The Effect of Probabilistic Information Threshold Forecasts

Susan Joslyn[1], Karla Schweitzer[1], David Jones[2], John Pyles[3] and Earl Hunt[1]

[1] Department of Psychology, University of Washington

[2] Department of Applied Physics, University of Washington

[3] Department of Cognitive Sciences, University of California, Irvine

Submitted to Weather and Forecasting

September dd 2005

Correspondence and requests for materials should be addressed to Susan Joslyn (e-mail: susanj@u.washington.edu) Department of Psychology, Box 351525

University of Washington

Seattle, Washington, 98125

Abstract

The study reported here asks whether the use of probabilistic information indicating forecast uncertainty improves the quality of deterministic weather decisions. Participants made realistic wind speed forecasts based on historical data in a controlled laboratory setting. They also decided whether it was appropriate to post an advisory for winds greater than 20 knots during the same time intervals and in the same geographic locations. On half of the forecasts each participant also read a chart showing color-coded probability of winds greater than 20 knots (23m s^{-1}). Participants had a general tendency to post too many advisories in the low to mid probability situations (0-50%) and too few advisories in very high probability situations (90-100%). However, the probability product attenuated these biases. When participants used the probability product they posted fewer advisories when the probability of high winds was low and they posted more advisories when the probability of high winds was high. The difference was due to the probability product alone because the within subjects design and counterbalancing of forecast dates ruled out those alternative explanations. The data suggest that the probability product improved threshold forecast decisions.

1. Introduction

Modern day weather forecasters rely heavily on numerical weather and climate models that make weather predictions by transforming present into future weather conditions according to the known principles of atmospheric physics. Due to uncertainties of the initial state of the atmosphere and the computational representation of the equations of motion, model predictions are accompanied by varying amounts of uncertainty. It is now possible, with ensemble forecasts in which multiple simulations of the atmosphere are made, to estimate and quantify the amount uncertainty in the model prediction (Anderson 1996; Grimit and Mass 2002). In theory, this is very useful information to both weather forecasters and to the general public. However, with the exception of probability of precipitation, forecast uncertainty is not usually communicated to the public. At present, most forecasts remain deterministic. In part this is due to a question about whether or not people can successfully make use of uncertainty information to improve deterministic forecast decisions.

As far as we know there is very little research that addresses this issue directly. There is some empirical evidence that people can make use of uncertainty information in simulated tasks to increase rewards and reduce exposure to risk (Roulston et al. manuscript accepted by *Wea. Forecasting*). However, claims that uncertainty information enhances weather forecast value are generally based on normative or prescriptive decision-making models (Richardson 2002; Palmer 2002). By contrast, there is ample research that suggests that experienced forecasters understand and reliably estimate credible intervals and probability forecasts (Murphy and Winkler, 1974a, 1974b, 1977). However some recent evidence suggests an over-forecasting bias when safety was an issue (Keith 2003). In sum, there is no research of which we are aware that attempts to gage the impact of uncertainty information on realistic deterministic forecasts.

The study presented here was conducted to determine whether wind speed or high wind advisory (wind speeds in excess of 20 knots (23m s^{-1})) forecasts differ as the result of reviewing charts indicating the probability of wind speeds exceeding 20 knots. Participants made the forecasts for marine areas in which small boaters would be affected. The hypothesis was that probabilistic information would impact a threshold warning forecast in this context. The reasoning was as follows: forecasters might decide to post an advisory even if the probability for high winds was small, to reduce the risk of boating accidents. Reviewing explicit probabilistic information would alert them to such situations.

2. Method

a. Participants

Ten male University of Washington atmospheric science students participated in this study. All participants had completed basic instruction in forecasting. Three of the participants were undergraduate students who had completed a course in atmospheric structure and analysis. Seven participants were graduate students. Participants were paid $40 for participating in the two-session study. They were paid $20 after their first session and $20 after their second session.

