This preprint (version 1) is currently under review and has not yet been peer reviewed.

**Tracking Accuracy of Strategic Intelligence Forecasts:
Findings from a Long-Term Canadian Study**

Running title: Tracking Accuracy

David R. Mandel^{a*} and Daniel Irwin^b

^aDefence Research and Development Canada

^bDepartment of National Defence

*Correspondence to: drmandel@gmail.com

Abstract

Forecasting plays a vital role in intelligence assessment and contributes to national security decision-making by improving strategic foresight. Remarkably, most intelligence organizations do not proactively track their forecasting accuracy and, therefore, do not know how accurate their forecasts are or what types of biases intelligence analysts (or organizations) might exhibit. We review research on geopolitical forecasting and a roughly decade-long program of research to assess the accuracy of strategic intelligence forecasts produced by and for the Government of Canada. This research is described in three phases corresponding to previously published research, following which novel analyses (drawing from the data used in the earlier phases) are reported. The findings reveal a high degree of forecasting accuracy as well as significant underconfidence. These results were evident regardless of whether analysts assigned numeric probabilities to their forecasts. However, the novel analyses clarified that there is a substantial cost to accuracy if end-users rely on their own interpretations of verbal probability terms used in the forecasts. We recommend that intelligence organizations proactively track forecasting accuracy as a means of supporting accountability and organizational learning. We also recommend that intelligence organizations use numeric probabilities in forecasts as a means of improving intelligence producer-consumer agreement in the interpretation of forecasts and other probabilistic assessments.

Keywords: forecasting, accuracy, strategic intelligence, probability

Better earlier warning than later mourning.

Jewish proverb

As the Jewish proverb conveys, accurate indications about consequential future events that arrive early enough can help decision-makers manage risk and avert trouble. It is, therefore, no surprise that forecasting plays a vital role in intelligence assessment. According to North Atlantic Treaty Organization (NATO) intelligence doctrine, “analysis does more than look at the current situation, it should be predictive and therefore should address what might happen next, based upon alternative assumptions regarding the actions and reactions of different actors (including the impact of any intervention)” (NATO, 2016, §3.38). An effective forecasting capability supports planning and decision-making at all levels, ranging from tactical to strategic. And, although the empirical research reported in this article focuses on efforts to track forecasting accuracy at the strategic level of intelligence production, the issues dealt with also apply to forecasting at the tactical and operational levels. Moreover, the methods later described for tracking forecast accuracy and forecasters’ skill could be applied at any level.

At the strategic level, geopolitical forecasting offers early warning indicators—or what in more recent times has been called *anticipatory intelligence*—to key decision makers, such as state leaders and other senior policymakers (Clapper, 2014; Kerbel, 2019). Timely, relevant, and accurate geopolitical forecasts promote transnational or national interests by reducing the probability of strategic surprises, mitigating security risks, and allowing policymakers to capitalize on opportunities in a dynamic, globalized world. As Sherman Kent (1964; 1994), a figure widely credited as the father of modern intelligence analysis, noted long ago, intelligence that provides the decision-maker with timely and accurate indications about the future are among

the most important of intelligence products, but they are also the products that often represent thankless tasks for intelligence agencies. As Kent (1994) explained, this is because forecasting involves judgment under uncertainty, and analysts do not have exclusive purview over judgments. Policymakers often believe (whether accurately or not) that they can predict future events better than analysts. When the intelligence community's judgments concur with their own, they may see in the concordance the goodness of their own judgment and the redundancy of the analysis. When the judgments are discordant, they may be less than thankful for dissenting views, especially if they have committed to a particular course of action or formed a stubborn bias towards a particular policy option.

Why Tracking Accuracy Matters

Although the intelligence community assigns great importance to the anticipatory function of intelligence, the quality of geopolitical forecasting remains largely unverified because intelligence organizations do not routinely and systematically monitor forecasting skill (Betts, 2007; Dhami et al., 2015; Friedman & Zeckhauser, 2016; Mandel, 2015a). The decision not to track forecasting skill on an ongoing basis with well-established quantitative scoring rules can have many deleterious consequences. Without credible monitoring processes, intelligence organizations simply cannot know how good their forecasts are, how they might be improved, and where improvement is most needed. In the absence of such knowledge and subsequent corrective actions that might have been implemented, intelligence organizations increase the risk of failure to detect strategic threats and opportunities. Without proper monitoring, intelligence organizations also remain unnecessarily susceptible to reactive pressure for institutional change—and in some cases, even institutional erasure (e.g., Moynihan, 1991)—after politicized intelligence failures (Johnson, 2007; Tetlock & Mellers, 2011).

Without adequate tracking, such organizations also forfeit opportunities for improving forecasting through calibration feedback (Rieber, 2004) or recalibration techniques aimed at mitigating judgment biases such as overconfidence or underconfidence (Mandel & Barnes, 2014) or violations of additivity in the assignment of probabilities to alternative hypotheses (Karvetski et al., 2020; Mandel et al., 2018). Such biases may be present—and detectable—at individual and organizational levels. Nor are intelligence organizations well poised to evaluate the effects of training or structured techniques to improve analytic qualities such as forecasting accuracy if such qualities are not carefully measured (Chang et al., 2018; Chang & Tetlock, 2016; Mandel, 2015b; Mellers et al., 2014).

Intelligence organizations that aim to track their forecasting accuracy should ensure that, at minimum, efforts to support this objective are ecologically valid and methodologically sound. That is, organizations should strive to ensure that (a) the tracking process represents the analytic environment in which real intelligence forecasts are made and (b) within that context, reasonable efforts are made to promote valid and reliable scoring of forecasting accuracy and its various skill components (e.g., discrimination and calibration). To achieve these objectives, monitoring systems should track the quality of real forecasting products that are made by real intelligence analysts who are working under normal conditions. As well, the monitoring systems should use well-established quantitative methods that have been applied in other domains of expert judgment under uncertainty (e.g., Åstebro & Koehler, 2007; Goodman-Delahunty et al. 2010; Keren, 1987; Lin & Bier, 2008; Murphy & Winkler, 1984).

Intelligence organizations may be less interested in the external validity of the monitoring process insofar as they will be more concerned about their own organizational performance than about how their performance generalizes to other organizations. Nevertheless, concern about the

external validity of tracking procedures should be of interest to the broader intelligence community, and especially to organizations (or organizational divisions) that are mandated to oversee, coordinate, or provide scientific advice to wider “intelligence communities.” For instance, in the US this would include the Analytic Standards and Integrity Group within the Office of the Director of National Intelligence (ODNI), and in the UK this would include the Professional Head of Intelligence Assessment (PHIA) within the Cabinet Office. It is important to know, for instance, whether (and if so, why) the organizations that comprise a national intelligence analytic capability exhibit small or large variance in forecasting skill. In short, does knowing the performance of one intelligence organization tell us much about the forecasting capability of another? As the research summarized later demonstrates, concerns about external validity can be amplified even within a single organization if its various divisions rely on different analytic methods for producing forecasts.

