

EVALUATING PROBABILITY FORECASTS

BY TZE LEUNG LAI¹, SHULAMITH T. GROSS² AND DAVID BO SHEN

Stanford University, Baruch College/CUNY and UBS

Probability forecasts of events are routinely used in climate predictions, in forecasting default probabilities on bank loans or in estimating the probability of a patient's positive response to treatment. Scoring rules have long been used to assess the efficacy of the forecast probabilities after observing the occurrence, or nonoccurrence, of the predicted events. We develop herein a statistical theory for scoring rules and propose an alternative approach to the evaluation of probability forecasts. This approach uses loss functions relating the predicted to the actual probabilities of the events and applies martingale theory to exploit the temporal structure between the forecast and the subsequent occurrence or nonoccurrence of the event.

1. Introduction. Probability forecasts of future events are widely used in diverse fields of application. Oncologists routinely predict the probability of a cancer patient's progression-free survival beyond a certain time horizon [Hari et al. (2009)]. Economists give the probability forecasts of an economic rebound or a recession by the end of a fiscal year. Banks are required by regulators assessing their capital requirements to predict periodically the risk of default of the loans they make. Engineers are routinely called upon to predict the survival probability of a system or infrastructure beyond five or ten years; this includes bridges, sewer systems and other structures. Finally, lawyers also assess the probability of particular trial outcome [Fox and Birke (2002)] in order to determine whether to go to trial or settle out of court. This list would not be complete without mentioning the field that is most advanced in its daily probability predictions, namely meteorology. In the past 60 years, remarkable advances in forecasting precipitation probabilities, temperatures, and rainfall amounts have been made in terms of breadth and accuracy. Murphy and Winkler (1984) provide an illuminating history of the US National Weather Service's transition from nonprobabilistic to probability predictions and its development of reliability and accuracy measures for these probability forecasts. Accuracy assessment is difficult to carry out directly because it requires comparing a forecaster's predicted probabilities with the actual but unknown probabilities of the events under study. Reliability is measured using "scoring rules,"

Received February 2011; revised May 2011.

¹Supported in part by NSF Grant DMS-08-05879.

²Supported in part by PSC-CUNY 2008 and 2009 grants and a 2008 Summer Research Support grant from Baruch Zicklin School of Business.

MSC2010 subject classifications. Primary 60G42, 62P99; secondary 62P05.

Key words and phrases. Forecasting, loss functions, martingales, scoring rules.

which are empirical distance measures between repeated predicted probabilities of an event, such as having no rain the next day, and indicator variables that take on the value 1 if the predicted event actually occurs, and 0 otherwise; see Gneiting and Raftery (2007), Gneiting, Balabdaoui and Raftery (2007) and Ranjan and Gneiting (2010) for recent reviews and developments.

To be more specific, a scoring rule for a sequence of n probability forecasts $\hat{p}_i, i = 1, \dots, n$, is the average score $n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$, where $Y_i = 1$ or 0 according to whether the i th event A_i actually occurs or not. An example is the widely used Brier’s score $L(y, \hat{p}) = (y - \hat{p})^2$ [Brier (1950)]. Noting that the Y_i are related to the actual but unknown probability p_i via $Y_i \sim \text{Bernoulli}(p_i)$, Cox (1958) proposed to evaluate how well the \hat{p}_i predict p_i by using the estimates of (β_1, β_2) in the regression model

$$(1.1) \quad \text{logit}(p_i) = \beta_1 + \beta_2 \text{logit}(\hat{p}_i)$$

and developed a test of the null hypothesis $(\beta_1, \beta_2) = (0, 1)$, which corresponds to perfect prediction. Spiegelhalter (1986) subsequently proposed a test of the null hypothesis $H_0: \hat{p}_i = p_i$ for all $i = 1, \dots, n$, based on a standardized form (under H_0) of Brier’s score. A serious limitation of this approach is the unrealistic benchmark of perfect prediction to formulate the null hypothesis, so significant departures from it are expected when n is large, and they convey little information on how well the \hat{p}_i predict p_i . Another limitation is the implicit assumption that the \hat{p}_i are independent random variables, which clearly is violated since \hat{p}_i usually involves previous observations and predictions.

Seillier-Moiseiwitsch and Dawid (1993) have developed a hypothesis testing approach that removes both limitations in testing the validity of a sequence of probability forecasts. The forecaster is modeled by a probability measure under which the conditional probability of the occurrence of A_i given the σ -field \mathcal{G}_{i-1} generated by the forecaster’s information set prior to the occurrence of the event is π_i . In this model, the forecaster uses $\hat{p}_i = \pi_i$ as the predicted probability of A_i . As pointed out earlier by Dawid (1982), this model fits neatly into de Finetti’s (1975) framework in which “the coherent subjectivist Bayesian can be shown to have a joint probability distribution over all conceivably observable quantities,” which is represented by the probability measure Π in the present case. To test if Π is “empirically valid” based on the observed outcomes Y_1, \dots, Y_n , Seillier-Moiseiwitsch and Dawid (1993) consider the null hypothesis H_0 that “the sequence of events is generated by the same joint distribution from which the forecasts are constructed.” Under this null hypothesis, $\sum_{i=1}^n \xi_i (Y_i - \pi_i), n \geq 1$, is a martingale with respect to the filtration $\{\mathcal{G}_i\}$ when ξ_i is \mathcal{G}_{i-1} -measurable for all i . Assuming certain regularity conditions on ξ_i , they apply the martingale central limit theorem to show that as $n \rightarrow \infty$,

$$(1.2) \quad \left\{ \sum_{i=1}^n \xi_i (Y_i - \pi_i) \right\} / \left\{ \sum_{i=1}^n \xi_i^2 \pi_i (1 - \pi_i) \right\}^{1/2} \implies N(0, 1)$$

under H_0 , where \implies denotes convergence in distribution. Since $\pi_i = \hat{p}_i$ in this model of a coherent forecaster, Seillier-Moiseiwitsch and Dawid (1993) have made use of (1.2) to construct various tests of H_0 . One such test, described at the end of their Section 6, involves another probability forecast \hat{p}'_i , which is “based on no more information” to define ξ_i , so that a significantly large value of the test statistic can be used to reject H_0 in favor of the alternative forecasting model or method.

Hypothesis testing has been extended from testing perfect prediction or empirical validity of a sequence of probability forecasts to testing equality of the predictive performance of two forecasts; see Redelmeier, Bloch and Hickam (1991) who extended Spiegelhalter’s approach mentioned above. Testing the equality of predictive performance, measured by some loss function of the predictors and the realized values, of two forecasting models or methods has attracted much recent interest in the econometrics literature, which is reviewed in Section 6.2. In this paper we develop a new approach to statistical inference, which involves confidence intervals rather than statistical tests of a null hypothesis asserting empirical validity of a forecasting model or method, or equal predictive performance for two forecasting models or methods. The essence of our approach is to evaluate probability forecasts via the average loss $L_n = n^{-1} \sum_{i=1}^n L(p_i, \hat{p}_i)$, where p_i is the actual but unknown probability of the occurrence of A_i . When L is linear in p_i , $L(Y_i, \hat{p}_i)$ is an unbiased estimate of $L(p_i, \hat{p}_i)$ since $E(Y_i | \hat{p}_i) = p_i$. We show in Section 2, where an overview of loss functions and scoring rules is also given, that even for L that is nonlinear in p_i there is a “linear equivalent” which carries the same information as L for comparing different forecasts. In Section 3 we make use of this insight to construct inferential procedures, such as confidence intervals, for the average loss L_n under certain assumptions and for comparing the average losses of different forecasts.