b. Task

Participants made four forecasts over two sessions. For each forecast they were asked to review historical data and predict wind speed and direction for four locations in the Puget Sound region: Smith Island, Destruction Island, West Point, and Tatoosh Island. As part of a single forecast, they predicted wind speed and direction for each location, every 6 hours, for a 48 hours period approximately 27 hours in the future. This resulted in 36 wind speeds and directions in all. The 48 hour forecasting period began at 4:00 P.M. (1600 local, 0000 UTC) the following day. Participants were also asked to decide whether they would issue a high wind advisory for any of the four forecast locations during the 48-hour period. If so, they were asked to indicate the hours during which the advisory should be posted. The concept of the wind advisory, its purpose and techniques for predicting it, had been covered in the completed coursework. For the purposes of this exercise, a high wind advisory was defined to be winds greater than 20 knots. Participants were asked to disregard wave height, which is usually a factor in small craft advisories.

c. Materials

1) WEATHER DATA

Participants were provided with historic weather information as data upon which to base the wind forecast. The selection of information sources was based on the results of an observational study of forecasters at a near-by Naval station who produced a similar forecast. Participants in the present study had several products from prominent models, including the Fifth-Generation Pennsylvania State University_National Center of Atmospheric Research Mesoscale Model (MM5), the Nested Grid Model (NGM), and the Aviation Model (AVN). The sources also included satellite and radar imagery, regional terminal airdrome forecasts (TAFs), meteograms, and observations taken from buoys in locations near the forecast sites (see Appendix A for a complete list).

On half the trials participants were also provided with a probability product. It was a chart, color-coded for the probability of 10-meter winds in excess of 20 knots (see Figure 1). The chart was based on the Centroid Mirroring Ensemble (ACME) in the 12-km and 36-km domains (Grimit & Mass, 2003). ACME consists of 17 individual forecasts (called ensemble members) all using the MM5 with an identical physics package, but different boundary and initial conditions drawn from variety of global models. Although most participants had been introduced to ensemble forecasting in their coursework, they were reminded that probability of winds in excess of 20 knots was estimated from the degree of agreement between ensemble members. The chart was color-coded and it divided probability of high winds into 6 categories. Areas in which 90-100% of the ensemble members predicted winds greater than 20 knots were color coded red and indicated 90 to 100% chance of winds greater than 20 knots. Similarly, yellow areas indicated 70-90%, the green areas indicated 50-70%, the blue areas indicated 30 - 50% and the purple areas indicated 10-30% chance. Any areas that were white had a 10% or less probability of winds over 20 knots.

All data was from approximately 1:30 P.M. (1330 local, 2130 UTC) on the day they were collected and participants were informed of this fact. Four days of historic data were used in the experiment (February 14, February 20, March 11, and March 26), one for each forecast. The historic data were selected to have some periods of high winds and some periods during which the winds were calm.

2) INTERFACE

Information, in web-page format was presented on a 1024 x 768 computer screen in High Color (16 bit) mode. The main difference between the presentation format in the study and that in use in a typical forecasting office was that some products (e.g. MM5, satellite) that are normally viewed in an animated loop were presented in a step through format using an image gallery web page design. Thumbnails of all the images from a product opened to a full size image in response to a mouse click. Navigation buttons moved the user forward and backward through the images simulating an animation loop. Other information sources (such as TAFs, meteograms, and buoy observations) were captured as complete web pages and presented in their original form and linked directly to the main link page. A web-gallery generator called

Express Thumbnail Creator (http://www.express-soft.com/etc/) was used to generate the thumbnails, thumbnail pages, and image pages with navigation links. All product-image web pages were connected to a simple main links page with a link that led to the main thumbnail page or information source.

Thus, the user saw the main page with text links to individual products such as the satellite, radar, and model products. Clicking on a text link to a product opened a page of thumbnail images for that product, which could then be clicked on to display full size images. Participants could then navigate through the images using forward and backward buttons. An up arrow navigation button reconnected with the thumbnails page, and a home button returned to the main links page. TAFs, meteograms, and buoy observations opened up in a separate window. When users closed the window they returned to the main links page. The experimenter demonstrated these procedures to participants.