Geopolitical Forecasting Accuracy

On the strategic level, intelligence forecasting mainly concerns events of geopolitical significance. Little is known about expert geopolitical forecasting skill, especially within intelligence organizations. Tetlock’s (2005) long-term study of close to 300 political experts found evidence of mediocre geopolitical forecasting skill. Experts were highly overconfident, and even the best expert forecasters—namely, the uncertainty-tolerant “foxes,” referring to Isaiah Berlin’s (1953) intellectually playful distinction between the metaphorical hedgehog and fox—performed substantially worse than the best statistical prediction models that Tetlock tested. Moreover, seasoned experts displayed comparable accuracy to those still wet behind the ears. Even more remarkably, experts who forecasted on topics in their areas of expertise

performed no better than dilettantes—namely, experts who forecasted on topics outside their areas of expertise.

If replicable and generalizable, the latter finding would have striking implications for intelligence organizations, which routinely organize their expertise along geopolitical fault lines. The presumption of such a system is that a China analyst will make better forecasts about China than an analyst of Iran, and vice versa. Tetlock's (2005) study challenges that bedrock assumption. However, it certainly does not invalidate the assumption because the study scores low on ecological validity for two important reasons. First, the experts in his study did not comprise a set of strategic intelligence analysts and, second, the forecasts were made as part of a research study, not as a tasking to produce finished intelligence for real intelligence consumers. Experts were neither incentivized nor under comparable degrees of accountability pressure as analysts in the wild. Whether such differences matter is unclear, but Tetlock's (2005) findings should prompt intelligence organizations to redouble efforts to answer such questions.

Many geopolitical topics in Tetlock's (2005) study required long-range forecasts that would only be resolved in several years. In contrast, most intelligence forecasts (excluding futures or foresight exercises) are short-to-medium range and resolve in less than 1 year. A geopolitical forecasting tournament sponsored by the US government's Intelligence Advanced Research Program Activity (IARPA) elicited forecasts that generally resolved in 1 year or less, thus increasing timeframe comparability to the strategic intelligence realm. The winners of IARPA's Aggregative Contingent Estimation (ACE) tournament found that elite "superforecasters" (i.e., those at or above the 98th percentile in forecasting accuracy) could be cultivated using a combination of effective sampling, elicitation, training, and aggregation methods (Mellers et al., 2014, 2015; Tetlock & Gardner, 2015). Superforecasters do not only

forecast more accurately, they exhibit superior cognitive abilities (Mellers et al., 2015). Superforecasters also better discriminate the meaning of verbal probability terms, are less susceptible to content effects on their interpretation of such terms, and are more coherent on other judgment tasks (Mellers et al., 2017). The findings of the ACE program and other IARPA programs on forecasting (such as the more recent Hybrid Forecasting Competition [HFC]) suggest several ways in which intelligence organizations could attempt to improve intelligence forecasting, but they do not bring us much closer to knowing how well intelligence organizations currently forecast because, once again, the forecasters are not sampled from intelligence analysts and the forecasts are elicited as part of applied research rather than on-the-job activities with their attendant accountability pressures and particular incentive structures.

Ecologically valid studies of geopolitical forecasting skill in intelligence organizations are exceedingly rare. In one example, Lehner et al. (2012) examined 187 geopolitical forecasts taken from unclassified or declassified intelligence reports and found poor discrimination but fairly good calibration. However, the methods the authors used cast doubt on the interpretability of the study's findings. For instance, descriptive statements (e.g., "Arab groups in Kirkuk continue to resist violently what they see as Kurdish encroachment" (see Lehner et al., 2012, p. 730) were rewritten as forecasts, thus changing their intended meaning (i.e., "Arab groups in Kirkuk will resist violently what they see as Kurdish encroachment in the January 2007 to July 2009 time frame" (see Lehner et al., 2012, p. 731). Another serious limitation of that study is that verbal probability qualifiers in the original assessments were omitted in redrafted forecasts, once again changing their meaning. Due to these and other methodological shortcomings, it is difficult to draw conclusions from this research about the quality of *actual* strategic intelligence forecasting. Stastny and Lehner (2018) summarize a more recent study of forecasting accuracy in

strategic intelligence reports. However, the interpretability of that study was also undermined by various methodological factors thoroughly described in Mandel (2019).

Tracking Strategic Intelligence Forecasting Accuracy in Canada

Phase I

The most comprehensive attempt to systematically track the forecasting accuracy of strategic intelligence was conducted in Canada. In the initial phase of this effort, Mandel and Barnes (2014; see also Mandel et al., 2014) conducted a long-term study of geopolitical forecasting accuracy in strategic intelligence, which examined forecasts extracted from a comprehensive review of 6 years of classified reports produced by the Middle East and Africa (MEA) division of the Intelligence Assessment Secretariat (IAS) under the Privy Council Office in the Government of Canada. As part of regular analytic practice in the MEA division, analysts recorded whether their assessments were forecasts or other types of judgments, such as explanations or purely descriptive statements (Barnes, 2016).

Analysts also assigned numeric probabilities to forecasts (but not to other judgments or statements). The nine permissible numeric probabilities, each expressed out of 10, were {0, 1, 2.5, 4, 5, 6, 7.5, 9, and 10}. These values were mapped to verbal probability terms following the numerically bound linguistic probability scheme described by Barnes (2016). For instance, for the probability level 2.5/10, analysts could select from the terms *low probability*, *probably not*, or *unlikely*, whereas for a probability level of 7.5/10, analysts could select from the terms *probably*, *probable*, or *likely*. Note that whereas the verbal probabilities appeared in finished intelligence reports, the numeric probabilities did not. For example, one forecast (edited to remove sensitive information) was, “It is very unlikely [1/10] that either of these countries will make a strategic decision to launch an offensive war in the coming six months.” The numeric

probability given in brackets was omitted from the finished intelligence reports. These probabilities were recorded for auditing and research purposes only.

Mandel and Barnes (2014) found that 75% of intelligence assessments were forecasts and just over two-thirds (68%) were expressed with sufficient clarity to be quantitatively scored for accuracy. The remainder was excluded from primary analyses either because the description of probability was too vague (i.e., using words such as *might* or *could*) or because the forecasted outcomes were expressed in ways that made them too difficult to code.

For the scorable forecasts, Brier scores were calculated. The Brier score, B , is a proper scoring rule equal to the squared deviation between probabilities assigned to forecasts and outcomes, which were coded 0 for non-occurrence and 1 for occurrence (Brier, 1950; Murphy, 1973). The Brier score ranges from 0-1, with 0 equalling a perfect score. A forecaster who simply gave a forecast of 0.5 (i.e., fifty-fifty) for every prediction would have a Brier score of 0.25, which can be regarded as a useful benchmark. That is, one would hope that expert judgment yields Brier scores much lower than 0.25, and as close to zero as possible. Among the 1,514 forecasts that were scored, forecasting accuracy based on the numeric probabilities that analysts assigned was very good. The mean Brier score reported by Mandel and Barnes (2014) was 0.074. If forecasts that were excluded due to the vagueness of the probability term used were included (i.e., by substituting the expected probability value of .5), B expectedly increased, but the mean value remained less than 0.10 in various tests of estimate sensitivity.