Note that we have used E to denote expectation with respect to the actual probability measure P , under which A_i occurs with probability p_i given the previous history represented by the σ -field \mathcal{G}_{i-1} , and that we have used Π to denote the probability measure assumed by a coherent Bayesian forecaster whose probability of occurrence of A_i given \mathcal{G}_{i-1} is π_i . Because $\pi_i = \hat{p}_i$ for a coherent Bayesian forecaster, Seillier-Moiseiwitsch and Dawid (1993) are able to use (1.2) to test the null hypothesis of empirical validity of Π in the sense that $E_{\Pi}(Y_i | \mathcal{G}_{i-1}) = \hat{p}_i$, where E_{Π} denotes expectation with respect to the measure Π . Replacing Π by P is much more ambitious, but it appears impossible to derive the studentized version of the obvious estimate $\hat{L}_n = n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$ and its sampling distribution under P to perform inference on L_n . We address this difficulty in several steps in Section 3. First we consider in Section 3.1 the case in which $L(p, \hat{p})$ is linear in p and make use of the martingale central limit theorem to prove an analog of (1.2) with p_i in place of π_i and $\xi_i = L(1, \hat{p}_i) - L(0, \hat{p}_i)$. Whereas $\pi_i = \hat{p}_i$ under Π , the p_i associated with P are unknown parameters that need to be estimated. Postponing their estimation to Section 3.4, we first use the simple bound $p_i(1 - p_i) \leq 1/4$ to obtain confidence intervals for L_n by making use of this analog of (1.2). In

Section 3.2 we consider the problem of comparing two probability forecasts via the difference of their average losses, and make use of the idea of linear equivalents introduced in Section 2 to remove the assumption of $L(p, \hat{p})$ being linear in p when we consider $\Delta_n = n^{-1} \sum_{i=1}^n \{L(p_i, \hat{p}'_i) - L(p_i, \hat{p}''_i)\}$. A variant of Δ_n , called Winkler's skill score in weather forecasting, is considered in Section 3.3. In Section 3.4, we return to the problem of estimating $p_i(1 - p_i)$. Motivated by applications in which the forecasts are grouped into "risk buckets" within which the p_i can be regarded as equal, Section 3.4 provides two main results on this problem. The first is Theorem 3, which gives consistent estimates of the asymptotic variance of $\hat{\Delta}_n$, or of \hat{L}_n when $L(p, \hat{p})$ is linear in p , in the presence of risk buckets with each bucket of size 2 or more. The second, given in Theorem 4, shows that in this bucket model it is possible to adjust the Brier score to obtain a consistent and asymptotically normal estimate of the average squared error loss $L_n = n^{-1} \sum_{i=1}^n (p_i - \hat{p}_i)^2$. Theorem 4 also provides a consistent estimate of the asymptotic variance of the adjusted Brier score when the bucket size is at least 3. In Section 3.5 we develop an analog of Theorem 3 for the more general setting of "quasi-buckets," for which the p_i within each bin (quasi-bucket) need not be equal. These quasi-buckets arise in "reliability diagrams" in the meteorology literature. Theorem 5 shows that the confidence intervals obtained under an assumed bucket model are still valid but tend to be conservative if the buckets are actually quasi-buckets. The proofs of Theorems 4 and 5 are given in Section 5.

Section 4 gives a simulation study of the performance of the proposed methodology, and some concluding remarks and discussion are given in Section 7. In Section 6 we extend the Y_i from the case of indicator variables of events to more general random variables by modifying the arguments in Section 5, and also show how the methods and results in Sections 3.2 and 3.4 can be used to address related problems in the econometrics literature on the expected difference in scores between two forecasts, after a brief review of that literature that has become a major strand of research in economic forecasts.

2. Scoring rules and associated loss functions. Instead of defining a scoring rule via L (which associates better forecasts with smaller values of L), Gneiting and Raftery (2007) and others assign higher scores to better forecasts; this is tantamount to using $-L$ instead of L in defining a scoring rule. More generally, considering p and its forecast \hat{p} as probability measures, they call a scoring rule S *proper* relative to a class \mathcal{P} of probability measures if $E_p S(Z, p) \geq E_p S(Z, \hat{p})$ for all p and \hat{p} belonging to \mathcal{P} , where Z is an observed random vector (generated from p) on which scoring is based. For the development in the subsequent sections, we find it more convenient to work with L instead of $-L$ and restrict to $Z = (Y_1, \dots, Y_n)$ so that $S(Z, (\hat{p}_1, \dots, \hat{p}_n)) = -n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$.

The function L in the scoring rule $n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$ measures the closeness of the probability forecast \hat{p}_i of event i before the indicator variable Y_i of the event is observed. We can also use L as a loss function in measuring the accuracy

of \hat{p}_i as an estimate of the probability p_i of event i . Besides the squared error loss $L(p, \hat{p}) = (p - \hat{p})^2$ used in Brier's score, another widely used loss function is the Kullback–Leibler divergence,

$$(2.1) \quad L(p, \hat{p}) = p \log(p/\hat{p}) + (1 - p) \log[(1 - p)/(1 - \hat{p})],$$

which is closely related to the log score introduced by Good (1952), as shown below. More general loss functions of this type are the Bregman divergences; see Section 3.5.4 of Grünwald and Dawid (2004) and Section 2.2 of Gneiting and Raftery (2007).

We call a loss function $\tilde{L}(p, \hat{p})$ a *linear equivalent* of the loss function $L(p, \hat{p})$ if $\tilde{L}(p, \hat{p})$ is a linear function of p and

$$(2.2) \quad L(p, \hat{p}) - \tilde{L}(p, \hat{p}) \quad \text{does not depend on } \hat{p}.$$

For example, $\tilde{L}(p, \hat{p}) = -2p\hat{p} + \hat{p}^2$ is a linear equivalent of the squared error loss $(p - \hat{p})^2$ used by Brier's score. A linear equivalent \tilde{L} of the Kullback–Leibler divergence (2.1) is given by $-\tilde{L}(p, \hat{p}) = p \log(\hat{p}) + (1 - p) \log(1 - \hat{p})$. This is the conditional expected value (given \hat{p}) of $Y \log(\hat{p}) + (1 - Y) \log(1 - \hat{p})$, which is Good's log score. Since the probability \hat{p}_i is determined before the Bernoulli random variable Y_i is observed,

$$(2.3) \quad E\{L(Y_i, \hat{p}_i) | \hat{p}_i, p_i\} = p_i L(1, \hat{p}_i) + (1 - p_i) L(0, \hat{p}_i).$$

Therefore the conditional expected loss of a scoring rule $L(Y, \hat{p})$ yields a loss function

$$(2.4) \quad \tilde{L}(p, \hat{p}) = \{L(1, \hat{p}) - L(0, \hat{p})\}p + L(0, \hat{p})$$

that is linear in p . For example, the absolute value scoring rule $L(Y, \hat{p}) = |Y - \hat{p}|$ is associated with $\tilde{L}(p, \hat{p}) = p(1 - \hat{p}) + (1 - p)\hat{p}$ that is linear in each argument. Using the notation (2.4), the scoring rule $L(Y, \hat{p})$ is proper if $\tilde{L}(p, p) \leq \tilde{L}(p, \hat{p})$ for all $p, \hat{p} \in [0, 1]$, and is strictly proper if $\min_{0 \leq \hat{p} \leq 1} \tilde{L}(p, \hat{p})$ is uniquely attained at $p = \hat{p}$. The scoring rule $|Y - \hat{p}|$, therefore, is not proper; moreover, $|p - \hat{p}|$ does not have a linear equivalent.

3. A new approach to evaluation of probability forecasts. In this section we first consider the evaluation of a sequence of probability forecasts $\hat{p}_1, \dots, \hat{p}_n$ based on the corresponding sequence of indicator variables Y_1, \dots, Y_n that denote whether the events actually occur or not. Whereas the traditional approach to evaluating $\hat{\mathbf{p}} = (\hat{p}_1, \dots, \hat{p}_n)$ uses the scoring rule $n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$, we propose to evaluate $\hat{\mathbf{p}}$ via

$$(3.1) \quad L_n = n^{-1} \sum_{i=1}^n L(p_i, \hat{p}_i),$$

where L is a loss function, and p_i is the actual probability of the occurrence of the i th event. Allowing the actual probabilities p_i to be generated by a stochastic

system and the forecast \hat{p}_k to depend on an information set \mathcal{G}_{k-1} that consists of the event and forecast histories and other covariates before Y_k is observed, the conditional distribution of Y_i given \mathcal{G}_{i-1} and p_i is Bernoulli(p_i), and therefore

$$(3.2) \quad P(Y_i = 1|\mathcal{G}_{i-1}, p_i) = p_i.$$

3.1. *Linear case.* In view of (3.2), an obvious estimate of the unknown p_i is Y_i . Suppose $L(p, \hat{p})$ is linear in p , as in the case of linear equivalents of general loss functions. Combining this linearity property with (3.2) yields

$$(3.3) \quad E\{L(Y_i, \hat{p}_i)|\mathcal{G}_{i-1}, p_i\} = L(p_i, \hat{p}_i),$$

