On measuring agreement with numerically bounded linguistic probability schemes: A re-analysis of data from Wintle, Fraser, Wills, Nicholson, and Fidler (2019)

Short Title: On measuring agreement with numerically bounded linguistic probability schemes

David R. Mandel^{1*}, Daniel Irwin²

¹Intelligence, Influence and Collaboration Section, Toronto Research Centre, Defence Research and Development Canada, Toronto, Ontario, Canada

²Department of National Defence, Toronto, Ontario, Canada

*Corresponding author

Email: drmandel66@gmail.com

Abstract

Across a wide range of domains, experts make probabilistic judgments under conditions of uncertainty to support decision-making. These judgments are often conveyed using linguistic expressions (e.g., *x is likely*). Seeking to foster shared understanding of these expressions between senders and receivers, the US intelligence community implemented a communication standard that prescribes a set of probability terms and assigns each term an equivalent numerical probability range. Wintle *et al.* (2019) tested whether access to the standard improves shared understanding and also explored the efficacy of various enhanced presentation formats. Notably, they found that embedding numeric equivalents in text (e.g., *x is likely [55-80%]*) substantially outperformed the status quo approach in terms of the percentage overlap between participants' interpretations of linguistic probabilities (defined in terms of the numeric range equivalents they provided for each term) and the numeric ranges in the standard. These results have important prescriptive implications, yet Wintle *et al.*'s percentage overlap measure of agreement may be viewed as unfairly punitive since it penalizes individuals for being *more* precise than the stipulated guidelines even when the individuals' interpretations fall perfectly within the stipulated ranges. Arguably, within-range precision is a positive attribute. Accordingly, we reanalyzed Wintle *et al.*'s data using an alternative measure of percentage overlap that does not penalize in-range precision. Predictably, we find that percentage overlap is elevated across conditions. More importantly, the effects of presentation format and probability level are highly consistent with the original study. By removing further ambiguity, these findings buttress the claim that the methods currently used by intelligence

organizations are ineffective at coordinating the meaning of uncertainty expressions between intelligence producers and intelligence consumers.

Introduction

Across a variety of domains, experts make probabilistic judgments under conditions of uncertainty to support decision-making (e.g., Murphy & Winkler, 1984; Morgan, 1998; Moxley et al., 2003; Budescu et al., 2009; Burgman et al., 2011; Ho et al., 2015). For instance, clinicians estimate the likelihood of certain treatment outcomes, meteorologists predict the weather, and intelligence analysts forecast geopolitical developments. Depending on the context, uncertainty may stem from the sense that critical information is missing or inconsistent with other pieces of information, doubts about the accuracy or completeness of one's causal model, awareness of the possibility of benign misinformation or malign deception, and even aversion to the risk of being proven wrong. In order to fulfill the decision-support function, expert assessors must not only be able to estimate their uncertainty to arrive at sound judgments, they must also effectively communicate those probabilistic judgments to their respective audiences. Without effective communication, even an accurate and timely judgment could potentially misinform the end-user.

Often, experts (and the professional or organizational bodies they may represent) communicate probabilistic judgments using linguistic expressions such as *likely* or *very likely* for probability levels that convey “more likely than not”; conversely, experts use terms such as *unlikely* or *very unlikely* to convey probability levels “less likely than not.” Political pundits, for example, are prone to using such language in their forecasts without specifying what the linguistic probabilities mean to them (Tetlock & Gardner, 2015). Although the free use of such language is still tolerated in some intelligence communities (Barnes, 2016), others have

implemented communication schemes aimed at fostering shared meaning between senders and receivers.