To use a perhaps more intuitive scoring rule, the correct classification rate was 94%. That is, in 94% of scorable cases, the forecast was on the correct side of fifty-fifty (i.e., higher for events that occurred and lower for events that did not occur). Using another standard metric, the normalized discrimination index (Yaniv et al., 1991), which computes discrimination over

uncertainty (i.e., a measure of the proportion of outcome variance explained by the forecasts, akin to R^2 in multiple linear regression), Mandel and Barnes (2014) found that 76% of outcome variance was explained by the forecasts. Senior analysts also showed better discrimination skill than junior analysts, contrary to what might be inferred from Tetlock (2005). Finally, forecasts, on average, were underconfident rather than overconfident, once again contrary to what Tetlock (2005) observed, but in line with forecasting poll results from the ACE forecasting tournament noted earlier (Atanasov et al., 2017).

Underconfidence was significantly greater for forecasts that expert coders identified as harder rather than easier, and that the expert coders identified as more rather than less important to intelligence consumers. Such findings contrast with the *hard-easy effect* in which harder forecasts tend to produce greater overconfidence than easier forecasts (Lichtenstein & Fischhoff, 1977). They suggest that analysts may be expressing uncertainty strategically as a means of preemptively deflecting accountability pressure. In support of this explanation, Tetlock and Kim (1987) found that when participants were told in advance that they would be accountable for their judgments, they expressed less confidence in their judgments than participants who were either told after their judgments that they would be accountable or who were not told they would be accountable. Couching one's estimates in a high degree of uncertainty can represent a form of insurance against reputational harm. If one happens to be right, there is a minor coverage cost in that one could have been bolder. However, if one turns out to be wrong, then he or she is shielded from blame by the blatant unfalsifiability of a highly uncertain estimate. Analysts—and the intelligence organizations they represent—have much more reputational credit to lose from overconfidence than from underconfidence. Perhaps where forecast accuracy is accompanied by high degrees of epistemic uncertainty, analysts may even have more to lose from being well

calibrated than underconfident. Although the issue deserves research attention, the preliminary fact pattern is consistent with a social functionalist model of how intelligence analysts are likely to behave when acting strategically as career-minding intuitive politicians (Mandel & Tetlock, 2016; Tetlock, 2002).

Mandel and Barnes (2014) further showed that the degree of miscalibration—and underconfidence, more specifically—could be substantially reduced through a simple recalibration procedure in which each probability issued was made one unit more extreme. The values of 0, 0.5, and 1 did not change, but all others became one unit more extreme. For instance, 0.25 became 0.1, just as 0.75 became 0.9. This simple, organizationally feasible procedure was as effective as more complicated mathematical procedures (Baron et al., 2014; Turner et al., 2014) for improving calibration by reducing underconfidence.

Phase II

Monitoring exercises such as the one just described have a peculiar feature in that the forecasting skill assessment is based on numeric probability estimates that end-users never received. Instead, as noted earlier, intelligence consumers received only the verbal probabilities, which they were free to interpret as they please. Therefore, a question that arises is how good are the forecasts when seen through the eyes of a typical consumer. To address this issue, Mandel (2015a) re-examined the accuracy of the forecasts reported by Mandel and Barnes (2014) from the perspective of a statistically-average intelligence consumer.

In the first stage of the research, university students and intelligence analysts were asked to provide their best numeric probability equivalent for each of the 20 probability terms in the MEA division probability scheme (Barnes, 2016). The two subsamples did not statistically differ and their results were combined. The median probabilities assigned were very close to the

probabilities stipulated in the scheme used by the MEA division at the time. In fact, the average deviation was less than 2%.

In the second stage of the research, participants' median estimates were substituted for analysts' numeric probabilities and the mean Brier score was recomputed. Using the median inferred probabilities, forecasting skill (from the average consumer's perspective) was slightly better ($B = 0.071$) than that reported by Mandel and Barnes (2014). This closeness of the two estimates was inevitable given that the scheme for mapping verbal to numeric probabilities used in the MEA division at the time was closely matched to participants' median estimates.

Phase III

The research of Mandel and Barnes (2014); Mandel et al. (2014); and Mandel (2015a) provides an unparalleled and detailed quantitative analysis of forecasting skill in strategic intelligence. The findings are informative, but also raise many questions. Why was forecasting skill so good? How generalizable are the results? One possibility is that the results were as positive as they were because, as Barnes (2016) described, several rather unusual analytic procedures were instituted in the MEA division to enhance analytic rigor. Perhaps those practices paid off in terms of better forecasting skill. If so, one might expect that attempts to generalize the findings to other strategic intelligence divisions that did not rely on such procedures but which produced similar products would flounder. Alternatively, as Arkes and Kajdasz (2011) posited, general environmental features of intelligence organizations, such as accountability to skeptical stakeholders or requirements to be explicit about judgments and supporting reasoning, might account for the degree of skill observed in the MEA division. In that case, one might expect strategic intelligence forecasting skill to show similar characteristics outside of the MEA division. Yet another hypothesis proposed by Tetlock and Mellers (2014) is that the forecasts

made by the MEA division were good because they were easy. Although Tetlock and Mellers (2014) did not provide empirical evidence in support of their facile-forecasting hypothesis, the hypothesis cannot simply be discounted.

To address these questions and the alternative hypotheses they raise, Mandel and Barnes (2018) examined a larger set of strategic intelligence forecasts. The majority of the sample, as in the earlier research, was drawn from the MEA division. However, strategic forecasts from other IAS divisions and from team products produced by multiple organizations were also assessed. These forecasts from outside of the MEA division were not generated with the same methods outlined by Barnes (2016). Analysts were not required to categorize their assessments into forecast and non-forecast categories and, critically, they were not required to assign numeric probabilities to forecasts or to use a fixed set of probability terms.

In total, Mandel and Barnes (2018) extracted 3,622 forecasts from available intelligence reports, and 73% (2,629) had events that could be unambiguously coded as either having occurred or having not occurred. Of the latter, 77% (2,013) had probability terms for which a numeric probability could be confidently assigned. That sample constituted the forecast dataset. Of the 2,013 forecasts, 1,735 (86%) were in the MEA subsample and 278 (14%) were in the non-MEA subsample. As well, 1,759 (87%) forecasts were from IAS intelligence memoranda and 254 (13%) forecasts were from interdepartmental committees. The vast majority (95%) of the probability terms used in forecasts were terms that appeared in the probability scheme Barnes had devised. This is not surprising given that 87% of the sample was from the MEA division. For these terms, the median numeric probability equivalents from Mandel (2015a) were substituted. In the remaining 5% of cases, the values were estimated based on similarity to terms in the standard (for details, see Mandel & Barnes, 2018).