and therefore $L(Y_i, \hat{p}_i) - L(p_i, \hat{p}_i)$ is a martingale difference sequence with respect to $\{\mathcal{F}_i\}$, where \mathcal{F}_{i-1} is the σ -field generated by \mathcal{G}_{i-1} and p_1, \dots, p_i . Let $d_i = L(Y_i, \hat{p}_i) - L(p_i, \hat{p}_i)$. Since $L(y, \hat{p})$ is linear in y , we can write $L(y, \hat{p}) = a(\hat{p})y + b(\hat{p})$. Setting $y = 0$ and $y = 1$ in this equation yields $a(\hat{p}) = L(1, \hat{p}) - L(0, \hat{p})$. Moreover, $d_i = a(\hat{p}_i)(Y_i - p_i)$. Since $Y_i|\mathcal{F}_i \sim \text{Bernoulli}(p_i)$ and \hat{p}_i is \mathcal{F}_{i-1} -measurable,

$$(3.4) \quad E(d_i^2|\mathcal{F}_{i-1}) = a^2(\hat{p}_i)p_i(1 - p_i).$$

By (3.4), $\sum_1^n E(d_i^2|\mathcal{F}_{i-1}) = \sum_1^n \{L(1, \hat{p}_i) - L(0, \hat{p}_i)\}^2 p_i(1 - p_i) = O(n)$ a.s., and therefore $n^{-1} \sum_{i=1}^n d_i \rightarrow 0$ a.s. by the martingale strong law [Williams (1991), Section 12.14] proving $\hat{L}_n - L_n \rightarrow 0$ a.s. Moreover, if $n^{-1} \sum_1^n E(d_i^2|\mathcal{F}_{i-1}) = \sigma_n^2$ converges in probability to a nonrandom positive constant, then $\sqrt{n}(\hat{L}_n - L_n)/\sigma_n$ has a limiting standard normal distribution by Theorem 1 of Seillier-Moiseiwitsch and Dawid (1993). Summarizing, we have the following.

THEOREM 1. *Suppose $L(p, \hat{p})$ is linear in p . Let $\hat{L}_n = n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$, and define L_n by (3.1). Letting*

$$(3.5) \quad \sigma_n^2 = n^{-1} \sum_{i=1}^n \{L(1, \hat{p}_i) - L(0, \hat{p}_i)\}^2 p_i(1 - p_i),$$

assume that $\sigma_n^2 = O(1)$ with probability 1. Then $\hat{L}_n - L_n$ converges to 0 with probability 1. If σ_n^2 converges in probability to some nonrandom positive constant, then $\sqrt{n}(\hat{L}_n - L_n)/\sigma_n$ has a limiting standard normal distribution.

To apply Theorem 1 to statistical inference on L_n , one needs to address the issue that σ_n^2 involves the unknown p_i . As noted in the third paragraph of Section 1, Seillier-Moiseiwitsch and Dawid (1993) have addressed this issue by using $p_i = E_\Pi(Y_i|\mathcal{G}_{i-1})$ under the null hypothesis H_0 that assumes the sequence of events are generated by the probability measure Π . This approach is related to the earlier work of Dawid (1982), who assumes a “subjective probability distribution” Π for the events so that Bayesian forecasts are given by $\hat{p}_i = \pi_i = E_\Pi(Y_i|\mathcal{G}_{i-1})$. Letting

$\xi_t = 1$ or 0 according to whether time t is included in the “test set” to evaluate forecasts, he calls the test set “admissible” if ξ_t depends only on \mathcal{G}_{t-1} , and uses martingale theory to show that

$$(3.6) \quad \left(\sum_{i=1}^n \xi_i Y_i - \sum_{i=1}^n \xi_i \hat{p}_i \right) / \sum_{i=1}^n \xi_i \rightarrow 0 \quad \text{a.s. } [\Pi] \text{ on } \left\{ \sum_{i=1}^n \xi_i = \infty \right\}.$$

From (3.6), it follows that for any $0 < x < 1$, the long-run average of Y_i (under the subjective probability measure) associated with $\hat{p}_i = x$ (i.e., $\xi_i = I_{\{\hat{p}_i=x\}}$) is equal to x provided that $\sum_{i=1}^n I_{\{\hat{p}_i=x\}} \rightarrow \infty$. Note that Dawid’s well-calibration theorem (3.6) involves the subjective probability measure Π . DeGroot and Fienberg (1983) have noted that well-calibrated forecasts need not reflect the forecaster’s “honest subjective probabilities,” that is, need not satisfy Dawid’s coherence criterion $\hat{p}_i = \pi_i$. They therefore use a criterion called “refinement” to compare well-calibrated forecasts.

In this paper we apply Theorem 1 to construct confidence intervals for L_n , under the actual probability measure P that generates the unknown p_i in (3.5). Whereas substituting p_i by Y_i in $L(p_i, \hat{p}_i)$ leads to a consistent estimate of L_n when L is linear, such substitution gives 0 as an overly optimistic estimate of $p_i(1 - p_i) = \text{Var}(Y_i | \mathcal{F}_{i-1})$. A conservative confidence interval for L_n can be obtained by replacing $p_i(1 - p_i)$ in (3.5) by its upper bound 1/4. In Section 3.4, we consider estimation of σ_n^2 and of $n^{-1} \sum_{i=1}^n L(p_i, \hat{p}_i)$ when L is nonlinear in p_i , under additional assumptions on how the p_i are generated.

3.2. *Application to comparison of probability forecasts.* Consider two sequences of probability forecasts $\hat{\mathbf{p}}' = (\hat{p}'_1, \dots, \hat{p}'_n)$ and $\hat{\mathbf{p}}'' = (\hat{p}''_1, \dots, \hat{p}''_n)$ of $\mathbf{p} = (p_1, \dots, p_n)$. Suppose a loss function $L(p, q)$ is used to evaluate each forecast, and let $\tilde{L}(p, q)$ be its linear equivalent. Since $L(p, q) - \tilde{L}(p, q)$ does not depend on q in view of (2.2), it is a function only of p , which we denote by $d(p)$. Hence

$$\begin{aligned} L(p_i, \hat{p}'_i) - L(p_i, \hat{p}''_i) &= \{\tilde{L}(p_i, \hat{p}'_i) + d(p_i)\} - \{\tilde{L}(p_i, \hat{p}''_i) + d(p_i)\} \\ &= \tilde{L}(p_i, \hat{p}'_i) - \tilde{L}(p_i, \hat{p}''_i) \end{aligned}$$

is a linear function of p_i , and therefore we can estimate $\Delta_n = n^{-1} \sum_{i=1}^n \{L(p_i, \hat{p}'_i) - L(p_i, \hat{p}''_i)\}$ by the difference $n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}'_i) - n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}''_i)$ of scores of the two forecasts. Application of Theorem 1 then yields the following theorem, whose part (ii) is related to (2.4).

THEOREM 2. *Let $\hat{\Delta}_n = n^{-1} \sum_{i=1}^n \{L(Y_i, \hat{p}'_i) - L(Y_i, \hat{p}''_i)\}$ and*

$$(3.7) \quad \begin{aligned} \delta_i &= \{L(1, \hat{p}'_i) - L(0, \hat{p}'_i)\} - \{L(1, \hat{p}''_i) - L(0, \hat{p}''_i)\}, \\ s_n^2 &= n^{-1} \sum_{i=1}^n \delta_i^2 p_i(1 - p_i). \end{aligned}$$

(i) Suppose L has a linear equivalent. Letting $\Delta_n = n^{-1} \sum_{i=1}^n \{L(p_i, \hat{p}'_i) - L(p_i, \hat{p}''_i)\}$, assume that $s_n^2 = O(1)$ with probability 1. Then $\hat{\Delta}_n - \Delta_n$ converges to 0 with probability 1. If furthermore s_n converges in probability to some non-random positive constant, then $\sqrt{n}(\hat{\Delta}_n - \Delta_n)/s_n$ has a limiting standard normal distribution.

(ii) Without assuming that L has a linear equivalent, the same conclusion as in (i) still holds with $\Delta_n = n^{-1} \sum_{i=1}^n \{\delta_i p_i + L(0, \hat{p}'_i) - L(0, \hat{p}''_i)\}$.

3.3. *Illustrative applications and skill scores.* As an illustration of Theorem 2, we compare the Brier scores B_k for the k -day ahead forecasts $\hat{p}_t^{(k)}$, $1 \leq k \leq 7$, for Queens, NY, provided by US National Weather Service from June 8, 2007, to March 31, 2009. Table 1 gives the values of B_1 and $B_k - B_{k-1}$ for $2 \leq k \leq 6$. Using $1/4$ to replace $p_i(1 - p_i)$ in (3.7), we can use Theorem 2(i) to construct conservative 95% confidence intervals for

$$\Delta(k) = n^{-1} \left\{ \sum_{t=1}^n (p_t - \hat{p}_t^{(k)})^2 - \sum_{t=1}^n (p_t - \hat{p}_t^{(k-1)})^2 \right\},$$

in which p_t is the actual probability of precipitation on day t . These confidence intervals, which are centered at $B_k - B_{k-1}$, are given in Table 1. The results show significant improvements, by shortening the lead time by one day, in forecasting precipitation $k = 2, 3, 4, 6$.