In the US, for instance, the Office of the Director of National Intelligence promulgated a communication standard under Intelligence Community Directive (ICD) 203 that prescribes a set of probability terms and assigns each term an equivalent numerical probability range (see Table 1). In practice, intelligence consumers are expected to reference a translation table to ensure that their interpretations of linguistic probabilities align with the prescribed meanings. NATO and the UK intelligence community follow the same basic approach, although the precise set of terms and numerical ranges assigned differ (Dhami & Mandel, forthcoming; Friedman, 2019; Irwin & Mandel, 2020). Numerically bounded linguistic probability (NBLP) schemes are used in many other domains as well. For instance, the Intergovernmental Panel on Climate Change (IPCC) uses this approach to communicate probabilistic assessments about global climate change (Mastrandrea *et al.*, 2011) and the European Commission recommends a similar system for communicating the risk of drug side effects (Webster *et al.*, 2017). NBLP schemes have also been proposed for use in extreme event attribution reports (Lewis *et al.*, 2019).

Table 1. Linguistic probability expressions and numeric range equivalents in US Intelligence Community Directive (ICD) 203. Note: ICD 203 provides two sets of linguistic probability expressions that are treated synonymously. In Wintle *et al.* (2019), participants were only shown terms from the first column and their numeric equivalents.

Almost no chance	Remote	1-5%
Very unlikely	Highly improbable	5-20%
Unlikely	Improbable (improbably)	20-45%
Roughly even chance	Roughly even odds	45-55%
Likely	Probable (probably)	55-80%
Very likely	Highly probable	80-95%
Almost certain(ly)	Nearly certain	95-99%

Since their adoption, a growing body of research has examined the extent to which NBLP schemes improve shared understanding, and such research has explored the efficacy of alternative presentation formats. Notably, Budescu *et al.* (2009, 2012, 2014) have examined the NBLP guideline used by the IPCC. This work has repeatedly shown that providing end-users access to a translation table only marginally improves compliance—namely, the shared agreement in meaning between the end-user and the standard. They have further demonstrated that including numeric range equivalents in text beside linguistic probabilities (e.g., *x is likely* [>66%]) yields better compliance than using the standard alone, but even this hybrid method of presenting linguistic probabilities with numeric probability ranges leaves considerable room for improvement. For instance, measuring percentage overlap from 0% (if the participant's interpretation fell completely outside the stipulated range) to 100% (if the participant's interpretation was fully contained), Budescu *et al.* (2014) found that mean percentage overlap was 18.5% among participants provided the translation table, compared to 33.6% among participants provided in-text numeric translations.

In an important extension of Budescu *et al.*'s (2009, 2012, 2014) work, Wintle *et al.* (2019) examined the efficacy of the US intelligence community guideline shown in Table 1. Aside from noteworthy differences between the two NBLP schemes (e.g., some numeric ranges overlap in the IPCC standard but are mutually exclusive in ICD 203), Wintle *et al.* compared the effect of four presentation formats (described in the next section) on the percentage of overlap between participants' interpretations of linguistic probabilities (defined in terms of the numeric range equivalents they provided for each term) and the numeric ranges in ICD 203. Compared to a control condition, they found that percentage overlap (i.e., a measure of compliance)

significantly improved only in the hybrid condition in which numeric range equivalents were shown next to the linguistic probability phrase in text.

Wintle *et al.*'s (2019) findings are important because they strongly generalize the findings of Budescu *et al.*'s (2009, 2012, 2014) studies from the domain of climate science to that of intelligence analysis. Given that the intelligence community, like many areas of government, is often slow to implement pan-organizational reform and to leverage relevant behavioral science (Dhami *et al.*, 2015; Mandel, 2020), it is important to demonstrate that the limitations of NBLP schemes do not only exist in other expert communities. Nevertheless, the practical value of this research might be minimized because the measure of percentage overlap (PO) that Wintle *et al.* (2019) used may be seen as unfairly punitive from a prescriptive perspective. Wintle *et al.* calculated PO between participants' interpreted ranges and the ICD 203-stipulated ranges as follows:

$$PO = \{[\min(U_e, U_s) - \max(L_e, L_s)] / [\max(U_e, U_s) - \min(L_e, L_s)]\} * 100,$$