Addressing the generalizability question first, recall that the mean Brier score was 0.074 in Mandel and Barnes (2014) based on analysts' numeric probabilities, and it was 0.071 in Mandel (2015a) using inferred numeric probabilities from median sample estimates of the best numeric equivalents to the verbal probability terms. Mandel and Barnes (2018) found no significant difference in Brier scores between MEA division and the non-MEA forecasts. Indeed, the two subsamples were virtually indistinguishable: $B = 0.059$ for the MEA subsample and $B = 0.058$ for the non-MEA subsample.

Examining the specific skill components of discrimination and calibration also showed no difference between the groups. Discrimination in both groups was very good. For instance, the slight difference in area under the receiver-operator characteristic curve (0.96 for MEA forecasts and 0.93 for non-MEA forecasts) was not statistically significant. Although the two subsamples did not differ in calibration, and both showed evidence of underconfidence, Mandel and Barnes (2018) found that the degree of underconfidence in the MEA subsample was greater than in the non-MEA subsample.

Overall, the findings indicate that the forecasting skill characteristics of the MEA division reported by Mandel and Barnes (2014) and Mandel (2015a) are generalizable to other organizational units producing strategic intelligence forecasts, at least within the Canadian context. The only notable difference (which was small) was the extent of underconfidence, which was greater in the MEA subsample. It is possible that the analytic standards introduced in the MEA division prompted analysts to be more cautious in their assessments, yet this is conjectural. It should be tested in controlled experiments on the effect of accountability processes on expert judgment.

To test the facile-forecasting hypothesis (Tetlock & Mellers, 2014)—namely, the suggestion that forecasts were good because they were easy—Mandel and Barnes (2018) first compared forecasts from key judgments to forecasts not in key judgments. Key judgments are the most authoritative and informative assessments made in intelligence reports. Therefore, they are less likely than other forecasts to address easy topics that would be uninformative to decision makers. Contrary to the facile-forecasting hypothesis, neither calibration nor discrimination significantly differed between key and non-key forecasts.

In a second test of the hypothesis, Mandel and Barnes (2018) compared the 40% of cases in which forecasts were assigned very high certainty (i.e., probabilities greater than or equal to .95 or less than or equal to .05) with the remaining 60% of cases that were less certain (i.e., probabilities greater than .05 but less than .95). The idea here—and one that Tetlock, Mellers and their colleagues have also recently expressed (Friedman et al., 2018)—is that easy forecasts tend to be made with more certainty and, hence, probability values closer to 0 or 1. Contrary to the facile-forecasting hypothesis, neither calibration nor discrimination significantly differed between the high-certainty and lower-certainty subsamples. Therefore, the facile-forecasting hypothesis was not supported by the findings of Mandel and Barnes (2018). If the facile-forecasting hypothesis were correct, it would call into question why governments spend vast sums of money tackling prediction problems that are easily or already known. Although the indirect tests of Mandel and Barnes (2018) did not support the hypothesis that good forecasting is due to easy topic selection, future research should examine intelligence analysts' forecasting skill under conditions where topic difficulty is either controlled or directly measurable, such as in prediction markets or forecasting tournaments like those sponsored by IARPA.

Novel Extensions

As noted earlier, Mandel (2015a) found that forecasting accuracy was as good from an “average consumer’s” perspective as from the analyst-sender’s perspective. In this comparison, the sender used numeric probabilities selected from a set of nine levels, whereas the consumer was represented by the median probabilities assigned to the verbal probabilities that were actually used in the finished intelligence reports. The median was based on a mixed sample of university students and intelligence analysts. Although the findings are not misleading, it would be incorrect to infer from them that the cost of using verbal probabilities to convey forecasts does not exact a cost to accuracy. Indeed, we expect that if forecasts were scored based on individual consumers’ interpretations of the verbal probabilities used in forecasts, accuracy would be significantly reduced due to the variability and unreliability of verbal-to-numeric probability mappings. For example, Mandel, Dhami, et al. (2021) found substantial unreliability in participants’ mappings of verbal probabilities to numeric probabilities even within a brief task in which repeated exposure to the same probability terms occurred without a change of context.

To test the effect of shifting from the statistically-average consumer to the average of individual consumers, we reanalyzed data from Mandel (2015a) and Mandel and Barnes (2018). Recall that in Mandel and Barnes (2018), 2,013 forecasts were scored for accuracy. Of those, we reanalyzed 1,922 forecasts that used probability terms matching those in the Barnes (2016) probability scheme. We further used data from the 19 intelligence analysts who provided their numeric probability equivalents twice for each of the terms in that probability scheme (once in a “success” context and once in a “failure” context, where each of these contexts was run as a block, the order of which was counterbalanced). Since these analysts were in fact also intelligence consumers, we restricted our re-analyses to this subsample (and hence excluded the university student subsample).

Using the aforementioned data from Mandel (2015a), we computed three forecast scores that varied the degree to which probability translations from verbal to numeric modes were averaged. In the averaged Brier score, which involved no averaging of probability translations, the 1,922 forecasts were scored for accuracy using the Brier score 38 times: twice using each of the 19 analysts' translations of the relevant probability terms. The two scorings per analyst were based on the translation values elicited in the success and failure blocks, respectively. The left violin plot of Figure 1 shows the distributional characteristics of these scores. In the averaged probability score, each analyst's numeric probability equivalents elicited under the two contexts (i.e., success and failure) were first averaged before a single vector of Brier scores was computed. We hypothesized that this scoring method would improve on the averaged Brier score because it harnesses the "wisdom of the internal crowd" (Herzog & Hertwig, 2014), providing more reliable (less noisy) numeric estimates of the verbal probabilities. Figure 1 shows the distributional characteristics of this scoring method in the middle violin plot. Compared to the scoring method that does not capitalize on averaging probability translations first (mean $B = .098$, $SD = .133$), mean accuracy of forecasts based on the averaged probability scores (mean $B = .087$, $SD = .133$) is significantly better, and the effect is very large ($t[1921] = 82.80$, $p < .001$, $r = .999$, $d_{\text{repeated_measures,pooled}} = 1.85$ [Lenhard & Lenhard, 2016]).