For another application of Theorem 2, we consider Winkler’s (1994) skill score. To evaluate weather forecasts, a skill score that is commonly used is the percentage improvement in average score over that provided by climatology, denoted by \hat{p}_i^c and considered as an “unskilled” forecaster, that is,

$$(3.8) \quad S_n = \left\{ n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i^c) - n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i) \right\} / n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i^c).$$

Climatology refers to the historic relative frequency, also called the base rate, of precipitation; we can take it to be $\hat{p}_i^c = (M + 1)^{-1} \sum_{t=-M}^0 Y_t$. Noting that (3.8) is not a proper score although it is intuitively appealing, Winkler (1994) proposed to replace the average climatology score in the denominator of (3.8) by individual

TABLE 1
Brier scores B_1 and 95% confidence intervals for $\Delta(k)$

B_1	$\Delta(2)$	$\Delta(3)$	$\Delta(4)$	$\Delta(5)$	$\Delta(6)$	$\Delta(7)$
0.125	0.021 ±0.010	0.012 ±0.011	0.020 ±0.012	0.010 ±0.011	0.015 ±0.011	0.007 ±0.010

weights $l(\hat{p}_i, \hat{p}_i^c)$, that is,

$$(3.9) \quad W_n = n^{-1} \sum_{i=1}^n \{L(Y_i, \hat{p}_i) - L(Y_i, \hat{p}_i^c)\} / l(\hat{p}_i, \hat{p}_i^c),$$

where $l(p, c) = \{L(1, p) - L(1, c)\}I_{\{p \geq c\}} + \{L(0, p) - L(0, c)\}I_{\{p < c\}}$. Theorem 2(i) can be readily extended to show that Winkler’s score W_n is a consistent estimate of

$$(3.10) \quad w_n = n^{-1} \sum_{i=1}^n \{L(p_i, \hat{p}_i) - L(p_i, \hat{p}_i^c)\} / l(\hat{p}_i, \hat{p}_i^c)$$

and that $\sqrt{n}(W_n - w_n) / \tilde{s}_n$ has a limiting standard normal distribution, where

$$(3.11) \quad \tilde{s}_n^2 = n^{-1} \sum_{i=1}^n \delta_i^2 p_i(1 - p_i) / l^2(\hat{p}_i, \hat{p}_i^c).$$

Winkler (1994) used the score (3.9), in which $L(p, \hat{p}) = (p - \hat{p})^2$, to evaluate precipitation probability forecasts, with a 12- to 24-hour lead time, given by the US National Weather Service for 20 cities in the period between April 1966 and September 1983. Besides the score (3.9), he also computed the Brier score and the skill score (3.8) of these forecasts and found that both the Brier and skill scores have high correlations (0.87 and 0.76) whereas (3.9) has a much lower correlation 0.44 with average climatology, suggesting that (3.9) provides a better reflection of the “skill” of the forecasts over an unskilled forecasting rule (based on historic relative frequency). Instead of using correlation coefficients, we performed a more detailed analysis of Winkler’s and skill scores to evaluate the one-day ahead probability forecasts of precipitation for six cities: Las Vegas, NV; Phoenix, AZ; Albuquerque, NM; Queens, NY; Boston, MA; and Portland, OR (listed in increasing order of relative frequency of precipitation), during the period January 1, 2005, to December 31, 2009. The period January 1, 2002, to December 31, 2004, is used to obtain the past three years’ climatology, which is used as the reference unskilled score in the calculation of the skill score and Winkler’s score (3.9). The left panel of Figure 1 plots Winkler’s score against the relative precipitation frequency taken from the period January 1, 2005, to December 31, 2009, which is simply the percentage of days with rain during that period and represents the climatology in (3.8). The dashed line in the right panel of Figure 1 represents linear regression of the skill scores (3.8) on climatology and has a markedly positive slope of 0.95. In contrast, the regression line of Winkler’s scores on climatology, shown in the left panel of Figure 1, is relatively flat and has slope 0.12. Unlike skill scores, Winkler’s scores are proper and provide consistent estimates of the average loss (3.10) involving the actual daily precipitation probabilities p_i for each city during the evaluation period. The vertical bar centered at the dot (representing Winkler’s score) for each city is a 95% confidence interval for (3.10), using a conservative

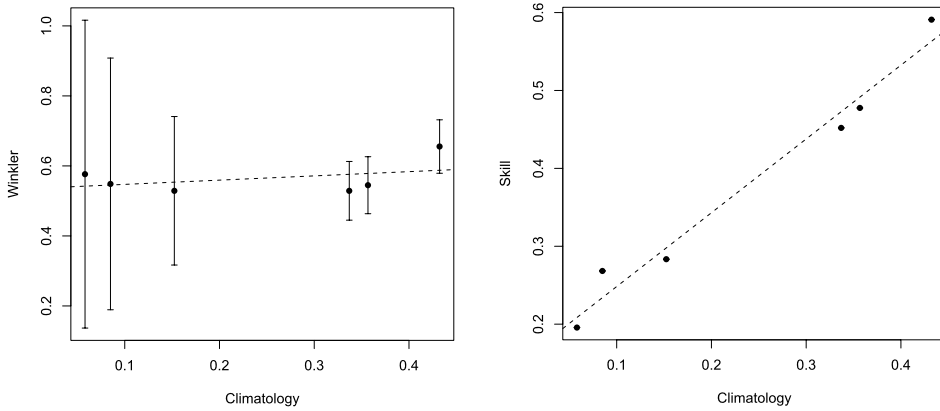


FIG. 1. The Winkler and skill scores versus climatology.

estimate of (3.11) that replaces $p_i(1 - p_i)$ by $1/4$. The confidence intervals are considerably longer for cities whose relative frequencies \hat{p}_i^c of precipitation fall below 0.1 because $\delta_i^2/l^2(\hat{p}_i, \hat{p}_i^c)$ tends to be substantially larger when \hat{p}_i^c is small.

3.4. Risk buckets and quadratic loss functions. Both (3.5) and (3.7) involve $p_i(1 - p_i)$, which is the variance of the Bernoulli random variable Y_i . It is not possible to estimate this variance based on a single observation unless there is some statistical structure on the p_i to make (3.5) or (3.7) estimable, and a conservative approach in the absence of such structure is to use the upper bound $1/4$ for $p_i(1 - p_i)$ in (3.5) or (3.7), as noted in Section 3.1. One such structure is that the p_i can be grouped into buckets within which they have the same value, as in risk assessment of a bank’s retail loans (e.g., mortgages, automobile loans and personal loans), for which the obligors are grouped into risk buckets within which they can be regarded as having the same risk (or more precisely, the same probability of default on their loans). According to the Basel Committee on Banking Supervision [(2006), page 91] each bank has to use at least seven risk buckets for borrowers who have not defaulted and at least one for those who have defaulted previously at the time of loan application.

A bucket model for risk assessment involves multivariate forecasts for events k , $1 \leq k \leq K_t$, at a given time t . Thus, identifying the index i with (t, k) , one has a vector of probability forecasts $(\hat{p}_{t,1}, \dots, \hat{p}_{t,K_t})$ at time $t - 1$ for the occurrences of K_t events at time t ; $K_t = 0$ if no forecast is made at time $t - 1$. The information set can then be expressed as \mathcal{G}_{t-1} that consists of event and forecast histories and other covariates up to time $t - 1$, and therefore conditional on \mathcal{G}_{t-1} and $p_{t,1}, \dots, p_{t,K_t}$, the events at time t can be regarded as the outcomes of K_t independent Bernoulli trials with respective probabilities $p_{t,1}, \dots, p_{t,K_t}$. The bucket model assumes that, conditional on \mathcal{G}_{t-1} and $p_{t,1}, \dots, p_{t,K_t}$, events in the same bucket at time t have the same probability of occurrence. That is, the $p_{t,k}$ are equal for all k belonging

to the same bucket. Let J_t be the number of buckets at time t and $n_{j,t}$ be the size of the j th bucket, $1 \leq j \leq J_t$, so that $n = \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t}$. Then the common $p_{t,k}$ of the j th bucket at time t , denoted by $p_t(j)$, can be estimated by the relative frequency $\bar{Y}_t(j) = n_{j,t}^{-1} \sum_{i \in I_{j,t}} Y_i$, where $I_{j,t}$ denotes the index set for the bucket. This in turn yields an unbiased estimate

$$(3.12) \quad \hat{v}_t(j) = n_{j,t} \bar{Y}_t(j) (1 - \bar{Y}_t(j)) / (n_{j,t} - 1)$$

of $p_i(1 - p_i)$ for $i \in I_{j,t}$, and we can replace $p_i(1 - p_i)$ in (3.5) or (3.7) by $\hat{v}_t(j)$ for $i \in I_{j,t}$ so that the results of Theorems 1 or 2 still hold with these estimates of the asymptotic variance, as shown in the following.