where L_e , B_e (used subsequently), and U_e refer to the participant's elicited lower-bound, best, and upper-bound estimates for a term in a particular statement, respectively, and the subscripts e and s refer to elicited and stipulated, respectively. This measure penalizes individuals who provide compliant "in-range" interpretations that are *more precise* than the stipulated numeric ranges, or what Benjamin and Budescu (2018) described as nested intervals. For instance, given that the stipulated range for *likely* is 55-80%, a participant interpreting *likely* as 60-75% would have a PO = 60% instead of 100% (i.e., $PO = \{[\min(75, 80) - \max(60, 55)] / [\max(75, 80) - \min(60, 55)]\} * 100 = 60$). However, such precision can (legitimately, we

argue) be seen as a positive communication attribute. Penalizing in-range precision where the receiver's interpreted range is a fully nested set of the stipulated range may also distort the intended use of the NBLP scheme under examination. That is, when an intelligence analyst selects a term such as *likely*, it only means that the analyst's credible interval falls *within* the stipulated range (55-80%), not necessarily that it *equals* the full range. Accordingly, we reanalyzed the data from Wintle *et al.* using an alternative measure of percentage overlap that does not penalize compliant participants for being more precise than the stipulated guidelines. Given that our measure of percentage overlap was less punitive, we expected to see an elevation of compliance overall. We further hypothesized that the relative differences between presentation formats would remain unchanged.

Materials and methods

Most of the relevant materials and methods are reported in full detail in Wintle *et al.* (2019). For ease of reference, we summarize key aspects of their experiment and then detail the procedures we used for re-analyzing their data. Wintle *et al.* recruited 924 adult participants from a pool of 4,122 people who had expressed interest in a larger research project. The experiment was administered online via Qualtrics. Participants were randomly assigned to one of four presentation conditions: Table, Tooltip, Brackets, or Control. In the Table condition, participants could click on a link, opening a separate tab/window containing the ICD 203 NBLP scheme. In the Tooltip condition, participants could position their cursor over the probability term appearing in a statement, revealing the term's numeric equivalent. In the Brackets condition, numeric range equivalents were embedded in text alongside verbal probability

expressions (e.g., “likely [55-80%]”). The brackets condition is identical to what we earlier referred to as the hybrid format. In the Control condition, participants received no guidance on the stipulated numeric equivalents of each probability expression (i.e., they were not shown the ICD 203 guideline in any form).

Across presentation conditions, participants were shown eight statements presented in random order that were extracted from US intelligence reports. Each statement contained one of four probability terms (very unlikely, unlikely, likely, very likely) drawn from ICD 203 (e.g., “ISIS is *unlikely* to announce that it is ending its self-declared caliphate...”). Each probability term was presented twice, each time in a different statement. In our subsequent analyses, we average results over these two instantiations for each term, as the specific statements are not of theoretical interest. After reviewing a statement, participants were asked, “What do you think the authors mean by [probability term]?” Using a 101-point percentage-chance slider scale ranging from 0 (No chance) to 100 (Certain), participants gave their minimum, best, and maximum numerical probability equivalents.

Analysis

Data and other supporting files associated with our re-analysis are available at

<https://osf.io/v9q82/>.

We treated missing data differently than Wintle *et al.* (2019). Where participants provided maximum estimates for the terms *unlikely* and/or *very unlikely* but were missing minimum estimates, Wintle *et al.* recoded the missing values as 0. This procedure was based on the assumption that these participants had intended to input 0 but did not realize that they had

to move the slider from its default position (this procedure is not described in the original article, but is outlined in their code: <https://osf.io/cp5rq/>). Although this explanation is plausible, there is in fact no way to ascertain the extent to which this reason accounts for those missing cases. Accordingly, we adopted the more conservative procedure of treating all missing cases as missing.

