

# The Value of Performance Weights and Discussion in Aggregated Expert Judgements

Anca M. Hanea,<sup>1</sup> Marissa F. McBride<sup>2,3</sup> Mark A. Burgman<sup>2,4</sup> Bonnie C. Wintle<sup>2,5</sup>

---

In risky situations characterized by imminent decisions, scarce resources and insufficient data, policy makers rely on experts to estimate model parameters and their associated uncertainties. Different elicitation and aggregation methods can vary substantially in their efficacy and robustness. While it is generally agreed that biases in expert judgements can be mitigated using structured elicitations involving groups rather than individuals, there is still some disagreement about how to best elicit and aggregate judgements. This mostly concerns the merits of using performance-based weighting schemes to combine judgements of different individuals (rather than assigning equal weights to individual experts), and the way that interaction between experts should be handled. This paper aims to contribute to, and complement the ongoing discussion on these topics.

---

**KEY WORDS:** structured expert judgement, performance based weighting schemes, elicitation protocol, aggregation, confidence

---

## 1. INTRODUCTION

Risk assessment and management are often characterised by imminent decisions and scarce resources. Usually, at least some requisite data for pressing decisions are unreliable, rudimentary or entirely absent. In these situations, policy makers rely on expert judgements to fill the knowledge gaps. Experts estimate quantities (parameters of models), decide the forms of cause-effect relationships, and predict the outcomes of management interventions, all of which are uncertain. We restrict our attention

to the elicitation of parameters and their associated uncertainties from experts. How these are best elicited and combined across experts is critical to a decision process, as differences in the efficacy and robustness of elicitation and aggregation methods can be substantial (e.g. <sup>(42,45,17,21)</sup>).

It is generally agreed that in scientific disciplines, experts are very sensitive to contextual and psychological frailties including anchoring, dominance effects, and overconfidence, which can severely distort technical judgements (e.g. <sup>(32,49,13,39)</sup>), and that many of these frailties can be mitigated using structured elicitations that involve groups rather than individuals (e.g. <sup>(21,50,15,42)</sup>). Since the 1990's, and building on developments in decision theory, mathematics and behavioural science, some researchers have argued that substantial improvement in the quality and reliability of aggregated expert judgements can be achieved by using performance-based weighting schemes of aggregation <sup>(12,24,1)</sup>. However, heated debate continues in the literature regarding the practical merits of performance-

<sup>1</sup>Centre of Excellence for Biosecurity Risk Analysis, University of Melbourne, Parkville, VIC, Australia, ahanea@unimelb.edu.au

<sup>2</sup>School of BioSciences, University of Melbourne, Parkville, VIC, Australia,

<sup>3</sup>Harvard Forest, Harvard University, Petersham, MA, USA, mfc.mcbride@gmail.com

<sup>4</sup>Centre for Environmental Policy, Imperial College London, London, UK, markab@unimelb.edu.au

<sup>5</sup>Centre for the Study of Existential Risk, University of Cambridge, Cambridge, UK, bondiewints@gmail.com

based weighting, with opponents arguing that stable individual differences between experts on which to base weights are too difficult to measure<sup>(6)</sup>.

Several (different) elicitation protocols developed over the last decades have been deployed successfully in political science, infrastructure planning and volcanology (e.g.<sup>(24,1,2,45,5)</sup>). Most follow thoroughly documented methodological rules, but they differ in several aspects, including the way interaction between experts is handled, and the way an aggregated opinion is obtained from individual experts. There is no single, best structured expert judgement (SEJ) protocol; each has strengths and weaknesses (e.g.<sup>(5,45)</sup>). A SEJ for uncertainty quantification can be roughly defined as a formalised procedure that:

1. Asks questions that have clear operational meanings;
2. Follows transparent methodological rules, i.e. is traceable, repeatable and open to review;
3. Anticipates and mitigates some of the most important psychological and motivational biases;
4. Is thoroughly documented;
5. Provides opportunities for empirical evaluation and validation.

These elements make SEJ accountable, transparent and repeatable<sup>6</sup>. If expert opinion is used as scientific data, then it should be subject to the same kind of methodological rules for quality assurance that are applied to other types of empirical data<sup>(23,21)</sup>.

This paper aims to contribute to, and complement the ongoing discussion about the reliability and validity of different aggregation schemes for quantitative expert judgements. The most recent debate presented in<sup>(6)</sup> and its accompanying commentaries<sup>(20)</sup>,<sup>(54)</sup>,<sup>(42)</sup>,<sup>(7)</sup> acted as a catalyst for writing this paper. We will discuss the following three fundamental conjectures:

1. Prior performance can be a useful guide to future performance.
2. There is no clear relationship between the various measurable qualities of experts judgements.

<sup>6</sup>Note that this definition includes the one formulated by Roger Cooke: “The qualifier *structured* means that expert judgment is treated as scientific data, albeit scientific data of a new type.” in<sup>(19)</sup>.

3. An extensive, facilitated discussion among experts, prior to eliciting their individual final estimates improves experts’ performance.

Various papers support the above conjectures and we will refer to them in due course, however this paper is not intended to be a literature review. We will complement the discussion on these three conjectures by analysing the data collected from a large scale elicitation exercise. The elicitation exercise used the IDEA protocol for SEJ<sup>(55,31)</sup>. A brief overview of IDEA, its development, and the data can be found in Section 2. Section 3.1 discusses the first conjecture. It introduces ideas about differential weighting that are somewhat in contrast to those espoused in<sup>(6)</sup> and<sup>(7)</sup>. It also touches on the paramount importance of validating expert opinion. Section 3.2 discusses the second topic. Measuring the quality of expert judgements, and the experts’ performance as uncertainty quantifiers, seems essential if one wants to treat expert data in the same way as scientific data. However, the various measures available for evaluating expert performance in making quantitative estimates reward different dimensions of performance, and may lead to differences in relative performance. Given that the question of whether and how to value the different dimensions of expert performance that exist arise regularly in discussions of expert performance evaluation, we believe an investigation into the question of whether any relationship tends to exist between these different measures in practice is of interest and worth discussing further. The discussion from<sup>(6)</sup>,<sup>(54)</sup>,<sup>(42)</sup>, and<sup>(7)</sup> offers little background on this topic, aside from a clear criticism of the relative use of these measures in the weighting scheme proposed in<sup>(21)</sup>. Because of this lack of context, some claims made in [6] and [51] may be misinterpreted. Section 3.2 makes statements which reflect the authors position on this topic and may seem in disagreement with certain claims made in<sup>(6)</sup> and<sup>(54)</sup>. Section 3.3 relates to the third topic. It follows and supports (to some extent) ideas presented in<sup>(6)</sup> and<sup>(7)</sup>. Arguments from Section 3 will be supported by and discussed in the light of data analysis. Section 4 concludes the paper.