3) ANSWER SHEETS

There were two answer sheets for each forecast, the wind speed and direction answer sheet and the wind advisory answer sheet. The wind speed and direction answer sheet provided four columns each headed by the name of one of the four locations for which a forecast was required. There was a row designated for each of the 8 hours of the forecast with a blank for the wind speed and direction. Although forecasters were asked to record wind speed and direction to make the task realistic, these data were not analyzed, as they were less likely to be influenced by the probability of wind speeds greater than 20 knots. Participants recorded the wind advisory on a separate sheet to reduce the influence of the wind speed forecast on the wind advisory. This procedure was used to discourage participants from simply posting a wind advisory for times periods during which the deterministic forecast was over 20 knots and encourage them to take the probabilistic information into account when making this decision. On the wind advisory answer sheet there were four questions asking whether participants would forecast a high wind advisory during the 48-hour forecast period for each of the four locations and if so, to indicate which hours.

4) MAP

Although the forecast locations were not marked on the probability charts, participants were provided with a separate map showing the four forecast locations as well as the locations of nearby airfields for which they had the TAFs.

d. Procedure

Participants were tested individually in two sessions each lasting approximately an hour and a half. In the first session, after informed consent procedures, the experimenter explained the forecasting task and how to fill out the answer sheets. The forecast locations were pointed out on the map that was posted above the workstation. The experimenter demonstrated how to access information on the computer and explained that participants could use whichever sources of information they wished.

Then the experimenter introduced the probability product and informed the participant that it would also be available on some forecasts. The experimenter demonstrated how to read the probability chart and how it was generated. In order to ensure that participants read and understood the probability information, in trials with the probability product, they were required to record the probability ranges for each of the four locations at each of the forecast times. This procedure was initiated after several of the pilot participants ignored the probability product altogether commenting that it was not useful for the forecast they were required to make.

Because the impact of the probability product was the focus, it was necessary to ensure that participants had encoded the information.

When all of the participants' questions had been answered, they made a practice forecast with the probability product to familiarize themselves with the procedure and the interface. Upon completing the practice forecast, the experimenter checked the answer sheets for completeness. Unless there were further questions, the participant then made two test forecasts one with and one without the probability product. The session was not time-limited but generally took between an hour and a half to two hours. Participants returned for a second session less than a week later to complete the remaining two weather forecasts.

e. Design

Probability information was manipulated within participants. Each participant had two trials with the probability product and two trials without the probability product. Half of the participants began the session with the probability product and half did not. Weather data date order was similarly counter balanced. The weather data dates were rotated through conditions so that each day was used in a forecast with the probability product on half of the trials and without it on the other half. Rotation ensured that weather conditions were equivalent across conditions. This was important because the weather on some dates might have been easier to forecast accurately than on other dates. However, no participant saw data from the same date twice.

3. Results

The study was designed to investigate the effect of the probabilistic information (the probability of wind speeds exceeding 20 knots) on decisions to post a high wind advisory. First, however, participants' ability to read the probability product was examined to ensure that participants had accurately encoded the probability information. Then the condition with the probability product was compared to the condition without the probability product to evaluate wind advisory decisions.

a. Reading probability product

In order to assess the consistency with which participants read the probability of winds greater than 20 knots, agreement between participants was calculated [1]. Agreement was the percentage of participants who recorded the same range of probabilities for a given location at a given time period divided by the total number of participants forecasting that date.[2] Summed across locations and time periods, participants agreed in 80% of the cases. In other words, participants disagreed in the interpretation of probability product in 20% of the cases. There are several possible explanations for this. The forecast sites were not labeled directly on the probability product so locations were inferred using a separate map. In addition, some of the colored areas were small and the boundaries between them were difficult to distinguish in the graphic. Cases in which participants disagreed on the interpretation of the probability product were omitted from subsequent analyses.

b. Wind advisory analyses

Next, the influence of the probability product on the wind advisory was examined. Accuracy for posting the wind advisory was defined in terms of the signal detection measure of

---

[1] This is an estimate, as the forecast locations were not marked on the probability product itself so there was no objective answer to the range of probabilities displayed. Only the cases in which more than half of the participants agreed on the interpretation of the probability product were included.