Our data-handling procedures also differed from Wintle *et al.* (2019) in another respect. In cases where participants violated the “logical” constraint $L_e < B_e < U_e$, Wintle *et al.* rearranged their responses to conform to that inequality. (Our use of scare quotes reflects the fact that the logical constraint should be the relaxed expression, $L_e \leq B_e \leq U_e$. A participant giving a best estimate equal to one of the bounds would be unable to satisfy the strict inequality.) Given that the bases of such incoherence are unknown, we do not believe this procedure is justified and Wintle *et al.* do not provide a justification for the intervention. Accordingly, we dropped 190 cases (2.6% of the sample) where participants provided a minimum estimate that exceeded their maximum estimate. We deleted these cases because they cannot be meaningfully used to calculate the percentage overlap measure subsequently described. Best estimates that fell outside the participant’s lower and upper bounds (354 or 5.3% of valid cases) were left unchanged because this form of incoherence does not prohibit the calculation of any measure used in the re-analysis.

In contrast to Wintle *et al.* (2019), we calculated percentage overlap as follows:

$$PO = \{1 - [\max(U_e - U_s, 0) + \max(L_s - L_e, 0)] / (U_e - L_e)\} * 100$$

This measure of PO is asymmetric, representing the proportion of the participant’s interval that overlaps with the stipulated interval. Note that if spread was equal to 0 (i.e., $L_e = U_e$), which

occurred in 95 cases, PO was equal to 100% if the value of the bounds was within the stipulated range and it was equal to 0% if the value of the bounds was outside the stipulated range.

To facilitate comparison with Budescu *et al.* (2009, 2012, 2014), Wintle *et al.* (2019) also analyzed the proportion of participants' best estimates falling within the stipulated ranges. To minimize the effect of outliers (Tukey, 1962) and ensure consistency with earlier studies (e.g., Budescu *et al.*, 2009), we Winsorized best estimates to focus on the central 90 percent of the distribution within each of the 32 cells defined by the 4 (probability term) $\times$ 2 (intelligence statement) $\times$ 4 (presentation format) design.

Results and discussion

Missing data

There was a maximum of 7,392 responses to each of the three questions (i.e., minimum, best, and maximum estimates) in the experiment; namely, 924 participants each of whom were asked to assess eight statements. Overall, 13.26% (2,941/22,176) of responses were missing. Table 2 shows the percentage of missing responses by estimate, probability term, and presentation format. Each of the three types of estimates was independent of presentation format, all $p > .2$. In contrast, each of the three types of estimates was not independent of probability term. Chi-square test values (with $df = 3$ and $N = 7,392$) for minimum, best, and maximum estimates are 115.09, 12.42, and 14.02, respectively, all $p < .01$. The pattern of missing responses for estimates of minimum and maximum values across probability term was comparable. In both cases, there was a greater percentage of missing data for the low

probability terms *unlikely* and *very unlikely* than for the high probability terms *likely* and *very likely*. We note that although the result for minimum values is consistent with Wintle *et al.*'s assumption that missing data reflected a failure to adjust the slider in order to set it to zero, this would not explain why the same pattern of missing data was observed for the estimates of maximum value. The effect of probability term on missing responses for best estimates, however, showed a different pattern in which the term *likely* had a higher percentage of missing responses.

Table 2. Percentage of missing responses by estimate, probability term, and presentation format.

Factor/Level	Estimate		
	Minimum	Best	Maximum
Format			
Table	17.32	10.61	12.61
Tooltip	16.33	9.61	11.72
Brackets	17.10	11.69	12.99
Control	17.64	10.28	11.23
Term			
Very Unlikely	23.76	10.44	11.96
Unlikely	18.83	10.12	14.45
Likely	14.61	12.55	11.53
Very Likely	11.20	9.09	10.61