## 2. IDEA: A DELPHI-LIKE PROTOCOL WITH A TWIST

There are two main ways in which experts’ judgements are pooled<sup>(17)</sup>. One method is usually referred

to as *behavioural aggregation*, and involves striving for consensus via discussion<sup>(45)</sup>. When experts (initially) disagree, the advocates of behavioural aggregation will advise on a discussion between the experts with divergent opinions, resulting in a “self-weighting” through consensus<sup>7</sup>. But this comes at the cost of verifiability and reproducibility. Moreover, while well-functioning behavioural groups offer participants the opportunity to share knowledge and correct misunderstandings, such interaction is frequently prone to biases including overconfidence, polarisation of judgements and groupthink<sup>(33)</sup>.

Mathematical rules used in *mathematical aggregation* provide a more explicit, auditable and objective approach. A weighted linear combination of opinions is one example of such a rule. Equal weighting is often used mostly because of its simplicity (no justification for weights is required). Evidence also shows that the equal weighting scheme frequently performs quite well relative to more sophisticated aggregation methods (e.g.<sup>(17)</sup>), but not always<sup>(24,20)</sup>. There are reasons to believe that differential weighting based on anything (e.g. self-ratings, peer-ratings, citation indices) other than performance should be avoided<sup>(56,15,22)</sup>. Probably the most well known and widely used version of a differential weighting scheme is the Classical Model or Cooke’s model (CM) for SEJ<sup>(21)</sup>. Perhaps its most distinguishing feature is the use of calibration variables<sup>8</sup> to derive performance based weights.

Mixed SEJ protocols combine behavioural and mathematical aggregation techniques<sup>(28)</sup>. The most common mixed protocol is the Delphi method<sup>(48)</sup>, in which experts receive feedback over successive questionnaire rounds, in the form of other group members’ judgements. Experts remain anonymous and do not interact with one another directly. Instead, a facilitator provides feedback between rounds. As originally conceived the Delphi method strives to reach consensus after a relatively small number of rounds<sup>(26)</sup>, though in modern usages achieving consensus is not necessarily the primary aim (e.g.<sup>(53)</sup>). While research supports a general

conclusion that Delphi method can improve accuracy over successive rounds, this is by no means guaranteed. Critical reviews suggest that even though individual judgements may converge after a number of rounds<sup>(53)</sup>, this convergence does not necessarily lead to greater accuracy (e.g.<sup>(44,8)</sup>). Increasingly it appears that for genuine improvement in judgement quality to take place between rounds, it is important that feedback include the reasoning behind the expert’s opinion (i.e.<sup>(8,9)</sup>). Moreover, while the Delphi permits sharing of rationales, it still places limitations on the degree to which information sharing can take freely place between experts, with the result that misunderstandings and linguistic ambiguity are more difficult to remove. For example, a recent study that attempted to employ the Delphi approach for use in eliciting probability distributions experienced considerable delays and difficulties possibly arising due to the lack of expert - expert interaction and expert - facilitator interaction, leading to a decreased ability to clarify reasoning, assumptions and the details of the survey instrument<sup>(4)</sup>.

## 2.1 IDEA Protocol

The IDEA protocol described in this section synthesizes specific elements from all the approaches described above. In doing so, it aims to minimize the disadvantages of existing approaches and optimise their advantages. The majority of elements that characterise IDEA are not new; its most important contribution is in the structured approach to the combination of these elements (a similar approach is also proposed in<sup>(7)</sup>, yet informally). The elements of the IDEA protocol, and the reasons behind these choices are discussed in great detail in<sup>(13,31)</sup>. In this paper we will reiterate the specific steps, without repeating the discussion in the above mentioned references. IDEA is so-called because it encourages experts to Investigate, Discuss, and Estimate, and concludes with a mathematical Aggregation of judgements. It is a Delphi-like protocol in that experts give individual judgements over subsequent rounds, and facilitators provide feedback. In contrast to the traditional Delphi, IDEA does not seek consensus and can not always ensure full anonymity. A diverse group of experts first answer questions without engaging in discussion. Experts are then provided with the judgements of their peers and have the opportunity to endorse agreements and discuss differences of opinion (unlike Delphi), allowing people to reconcile the meanings

<sup>7</sup>However, where a group consensus judgement cannot be reached, individual expert distributions can be elicited and combined using a mathematical aggregation technique. Or alternatively, where consensus is not the aim, the resulting spread of expert viewpoints following discussion can be maintained and presented to decision-makers<sup>(42)</sup>.

<sup>8</sup>Calibration variables are variables taken from the experts’ domain for which the true values are known, or will become known, within the time frame of the study<sup>(1)</sup>.

of words and context (e.g.<sup>(16)</sup>). The discussion may be remote (on an on-line platform or over email) or face-to-face. Facilitators encourage and moderate discussion between rounds. In typical face-to-face discussions, people are often forced to declare the identity of their estimates to others, exposing them to dominant individuals, halo effects and potential pressure to conform. Conversely, people may refrain from defending their estimate in order to avoid identification. In IDEA, while the complete anonymity that characterizes a traditional Delphi process is lost, making the second estimate strictly anonymous largely mitigates these phenomena. This promotes the benefits of behavioural aggregation, while guarding against some of the most debilitating elements of group elicitation<sup>(14,39)</sup>. Expert estimates obtained through the second round (post-discussion) are mathematically aggregated. The method for mathematical aggregation can be chosen by the practitioners and it is not dictated by the IDEA protocol. Because IDEA does not strive for consensus, on the contrary encourages independent anonymous final estimates, the dangers of reaching an artificial consensus are mitigated against. Differences of opinion are recorded and can be easily communicated to decision makers.