[2] In some cases, even when directed to write down a range, the participants wrote down a single number. In these cases the range was inferred from the number written.

sensitivity[3] (Green and Swets, 1966). Sensitivity is the degree to which the participant can discriminate between a high wind event and a non-high wind event eliminating the influence of response bias. Response bias is the participants overall willingness to post an advisory. In order to compute d´ (sensitivity) the hours during which the participant posted an advisory were compared to the hours during which the observed wind speeds exceeded 20 knots[4]. Then, forecast hours were divided into the following four categories. Hits were defined as cases in which the winds were greater than 20 knots and the participant posted a wind advisory. Misses were cases in which the winds were greater than 20 knots but the participant did not post an advisory. False alarms were cases in which the participant posted an advisory and the winds were less than 20 knots. Correct rejections were cases in which the winds were less than 20 knots and the participant did not post an advisory (see Table 1). For the sensitivity measure (d') higher scores indicate greater ability to discriminate between a high wind event and a non-high wind event. The mean d-prime was greater in the condition with the probability product (d'_{with}=. 1.25) than in the condition without the probability product ($d'_{without}$=. 92). Moreover, when using the probability product, participants had a more conservative response bias (C= .11) than they did without it (C = - .19)[5]. That is, they tended to post <u>fewer</u> advisories with the probability product (38% of the time) than without it (45% of the time). Although none of these differences quite reached significance the implications were intriguing.

---

[4] We examined only the 120 hours per participant for which a frame of the probability product were provided and during which participants agree on the value represented by the product for that location.

[5] $C = .5[z(H)+z(F)]$

In order to further investigate the impact of the probability product on the frequency of posting a wind advisory, the percentage of times participants posted an advisory for each of the probability ranges displayed in the product (0-10%, 10-30%, 30-50%, 50-70%, 70-90%, 90-100%) was calculated. The number of cases in which participants issued an advisory in a given range was divided by the total number of cases in which that range was identified to determine the percent advisories issued. Then, the percent advisories issued in the conditions with and without the probability product were plotted over probability ranges (See Figure 2). For reference, a line matching probability to frequency (expected response) was included. The latter can be interpreted as the hypothetical pattern of responses that perfectly reflects the probability product's forecast.

There were some similarities between the forecasts with and without the probability product. In both conditions participants posted more advisories as the probability of high winds increased. This is not surprising in that participants had the model-produced deterministic prediction for all forecasts. In general, model predicted wind speeds increase as the likelihood of high wind increases. When comparing the human forecasters to the model (expected response), note that participants tended to post advisories in a larger percentage of cases than was indicated in the lower ranges (0-30%). However, this bias was reversed when the probability of high winds was very high. In the very highest probability category (90-100%) participants issued a smaller percentage of advisories than was indicated. This suggests that in general human forecasters are too liberal in their willingness to issue a wind advisory when likelihood is low and too conservative when likelihood is high.

Importantly, these two tendencies, a liberal bias at lower probabilities and conservative bias at higher probabilities, were attenuated by the probability product. Those in the condition with the probability product posted fewer advisories than those in the control condition in the lower ranges (0-50%) and more advisories in the very highest range (70-100%). This effect was statistically significant. Forecasts were divided into two categories, one in which probability of high winds was high (70-90%) and one in which the probability of high winds was low (0-50%)[6]. Then, a 2 (Probability of High Winds: High vs. Low) X 2 (Probability Product: With vs. Without) repeated measures Analysis of Variance (ANOVA) was conducted. The ANOVA yielded a main effect for Probability of High Winds, F (1, 9) = 177.68, MSE = 2.69, $p < .001$. Not surprisingly, people posted significantly more advisories when the likelihood of high winds was high (*M= 81% =*, *SD* =16%) than when it was low (*M* =29%, *SD* =24%). Moreover the interaction of the probability of high winds and the use of the product was also significant (F (1, 9) = 7.05, MSE = .13, p <. 05). The probability product had a significantly different effect on posting decisions, depending on the likelihood of high winds. This suggests that the probability product discouraged participants from posting advisories when the likelihood of high winds was low and encouraged them to post more advisories when likelihood was high.

4. Conclusion

These results suggest that the probability product significantly improved the threshold forecast: posting high wind advisories. It appears to have had its effect by counteracting natural biases in high and low likelihood situations. It has long been known that people do not treat probability linearly (Kahneman and Tversky 1979; Gonzalez and Wu 1999). In this study,

---

[6]Two missing data points for two participants in one category range (90-100) were estimated by calculating the average percent advisories posted for that group.

participants had a liberal bias in the lower probability ranges and a conservative bias in the very highest range. A similar pattern was observed in the probability estimates of experienced forecasters over extended forecast periods (Baars and Moss 2004) and when safety is an issue (Keith 2003).