Compliance as percentage overlap

We conducted a mixed (Presentation Format × Probability Term) analysis of variance (ANOVA) on percentage overlap. The effect of probability term was significant, $F(3, 472) = 25.07, p < .001, \eta_p^2 = .137$. There was also a significant effect of presentation format, $F(3, 474) = 41.30, p < .001, \eta_p^2 = .207$. However, the interaction effect was not significant, $F(9, 1422) = 1.08, p = .37, \eta_p^2 = .007$. Tables 3 and 4 show the mean percentage overlap as a function of presentation format and probability level, respectively, and Figure 1 plots percentage overlap

by these factors. As Table 3 shows, all pairwise comparisons between levels of presentation format were significant with the exception of the comparison between the Table and Tooltip conditions. As Table 4 shows, all pairwise comparisons among levels of probability term were significant with the exception of the comparison between the *likely* and *very likely*. Replicating Wintle *et al.*'s (2019) findings, overlap was markedly worse for the phrase *unlikely*. Overall, the asymmetric measure of percentage overlap was less punitive than the symmetric measure used by Wintle *et al.*, but it revealed a highly consistent pattern of results regarding the effects of the independent variables (see Figure 1).

Table 3. Mean percentage overlap by presentation format. Note: superscripts in the "Mean" column denote a significant difference at the $\alpha = .05$ probability level.

Format	Mean	95% confidence interval	
		LB	UB
Table	60.10 ^a	55.79	64.32
Tooltip	60.46 ^a	56.14	64.79
Brackets	81.73 ^b	77.89	85.58
Control	51.12 ^c	46.90	55.33

Table 4. Mean percentage overlap by probability term. Note: superscripts in the mean column denote a significant difference at the $\alpha = .05$ level.

Term	Mean	95% confidence interval	
		LB	UB
Very Unlikely	63.73 ^a	60.95	66.50
Unlikely	55.79 ^b	52.61	58.97
Likely	67.43 ^c	64.91	69.94
Very Likely	66.43 ^c	64.06	68.80

[FIGURE 1 HERE]

Figure 1. Percentage overlap by presentation format and probability term.

Compliance as percentage of in-range best estimates

We ran a mixed (Presentation Format $\times$ Probability Term) ANOVA on the percentage of compliant best estimates; namely, best estimates that fell within the ranges specified in ICD 203. The effect of probability term on the percentage of compliant best estimates was significant, $F(3, 638) = 49.51, p < .001, \eta_p^2 = .189$. There was also a significant effect of presentation format, $F(3, 640) = 25.76, p < .001, \eta_p^2 = .108$. However, the interaction effect was not significant, $F(9, 1920) = 1.03, p = .41, \eta_p^2 = .005$. Tables 5 and 6 show the mean percentage of compliant best estimates as a function of presentation format and probability level, respectively, and Figure 2 plots the percentage of compliant best estimates by these factors. As Table 5 shows, the pairwise differences perfectly paralleled those obtained on the PO measure of compliance; namely, all conditions significantly differed except for the Table and Tooltip

conditions. The pairwise differences between levels of probability term showed some discrepancies with those reported for the PO measure. Once again, compliance was poorest for the phrase *unlikely* (see Table 6). However, compliance was significantly better for the term *very likely* than for the remaining terms. Figure 2 shows how our results closely parallel those of Wintle *et al.* (2019). This is to be expected given that our analyses differ mainly in terms of the data exclusion rules we applied and also the fact that we Winsorized the data.

Table 5. Mean percentage of compliant best estimates by presentation format. Note: superscripts in the mean column denote a significant difference at the $\alpha = .05$ level.

Format	95% confidence interval		
	Mean	LB	UB
Table	66.50 ^a	62.45	70.56
Tooltip	67.71 ^a	63.70	71.73
Brackets	83.36 ^b	79.53	87.18
Control	59.74 ^c	55.81	63.67

Table 6. Mean percentage of compliant best estimates by probability term. Note: superscripts in the mean column denote a significant difference at the $\alpha = .05$ level.

Format	95% confidence interval		
	Mean	LB	UB
Very Unlikely	73.56 ^a	70.76	76.36
Unlikely	55.41 ^b	52.14	58.67
Likely	70.88 ^a	68.00	73.77
Very Likely	77.46 ^c	74.86	80.05

[FIGURE 2 HERE]

Figure 2. Percentage of participant's best estimates falling inside the ranges stipulated by ICD 203. Results are separated by presentation format and probability term.