IDEA uses a structured procedure for questioning experts about uncertain variables, described in detail in<sup>(14)</sup>. When used to elicit continuous quantities this procedure uses four questions to elicit the values of variables (corresponding to different quantiles), termed a four-step format. A slightly different version of this procedure (for fixed quantiles) corresponds to the way questions are asked in the CM, so the mathematics used in the CM for scoring and aggregating expert judgements can be easily used in the last step of IDEA.

When eliciting probabilities of event occurrences, IDEA uses three questions, termed a three-step format, one for the expert's *best estimate* about the probability of occurrence, and the other two for an interval that captures uncertainty around it. Other approaches, including CM, ask the experts to assign events to probability bins  $b_i = (p_i; 1 - p_i)$ , where  $p_i$  corresponds to the probability of occurrence. Bins can have the following form:  $b_1 = (0.1; 0.9)$ ,  $b_2 = (0.2; 0.8)$ ,  $b_3 = (0.3; 0.7)$ , etc. if the probability of occurrence scale is discretised into 10 intervals. An expert assigns an event to the  $b_2$  bin if their best estimate (about the probability of occurrence) is anywhere between 0.1 and 0.2. So, these approaches

only ask for best estimates, acknowledging the imprecision in the experts' judgements by allowing a fixed interval around them (equal to the respective bin's length). IDEA allows the expert to choose their own interval. The probabilities of binary variables can sometimes be interpreted in terms of relative frequencies. It is then legitimate to ask experts to quantify their degree of belief using a subjective distribution. In this case the upper and lower bounds asked for in the three-step format may be thought of as quantiles of this subjective probability distribution. However, when the relative frequency interpretation is not appropriate (i.e. when the probability of a unique event is elicited) the three-step format may be criticised for lacking operational definitions for the upper and lower bounds, in a probabilistic framework. Then, the elicited bounds can be thought of as an analogue of the fixed bounds of the bins, when the probability of occurrence scale is discretised into bins (see above). The imprecision in the experts' judgements is decided and quantified by the experts themselves, rather than chosen by the analyst.

The main reason to elicit bounds in such cases is to improve thinking about the best estimates. Another reason is to provide an indication of relative participant uncertainty. Because the data analysed in this paper consists of elicited unique events' probabilities, most measures of performance that we consider and use to compare expert performances are calculated using the elicited best estimates, rather than the bounds.

Irrespective of the way the probabilities are interpreted, the mathematical apparatus of the CM and its specific feature of using calibration variables can be used in the last step of IDEA.

## 2.2 Data

The IDEA protocol was refined and tested as part of a forecasting "tournament" that started in 2011 as an initiative of the US Intelligence Advanced Research Projects Activity (IARPA)<sup>9</sup>. Five university-based research teams were involved in predicting hundreds of geopolitical, economic and military events, with the goal of finding the key characteristics of efficient protocols for eliciting and aggregating accurate probabilistic judgements. The project used real events that resolved in the near-

<sup>9</sup><http://www.iarpa.gov/index.php/research-programs/ace>

future to test the accuracy of forecasts. Thousands of forecasters made over a million forecasts on hundreds of questions<sup>(52,36)</sup>. The Good Judgment Project (GJP) team who won the tournament<sup>(36)</sup> encouraged similar *think again* and *estimate lower and upper bounds* style practices e.g.<sup>(37)</sup>, though to our knowledge, IDEA was unique in eliciting multiple, uncertainty judgements for each estimate. As opposed to other protocols, the IDEA protocol combines both feedback and (facilitated) interaction. The benefits of mathematical aggregation for combining forecasters' opinions were also investigated by other teams e.g.<sup>(3)</sup>, but none considered performance weighting.

The results shown in this paper are based on the data gathered through this tournament. The data elicited with the IDEA protocol represent the answers to a subset of the questions developed by IARPA. All questions considered correspond to Bernoulli variables of the following sort: "Will the Turkish government release imprisoned Kurdish rebel leader Abdullah Ocalan before 1 April 2013?", which were answered using the three-step format outlined above. All questions usually resolved within 12 months, hence they were suited for empirical validation studies. The elicitation took place remotely, initially via email, and from the second year of the tournament through a dedicated website<sup>10</sup> which was set up for the participants to answer the questions and upload/download necessary materials. This website facilitated discussion via a discussion board that could be used to share relevant resources from outside the platform (e.g. web links), and to comment on and rate the quality of the information shared by others.

In answering the questions, experts followed the three-step format (see Figure 1). The best estimate of each expert represents their subjective degree of belief in the probability of the occurrence of the event. Accepting that the judgements of experts are inherently imprecise and trying to formalise this imprecision, the bounds (about which experts are asked prior to asking the best estimate in order to avoid anchoring effects) are thought of as an uncertainty interval. The width of this interval provides an indication of how sure the expert is of their estimate.

The tournament operated on a yearly basis, over the course of four years. Each year, new participants

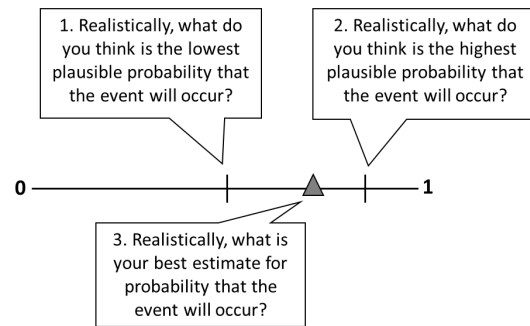


Fig. 1. A three-step format question

joined the IDEA groups, and other participants dropped out. There were 150 participants (over the four years) who answered at least one question (both rounds). Eight of these participants returned each year. The level of participants' expertise covered a very wide range from self-taught individuals with specialist knowledge to intelligence analyst<sup>11</sup>. A total of 155 questions were answered by at least one participant. However, no participant answered more than 96 questions. The participants were divided into groups and the number of groups varied across years to keep the number of participants per group fairly constant (typically ten). Starting from the third year *Super-groups* were formed<sup>12</sup> composed of the best performing participants from the previous year<sup>13</sup>. The participants composing the Super-group were unaware of the fact that their group was in any way special. The number of participants composing the Super-group was equal to the number of participants from any other group. This construction was meant to help us accumulate evidence around whether prior performance is a useful guide to future performance on forecasting tasks.