The tendencies observed in the study reported here may be related to the warning task assigned to participants. Participants may have been sensitive to different errors when the likelihood of high winds was very high than they were when the likelihood was low. Perhaps participants attempted to minimize misses in the low probability situations leading them to post more advisories than were warranted. In high probability situations they may have shifted their focus to false alarms, causing a reduction in the number of advisories posted. Although this is mere speculation in context of the present data, there is evidence that severity of outcome and sensitivity to loss affect the interpretation of even precisely quantified uncertainty (Weber 1994; Windschitl and Weber 1999).

From a practical standpoint, a liberal bias makes sense in the context of the high wind warning task studied here. The purpose of the wind advisory was to prevent boaters from setting sail in conditions of dangerously high winds. Participants may have chosen to err on the side of caution in the lower probability ranges by posting an advisory even when the chance of high winds was small. However, in real life situations an overly liberal bias could lead to problems. Boaters may begin to disregard the advisory if it proves wrong too often and high winds fail to materialize. The user's false alarm tolerance is thought to be critical to the success of warnings such as this (Roulston and Smith 2002). Thus, for situations like this, the use probabilistic information by the forecaster may be especially important. In the study reported here, participants posted fewer advisories in the lower probability ranges when using the product than they did without it, reducing the overall number of false alarms (15% false alarms with the probability product versus 28% false alarms without). This improvement could be critical in real life situations in which trust in the advisory system is crucial for boater safety.

Participants were reluctant to post advisories as often as was warranted in the very highest category (90-100%). This tendency is also problematic in a real life situation in which small boaters could be endangered by setting sail in high wind situations of which they were not warned. Again the probability product attenuated this effect. When participants used the probability product they posted more advisories when high winds were very likely than they did when they did not use it.

It is important to remember that the same participants and the same weather data were used in both conditions. The only difference between conditions was the probability product itself. Thus the probability product had an important positive impact counteracting two problematic biases and improving threshold forecast decisions. There is now strong evidence that probabilistic information is beneficial for realistic deterministic forecast decisions.

Acknowledgements. This research was supported by the DOD Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under Grant N00014-01-10745.

Appendix A: complete list of products available to participants

- MM5 (12UTC run)
    - SLP, 10m winds, topography or 925mb temp, 36km domain (72hrs, 25 frames)
    - SLP, 10m winds, topography or 925mb temp, 12km domain (72hrs, 25 frames)
    - SLP, 10m winds, topography or 925mb temp, 4km domain (48hrs, 17 frames)
    - 850mb heights, temperature, winds, 12km domain (72hrs, 25 frames)
    - Subdomain SLP, 10m winds, 925mb temp, 4km domain (48hrs, 17 frames)
    - Surface wind speed, 4km domain (48hrs, 17 frames)
    - 500mb heights, temperatures, winds, 12km domain (72hrs, 25 frames)
    - 3 hour precipitation, 12km domain (72hrs, 23 frames)
    - Meteograms
        - NWS Seattle
        - Port Angeles
        - Quillayute
        - Victoria, BC
- Buoys
    - Smith Island
    - Destruction Island
    - Tatoosh Island
    - West Point
- Satellite Imagery
    - Enhanced 4km
    - Infrared 4km
    - Infrared Enhanced 4km
    - Visible 4km

- TAFs and current METARs
    - Whidbey Island
    - McCord
    - Sea-Tac
    - Portland
    - Hoquiam
    - Bremerton
    - Everett
    - Bellingham
    - Port Angeles
    - Fairchild
    - Moses Lake
    - Pasco
    - Friday Harbor
    - Victoria
- Radar
    - Base Reflectivity Elevation 1
    - Base Radial Velocity Elevation 1
- AVN (00UTC run)
    - 850mb winds, heights, temperatures
- NGM (00UTC run)
    - 850mb winds, heights, temperatures

PROBABILITY STIMULUS

- MM5 Ensemble Probability of winds greater than 21kts
    - ACME 36km (48hrs, 8 frames)
    - ACME 12km (48hrs, 16 frames)
    - ACMEcore 36km (48hrs, 8 frames)

- ACMEcore 12km (48hrs, 16 frames)

References

Anderson, O., 1996: A method for producing and evaluating probabilistic forecasts from ensemble model integrations. *J. Climate*, **9**, 1518-1530.

