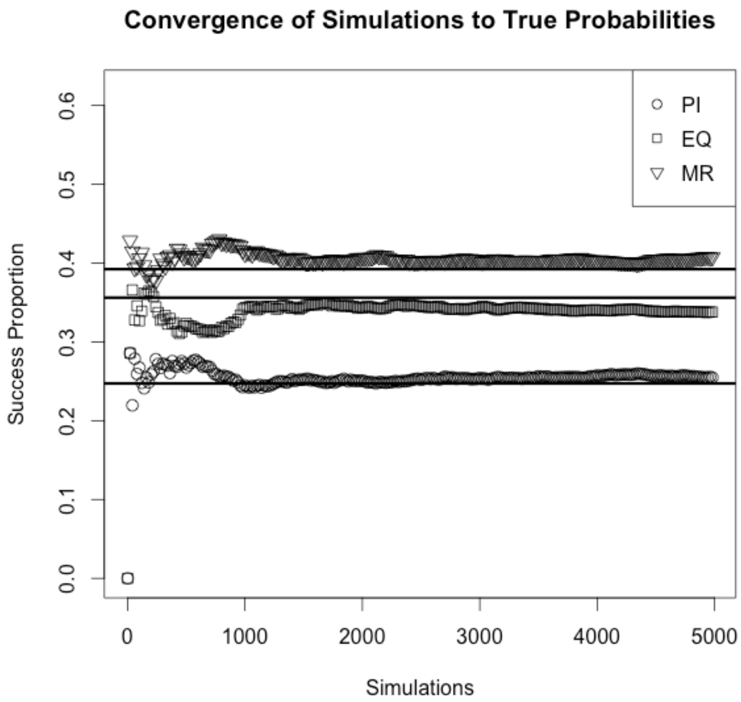

Fig. E.4: Success proportions/the true probability in a Normal-Normal case.

Appendix F

DV for Multiple Simultaneous Returns

In Section 3.8 we discussed how the PI approach could be extended from a setting where it was assumed that all DMs witnessed a common return at each epoch (*e.g.*, a stock price or a particular horse winning a race) to one in which all DMs observed distinct returns. This more complex environment, entailing multiple simultaneous returns, increased the applicability of this technique and further differentiated it from the classical setting of Cooke (1991). Here we construct an analogous extension for the DV approach derived in Section 7.1.

Each DM should update their opinion in light of all the information witnessed as in Equations (3.30) and (3.31), leading to an augmented belief of the form given in Equation (3.33). Our topic of interest regards how DV weights should be calculated given multiple returns. Using notation from Section 3.8, how should P_i, with utility function $u_i(r)$, calculate an analogy of the DV weight in Equation (7.2) given the stream of decisions $\underline{d_1}$ and corresponding returns $\underline{r_1}$ at the first epoch? We propose that this be done by pairwise summation across all sets of returns observed, *i.e.*,

$$v_{i,j} = \sum_{k=1}^{n} |u_i(r_{1,j}) - \mathbb{E}_{i|j}[u_i(d_{1,k})| \tag{F.1}$$

This involves consideration of all available information, with a DM comparing the predictions given by the information source to the outcomes resulting from the various decisions made by neighbours and contrasting these with the corresponding returns. If a DM calculated weights simply using Equation (7.2), *i.e.*, the return she witnessed, then she is ignoring the rest of the data she has access to. Her neighbour may appear very accurate but this may be a result of the outcome witnessed by the DM being a "fluke", *i.e.*, one with a probability of occurring that is in the tails of the true data generating mechanism. If this was the case, and Equation ($F.1$) was used the neighbour would score poorly for the other (more likely) returns witnessed, and hence receive a low weight. Note that if a DM makes a "trivial" decision (*i.e.*, one leading to no relevant return being witnessed, such as a decision to not enter a transaction) then this is not considered in the summation in Equation ($F.1$).