<sup>11</sup>Because the study was conducted with people who may be considered in some contexts to be non-expert participants in the remainder of this paper, whenever data from this study are concerned, we will use the term *participant* instead of *expert*.

<sup>12</sup>The GJP team tested an idea similar to the super-group of top forecasters, but on a considerably greater scale, i.e. with the benefit of many more participants.

<sup>13</sup>Performance was measured using the average Brier score. This measure (defined later in this paper) was imposed by the forecasting tournament rules and all participating team had to use it.

<sup>10</sup><http://intelgame.acera.unimelb.edu.au/>

### 3. PERFORMANCE MEASURES AND THE VALUE OF DISCUSSION

In some protocols using mathematical aggregation, the interaction between experts is limited to training and briefing (e.g.<sup>(21)</sup>). It is generally believed that interaction between experts may harm aggregation because it induces correlation between judgements (e.g.<sup>(45)</sup>). Nevertheless, when experts share and debate their knowledge during interaction, they establish a common and better understanding of what is being asked of them (e.g.<sup>(33)</sup>). We postulate that, in general any additional dependence between judgements introduced through the discussion between rounds is justified by the increase in information resulting from discussion and by the reduction of misunderstandings and linguistic uncertainty. Results based on the dataset described above support this postulate and are presented in<sup>(30)</sup>. The dependence structures within and between the groups (before and after discussion) are investigated in<sup>(30)</sup> and found to be effectively the same<sup>14</sup>. Due to these findings we feel comfortable in pooling the expert data from all groups and across all years to form a larger dataset and hence permit more powerful statistical testing.

In evaluating the performance of probability forecasters, we assess their accuracy and their confidence. While alternative definitions of *accuracy* exist in the literature on evaluating subjective probabilities (e.g.<sup>(34), (57), (41)</sup>), we will hereafter consider accuracy to be a measure of the degree of correspondence between the participants' predictions (which are the probabilities of event occurrences) and the observed outcomes. The average Brier score<sup>(10)</sup> is used to evaluate and compare participants' long term accuracy. The Brier score for binary variables ranges from 0 for a perfect forecast to 2 for the worst possible forecast. Although the score can be computed based on a single forecast, performance on many questions is necessary to provide a reliable guide to accuracy. In our case, long term performance is calculated as an average Brier score, using the participants best estimates of the event probabilities (i.e. Figure 1, box 3). The average Brier score can be further decomposed into two additive components called calibration and refinement<sup>(43)</sup>. The calibration

component of the Brier score is one of the measures thoroughly investigated in<sup>(30)</sup> as a measure of performance. In this paper we will also use it in Section 3.2. We treat the length of the uncertainty intervals provided by participants around their best estimates as a measure of their confidence. Under this interpretation, a wider interval denotes decreased confidence, while a narrower interval denotes an increased level of confidence. The length of these intervals can also be thought of as a measure of informativeness: wide intervals are uninformative.

A more extensive discussion about different concepts and measures of accuracy, calibration and confidence, and different uses of these terms can be found in<sup>(31)</sup>. A description of the calibration term of the Brier score and other measures of confidence and informativeness calculated for the participants' assessments are discussed and analysed in<sup>(30)</sup>. For the sake of conciseness we shall not repeat these discussions here, but merely refer the reader to the two mentioned references.

#### 3.1 Prior and future performance

Each year we compared the performance of the different groups. When we consider groups of participants, we are interested in the group accuracy as measured by the average Brier score of the group aggregated estimates. Specifically, for each group, the best estimates of the participants for a given question are averaged, and a Brier score per question, per group is calculated. The average Brier score (across all questions answered by at least one participant from that group) is reported as the accuracy measure of the group. Each year we compared the accuracy of the equally weighted opinions of the groups after discussion, using a within-subject design.

In one year alone (2013 - 2014) we had sufficient data to calculate differential weights using the CM<sup>(21)</sup>. We used the best estimate responses of the participants in the IDEA Super-group 1 to calculate calibration and information scores for discrete variables defined in<sup>(21)</sup> and reiterated in<sup>(30)</sup>. Performance-based weights were then derived based on these calibrations scores, and the best estimates of the participants were weighted accordingly to obtain the differentially weighted response per question, per group. Brier scores were then calculated per question, per group, for the new performance weighted responses. Figure 2 shows the average Brier scores of the equally weighted combination of judgements for each group after discussion, together with standard

<sup>14</sup>The dependence between experts' answers is measured using the average correlation coefficient of experts within each group and conditional on the outcomes. For details of the calculations we refer to<sup>(30)</sup>.

error bars, for 2013 - 2014. The performance-based differentially weighted combination of the Super-group participants' judgements is shown in the same figure.

Although not statistically significant (as indicated by the t-test), the Super-group G1 outperformed the other groups of participants, suggesting that prior performance is a useful guide to future performance on similar estimation tasks (out of sample). The same was observed for all the other years of the tournament (except the first year, where no super-group was used.). This finding is in agreement with the findings of<sup>(38)</sup>.

Employing the CM, a slight improvement in performance is achieved using this weighting scheme (when compared to equal weighting). Even though not significant, this improvement provides additional evidence to support prior performance as a useful guide to future performance (in - sample).

Asking questions that resolve in the near future and using performance measures allows for validation of the utility of predictions from the particular group of experts being consulted, which is an essential step for any model. Several methods for validating expert opinion have now been used and proposed<sup>(27,25)</sup> mainly by the advocates of the CM<sup>15</sup>. These methods include in-sample validation and cross-validation, if variables of interest (the variables for which expert judgement is required) are unobservable on relevant time scales. We conjecture that well-designed cross-validation studies offer a credible approach to validating expert judgement and the utility of performance based weighting<sup>(25,27)</sup>, certainly preferable to no validation at all<sup>16</sup>. The most recent cross-validation study performed by the authors of<sup>(18)</sup> offers a high level of support for the idea that prior performance a useful guide to future performance: for 54 out of 73 professional expert elicitations, the performance based combination of experts' judgements outperformed the equally weighted combination.