The best accuracy (mean $B = .057$, $SD = .159$), however, is seen in the rightmost violin plot, which recapitulates the method used in Mandel (2015a); namely, averaging estimates not only within subjects by averaging the two estimates across contexts, but also taking the median value of those estimates across translators. This score yielded a very large improvement in accuracy over the averaged probability scoring method ($t[1921] = 43.71$, $p < .001$, $r = .993$, $d_{\text{repeated_measures,pooled}} = 1.73$), and the effect size was even larger if compared to the scoring method

that does not average translations at all ($d_{\text{repeated_measures,pooled}} = 2.21$). In sum, the preceding analyses show a massive benefit of averaging consumers' numeric translations of verbal probabilities, both internally (within subjects) and externally (within and between subjects), when examining how accurately strategic intelligence forecasts are likely to be when viewed through the lens of the intelligence consumer. As we discuss below, these findings have implications for current methods that intelligence organizations use to communicate uncertainty in intelligence estimates.

Figure 1. Box and violin plots of accuracy using three scoring methods.

Remarks

A roughly decade-long investigation into the forecasting accuracy of Canadian strategic intelligence forecasts has yielded several important findings. First, the study pioneered by Barnes and Mandel (e.g., Barnes, 2016; Mandel & Barnes, 2014) showed that intelligence organizations could take tangible steps toward implementing an outcome-based accountability system that proactively tracks the accuracy of intelligence forecasts. Such a system can benefit from the use

of numeric probabilities. However, as the later work in this research program shows (Mandel, 2015a; Mandel & Barnes, 2018), the use of numeric probabilities is not essential for tracking accuracy. Moreover, forecasts that were made without the use of numeric probabilities did not differ substantially in terms of accuracy or the type of biases recorded. For instance, forecasts made using numeric probabilities and those made using only verbal probabilities both showed significant underconfidence and comparable discrimination skill.

The advantages of proactive tracking revealed by this study are evident by the many discrepancies observed with other related research that has not been conducted on real intelligence products. For instance, rather than finding the typical hard-easy effect demonstrated in laboratory experiments (Lichtenstein & Fischhoff, 1977), there was evidence of an “easy-hard effect” in which forecasting topics rated by expert coders as more difficult generated more underconfidence than those rated as easier. This may be due to the fact that harder topics were generally rated as more important and, therefore, analysts may have been prone to be more cautious in their assessments to avert unwarranted risk taking by decision-makers or to mitigate blame should the forecasts be proven wrong.

The findings from this program of research also yielded several discrepancies with Tetlock’s (2005) study of expert political judgment. Perhaps the most obvious difference was that the two studies produced evidence of miscalibration in opposing directions: overconfidence in Tetlock (2005) and underconfidence in Barnes and Mandel (2014, 2018). Whereas Tetlock (2005) found no benefit of experience, Barnes and Mandel (2014) found that senior analysts had significantly better discrimination than junior analysts. In short, the research reviewed here on the tracking of forecasting accuracy in Canadian strategic intelligence reports underscores the risks associated with organizational learning that is based largely on knowledge extraction from

the social and behavioral sciences. The cost of implementing a tracking system such as Mandel and Barnes (2014) implemented is minuscule not only relative to yearly intelligence budgets, but even relative to IARPA programs such as ACE and HFC, which although informative, have not shed direct light on analytic performance.

Intelligence organizations could require forecasts to be identified, logged and scored for accuracy once outcome data was available. Analysts who use verbal probabilities to convey forecast probabilities could be asked to provide numeric range equivalents for the terms, the midpoints of which could be used as a numeric probability equivalents for scoring purposes. Or in cases where the probability terms are already mapped to numeric ranges through a numerically bound linguistic probability scheme adopted by the intelligence organization (e.g., see Dhami & Mandel, 2020; Friedman, 2019; Mandel & Irwin, 2020), those ranges could be used. A useful first step in this approach might be to apply it to key judgments in intelligence reports—namely, those judgments deemed to be of highest importance that are usually flagged in executive summaries. As noted earlier, Mandel and Barnes (2018) found there was no significant difference in forecasting skill characteristics between key and non-key judgments. Therefore, analyses of key judgments might be an economical method of gaining insight into an organization's forecasting skill. The effect of interventions aimed at improving forecasting accuracy could also be tracked in one manner or another (e.g., as pre-post interventions or in [quasi-]randomly controlled trials) to expedite organizational understanding of which interventions work in practice and which are worth the investment of required resources.

The present extension of Mandel (2015a) and Mandel and Barnes (2018) showed unambiguously that the accuracy of forecasts scored from the perspective of individual intelligence consumers interpreting the verbal probabilities, on average, is substantially reduced

relative to the accuracy of such forecasts when they are scored based on the average interpretation of probability terms across a group of intelligence consumers. These findings provide a counterpoint to those reported by Mandel (2015a), which showed that forecast accuracy was comparable when scored on the basis of analysts' numeric probabilities or the median probability equivalents assigned by independent "consumers." Such findings could be misinterpreted as indicating that verbal probabilities are about as effective as numeric probabilities when decoded by intelligence consumers. This conclusion is not supported by the evidence. The accuracy of forecasts was substantially worse when individual consumers' interpretations of the verbal terms were used as a basis for scoring forecasts than when the median interpretation of the same consumers was used. Indeed, they were substantially worse than when two numeric equivalents produced by a given consumer were averaged first.

In practice, policymakers, military commanders, and other intelligence consumers rely upon their own interpretations of verbal probabilities as they process intelligence. They do not first approach a group of their peers or subordinates and ask them what the words such as *unlikely* or *almost certain* mean to them and then find the average interpretation to then adopt as their own. Therefore, the novel results we reported in this article reflect a more accurate view of the real-world cost that the unreliability in the interpretation of verbal probabilities exacts on forecast accuracy from the intelligence consumer's perspective. Unfortunately, although the costs may be substantial, opposition to communicating probabilities with the use of numeric ranges has not waned much, if at all, in the many decades since Kent (1964) expressed his frustration with the "poets" of intelligence analysis and used numeric ranges to at least stipulate an organizational meaning for probability terms (Barnes, 2016; Dhami & Mandel, 2020; Friedman & Zeckhauser, 2016; Mandel & Irwin, 2020; Marchio, 2014; Spielmann, 2016).

Objections to the quantification of probability estimates is far from unfathomable. It is well established that, on average, information senders prefer to encode probabilities in utterances or written statements using verbal probabilities that are easier and “more natural” to produce, whereas information receivers prefer to decode numeric probability, which they find more informative and less vague and ambiguous (Brun & Teigen, 1988; Erev & Cohen, 1990; Juanchich & Sirota, 2020; Wallsten et al., 1993). Aside from greater ease, verbal probabilities (due to what is often described as their “directionality”) also function more effectively as pragmatic devices to signal preference and recommendations for subsequent action (Brun & Teigen, 1988; Sher & McKenzie, 2006; Teigen & Brun, 1995, 1999). For instance, Collins and Mandel (2019) found that, on average, whereas subjects rated numeric probabilities as more informative than verbal probabilities about the actual degree of probability communicated in a forecast by a financial advisor (described in a vignette), subjects rated verbal probabilities as better than numeric probabilities at conveying the advisor’s *implicit* recommendation for investing (in fact, the advisor offered no explicit recommendation in the vignette). Such “reading between the lines” may be of value in some contexts. For instance, if probability communicators want to “nudge” decision-makers towards a particular option or plan of action, verbal probabilities might be more effective. However, if the principal aim of communicating probabilities is to inform end-users of the degree of probability and uncertainty surrounding focal events, as is routinely the case in the intelligence domain—and further, if the mandate of intelligence is not to “nudge” but to remain policy neutral, as has long been the case (Kent, 1951)—then the empirical evidence suggests that intelligence organizations should quantify probabilities in intelligence assessments, and especially in forecasts which invariably require probability descriptions.