Conclusions

Breakdowns in the communication of uncertainty have played a significant role in major intelligence failures in recent history (US Central Intelligence Agency, 1983; UK House of Commons Foreign Affairs Committee, 2003; UK Intelligence and Security Committee, 2003; US Congressional Select Committee on Intelligence, 2004). Wintle *et al.* (2019) made an important contribution to the evaluation of the current standard used by the US intelligence community for communicating uncertainty in intelligence products. Their research showed that the NBLP scheme currently used did not promote a level of shared understanding one would expect for communications of great importance that could affect national security. We reanalyzed the

data of Wintle *et al.* by using more conservative data exclusion rules and a less punitive overlap measure of compliance. We did so not because we view Wintle *et al.*'s method as wrong, but rather because we believe it is fair to construe it as overly punitive from a prescriptive vantage point. We further believe that defenders of the status quo within the intelligence community might draw upon this characteristic and dismiss the validity of the research on that basis. Indeed, the in-range precision that Wintle *et al.*'s PO measure scores as non-overlap might be regarded as a beneficial attribute given that receivers are often willing to trade accuracy for informativeness in the form of greater precision (Yaniv & Foster, 1995, 1997). The measure that we used instead is not susceptible to the same critique. Although our asymmetric measure of percentage overlap does not reward precision within the numeric probability bounds stipulated in ICD 203, it does not penalize it either.

Our reanalysis resulted in an elevated percentage overlap across conditions. This is to be expected given that our measure is less punitive than Wintle *et al.*'s (2019) measure. However, critically, the effect of presentation format and probability level was highly consistent with that reported in Wintle *et al.* (2019). We found that the Table and Tooltip methods outperformed the control condition, and these methods were outperformed by the Brackets method. This ordering suggests that the more numeric representations are present at the time of interpretation, the more faithfully the sender's interpretation is decoded. These results further suggest that the methods intelligence organizations currently use, which are closest to the Table condition in the present research, are unlikely to effectively sync the meaning of uncertainty expressions assigned by intelligence producers and inferred by intelligence consumers, such as policymakers who must rely on uncertain estimates to support critical

decision-making about national security matters (Friedman, 2019). Even by our less punitive measurement standards, percentage overlap was about 60% in the Table condition, a mere 9% above the control condition, which would represent the intelligence community's approach prior to the adoption of the ICD 203 standard. Similarly, if we compare the Table condition to the control group in terms of the compliance of participants' best estimates, there is only a small increase from approximately 60% to 67%.

By comparison, the improvement in communication fidelity between the current intelligence community method and the enhanced method of including stipulated numeric probability ranges directly in assessments was far more substantial by both compliance metrics. This is consistent with the findings of Budescu *et al.* (2009, 2012, 2014). Although these findings suggest that use of bracketed numeric ranges in estimates would improve communication fidelity, we urge caution in our recommending this strategy. One potential problem with the hybrid approach is that end-users might interpret numeric ranges as credible intervals on the substantive assessment given in an intelligence report. For example, imagine an intelligence analyst states, "It is likely [55%-80%] that the Blanks will attack our country in the next week." According to the hybrid method, the range given is meant to clarify what the term *likely* means, in general—namely, as outlined in ICD 203. However, it is plausible that end-users would interpret the range to represent the credible interval the analyst assigned to the specific event, "the Blanks will attack our country in the next week." However, the analyst might have had in mind a different credible interval. Indeed, the analyst's credible interval might span more than one of the stipulated ranges. For instance, if pressed for a credible interval, the analyst might have assigned a range of 75%-85%. This range spans the stipulated ranges associated with the

terms *likely* and *very likely*. Some analysts with such a range in mind might use the more extreme term (i.e., *very likely*), but perhaps more often than not they would choose the less extreme term. At least, this is suggested by work showing substantial underconfidence in analysts' probabilistic judgments (Mandel & Barnes, 2014, 2018). An alternative approach would be to report only the numeric probability ranges (e.g., Dhimi & Mandel, forthcoming; Friedman, 2019; Irwin & Mandel, 2019). If so, the ranges could unambiguously refer to the analyst's credible interval. To the best of our knowledge, no research has yet compared the effectiveness of a purely numerical format with the ones examined here. This could be a promising direction for future research with implications for public policy.