We acknowledge and share the concerns of the authors of<sup>(6,7)</sup> about true validation studies (i.e. where the value of expert performance-based weighting for predicting the target variables of interest in a genuine application of expert knowledge

is able to be evaluated post-hoc). To this end, we hope that the out-of-sample validation results obtained with IDEA and presented here and in<sup>(31,30)</sup> strengthen the arguments discussed in<sup>(27,25,20,18)</sup>.

Moreover, when judgements conflict, one will (usually) turn out to be "better" than the other. How can we know, a priori, whose opinion is valid (or more reliable) and should assigned more weight in an aggregation? Previous studies (e.g.<sup>(15,22)</sup>) assure us that we cannot, and the only way of knowing is by assessing prior performance on similar tasks.<sup>(42)</sup> argues that the "outlier's views should not get masked by combining multiple experts". While we agree with this statement in terms of the importance of retaining experts' un-aggregated opinions along with any group aggregate judgement in order to avoid masking individual differences of opinion, we argue that aggregation is often necessary, and that where information on past performance is available it can provide a valuable way to discriminate between (and aggregate accordingly) situations when an outlier's views are likely to be worthwhile, and those where they are not.

The topics of measuring performance and differential weighting are obviously interrelated and the above discussion suggests that there are good validated reasons to base differential weighting schemes on measures of prior performance on similar tasks. The question that follows naturally is which measures of performance should one use. If experts' answers have one quality (e.g. accuracy) would this guarantee other qualities (e.g. informativeness)? Does this relationship depend on the way we define and measure these qualities? These questions are not new, and such relationships between various measures of performance have been investigated before in the context of subjective probability (distributions) elicitations. The next section adds yet another case study to the existing literature and the ongoing debate.

### 3.2 The correlation between performance measures

Ideally, any increase in an expert's confidence would be matched with an increase in their accuracy or calibration. However research across many different measures of confidence, accuracy and calibration has found that these measures are often poorly correlated with one another (e.g.<sup>(34)</sup>,<sup>(29)</sup>,<sup>(40)</sup>,<sup>(2)</sup>).

In asking experts to quantify parameters and their uncertainty around these estimates, SEJ meth-

<sup>15</sup>The data used in many of these studies are freely available online at <http://rogermcooke.net/>

<sup>16</sup>In other words (using the terminology in<sup>(7)</sup>) we argue that starting from *data* in our journey towards data will get us further than starting from a void.




Fig. 2. Forecasting tournament 2013-2014

ods rely on each expert's ability to realistically quantify the limitations of their knowledge. When experts underestimate the uncertainty in their judgements, apparent overconfidence occurs, leading to faulty assessments on which hazardous decisions may be based. Even though less uncertainty around a judgement (more confidence) often translates to greater informativeness, we believe that certainty should only be valued when associated with well-calibrated assessments. This contrasts with our interpretation of the views expressed in<sup>(54)</sup>, who seemingly views narrow (*informative*) distributions as more important than calibrated ones<sup>17</sup>.

This being said, we would like to emphasize that both measures are necessary in evaluating performance, and they should be evaluated in tandem whenever possible to avoid rewarding high performance in one measure at the expense of the other.

In this section we investigate the relationship between accuracy (measured by the average Brier score) and confidence (as measured by uncertainty interval length) using the dataset described in

<sup>17</sup>“I feel that in most cases, informativeness is more important than calibration”

Section 2.2. Because the average Brier score can be decomposed such that we can measure calibration as well, and the calibration component of the Brier score is one of the measures thoroughly investigated in<sup>(30)</sup> we will extend the analysis and investigate the relationship between calibration and confidence using the same data set.

Figure 3 looks at these two relationships in terms of rank correlations. We have selected the participants who answered at least 20 questions, to reduce the variance of the average measures of performance estimators.

The rank correlation between the average Brier score and the average confidence is represented by stars. The rank correlation between the calibration term of the average Brier score and the average confidence is represented by circles. These correlations are shown as a function of the minimum number of questions ( $n$ ) answered by participants. We denote  $m$  the number of participants who answered at least  $n$  questions. We added  $n$ 's corresponding  $m$  just above the  $X$ -axis (e.g. for  $n = 20$ ,  $m = 53$ ). Each rank correlation is calculated based on  $m$  samples, and the value of  $m$  determines if the correlation is statistically different than zero. The only values found statistically different than zero at




Fig. 3. Confidence vs. accuracy and calibration

0.1 significance level, when performing a two-tailed statistical test<sup>(46)</sup> are the ones enclosed by the ellipse from Figure 3. However, *none* of the correlations shown in Figure 3 are statistically different than zero at 0.05 significance level.

Based on the above analysis, there is very little reason to believe that an increase in an expert’s confidence would correspond to an increase in their accuracy or calibration. This warns us against relying too heavily on any one dimension of expert performance when evaluating experts under the assumption that adequate performance according to one metric of expert performance (e.g. informativeness) would suggest it for others (e.g. calibration). Moreover, if these evaluations are used to form weights, the weights should reflect this dual importance and evaluate for adequate performance across multiple dimensions of expert performance.


### 3.3 Does discussion help?

The authors of<sup>(7)</sup> recommend behavioural aggregation as an alternative to the CM. We agree with only one aspect of this recommendation, the possibility of facilitated discussion. IDEA assimilated the extensive discussion aspect of the behavioural aggregation, but strongly advises against striving for consensus and the “self-weighting” associated with it. The reasons for this advice are briefly addressed in

this paper (see section 2.1) and thoroughly discussed in<sup>(14)</sup>.

We consistently find (throughout our studies) that sharing information between rounds appears to improve on average both accuracy and confidence. Figure 4 is one more example of this phenomenon. The average Brier scores of all participants before discussion are plotted against their average Brier scores after discussion. These are represented by stars. If participants do not change their minds (about any question) in the second round, their average Brier score does not change (so the point representing this pair of scores will fall on the main diagonal). The majority of points in this scatter plot fall under the main diagonal, indicating that the majority of participants who change their mind, changed in the direction of the realised outcome. The average confidence of the participants before discussion is plotted against their average confidence after discussion in the same plot. These are represented by plus (+) signs. The same general pattern can be observed. Confidence was greater after discussion.