Intelligence forecasts do not only support decision-makers—they also provide input into other intelligence products such as threat and risk assessments (Friedman, 2019). Most national security threats require a set of individually necessary and jointly sufficient conditions to occur for the threat to manifest. Estimating such threats, even in the simplest cases in which the necessary conditions are independent, requires computing the product of the probabilities. Likewise, when a decision-maker receives intelligence from multiple sources, it may be helpful to be able to compute an average estimate of chances, as was highlighted by Friedman and Zeckhauser (2015) in President Obama’s decision to capture or kill Osama bin Laden. Recent research examining the computability of products and averages of probabilities shows that these values can be estimated more accurately and coherently using numeric point estimates than using verbal probabilities or numeric ranges (Mandel, Dhimi et al., 2021). Although numeric ranges did not support better computation than verbal probabilities, ranges (unlike verbal probabilities) can unambiguously be converted into point estimates with margins of error (Moore et al., 2009).

Intelligence organizations might counter that they have, in fact, implemented schemes for improving the communication of probability. Most of the current approaches are of the type described by Mandel, Wallsten et al. (2021) as *numerically bounded linguistic probability* (NBLP) *schemes*. Recent reviews of such schemes in intelligence can be found in Friedman (2019), Dhimi and Mandel (2020), and Mandel and Irwin (2020). The key characteristic of NBLP schemes is that a small set of verbal probabilities are selected for use in intelligence reports and each of these is defined in terms of a numeric probability range. However, most of the definitions are by fiat and recent research has shown that they do not optimize the agreement between end-users’ understanding of the terms’ meaning and the stipulated meaning (Ho et al., 2015; Wintle et al., 2019). Moreover, studies examining how well such schemes foster

agreement in the interpretation of probabilities in assessments has yielded disappointing results in climate science (Budescu et al., 2009; 2014) and intelligence analysis (Wintle et al., 2019; Mandel & Irwin, 2021a, 2021b). These studies indicate quite clearly that people do not simply accept and keep track of stipulated meaning of terms that are frequently encountered in everyday language use. Even when asked what “the sender meant” (as opposed to what they personally understand the term to mean) by such and such a term, participants appear to default to their own interpretations. Should we expect anything different of elite, time-strapped and account-pressured decision-makers?

Conclusion

The program of research conducted in Canada to verify the accuracy of strategic intelligence forecasts is not only rare but unique. To the best of our knowledge, there are no comparable studies of the forecast accuracy of real intelligence assessments that have conducted, let alone published in the open scientific literature. The preconditions for such research are difficult to attain. Most intelligence assessments are classified and off limits for research purposes. In order to analyze the forecasts used by Mandel and Barnes (2014, 2018), the content of the forecasts had to be separated from the forecast probabilities, outcomes, and other variables used in the analyses. Had an intelligence director (Barnes) and a defence scientist (Mandel) not teamed up, the project would not have been possible. We suspect that future studies like this will be rare precisely because of the improbable preconditions required for success. However, prospects for internal tracking of forecasting accuracy by intelligence organizations as part of an outcome-based accountability approach are much more feasible. The research reviewed in this article demonstrates how intelligence organizations could implement such an approach, the results of which could remain classified and shared on a need-to-know basis. Furthermore, the

research indicates that forecast accuracy could be improved by using numeric probabilities to communicate uncertainties in assessments to intelligence consumers.

Over a decade ago, after having been briefed by the first author of this article on the early phase of the study, the US National Academies' National Research Council Committee on Behavioral and Social Science Research to Improve Intelligence Analysis for National Security stated in its report to ODNI the following preamble, recommendation and immediate actions:

We recognize that there has historically been resistance to numeric probability estimates from analysts who believe they imply artificial precision. However, as discussed in Chapter 2, the scientific evidence, including Canada's real-world success with numeric probabilities in intelligence analysis (Mandel, 2009), suggest that, with proper training and feedback, such judgments could substantially improve analytic products and customer understanding of them. Proper incentives seek to encourage learning, not to determine culpability. They reward positive performance and cultivate the natural desire to do well, a desire that is especially prevalent in the IC [intelligence community]. In addition, numeric probabilities allow feedback that is essential to learning. Proper incentives discourage both overconfidence (intended perhaps to carry an argument) and underconfidence (intended perhaps to avoid responsibility). They encourage good calibration: being as confident as one's understanding warrants. Thus, DNI [Director of National Intelligence] must ensure that numeric probabilities are implemented in a constructive way, using them for useful feedback, not destructive criticism.

Recommendation 2

The Director of National Intelligence should ensure that the intelligence community adopts scientifically validated analytical methods and subjects all of its methods to

performance evaluation. To that end, each analytical product should report, in a standardized format, the elements necessary for such evaluation, including its analytical method, domain, conclusions, analysts' background, and the collaboration that produced it. Analyses must include quantitative judgments of the probability and uncertainty of the events that they forecast. These reports should be archived in a database that is routinely used to promote institutional learning and individual training and as input to the Director of National Intelligence's ongoing review efforts of analytic shortfalls and plans to address them.

Immediate Actions

1. Institutionalize an "Analytical Olympics," with analysts and analytical methods competing to provide the best calibrated probabilities (i.e., showing appropriate levels of confidence) in assessments and predictions made for well-specified outcomes that have occurred or will occur in the near future.
 2. Begin assessing how well-calibrated individual analysts are, using the results as personal feedback that will allow analysts to improve their own performance and IC to learn how this performance is related to workforce factors, such as personal capabilities, training, and incentives.
 3. Create a research program that reviews current and historic assessments, looking for correlates of accuracy and calibration, considers properties such as method used, collaboration process, classification level, substantive domain, and team composition.
- (National Research Council, 2011, pp. 86–87)

Unfortunately, it does not appear that ODNI has taken this advice, although it has invested many millions into IARPA programs aimed at studying how to improve forecasting accuracy. Such programs have stimulated important research on forecasting methods, but they cannot replace self-monitoring of forecast accuracy in real products by intelligence organizations. These two types of activities should be mutually reinforcing. In the absence of proactive tracking of forecast accuracy by intelligence organizations, it is unclear whether the scientific findings of IARPA programs like ACE and HFC will appear on the radar of day-to-day intelligence producing directors and analysts.