THEOREM 3. *Using the same notation as in the preceding paragraph, suppose $n_{j,t} \geq 2$ for $1 \leq j \leq J_t$ and define $\hat{v}_t(j)$ by (3.12).*

(i) *Under the same assumptions as in Theorem 1, define*

$$\hat{\sigma}_n^2 = n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} \sum_{i \in I_{j,t}} \{L(1, \hat{p}_i) - L(0, \hat{p}_i)\}^2 \hat{v}_t(j).$$

Then $\hat{\sigma}_n^2 - \sigma_n^2$ converges to 0 with probability 1.

(ii) *Under the same assumptions as in Theorem 2, $\hat{s}_n^2 - s_n^2$ converges to 0 with probability 1, where $\hat{s}_n^2 = n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} \sum_{i \in I_{j,t}} \delta_i^2 \hat{v}_t(j)$.*

PROOF. Let \mathcal{F}_{t-1} be the σ -field generated by \mathcal{G}_{t-1} and $p_{s,1}, \dots, p_{s,K_s}$ for $s \leq t$. Note that $\hat{v}_t(j) = \sum_{i \in I_{j,t}} (Y_i - \bar{Y}_t(j))^2 / (n_{j,t} - 1)$ and that

$$(3.13) \quad E(\hat{v}_t(j) | \mathcal{F}_{t-1}) = p_t(j)(1 - p_t(j)),$$

which is the variance of Y_i associated with $I_{j,t}$. Therefore

$$\sum_{j=1}^{J_t} \{ \hat{v}_t(j) - p_t(j)(1 - p_t(j)) \} \left\{ \sum_{i \in I_{j,t}} [L(1, \hat{p}_i) - L(0, \hat{p}_i)]^2 \right\}$$

is a martingale difference sequence with respect to $\{\mathcal{F}_t\}$. Hence we can apply the martingale strong law as in the proof of Theorem 1 to show that $\hat{\sigma}_n^2 - \sigma_n^2$ converges a.s., and the same argument also applies to $\hat{s}_n^2 - s_n^2$. \square

The preceding proof also shows that for the squared error loss $L(p, \hat{p}) = (p - \hat{p})^2$, we can estimate (3.1) in the bucket model by the *adjusted Brier score*

$$(3.14) \quad \hat{L}_n - n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} \hat{v}_t(j),$$

since $\hat{L}_n = n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$ is a consistent estimate of the linear equivalent $n^{-1} \sum_{i=1}^n (\hat{p}_i^2 - 2p_i \hat{p}_i + p_i)$, and $n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} \hat{v}_t(j)$ is a consistent estimate of $n^{-1} \sum_{i=1}^n p_i(1 - p_i)$. Consistency of an estimate \hat{l}_n of l_n means that $\hat{l}_n - l_n$ converges to 0 in probability as $n \rightarrow \infty$. Moreover, the following theorem shows that $\sqrt{n}(\hat{L}_n - n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} \hat{v}_t(j) - L_n)$ has a limiting normal distribution in the bucket model and can be studentized to give a limiting standard normal distribution. Its proof is given in Section 5.

THEOREM 4. *Suppose $n_{j,t} \geq 2$ for $1 \leq j \leq J_t$. Letting $L(p, \hat{p}) = (p - \hat{p})^2$, define L_n by (3.1) and the adjusted Brier score by (3.14). Let $v_t(j) = p_t(j)(1 - p_t(j))$,*

$$\begin{aligned}
 \beta_n^2 = n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} & \left\{ v_t(j) \sum_{i \in I_{j,t}} (1 - 2\hat{p}_i)^2 \right. \\
 (3.15) \qquad & - 2v_t(j)(1 - 2p_t(j)) \sum_{i \in I_{j,t}} (1 - 2\hat{p}_i) \\
 & \left. + n_{j,t} v_t(j)(1 - 4v_t(j)) + 2n_{j,t} v_t^2(j)/(n_{j,t} - 1) \right\}.
 \end{aligned}$$

If β_n converges in probability to some nonrandom positive constant, then $\sqrt{n}(\hat{L}_n - n^{-1} \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} \hat{v}_t(j) - L_n)/\beta_n$ has a limiting standard normal distribution. Moreover, if $n_j \geq 3$ for all $1 \leq j \leq J_t$, then $\hat{\beta}_n - \beta_n$ converges to 0 with probability 1, where

$$\begin{aligned}
 \hat{\beta}_n^2 = \frac{1}{n} \sum_{t=1}^T \sum_{j=1}^{J_t} & \left\{ \hat{v}_t(j) \sum_{i \in I_{j,t}} (1 - 2\hat{p}_i)^2 \right. \\
 (3.16) \qquad & - \frac{2n_{j,t}^2}{(n_{j,t} - 1)^3} \left[\sum_{i \in I_{j,t}} (1 - 2\hat{p}_i) \right] \left[\sum_{i \in I_{j,t}} (Y_i - \bar{Y}_t(j))^3 \right] \\
 & + \frac{4n_{j,t}(n_{j,t} - 1)}{(n_{j,t} - 2)^2} \\
 & \left. \times \sum_{i \in I_{j,t}} \left[\frac{1}{2(n_{j,t} - 1)} \sum_{k \in I_{j,t}, k \neq i} (Y_i - Y_k)^2 - \hat{v}_t(j) \right]^2 \right\}.
 \end{aligned}$$

3.5. Quasi-buckets and reliability diagrams. When the actual $p_{t,i}$ in a bin with index set $I_{j,t}$ are not the same for all $i \in I_{j,t}$, we call the bin a “quasi-bucket.” These quasi-buckets are the basic components of *reliability diagrams* that

are widely used as graphical tools to evaluate probability forecasts. In his description of reliability diagrams, Wilks [(2005), Sections 7.1.2, 7.1.3] notes that reliability, or calibration, relates the forecast to the average observation, “for specific values of (i.e., conditional on) the forecast.” A widely used approach to “verification” of forecasts in meteorology is to group the forecasts \hat{p}_i into bins so that “they are rounded operationally to a finite set of values,” denoted by $\hat{p}(1), \dots, \hat{p}(J)$. Corresponding to each $\hat{p}(j)$ is a set of observations $Y_i, i \in I_j$, taking the values 0 and 1, where $I_j = \{i : \hat{p}_i = \hat{p}(j)\}$. The reliability diagram plots $\bar{Y}(j) = (\sum_{i \in I_j} Y_i)/n_j$ versus $\hat{p}(j)$, where n_j is the size of I_j ; see Figure 3 in Section 4. Statistical inference for reliability diagrams has been developed in the meteorology literature under the assumption of “independence and stationarity,” that is, that (\hat{p}_i, Y_i) are i.i.d. samples from a bivariate distribution of forecast and observation; see Wilks [(2005), Section 7.9.3] and Bröcker and Smith (2007). Under this assumption, the index sets I_j define a bucket model and a $(1 - \alpha)$ -level confidence interval for the common mean $p(j)$ of the Y_i for $i \in I_j$ is

$$(3.17) \quad \bar{Y}(j) \pm z_{1-\alpha/2} \{ \bar{Y}(j)(1 - \bar{Y}(j))/n_j \}^{1/2},$$

where z_q is the q th quantile of the standard normal distribution.