Because our study does not include a control group that provides both first and second round judgements but do not undertake discussion in-between, we cannot rule out the possibility that any changes in performance observed are due to iteration alone, rather than the discussion that took place during each iteration. However, this



**Fig. 4.** Average Brier Scores & Average Confidence before and after discussion. Both measures have similar possible ranges allowing for presentation on the same plot. Averages are calculated for each participant across all questions.

within-subject quasi-experimental approach is not unusual in studying the performance of Delphi-related methods (e.g.<sup>(47)</sup>). Our result complements other work which reveals the potential importance that any feedback provided include the reasoning behind the expert’s opinion (e.g.<sup>(8,9)</sup>) in order to contribute to improvements in performance between different rounds on estimation.

In the last three years of the tournament the data was collected through a dedicated website. This gave us the opportunity to collect information about the participants’ level of engagement in the discussion phase. During the discussion phase, the participants contributed comments and shared resources (i.e. additional relevant information such as links to news articles, scientific papers, blog posts, and additional data) with the group. After resources were added on the group discussion board, they could be viewed and rated. We recorded the number of these actions. These activities define to some extent the engagement of participants in the discussion between the rounds, and we can look at the relationship between measures of performance and these activity measures. We have (again) selected the participants who answered at least 20 questions, to reduce the variance of the average measures of performance and activity estimators. We first look at the rank correlation between the average Brier score and:

1) the average number of comments (star), 2) the average number of added resources (circle), 3) the average number of viewed resources (plus), and 4) the average number of rated resources (diamond). These correlations are shown as a function of the minimum number of questions ( $n$ ) answered by ( $m$ ) participants. Just like in the previous section, each rank correlation is calculated based on  $m$  samples, and the value of  $m$  determines if the correlation is statistically different than zero. Figure 5 shows these correlations. All correlations below the horizontal line are significantly different than zero at 0.05 significance level.

There is strong negative correlation between the average activity level (as defined by the number of comments made, and by the number of added and viewed resources) and the average Brier score, which means that higher levels of activity correspond to lower, hence better scores. This result supports our conjecture that discussion improves (at least one aspect) of performance.

However, when looking at the relationships between activity and the length of the uncertainty intervals, the rank correlations are much smaller and none is significantly different than zero (see Figure 6). Given this apparent weaker relationship observed between activity and confidence we speculate that the shift (increase) in average participant confidence following discussion (Figure 4) may be the result of receiving feedback (i.e. exposure to the intervals of the other group members) rather than discussion, as this is the element of the discussion phase that would not necessarily relate to level of on-line activity and interaction with other group members.

#### 4. DISCUSSION

The aggregation of expert judgements into a single, transparently constructed estimate is one of the most important steps in eliciting and using expert judgements. Equal weights are intuitively appealing because there is no need to defend differential weights, they are easy to understand, and all experts feel that they contribute equally. However, one of the most important lessons of empirical studies over the last decade is that an expert’s performance on technical questions may be predicted to some extent by the history of their performance on similar questions previously<sup>(18,37,11)</sup>. Taking advantage of this phenomenon, Cooke’s approach to differential weighting circumvents the difficult issue of the different scales associated with

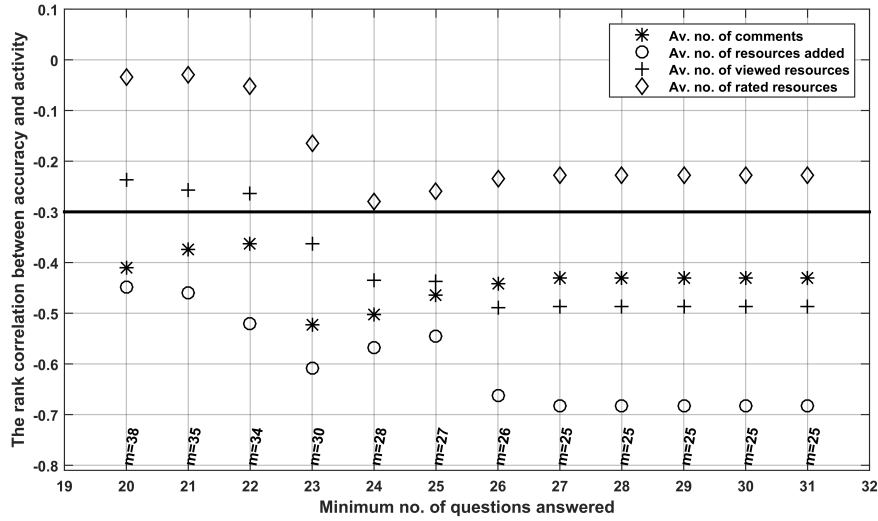


Fig. 5. Average Brier Scores & Average Level of Activity




Fig. 6. Average Confidence & Average Level of Activity

different kinds of questions, and assimilates each expert's confidence and statistical accuracy into a single weight. The result is the potential for an improvement in group performance over equal-weighted based aggregation methods. Our results suggest that even in the relatively difficult conditions imposed in answering binary questions on the outcomes of geopolitical events, performance based differential weights calculated using Cooke's method

may improve upon the performance of groups, even those composed of comparatively higher performing forecasters, such as those in to the super-groups.

We agree with<sup>(6)</sup> that the way forward for expert knowledge research involves additional experiments and analyses to further test the benefits of performance weighting schemes against equal weighting and behavioural aggregation. Our paper provides one such contribution, by drawing on recent experimental

findings. The discussion papers<sup>(6)</sup>,<sup>(54)</sup>,<sup>(42)</sup>, and<sup>(7)</sup> were published in 2015, and arguably drew more heavily on an older body of literature, findings and established positions around the relative merits of performance weighting. We have endeavoured to support this dialogue by contributing more recent findings and methodologies for drawing on the benefits of both behavioural aggregation (here via facilitated discussion) and performance weighting. Both our study, and that of the GJP present empirical evidence incorporating out-of-sample testing indicating that prior performance predicts future performance. We argue that this provides grounds for the use and the continued investigation of performance-based weighting.