Baars, J. A., C. Mass, and M. Albright, 2004: Performance of national weather service forecasts versus model output statistics. *Proc. 20th AMS Conf. On Weather Analysis and Forecasting/ 16th Conf. On Numerical Weather Prediction, Seattle, WA,* American Meteorological Society, np.

Gonzalez, R., and G. Wu, 1999: On the form of the probability weighting function. *Cognitive Psychology*, **38,** 129-166.

Green, D. M., and J. Z. Swets, 1966: *Signal detection theory and psychophysics*. Wiley, 455 pp.

Grimit, E.P., and C. F. Mass, 2003: Toward short-range ensemble prediction of mesoscale forecast skill. *Proc. 10th AMS Conf. on Mesoscale Processes*, Portland, OR, American Meteorological Society, 22-27.

Grimit, E. P., and C. F. Mass, 2002: Initial results of a mesoscale short-range ensemble forecasting system over the Pacific Northwest. *Wea. Forecasting*, **17,** 192-205.

Kahneman, D., and A. Tversky, 1979: Prospect theory: An analysis of decision under risk. *Econometrica,* **47,** 2 263-291.

Keith, R., 2003: Optimization of value of aerodrome forecasts. *Wea. Forecasting*, **18,** 5 808-824.

Murphy, A. H., and R. L. Winkler, 1974a: Credible interval temperature forecasting: some experimental results. *Mon. Wea. Rev.,* **162,** 11 784-794.

Murphy, A. H., and R. L. Winkler, 1974b: Subjective probability forecasting experiments in meteorology: some preliminary results. *Bull. Amer. Meteor. Soc.,* **55,** 10 1206-1216.

Murphy, A. H., and R. L. Winkler, 1977: Can weather forecasters formulate reliable probability forecasts of precipitation and temperature?. *Natl. Wea. Dig.* **2,** 2 2-9.

Palmer, T. N., 2002: The economic value of ensemble forecasts as a tool for risk assessment: From days to decades. *Quart. J. Roy. Meteor. Soc.*, **128,** 747-774.

Richardson, D. S., 2002: Measures of skill and value of ensemble prediction systems, their interrelationship and the effect of ensemble size. *Quart. J. Roy. Meteor. Soc.*, **127,** 2473-2489.

Roulston, M. S., and L. A. Smith, 2004: The boy who cried wolf revisited: The impact of false alarm intolerance on cost-loss scenarios. *Wea. Forecasting*, **19,** 2 391-397.

Roulston, M.S., G.E. Bolton, A. N. Kleit, and A. L. Sears-Collins 2005: A laboratory study of the benefits of including uncertainty information in weather forecasts. *Wea. Forecasting*, in press.

Windschitl, P. D., and E. Weber, 1999: The interpretation of "likely" depends on the context, but "70%" is 70%----right? The influence of associative processes on perceived certainty. *Journal of Experimental Psychology: Learning, Memory & Cognition,* **25,** 6 1514-1533.

Weber, E. U., 1994: From subjective probabilities to decision weights: The effect of asymmetric loss function. *Psychological Bulletin*, **115,** 2 228-242.

Figure 1: ACME Ensemble Winds Greater than 20 Knots

Figure 2: Percent of advisories posted and percent of observed wind speeds over probability range given by the model.

Table 1. Hits, False Alarms (FA), Misses and Correct Rejections (CR) for each forecaster. These were ascertained by comparing the hours for which the forecaster posted an advisory to the hours during which observed winds exceeded 20 knots. Slight differences in the total number of hours occur because of a few missing verification data.