The assumption of i.i.d. forecast-observation pairs is clearly violated in weather forecasting, and this has led to the concern that the confidence intervals given by (3.17) “are possibly too narrow” [Wilks (2005), page 331]. The temporal dependence between the forecast-observation pairs can be handled by incorporating time as in Section 3.4. To be specific, let $\hat{p}_{t,k}, k \leq K_t$, be the probability forecasts, at time $t - 1$, of events in the next period. We divide the set $\{\hat{p}_{t,k} : k \leq K_t, 1 \leq t \leq T\}$ into bins B_1, \dots, B_J , which are typically disjoint sub-intervals of $[0,1]$. Let

$$(3.18) \quad \begin{aligned} I_{j,t} &= \{k : \hat{p}_{t,k} \in B_j\}, \\ \bar{Y}_t(j) &= \sum_{i \in I_{j,t}} Y_i/n_{j,t}, \\ \bar{Y}(j) &= \sum_{t=1}^T \sum_{i \in I_{j,t}} Y_i/n_j, \end{aligned}$$

where $n_{j,t}$ is the cardinality of $I_{j,t}$ and $n_j = \sum_{t=1}^T n_{j,t}$. Note that $n_{j,t}$ and $\bar{Y}_t(j)$ have already been introduced in Section 3.4 and that $\bar{Y}(j)$ is the average of the observations in the j th bin, as in (3.17). In the absence of any assumption on p_i for $i \in I_{j,t}$, these index sets define quasi-buckets instead of buckets. We can extend the arguments of Section 3.4 to the general case that makes no assumptions on the p_i and thereby derive the statistical properties of $\bar{Y}(j)$ without the restrictive assumption of i.i.d. (\hat{p}_i, Y_i) . With the same notation as in Section 3.4, note that the index sets $I_{j,t}$ defined in (3.18) are \mathcal{G}_{t-1} -measurable since the $\hat{p}_{t,k}$ are \mathcal{G}_{t-1} -measurable.

Whereas $\bar{Y}_t(j)$ is used to estimate the common value of p_i for $i \in I_{j,t}$ and $\hat{v}_t(j)$, defined in (3.12), is used to estimate the common value of $p_i(1 - p_i)$ in Section 3.4, the p_i in quasi-buckets need no longer be equal. Replacing Y_i by p_i in $\bar{Y}(j)$ and taking a weighted average of $\hat{v}_t(j)$ over t , we obtain

$$(3.19) \quad \bar{p}(j) = \frac{\sum_{t=1}^T \sum_{i \in I_{j,t}} p_i}{n_j}, \quad \hat{v}(j) = \frac{\sum_{t=1}^T n_{j,t} \hat{v}_t(j)}{n_j}.$$

Instead of (3.17) that is based on overly strong assumptions, we propose to use

$$(3.20) \quad \bar{Y}(j) \pm z_{1-\alpha/2} \{ \hat{v}(j) / n_j \}^{1/2}$$

as a $(1 - \alpha)$ -level confidence interval for $\bar{p}(j)$. Part (iii) of the following theorem, whose proof is given in Section 5, shows that the confidence interval tends to be conservative. Parts (i) and (ii) modify the estimates in Theorem 3 for σ_n^2 and s_n^2 when the p_i in the assumed buckets turn out to be unequal.

THEOREM 5. *With the same notation as in Theorem 3, remove the assumption that p_i are all equal for $i \in I_{j,t}$ but assume that $I_{j,t}$ is \mathcal{G}_{t-1} -measurable for $1 \leq j \leq J_t$.*

(i) *Under the assumptions of Theorem 1, let*

$$(3.21) \quad \tilde{\sigma}_n^2 = \frac{1}{n} \sum_{t=1}^T \sum_{j=1}^{J_t} \sum_{i \in I_{j,t}} \{L(1, \hat{p}_i) - L(0, \hat{p}_i)\}^2 (Y_i - \bar{Y}_t(j))^2 \frac{n_{j,t}}{n_{j,t} - 1}.$$

Then $\tilde{\sigma}_n^2 \geq \sigma_n^2 + o(1)$ a.s. Moreover, if the p_i are equal for all $i \in I_{j,t}$ and $1 \leq j \leq J_t$, then $\tilde{\sigma}_n^2 - \sigma_n^2$ converges to 0 a.s.

(ii) *Under the assumptions of Theorem 2, $\tilde{s}_n^2 \geq s_n^2 + o(1)$ a.s., where*

$$(3.22) \quad \tilde{s}_n^2 = \frac{1}{n} \sum_{t=1}^T \sum_{j=1}^{J_t} \sum_{i \in I_{j,t}} \delta_i^2 (Y_i - \bar{Y}_t(j))^2 \frac{n_{j,t}}{n_{j,t} - 1}.$$

(iii) *Suppose $J_t = J$ for all $t = 1, \dots, T$. For $1 \leq j \leq J$, define $\bar{Y}(j)$ by (3.18), and $\bar{p}(j)$ and $\hat{v}(j)$ by (3.19), in which $n_j = \sum_{t=1}^T n_{j,t}$. Let*

$$(3.23) \quad v(j) = n_j^{-1} \sum_{t=1}^T \sum_{i \in I_{j,t}} p_i(1 - p_i),$$

and let $n = n_1 + \dots + n_J$ be the total sample size. Suppose n_j/n and $v(j)$ converge in probability to nonrandom positive constants as $n \rightarrow \infty$. Then $(n_j/v(j))^{1/2} \{ \bar{Y}(j) - \bar{p}(j) \}$ has a limiting standard normal distribution as $n \rightarrow \infty$. Moreover, $\hat{v}(j) \geq v(j) + o_p(1)$ and equality holds if the p_i are equal for all $i \in I_{j,t}$.

Note that the numerator of (3.12) is equal to $\sum_{i \in I_{j,t}} (Y_i - \bar{Y}_t(j))^2$. The estimate (3.21) or (3.22) essentially replaces this sum by a weighted sum, using the weights associated with $p_i(1 - p_i)$ in the sum (3.5) or (3.7) that defines σ_n^2 or s_n^2 . The term $n_{j,t}/(n_{j,t} - 1)$ in (3.21) and (3.22) corresponds to the bias correction factor in the sample variance (3.12). Theorem 5 shows that (3.21) [or (3.22)] is still a consistent estimator of σ_n^2 (or s_n^2) if the bucket model holds, and that it tends to over-estimate σ_n^2 (or s_n^2) otherwise, erring only on the conservative side.

4. Simulation studies. The risk buckets in Section 3.4 and the forecasts are usually based on covariates. In this section we consider $T = 2$ in the case of discrete covariates so that there are J_t buckets of various sizes for $n = \sum_{t=1}^2 \sum_{j=1}^{J_t} n_{j,t} = 300$ probability forecasts prior to observing the indicator variables Y_1, \dots, Y_n of the events. We use the Brier score and its associated loss function $L(p, \hat{p}) = (p - \hat{p})^2$ to evaluate the probability forecasts and study by simulations the adequacy of the estimates $\hat{\beta}_n^2$ and \hat{s}_n^2 and their use in the normal approximations. The simulation study covers four scenarios and involves 1,000 simulation runs for each scenario. Scenario 1 considers the Brier score of a forecasting rule, while Scenarios 2–4 consider the difference of Brier scores of two forecasts. The bucket sizes and how the p_i and \hat{p}_i are generated in each scenario are described as follows.

SCENARIO 1. There are ten buckets of size 15 each for each period. The common values $p_t(j)$ in the buckets are 0.1, 0.25, 0.3, 0.35, 0.4, 0.5, 0.65, 0.7, 0.75 and 0.8, respectively, for $t = 1, 2$. The probability forecast $\hat{p}_{t,k}$, $1 \leq k \leq 150$, made at time $t - 1$, uses covariate information to identify the bucket j associated with the k th event at time t and predicts that it occurs with probability $\bar{Y}_{t-1}(j)$, assuming that 150 indicator variables at time 0 are also observed so that $\bar{Y}_0(j)$ is available.

SCENARIO 2. For each period, there are nine buckets, three of which have size 2 and two of which have size 5; the other bucket sizes are 24, 30, 35 and 45 (one bucket for each size). The bucket probabilities $p_t(j)$ are i.i.d. random variables generated from Uniform (0,1). The forecast $\hat{p}_{t,k}$ is the same as that in Scenario 1, and there is another forecast $\hat{p}'_{t,k} = \bar{Y}_{t-1}$ that ignores covariate information.

SCENARIO 3. For each period, there are five buckets of size 30 each, and $p_t(j) = -0.1 + j/5$ for $j = 1, \dots, 5$. The two forecasts are the same as in Scenario 2.

SCENARIO 4. This is the same as Scenario 3, except that p_i is uniformly distributed on $[(j - 1)/5, j/5]$ for $i \in I_{j,t}$, that is, the bucket assumption is only approximately correct.