Our paper also explored the contested question of whether social interaction between group members erodes judgement quality<sup>(35,51)</sup>. As with IDEA, the GJP also made use of structured discussion between forecasters and found that it improved performance<sup>(38)</sup>. In their case, group members were known to each other via their online identifiers only. They could offer rationales and critiques, and share information, including their predictions (although, unlike in IDEA, formal feedback on initial group member predictions was not provided). The results of both approaches illuminate the value of facilitated conversations between group participants in reconciling language-based misunderstandings and interpretations of evidence. They also both suggest that controlled interaction and feedback can enhance performance, without overwhelming these benefits with dysfunctional group dynamics. Whether or not these findings hold up in different settings, such as those where group members are motivated to seek status or guard information, remains to be seen.

We conclude by reiterating the recommendation for risk analysts and decision makers to use consistent, transparent and repeatable approaches to SEJ. As data accumulate over time, the strengths and weaknesses of these new methodologies for aggregating expert judgements will become clearer and help to ensure better judgements in the future.

## ACKNOWLEDGMENT

The authors would like thank the reviewers for their thorough and numerous comments.

The authors declare no conflict of interest.

## REFERENCES

1. W. Aspinall. A route to more tractable expert advice. *Nature*, 463:294–295, 2010.
2. W. Aspinall and R.M. Cooke. Quantifying scientific uncertainty from expert judgement elicitation. In J. Rougier, S. Sparks, and L. Hill, editors, *Risk and uncertainty assessment for natural hazards*, chapter 10, pages 64–99. Cambridge University Press, Cambridge, 2013.
3. J. Baron, B.A. Mellers, P.E. Tetlock, E. Stone, and L. Ungar. Two reasons to make aggregated probability forecasts more extreme. *Decision Analysis*, 11:451–468, 2014.
4. F. Bolger, A. Hanea, A. O’ Hagan, O. Mosbach-Schulz, J. Oakley, G. Rowe, and M. Wenholt. Case Study in Plant Health. *EFSA Journal*, Guidance on Expert Knowledge Elicitation in Food and Feed Safety Risk Assessment(Parma):Italy, 2014.
5. F. Bolger, A. Hanea, A. O’ Hagan, O. Mosbach-Schulz, J. Oakley, G. Rowe, and M. Wenholt. Guidance on Expert Knowledge Elicitation in Food and Feed Safety Risk Assessment. *EFSA Journal*, 12(6):Parma, Italy, 2014.
6. F. Bolger and G. Rowe. The aggregation of expert judgment: Do good things come to those who weight? *Risk Analysis*, 35:5–11, 2015.
7. F. Bolger and G. Rowe. There is Data, and then there is *Data*: Only Experimental Evidence will Determine the Utility of Differential Weighting of Expert Judgment. *Risk Analysis*, 35:21–26, 2015.
8. F. Bolger, A. Stranieri, G. Wright, and J. Yearwood. Does the Delphi process lead to increased accuracy in group-based judgmental forecasts or does it simply induce consensus amongst judgmental forecasters? *Technological forecasting and social change*, 78(9):16711680, 2011.
9. F. Bolger and G. Wright. Improving the Delphi process: lessons from social psychological research. *Technological forecasting and social change*, 78:15001513, 2011.
10. G.W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:1–3, 1950.
11. D.V. Budescu and E. Chen. Identifying expertise to extract the wisdom of crowds. *Management Science*, 61(2):267 – 280, 2015.
12. M. Burgman, A. Carr, L. Godden, R. Gregory, M. McBride, L. Flander, and L. Maguire. Redefining expertise and improving ecological judgment. *Conservation Letters*, 4:81–87, 2011.
13. M.A. Burgman. *Risks and Decisions for Conservation and Environmental Management*. Cambridge University Press, Cambridge, 2005.
14. M.A. Burgman. *Trusting judgements: how to get the best out of experts*. Cambridge University Press, 2015.
15. M.A. Burgman, M. McBride, R. Ashton, A. Speirs-Bridge, and L. et al. Flander. Expert status and performance. *PLoS ONE*, 6:e22998, 2011.
16. J. Carey and M. Burgman. Linguistic uncertainty in qualitative risk analysis and how to minimize it. *Annals of the New York Academy of Sciences*, 1128:13–17, 2008.
17. R. Clemen and R. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19:187–203, 1999.
18. A. R. Colson and R. M. Cooke. Cross validation for the classical model of structured expert judgement. *Reliability Engineering and System Safety*, 163:109–120, 2017.
19. R. M. Cooke. Special issue on expert judgment. *Reliability Engineering & System Safety*, 93(5):655–656, 2008.
20. R. M. Cooke. The aggregation of expert judgment: Do good things come to those who weight? *Metron*, 35:12–