Forecaster	1	2	3	4	5	6	7	8	9	10	Total

WITH

	1	2	4	5	6	7	8	9	10	11	TOTAL
Correct	31	27	26	31	23	21	27	22	14	46	268
False	6	6	7	6	9	12	10	10	6	4	76
Hit	13	8	9	16	6	9	15	18	13	0	107
Miss	8	4	3	3	7	3	5	1	19	1	54

WITHOUT

	1	2	4	5	6	7	8	9	10	11	TOTAL
Correct	29	35	20	22	19	14	8	17	31	10	205
False	4	2	17	8	18	23	25	20	18	11	146
Hit	5	8	18	13	19	21	8	11	1	18	122
Miss	7	13	3	2	2	0	4	4	1	13	49

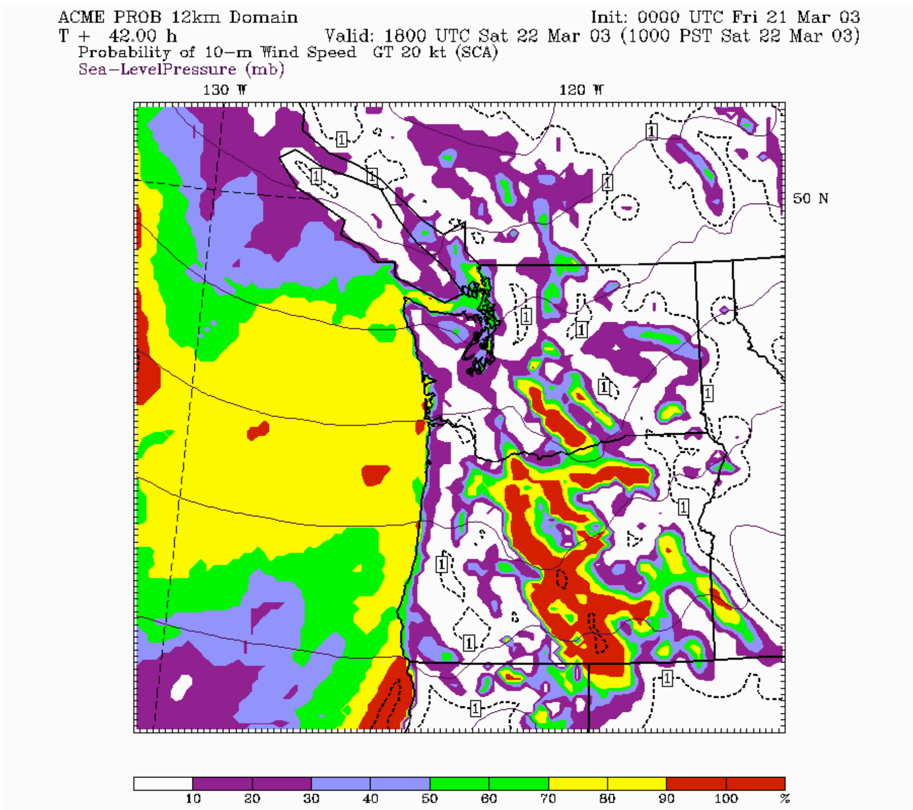

Figure 1: ACME Ensemble Winds Greater than 20 Knots is color coded for probability with warmer colors indicating higher probability of winds greater than 20 knots.

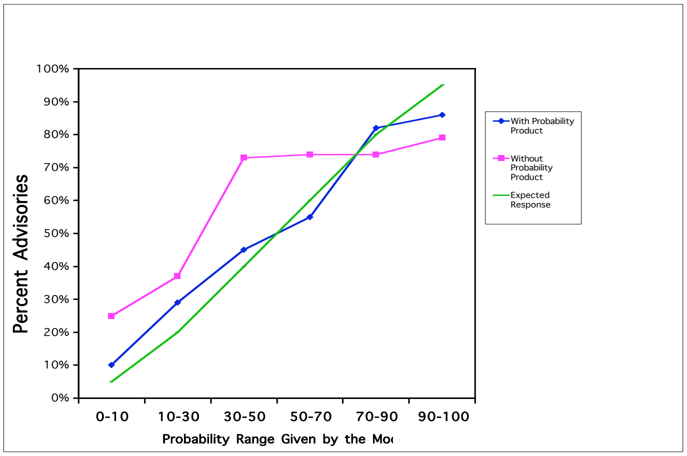

Figure 2. Percent of advisories posted, with and without the probability product, for each probability range given by the model.