References

- Arkes, H.R., & Kajdasz, J. (2011). Intuitive theories of behavior. In B. Fischhoff & C. Chauvin (Eds.), *Intelligence Analysis: Behavioral and Social Scientific Foundations* (pp. 143-168). Washington, DC: The National Academies Press.
- Åstebro, T., & Koehler, D. J. (2007). Calibration accuracy of a judgmental process that predicts the commercial success of new product ideas. *Journal of Behavioral Decision Making*, 20, 381-403.
- Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., Tetlock, P., Ungar, L., & Mellers, B. (2017). Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science*, 63, 691-706.
- Barnes, A. (2016). Making intelligence analysis more intelligent: Using numeric probabilities. *Intelligence and National Security*, 31, 327-344.
- Baron, J., Mellers, B. A., Tetlock, P. E., Stone, E., & Ungar, L. H. (2014). Two reasons to make aggregated probability forecasts more extreme. *Decision Analysis*, 11, 133-145.
- Berlin, I. (1953). *The Hedgehog and the Fox: An Essay on Tolstoy's View of History*. New York, NY: Weidenfeld & Nicolson.
- Betts, R. K. (2007). *Enemies of Intelligence: Knowledge and Power in American National Security*. New York, NY: Columbia University Press.
- Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78, 1-3.
- Brun, W., & Teigen, K. H. (1988). Verbal probabilities: Ambiguous, context-dependent, or both? *Organizational Behavior and Human Decision Processes*, 41, 390-404.
- Budescu, D. V., Broomell, S., & Por, H. (2009). Improving communication of

uncertainty in the reports of the intergovernmental panel on climate change. *Psychological Science*, 20(3), 299-308.

Budescu, D. V., Por, H., Broomell, S. B., & Smithson, M. (2014). The interpretation of IPCC probabilistic statements around the world. *Nature Climate Change*, 4, 508-512.

Chang, W., & Tetlock, P. E. (2016). Rethinking the training of intelligence analysts. *Intelligence and National Security*, 31, 903-920.

Chang, W., Berdini, E., Mandel, D. R., & Tetlock, P. E. (2018). Restructuring structured analytic techniques in intelligence. *Intelligence and National Security*, 33, 337-356.

Clapper, J. (2014). *The National Intelligence Strategy of the United States of America: 2014*. Washington DC: Office of the Director of National Intelligence.

Collins, R. N., & Mandel, D. R. (2019). Cultivating credibility with probability words and numbers. *Judgment and Decision Making*, 14(6), 683-695.

Dhami, M. K., & Mandel, D. R. (2020). Words or numbers? Communicating probability in intelligence analysis. *American Psychologist*. <http://dx.doi.org/10.1037/amp0000637>

Dhami, M. K., Mandel, D. R., Mellers, B. A., & Tetlock, P. E. (2015). Improving intelligence analysis with decision science. *Perspectives on Psychological Science*, 10, 753-757.

Erev, I., & Cohen, B. L. (1990). Verbal versus numerical probabilities: Efficiency, biases, and the preference paradox. *Organizational Behavior and Human Decision Processes*, 45(1), 1-18.

Friedman, J. A. (2019). *War and Chance: Assessing Uncertainty in International Politics*. New York, NY: Oxford University Press.

Friedman, J. A., & Zeckhauser, R. (2015). Handling and mishandling estimative probability: Likelihood, confidence, and the search for bin Laden. *Intelligence and National*

Security, 30(1), 77-99.

Friedman, J. A., & Zeckhauser, R. (2016). Why assessing estimative accuracy is feasible and desirable. *Intelligence and National Security*, 31, 178-200.

Friedman, J. A., Baker, J. D., Mellers, B. A., Tetlock, P. E., & Zeckhauser, R. (2018). The value of precision in probability assessment: Evidence from a large-scale geopolitical forecasting tournament. *International Studies Quarterly*, 62, 410-422.

Goodman-Delahunty, J., Granhag, P. A., Hartwig, M., & Loftus, E. F. (2010). Insightful or wishful: Lawyers' ability to predict case outcomes. *Psychology, Public Policy and Law*, 16, 133-157.

Herzog, S. M., & Hertwig, R. (2014). Harnessing the wisdom of the inner crowd. *Trends in Cognitive Sciences*, 18, 504-506.

Ho, E., Budescu, D. V., Dhami, M. K., & Mandel, D. R. (2015). Improving the communication of uncertainty in climate science and intelligence analysis. *Behavioral Science & Policy*, 1(2), 43-55.

Johnson, L. K. (2007). A shock theory of congressional accountability for intelligence. In: *Handbook of Intelligence Studies*, Johnson, L. K. (Ed.), 343-360. New York, NY: Routledge.

Juanchich, M., & Sirota, M. (2020). Most family physicians report communicating the risks of adverse drug reactions in words (vs. numbers). *Applied Cognitive Psychology*, 34(2), 526-534.

Karvetski, C. W., Mandel, D. R., & Irwin, D. (2020). Improving probability judgment in intelligence analysis: From structured analysis to statistical aggregation. *Risk Analysis*, 40(5), 1040-1057. <https://doi.org/10.1111/risa.13443>

Kent, S. (1951). *Strategic intelligence for American world policy*. Princeton, NJ:

Princeton University Press.

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

Kent, S. (1994). Estimates and influence. In: *Sherman Kent and the Board of National Estimates: Collected Essays*, Steury, D. P. (Ed.), 51-59. Washington DC: Center for the Study of Intelligence.

Keren, G. (1987). Facing uncertainty in the game of bridge: A calibration study. *Organizational Behavior and Human Decision Processes*, 39(1), 98-114.

Lehner, P., Michelson, A., Adelman, L., & Goodman, A. (2012). Using inferred probabilities to measure the accuracy of imprecise forecasts. *Judgment and Decision Making*, 7(6), 728-740.

Lenhard, W., & Lenhard, A. (2016). *Calculation of Effect Sizes*. Retrieved from: https://www.psychometrica.de/effect_size.html. Dettelbach, Germany: Psychometrica. <https://doi.org/10.13140/RG.2.2.17823.92329>

Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? *Organizational Behavior and Human Performance*, 20(2), 159-183.

Lin, S.-W., & Bier, V. M. (2008). A study of expert overconfidence. *Reliability Engineering and System Safety*, 93(5), 711-721.

Mandel, D.R. (2009). *Canadian perspectives: Applied behavioral science in support of intelligence analysis*. Invited paper presented at the Public Workshop of the National Academy of Sciences Committee on Behavioral and Social Science Research to Improve Intelligence Analysis for National Security, Washington, D.C.

Mandel, D. R. (2015a). Accuracy of intelligence forecasts from the intelligence consumer's perspective. *Policy Insights from the Behavioral and Brain Sciences*, 2(1), 111-120.

Mandel, D. R. (2015b). Instruction in information structuring improves Bayesian judgment in intelligence analysts. *Frontiers in Psychology*, 6, 387.