- 15, 2015.
21. R.M. Cooke. *Experts in uncertainty: Opinion and subjective probability in science*. Environmental Ethics and Science Policy Series. Oxford University Press, 1991.
  22. R.M. Cooke, S. ElSaadany, and X. Huang. On the performance of social network and likelihood-based expert weighting schemes. *Reliability Engineering and System Safety*, 93(5):745–756, 2008.
  23. R.M. Cooke and L.H.J. Goossens. Procedures guide for structural expert judgement in accident consequence modelling. *Radiation Protection Dosimetry*, 90(3):303–309, 2000.
  24. R.M. Cooke and L.H.J. Goossens. TU Delft expert judgment data base. *Reliability Engineering and System Safety*, 93(5):657–674, 2008.
  25. R.M. Cooke, M.E. Wittmann, D.M. Lodge, J.D. Rothlisberger, E.S. Rutherford, H. Zhang, and D.M. Mason. Out-of-sample validation for structured expert judgment of Asian carp establishment in Lake Erie. *Integrated Environmental Assessment and Management*, 10(4):522–528, 2014.
  26. N. Dalkey. An experimental study of group opinion: the Delphi method. *Futures*, 1:408 – 426, 1969.
  27. J.W. Eggstaff, T.A. Mazzuchi, and S. Sarkani. The effect of the number of seed variables on the performance of Cooke’s classical model. *Reliability Engineering and System Safety*, 121:72–82, 2014.
  28. W. Ferrell. Discrete subjective probabilities and decision analysis: elicitation, calibration and combination. In *Subjective probability*, volume Eds. Wright, G. and Ayton, P. Cambridge Press, New York, 1994.
  29. D.W. Griffin and A. Tversky. The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24:411–435, 1992.
  30. A. Hanea, M. McBride, M. Burgman, and B. Wintle. Classical meets modern in the idea protocol for structured expert judgement. *Journal of Risk Research*, 2016.
  31. A. Hanea, M. McBride, M. Burgman, B. Wintle, F. Fidler, L. Flander, S. Mascaro, and B. Manning. *Investigate Discuss Estimate Aggregate* for structured expert judgement. *International Journal of Forecasting*, 33(1):267–279, 2016.
  32. D. Kahneman and A. Tversky. Choices, values, and frames. *American Psychologist*, 39(4):341–350, 1984.
  33. N.L. Kerr and R.S. Tindale. Group-based forecasting?: a social psychological analysis. *International Journal of Forecasting*, 27:14–40, 2011.
  34. S. Lichtenstein, B. Fischhoff, and L.D. Phillips. Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, and A. Tversky, editors, *Judgment under uncertainty: Heuristics and biases*, pages 306–334. Cambridge University Press, Cambridge, 1982.
  35. J. Lorenz, H. Rauhut, F. Schweitzer, and D. Helbing. How social influence can undermine the wisdom of crowd effect. volume 108, pages 9020–9025. Proceedings of the National Academy of Sciences of the United States of America, 2011.
  36. B. Mellers, E. Stone, P. Atanasov, N. Rohrbaugh, S.E. Metz, L. Ungar, M.M. Bishop, M. Horowitz, E. Merkle, and P. Tetlock. The psychology of intelligence analysis: Drivers of prediction accuracy in world politics. *Journal of Experimental Psychology: Applied*, 21:1–14, 2015.
  37. B. Mellers, E. Stone, T. Murray, A. Minster, N. Rohrbaugh, M. Bishop, and E. Chen. Identifying and cultivating superforecasters as a method of improving probabilistic predictions. *Perspectives on Psychological Science*, 10(3):267–281, 2015.
  38. B.A. Mellers, L. Ungar, J. Baron, J. Ramos, B. Gurcay, K. Fincher, and P.E. Tetlock. Psychological strategies for winning a geopolitical forecasting tournament. *Psychological Science*, 25(4), 2014.
  39. G. Montibeller and D. von Winterfeldt. Cognitive and motivational biases in decision and risk analysis. *Risk Analysis*, 35 (7):1230–1251, 2015.
  40. D. A. Moore, S.A. Swift, A. Minster, B. Mellers, L. Ungar, P. Tetlock, and E.R. Tenney. Confidence calibration in a multi-year geopolitical forecasting competition. *Working Paper*, 2015.
  41. D.A. Moore, E.R. Tenney, and U. Haran. Overprecision in judgment. In G. Keren and W. Wu, editors, *The Wiley Blackwell handbook of judgment and decision making*, pages 182–212. John Wiley & Sons, Ltd., Chichester, 2015.
  42. M. G. Morgan. Our Knowledge of the World is Often Not Simple: Policymakers Should Not Duck that Fact, But Should Deal with It. *Risk Analysis*, 35:19–20, 2015.
  43. A.H. Murphy. A new vector partition of the probability score. *Journal of Applied Meteorology*, 12(4):595–600, 1973.
  44. M.K. Murphy, N. Black, D. Lamping, C. Mckee, and C. Sanderson. Consensus development methods and their use in clinical guideline development. *Health Technology Assessment*, 2(3):1–88, 1998.
  45. A. O’Hagan, C.E. Buck, A. Daneshkhah, J.R. Eiser, P.H. Garthwaite, D.J. Jenkinson, J.E. Oakley, and T. Rakow. *Uncertain judgements: Eliciting experts’ probabilities*. Wiley, London, 2006.
  46. P.H. Ramsey. Critical values for Spearman’s rank order correlation. *Journal of educational statistics*, 14(3):245–253, 1989.
  47. G. Rowe and G. Wright. The Delphi technique as a forecasting tool: issues and analysis. *International journal of forecasting*, 15 (4):353–375, 1999.
  48. G. Rowe and G. Wright. Expert opinions in forecasting: the role of the Delphi technique. In *Principles of forecasting: A handbook for researchers and practitioners*, Norwell: Kluwer Academic Publishers:125–144, 2001.
  49. P. Slovic. Trust, emotion, sex, politics, and science: surveying the risk-assessment battle field. *Risk Analysis*, 19:689–701, 1999.
  50. C.R. Sunstein. Group judgments: Deliberation, statistical means, and information markets. *John M. Olin Law & Economics Working Paper No. 219, University of Chicago Law School*, 2004.
  51. R.S. Tindale and K. Kluwe. Decision making in groups and organizations. *The Wiley Blackwell Handbook of Judgment and Decision Making*, G. Wu, and G. Keren (Eds.):849–874, 2015.
  52. L.H. Ungar, B. Mellers, V.A. Satopaa, J. Baron, P.E. Tetlock, J. Ramos, and S. Swift. The good judgment project: A large scale test of different methods of combining expert predictions. AAI Fall Symposium Series, (AAAI Technical Report FS-12-06), 2012.
  53. H.A. von der Gracht. Consensus measurement in Delphi studies: Review and implications for future quality assurance. *Technological Forecasting & Social Change*, 79:15251536, 2012.
  54. R.L. Winkler. Equal Versus Differential Weighting in Combining Forecasts. *Risk Analysis*, 35:16 – 18, 2015.
  55. B. Wintle, M. Mascaro, F. Fidler, M. McBride, M. Burgman, L. Flander, G. Saw, C. Twardy, A. Lyon, and B. Manning. The Intelligence Game: Assessing Delphi groups and structured question formats. In *Proceedings of the 5th Australian Security and Intelligence Conference*, 2012.
  56. F. Woudenberg. An evaluation of Delphi. *Technological Forecasting and Social Change*, 40:131–150, 1991.
  57. I. Yaniv and D.P. Foster. Precision and accuracy of

judgmental estimation. *Journal of Behavioral Decision Making*, 10(1):21–32, 1997.

Author Manuscript