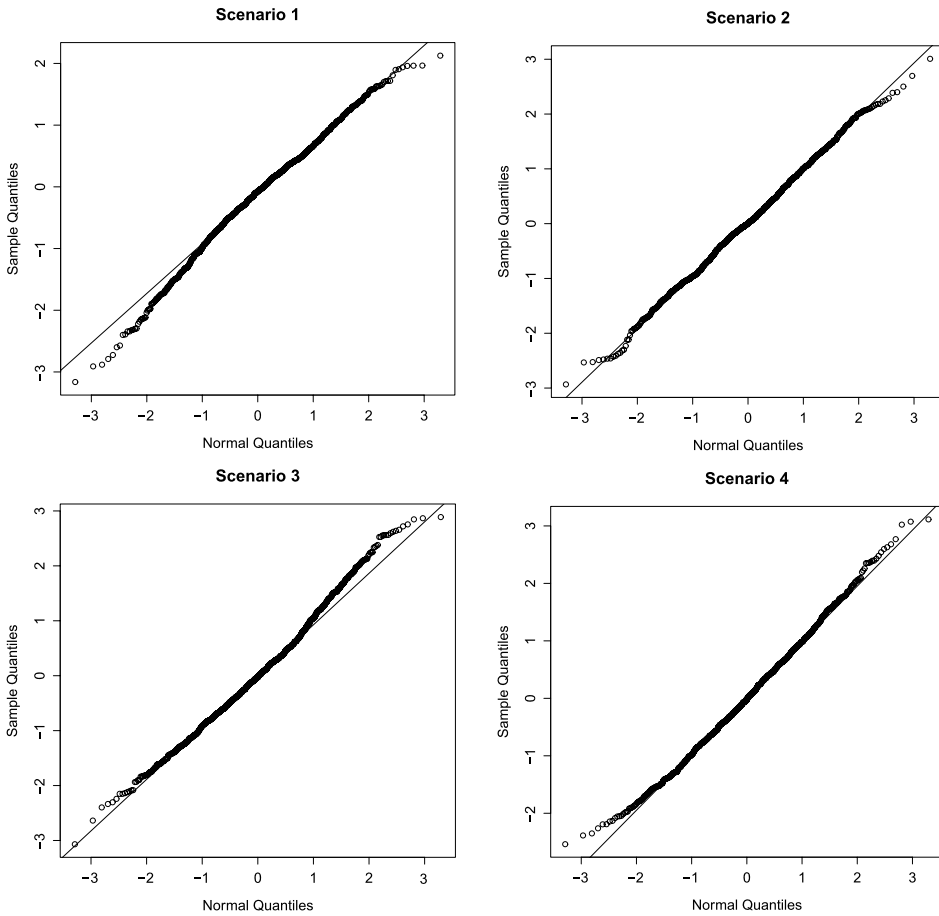


FIG. 2. *Q–Q plots for Scenarios 1–4.*

Figure 2 gives the Q–Q plots of $\sqrt{n}(\hat{L}_n - n^{-1} \sum_{t=1}^2 \sum_{j=1}^{J_t} n_{j,t} \hat{v}_t(j) - L_n) / \hat{\beta}_n$ for Scenario 1 and $\sqrt{n}(\hat{\Delta}_n - \Delta_n) / \hat{s}_n$ for Scenarios 2–4. Despite the deviation from the assumed bucket model in Scenario 4, the Q–Q plot does not deviate much from the 45° line. Table 2 gives the means and 5-number summaries (minimum, maximum, median, 1st and 3rd quartiles) of \hat{s}_n/s_n for Scenarios 2–4 and $\hat{\beta}_n/\beta_n$ for Scenario 1.

To illustrate the reliability diagram and the associated confidence intervals (3.20) in Section 3.5, we use one of the simulated data sets in Scenario 4 to construct the reliability diagram for the forecasts $\hat{p}_{t,k}$ ($t = 1, 2; k = 1, \dots, 5$), grouping the forecasts over time in the bins $[(j - 1)/5, j/5]$, $j = 1, \dots, 5$, that are natural for this scenario. The diagram is given in Figure 3. Table 3 gives the means, standard deviations (SD), and 5-number summaries of $\bar{Y}(j)$, $\bar{p}(j)$, $\hat{v}(j)$ and $v(j)$ defined in (3.18), (3.19) and (3.23) based on the 1,000 simulations. In particular,

TABLE 2
Simulation results for $\hat{\beta}_n/\beta_n$ (Scenario 1) and \hat{s}_n/s_n

	Min.	1st qrt.	Median	3rd qrt.	Max.	Mean
Scenario 1	0.6397	1.0840	1.1810	1.2830	1.6520	1.1780
Scenario 2	0.7442	0.9647	1.0060	1.0490	1.1970	1.0050
Scenario 3	0.7586	0.9506	1.0060	1.0570	1.2070	1.0010
Scenario 4	0.7420	0.9661	1.0180	1.0730	1.2240	1.0160

it shows that $\hat{v}(j)$ tends to over-estimate $v(j)$. Moreover, the probability of coverage of the 95% interval (3.20) for $\bar{p}(j)$, evaluated by averaging over the 1,000 simulations, is 0.949, 0.947, 0.944, 0.940 and 0.928, for $j = 1, \dots, 5$, respectively, suggesting that the results of Theorem 5 still apply even for moderate sample sizes. We do not consider the second forecast $\hat{p}'_{t,k} = \bar{Y}_{t-1}$ to illustrate reliability diagrams because by the central limit theorem, the $\hat{p}'_{t,k}$ are concentrated around 0.5 and nearly all of the forecasts lie in the bin $[0.4, 0.6]$.

5. Proofs of Theorems 4 and 5. Re-labeling the Y_i as $Y_{t,1}, \dots, Y_{t,K_t}$, we note that conditional on \mathcal{F}_{t-1} , $\{Y_{t,k} : 1 \leq k \leq K_t\}$ is a set of independent Bernoulli random variables with respective parameters $p_{t,1}, \dots, p_{t,K_t}$. This point, which has been noted in the second paragraph of Section 3.4 and will be discussed further in Section 7, explains why we can first derive the result for the special case in which $Y_{t,k}$ are independent and then modify the argument by conditioning on \mathcal{F}_{t-1} and appealing to martingale theory. As an illustration, note that if $Y_i, i \in I_{j,t}$, are




FIG. 3. Reliability diagram for the forecasts $\hat{p}_{t,k}$. At the midpoint of each of the five bins $[(j - 1)/5, j/5]$, $j = 1, \dots, 5$, a 95% confidence interval, centered at $\bar{Y}(j)$, for $\bar{p}(j)$ is shown; only the upper (or lower) half of the interval is shown at $j = 1$ (or 5) to keep the range of the vertical axis between 0.1 and 0.9.

TABLE 3
Simulation results for $\bar{p}(j), \bar{Y}(j), v(j)$ and $\hat{v}(j)$

	Min	1st qrt.	Median	3rd qrt.	Max	Mean	SD
$\bar{p}(1)$	0.101	0.101	0.101	0.168	0.234	0.121	0.033
$\bar{Y}(1)$	0.017	0.083	0.117	0.156	0.350	0.123	0.051
$v(1)$	0.087	0.087	0.087	0.127	0.167	0.100	0.020
$\hat{v}(1)$	0.067	0.089	0.106	0.132	0.233	0.106	0.037
$\bar{p}(2)$	0.101	0.300	0.300	0.355	0.515	0.320	0.049
$\bar{Y}(2)$	0.050	0.267	0.317	0.378	0.633	0.319	0.089
$v(2)$	0.087	0.207	0.207	0.207	0.247	0.209	0.015
$\hat{v}(2)$	0.048	0.208	0.221	0.239	0.259	0.213	0.034
$\bar{p}(3)$	0.300	0.515	0.515	0.577	0.701	0.527	0.058
$\bar{Y}(3)$	0.217	0.467	0.533	0.589	0.833	0.529	0.096
$v(3)$	0.206	0.233	0.247	0.247	0.247	0.239	0.011
$\hat{v}(3)$	0.144	0.240	0.249	0.254	0.259	0.244	0.015
$\bar{p}(4)$	0.515	0.659	0.701	0.701	0.906	0.690	0.052
$\bar{Y}(4)$	0.367	0.633	0.689	0.750	1.000	0.687	0.090
$v(4)$	0.082	0.206	0.206	0.206	0.247	0.204	0.021
$\hat{v}(4)$	0.063	0.207	0.217	0.236	0.259	0.211	0.035
$\bar{p}(5)$	0.769	0.906	0.906	0.906	0.906	0.895	0.026
$\bar{Y}(5)$	0.733	0.867	0.900	0.933	1.000	0.892	0.049
$v(5)$	0.082	0.082	0.082	0.082	0.164	0.088	0.016
$\hat{v}(5)$	0.077	0.084	0.093	0.120	0.202	0.096	0.037