Acknowledgments

This research was supported by Canadian Safety and Security Program project CSSP-2018-TI-2394. We thank Robert Collins for lending his data science expertise.

Author contributions

Conceptualization: David R. Mandel

Data Curation: Daniel Irwin

Formal Analysis: David R. Mandel, Daniel Irwin

Methodology: David R. Mandel

Validation: David R. Mandel, Daniel Irwin

Visualization: Daniel Irwin

Writing – Original Draft Preparation: David R. Mandel, Daniel Irwin

Writing – Review & Editing: David R. Mandel, Daniel Irwin

References

Barnes A. Making intelligence analysis more intelligent: Using numeric probabilities. *Intelligence and National Security*. 2016; 31(3): 327–44. doi: 10.1080/02684527.2014.994955

Benjamin DM, Budescu DV. The role of type and source of uncertainty on the processing of climate models projections. *Frontiers. Psychol.* 2018; 9(403). doi: 10.3389/fpsyg.2018.00403

Budescu DV, Broomell SB, Por HH. Improving communication of uncertainty in the reports of the Intergovernmental Panel on Climate Change. *Psych Sci.* 2009; 20(3): 299–308. doi: 10.1111/j.1467-9280.2009.02284.x

Budescu DV, Por HH, Broomell SB. Effective communication of uncertainty in the IPCC reports: A nationally representative survey. *Climatic Change*. 2012; 113(2): 181–200. doi: 10.1007/s10584-011-0330-3

Budescu DV, Por HH, Broomell SB, Smithson M. The interpretation of IPCC probabilistic statements around the world. *Nature Climate Change*. 2014; 4(6): 508–12. doi: 10.1038/nclimate2194

Burgman MA, McBride M, Ashton R, Speirs-Bridge A, Flander L, Wintle B., et al. Expert status and performance. *PLOS ONE*. 2011; 6(7): e22998. doi: 10.1371/journal.pone.0022998

Central Intelligence Agency

Dhami MK, Mandel DR. Words or numbers? Communicating probability in intelligence analysis. *American Psychologist*. Forthcoming 2020.

Dhami MK, Mandel DR, Mellers BA, Tetlock PE. Improving intelligence analysis with decision science. *Perspectives on Psychological Science*. 2015; 106(6): 753–757. doi: 10.1177/1745691615598511

Friedman JA. *War and chance: Assessing uncertainty in international politics*. 1st ed. New York: Oxford University Press; 2019.

Ho EH, Budescu DV, Dhami MK, Mandel DR. Improving the communication of uncertainty in climate science and intelligence analysis. *Behavioral Science & Policy*. 2015; 1(2): 43–55. doi: 10.1353/bsp.2015.0015

Irwin D, Mandel DR. Variants of vague verbiage: Intelligence community methods for communicating probability. In: Mandel DR, editor. *Assessment and communication of uncertainty in intelligence to*

support decision making: Final report of Research Task Group SAS-114. Brussels: NATO Science and Technology Organization; 2020. p. 319-48.

Kent S. Words of estimative probability. *Studies in Intelligence*. 1964; 8(4): 49-65.

