Rolling Forecasts

Don A. Moore and Thomas S. Wallsten

Draft of November 6, 2014¹

Rolling forecasts will provide continuously updated probability forecasts of when an event will occur. Such forecasts could be of considerable interest to the intelligence community because they will allow policy or decision makers (DMs) to infer probabilities associated with specific timelines or intervals at any time of their choosing. For example, probability forecasts regarding event occurrences over a given 10 day period might be useful in making decisions regarding deployment of military forces or intelligence assets.

IFP 1112 asked, “When will Bashar al-Assad resign or otherwise vacate the office of President of Syria?” Forecasters estimated the probability that he would leave by 1 April 2013. But the question remains interesting despite the IFP having closed. And more to the point, DMs might have an ongoing interest in the likelihood of that event within specific 30, 60 or 90 day intervals.

Of the various options one might consider for obtaining rolling forecasts, we present one here that is relatively easy to implement and that requires only moderate additional effort on the part of the forecasters. It involves structuring the IFP with a fixed number of time bins, say 5 bins of one month each (or of whatever interval the question generators deem appropriate), adding a new bin at the distant end of the series when the most proximal bin expires and modeling the consequent forecasts based on the procedures we have been developing (Tidwell, Wallsten, Yang, & Moore, 2014).

Forecasters will be asked to update their forecasts every time a bin expires and a new one is added. We will treat these updated forecasts as new distributions for all computational purposes, including scoring and aggregation.² Of course, forecasters can update their judgments any other time they wish to as well. To be clear, when forecasters update their estimates in the absence of a bin expiring, this does not create a new distribution; rather, it simply modifies the probabilities assigned to the current set of bins. New distributions are created only as old bins expire and new ones are formed.

The following modification of the al-Assad question illustrates the approach and provides a concrete example for discussing the modeling. Assume a release date of 20 August 2014.

What is the probability that al-Assad will vacate office...

- a. By 30 Nov 2014
- b. 1 Dec 2014 – 31 Jan 2015
- c. 1 Feb – 31 Mar 2015
- d. 1 April – 31 May 2015

¹ Don Moore and Tom Wallsten drafted this memo, with input from Barbara Mellers, Welton Chang, and Joe Tidwell. The Daily Updating and Aggregation summarizes work done primarily by Mr. Tidwell.

² If they fail to update, we will use a model to estimate their new probability distribution, but forecaster updates will be preferable.

e. Not before 1 June 2015

When the first (“By 30 November”) bin expires, we automatically convert the final open bin (“Not before 1 June 2015”) to a closed bin, (“1 June – 31 July 2015”) and create a new final open bin (“Not before 1 Aug 2015”). The process repeats as early bins expire and new ones are required. Note that this example utilizes 2-month bins, but the question generators may prefer longer or shorter intervals. Bin lengths should be based on (1) what the question formulator considers is the easiest or most natural time bin to think about for the issue at hand, which could be days, weeks or months, (2) the preferred granularity of the forecasts, (3) the volatility of the underlying question, and (4) the natural calendar of events that may increase or decrease the probability of occurrence. Bins need not all be of the same size.

Scoring Rolling Questions

Each forecast includes estimated probabilities over a mutually exclusive and collectively exhaustive set of outcome bins. We refer to this set of probability estimates as a distribution. Overall, each forecaster will have a number of distributions for each rolling IFP equal to $k+1$, where k is the number of bins that expired (or, equivalently, were created) during the IFP’s life. Thus, if an IFP closes before the first bin expires, the number of scorable distributions is one. If five bins have expired before the IFP closes, the number of scorable distributions is six.

We propose scoring each bin of each distribution, as it resolves, using the Brier scoring rule for binary events. In this manner, every bin, including the final “Not before” bin of each distribution is scored until an IFP is closed, either because the event occurred or the IFP was terminated for some other reason. The daily score for each distribution, then, is the average of its separate daily bin scores. The score for each distribution is its average over all the days it is open. The overall IFP score is the sum of the individual distribution scores.

The longer a rolling IFP is open, the higher the forecaster’s score. Because the accumulation of error scores could be demoralizing, we propose using “Brier points”: 2 minus the Brier score, such that higher scores reflect greater accuracy.

To illustrate, consider the question described above: What is the probability that al-Assad will vacate office...

- a. By 30 Nov 2014
- b. 1 Dec 2014 – 31 Jan 2015
- c. 1 Feb – 31 Mar 2015
- d. 1 April – 31 May 2015
- e. Not before 1 June 2015

When the 30 November bin expires without the event having occurred, we score each day’s forecast up until 30 Nov as if it were its own binary IFP (“Will Assad vacate office by 30 Nov 2014?” – Yes, No). When the bin “1 Dec 2014 – 31 Jan 2015” expires with Assad still in office, we can score it similarly. We score using Brier points, such that higher scores reflect more accurate forecasts. On 1 June 2015, all five of

the bins in that distribution will have resolved, and each bin's Brier score can be averaged within each day. The score for the IFP then averages bins across each day's distribution and sums the over days.

Daily Updating and Aggregation

As forecasts to the newly released rolling IFP come in, we model them in the manner we have been doing for our continuous elicitations, which involves fitting a continuous probability distribution to each forecaster's judgments. The gamma distribution has proven useful when modeling forecasts of event dates. Forecasters will not update their estimates on a daily basis, but we will update their modeled values daily. Survival analysis provides one method for doing this, but we have found a modification of that method to work very well and be preferable in this context. Namely, as days expire, for each forecaster, we use the initial fitted distribution for the forecast to assign probability estimates for a large number of bins, with the first bin always lower-bounded by the current day. We then rescale these bins to sum to one and then estimate the continuous model (gamma for date judgments) for this large number of bins. The advantage of this method over straight survival analysis is that it maintains the functional form we are using to represent the forecasts.

We can then aggregate on a daily basis across individuals by one of two methods (Tidwell, Wallsten, Yang, & Moore, 2014). One method is to define the aggregated gamma distribution by taking the median parameters of the individual fitted distributions. The other method is to take the median fitted cumulative distribution at each of numerous points across the domain. Thus far in our experience, the results of the two methods are close to indistinguishable because they basically accomplish the same thing. We recognize the importance of the specific distribution in estimating tail probabilities and expect to continue working on this issue.

The result of the aggregation is a single probability distribution. Among other things, this allows a DM to query the aggregate distribution for probabilities associated with any time interval she or he is interested in. Going back to al-Assad example, a DM might want to know the forecast probability that Assad will vacate office by 25 December 2014 or by 1 March 2015. Note that these dates fall within intervals, but are no problem for our methodology.

We close by emphasizing that even if forecasters update their forecasts only occasionally, the aggregated distribution provided for DMs is updated on a daily basis.

References

- Haran, U., Moore, D. A., & Morewedge, C. K. (2010). A simple remedy for overprecision in judgment. *Judgment and Decision Making*, 5(7), 467–476.
- Matheson, J. E. & Winkler, R. L. (1976). Scoring rules for continuous probability distributions. *Management Science*, 22, 1087-1096.
- Tidwell, J. W., Wallsten, T. S., Yang, H. & Moore, D. (2014). Continuous Elicitation. Chapter 9.A. in *Exploring the Optimal Forecasting Frontier: How Much Room Is There to Improve Subjective Forecasting Accuracy? Annual Report of the Good Judgment Project, Option Year 2, June 3, 2013 – June 2, 2014*.
- Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, 101(4), 547–567.