Mandel, D. (2019). Too soon to tell if the US intelligence community prediction market is more accurate than intelligence reports: Commentary on Stastny and Lehner (2018). *Judgment and Decision Making*, 14(3), 288–292.

Mandel, D. R., & Barnes, A. (2014). Accuracy of forecasts in strategic intelligence. *Proceedings of the National Academy of Sciences of the United States of America*, 111(30), 10984-10989.

Mandel, D.R., & Barnes, A. (2018). Geopolitical forecasting skill in strategic intelligence. *Journal of Behavioral Decision Making*, 31(1), 127-137.

Mandel, D. R., & Irwin, D. (2020). Uncertainty, Intelligence, and National Security Decisionmaking. *International Journal of Intelligence and CounterIntelligence*, 1–25.

Mandel, D. R., & Irwin, D. (2021a). Facilitating sender-receiver agreement in communicated probabilities: Is it best to use words, numbers or both? *Judgment and Decision Making*, 16(2), 363-393.

Mandel, D. R., & Irwin, D. (2021b). On measuring agreement with numerically bounded linguistic probability schemes: A re-analysis of data from Wintle, Fraser, Wills, Nicholson, and Fidler (2019). *PLoS ONE*, 16(3), e0248424.

Mandel, D. R., Karvetski, C. & Dhami, M. K. (2018). Boosting intelligence analysts' judgment accuracy: what works, what fails? *Judgment and Decision Making*, 13, 607-621.

Mandel, D. R., & Tetlock, P. E. (2016). Debunking the myth of value-neutral virginity: Toward truth in scientific advertising. *Frontiers in Psychology* 7, 451.

Mandel, D.R., Barnes, A., & Richards, K. (2014). *A Quantitative Assessment of the*

Quality of Strategic Intelligence Forecasts. DRDC Toronto Technical Report 2013-036.

Toronto, ON: DRDC.

Mandel, D. R., Dhimi, M. K., Tran, S., & Irwin, D. (2021). Arithmetic computation with probability words and numbers. *Journal of Behavioral Decision Making*.

<https://doi.org/10.1002/bdm.2232>

Mandel, D. R., Karvetski, C., & Dhimi, M. K. (2018). Boosting intelligence analysts' judgment accuracy: what works, what fails? *Judgment and Decision Making*, 13(6), 607-621.

Mandel, D. R., Wallsten, T. S., & Budescu, D. V. (2021). Numerically bounded linguistic probability schemes are unlikely to communicate uncertainty effectively. *Earth's Future*, 9:e2020EF001526.

Marchio, J. (2014). "If the weatherman can...": The intelligence community's struggle to express analytic uncertainty in the 1970s. *Studies in Intelligence*, 58(4), 31-42.

Mellers, B. A., Baker, J. D., Chen, E., Mandel, D. R., & Tetlock, P. E. (2017). How generalizable is good judgment? A multi-task, multi-benchmark study. *Judgment and Decision Making*, 12(4), 369-381.

Mellers, B. A., Stone, E., Murray, T., Minster, A., Rohrbaugh, N., Bishop, M., Chen, E., Baker, J., Hou, Y., Horowitz, M., Ungar, L., & Tetlock, P. E. (2015). Identifying and cultivating "Superforecasters" as a method of improving probabilistic predictions. *Perspectives in Psychological Science*, 10(3), 267-281.

Mellers, B. A., Ungar, L., Baron, J., Ramos, J., Gurcay, B., Fincher, K., Scott, S. E., Moore, D., Atanasov, P., Swift, S. A., Murray, T., Stone, E., & Tetlock, P. E. (2014). Psychological strategies for winning geopolitical forecasting tournaments. *Psychological Science*, 25(5), 1106-1115.

- Moore, R. E., Kearfott, R. B., & Cloud, M. J. (2009). *Introduction to Interval Analysis*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Moynihan, D. P. (1991, May 19). Do we still need the C.I.A.? The State Dept. can do the job. *The New York Times*, E17.
- Murphy, A. H. (1973). A new vector partition of the probability score. *Journal of Applied Meteorology*, 12, 595-600.
- Murphy, A. H., & Winkler, R. L. (1984). Probability forecasting in meteorology. *Journal of the American Statistical Association*, 79(387), 489-500.
- National Research Council. (2011). *Intelligence analysis for tomorrow: Advances from the behavioural and social sciences*. Washington, DC: National Academies Press.
- North Atlantic Treaty Organization. (2016). *AJP-2.1, edition B, version 1: Allied joint doctrine for intelligence procedures*. Brussels: NATO Standardization Office.
- Rieber, S. (2004). Intelligence analysis and judgmental calibration. *International Journal of Intelligence and Counter-Intelligence*, 17(1), 97-112.
- Spielmann, K. (2016). I got algorithm: Can there be a Nate Silver in intelligence? *International Journal of Intelligence and CounterIntelligence*, 29(3), 525-544.
- Stastny, B., & Lehner, P. (2018). Comparative evaluation of the forecast accuracy of analysis reports and a prediction market. *Judgment and Decision Making*, 13(2), 202–211.
- Tetlock, P. E. (2002). Social functionalist frameworks for judgment and choice: Intuitive politicians, theologians, and prosecutors. *Psychological Review*, 109(3), 451-471.
- Tetlock, P. E. (2005). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.
- Tetlock, P. E., & Gardner, D. (2015). *Superforecasting: The Art and Science of*

Prediction. New York, NY: Crown Publishing Group.

Tetlock, P. E., & Kim, J.I. (1987). Accountability and judgment processes in a personality prediction task. *Journal of Personality and Social Psychology*, 52(4), 700-709.

Tetlock, P. E., & Mellers, B. A. (2011). Intelligent management of intelligence agencies: Beyond accountability ping-pong. *American Psychologist*, 66(6), 542-554.

Tetlock, P. E., & Mellers, B. A. (2014). Judging political judgment. *Proceedings of the National Academy of Sciences of the United States of America*, 111(32), 11574-11575.

Turner, B. M., Steyvers, M., Merkle, E. C., Budescu, D. V., & Wallsten, T. S. (2014). Forecast aggregation via recalibration. *Machine Learning*, 95(3), 261-289.

Wallsten, T. S., Budescu, D. V., Zwick, R., & Kemp, S. M. (1993). Preferences and reasons for communicating probabilistic information in numeric or verbal terms. *Bulletin of the Psychonomic Society*, 31(2), 135-138.

Wintle, B. C., Fraser, H., Wills, B. C., Nicholson, A. E., & Fidler, F. (2019). Verbal probabilities: Very likely to be somewhat more confusing than numbers. *PloS One*, 14(4), e0213522-e0213522.

Yaniv, I., Yates, J. F., & Smith, J. E. K. (1991). Measures of discrimination skill in probabilistic judgment. *Psychological Bulletin*, 110(3), 611-617.