Lewis SC, King AD, Perkins-Kirkpatrick SE, Wehner MF. Toward calibrated language for effectively communicating the results of extreme event attribution studies. *Earth's Future*. 2019; 7(9): 1020-26. doi: 10.1029/2019EF001273

Mandel DR. The occasional maverick of analytic tradecraft. *Intelligence and National Security*. 2020; 35(3): 438-443. doi: 10.1080/02684527.2020.1723830. Mastrandrea MD, Mach KJ, Plattner GK, Edenhofer O, Stocker TF, Field CB, Ebi KL, Matschoss PR. The IPCC AR5 guidance note on consistent treatment of uncertainties: A common approach across the working groups. *Climatic Change*. 2011; 108(4): 675-91. doi: 10.1007/s10584-011-0178-6

Mandel DR, Barnes A. Accuracy of forecasts in strategic intelligence. *Proceedings of the National Academy of Sciences*. 2014; 111(30): 10984-10989. doi: 10.1073/pnas.1406138111

Mandel DR, Barnes A. Geopolitical forecasting skill in strategic intelligence. *Journal of Behavioral Decision Making*. 2018; 31(1): 127-137. doi: 10.1002/bdm.2055

Morgan MG. Commentary: Uncertainty analysis in risk assessment. *Human and Ecological Risk Assessment*. 1998; 4(1): 25-39. doi: 10.1080/10807039.1998.11009680

Moxey A, O'Connell D, McGettigan P, Henry D. Describing treatment effects to patients: How they are expressed makes a difference. *Journal of General Internal Medicine*. 2003; 18(11): 948-59. doi: 10.1046/j.1525-1497.2003.20928.x

Murphy AH, Winkler RL. Probability forecasting in meteorology. *Journal of the American Statistical Association*. 1984; 79(387): 489-500. doi: 10.1080/01621459.1984.10478075

Tetlock PE, Gardner D. *Superforecasting*. 1st ed. New York: Crown Publishing Group; 2015.

Tukey JW. The future of data analysis. *The Annals of Mathematical Statistics*. 1962; 33(1): 1-67. doi: 10.1214/aoms/1177704711

UK House of Commons Foreign Affairs Committee. Ninth Report of the Foreign Affairs Committee Session 2002-03: The Decision to Go to War in Iraq. 2003. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/272087/6062.pdf

UK Intelligence and Security Committee. Iraqi Weapons of Mass Destruction – Intelligence and Assessments. 2003. Available from: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/272079/5972.pdf

US Central Intelligence Agency. Report on a Study of Intelligence Judgments Preceding Significant Historical Failures: The Hazards of Single-outcome Forecasting, 16 December 1983. 1983. Available from: <https://www.cia.gov/library/readingroom/docs/CIA-RDP86B00269R001100100010-7.pdf>

US Congressional Select Committee on Intelligence. Report on the U.S. Intelligence Community's Prewar Intelligence Assessments on Iraq. 2004. Available from: https://fas.org/irp/congress/2004_rpt/ssci_iraq.pdf

US Office of the Director of National Intelligence. Intelligence Community Directive (ICD) 203, Analytic Standards, 2 January 2015. 2015. Available from: <https://www.dni.gov/files/documents/ICD/ICD%20203%20Analytic%20Standards.pdf>

Webster RK, Weinman J, Rubin GJ. People's understanding of verbal risk descriptors in patient information leaflets: A cross-sectional national survey of 18- to 65-year-olds in England. *Drug Safety*. 2017; 40(8): 743-54. doi: 10.1007/s40264-017-0542-1

Wintle BC, Fraser H, Willis BC, Nicholson AE, Fidler F. Verbal probabilities: Very likely to be somewhat more confusing than numbers. *PLOS ONE*. 2019; 14(4): e0213522. doi: 10.1371/journal.pone.0213522

Yaniv I, DP Foster. Graininess of judgment under uncertainty: An accuracy-informativeness trade-off. *Journal of Experimental Psychology: General*. 1995; 124(4): 424-32. doi: 10.1037/0096-3445.124.4.424

Yaniv I, DP Foster. Precision and accuracy of judgmental estimation. *Behavioral Decision Making*. 1997; 10(1): 21-32. doi: 10.1002/(SICI)1099-0771(199703)10:1<21::AID-BDM243>3.0.CO;2-G