i.i.d. Bernoulli random variables with common parameter $p_t(j)$, then $\hat{v}_t(j)$ defined in (3.12) is an unbiased estimate of $p_t(j)(1 - p_t(j))$ and one can use the classical strong law of large numbers to derive the result. The proof of Theorem 3 basically shows that $\hat{v}_t(j)$ is “conditionally unbiased” given \mathcal{F}_{t-1} in the sense of (3.13) and then uses the martingale strong law to derive the result. To prove Theorem 4, we extend this idea to obtain a conditionally unbiased estimate of $\text{Var}(\hat{v}_t(j))$ by first considering the i.i.d. case: let X_1, \dots, X_m be i.i.d. random variables. As is well known, the sample variance $\hat{v} = \sum_{i=1}^m (X_i - \bar{X})^2 / (m - 1)$ is a U -statistic of order 2, with kernel $h(X_i, X_k) = (X_i - X_k)^2 / 2$. Arvesen (1969) has shown that an unbiased estimate of the variance of the U -statistic is the jackknife estimate

$$(5.1) \quad \frac{4(m - 1)}{m(m - 2)^2} \sum_{i=1}^m \left\{ \frac{1}{m - 1} \sum_{\substack{k=1 \\ k \neq i}}^m h(X_i, X_k) - \hat{v} \right\}^2.$$

PROOF OF THEOREM 4. Use $L(p, \hat{p}) = (p - \hat{p})^2$ to express $n\{\hat{L}_n - n^{-1} \times \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} \hat{v}_t(j) - L_n\}$ as

$$(5.2) \quad \sum_{t=1}^T \sum_{k=1}^{K_t} (1 - 2\hat{p}_{t,k})(Y_{t,k} - p_{t,k}) - \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} [\hat{v}_t(j) - p_t(j)(1 - p_t(j))],$$

which is the difference of two martingales and is therefore a martingale. To compute the conditional variance (or predictable variation) of (5.2), we can use the ‘‘angle bracket’’ notation and formulas for predictable variation and covariation [Williams (1991), Section 12.12] to obtain

$$(5.3) \quad \begin{aligned} & \left\langle \sum_{t=1}^T \sum_{k=1}^{K_t} (1 - 2\hat{p}_{t,k})(Y_{t,k} - p_{t,k}) \right\rangle \\ &= \sum_{t=1}^T \sum_{k=1}^{K_t} (1 - 2\hat{p}_{t,k})^2 E((Y_{t,k} - p_{t,k})^2 | \mathcal{F}_{t-1}) \\ &= \sum_{t=1}^T \sum_{j=1}^{J_t} \left(\sum_{i \in I_{j,t}} (1 - 2\hat{p}_i)^2 \right) p_t(j)(1 - p_t(j)), \end{aligned}$$

$$(5.4) \quad \begin{aligned} & \left\langle \sum_{t=1}^T \sum_{k=1}^{K_t} (1 - 2\hat{p}_{t,k})(Y_{t,k} - p_{t,k}), \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} [\hat{v}_t(j) - p_t(j)(1 - p_t(j))] \right\rangle \\ &= \sum_{t=1}^T \sum_{j=1}^{J_t} \left[\sum_{i \in I_{j,t}} (1 - 2\hat{p}_i) \right] p_t(j)(1 - p_t(j))(1 - 2p_t(j)), \end{aligned}$$

$$(5.5) \quad \begin{aligned} & \left\langle \sum_{t=1}^T \sum_{j=1}^{J_t} n_{j,t} [\hat{v}_t(j) - p_t(j)(1 - p_t(j))] \right\rangle \\ &= \sum_{t=1}^T \sum_{j=1}^{J_t} \{ n_{j,t} p_t(j)(1 - p_t(j)) [1 - 4p_t(j)(1 - p_t(j))] \\ & \quad + 2n_{j,t} p_t^2(j)(1 - p_t(j))^2 / (n_{j,t} - 1) \}. \end{aligned}$$

Combining (5.3), (5.4) and (5.5) yields formula (3.15) for the conditional variance of (5.2) divided by n .

In view of (3.13), $\hat{v}_t(j)$ is a conditionally unbiased estimate of $p_t(j)(1 - p_t(j))$ given \mathcal{F}_{t-1} . If $Y \sim \text{Bernoulli}(p_i)$, then $E(Y - p)^3 = p(1 - p)(1 - 2p)$. Hence a conditionally unbiased estimate of $p_t(j)(1 - p_t(j))(1 - 2p_t(j))$ given \mathcal{F}_{t-1} is

$$(5.6) \quad [n_{j,t}^2 / (n_{j,t} - 1)^3] \sum_{i \in I_{j,t}} (Y_i - \bar{Y}_t(j))^3,$$

analogous to (3.13). Replacing $p_t(j)(1 - p_t(j))(1 - 2p_t(j))$ in (5.4) by (5.6) and multiplying (5.4) by $-2/n$ gives the second summand of (3.16). Note that the first summand of (3.16) corresponds to replacing $p_t(j)(1 - p_t(j))$ in (5.3) by $\hat{v}_t(j)$. The last summand of (3.16) corresponds to using the jackknife estimate (5.1) to estimate the conditional variance of $\hat{v}_t(j)$ given \mathcal{F}_{t-1} . Since $\{Y_i, i \in I_{j,t}\}$ is a set of i.i.d. random variables conditional on \mathcal{F}_{t-1} , the jackknife estimate is conditionally unbiased given \mathcal{F}_{t-1} ; see the paragraph preceding the proof of this theorem. The rest of the argument is similar to that of Theorem 3. \square

PROOF OF THEOREM 5. We first prove (iii). Using the notation in the paragraph preceding the proof of Theorem 4, recall that conditional on \mathcal{F}_{t-1} , the $Y_{t,k}$ are independent Bernoulli($p_{t,k}$) random variables. Since $I_{j,t}$ is \mathcal{F}_{t-1} -measurable, it follows that $\sum_{i \in I_{j,t}} (Y_i - p_i)$ is a martingale difference sequence with respect to $\{\mathcal{F}_t\}$ and $E\{\sum_{i \in I_{j,t}} (Y_i - p_i)^2 | \mathcal{F}_{t-1}\} = \sum_{i \in I_{j,t}} p_i(1 - p_i)$. Since $n^{-1} \sum_{t=1}^T \sum_{i \in I_{j,t}} p_i(1 - p_i)$ converges in probability to a nonrandom positive constant as $n \rightarrow \infty$, we can apply the martingale central limit theorem as in the proof of Theorem 1 to conclude that

$$\frac{\sum_{t=1}^T \sum_{i \in I_{j,t}} (Y_i - p_i)}{\{\sum_{t=1}^T \sum_{i \in I_{j,t}} p_i(1 - p_i)\}^{1/2}} \implies N(0, 1)$$

proving the first part of (iii).

To prove the second part of (iii), and also (i) and (ii), we first show that for any nonnegative \mathcal{F}_{t-1} -measurable random variables $w_{t,1}, \dots, w_{t,K_t}$,

$$\begin{aligned} (5.7) \quad & E\left\{ \sum_{k \in I_{j,t}} w_{t,k} (Y_{t,k} - \bar{Y}_t(j))^2 | \mathcal{F}_{t-1} \right\} \\ & \geq \sum_{k \in I_{j,t}} (1 - n_{j,t}^{-1}) w_{t,k} p_{t,k} (1 - p_{t,k}), \end{aligned}$$

in which $\sum_{k \in I_{j,t}}$ means $\sum_{i \in I_{j,t}}$ when i is represented as (t, k) ; see the second paragraph of Section 3.4. Define $\bar{Y}_t(j)$, $\bar{Y}(j)$ and $\bar{p}(j)$ as in (3.18) and (3.19), and let $\bar{p}_t(j) = (\sum_{k \in I_{j,t}} p_{t,k})/n_{j,t}$. From the decomposition

$$(5.8) \quad Y_{t,k} - \bar{Y}_t(j) = (Y_{t,k} - p_{t,k}) + (p_{t,k} - \bar{p}_t(j)) + (\bar{p}_t(j) - \bar{Y}_t(j)),$$

it follows that the left-hand side of (5.7) is equal to

$$\begin{aligned} (5.9) \quad & \sum_{k \in I_{j,t}} w_{t,k} E[(Y_{t,k} - p_{t,k})^2 | \mathcal{F}_{t-1}] + \sum_{k \in I_{j,t}} w_{t,k} (p_{t,k} - \bar{p}_t(j))^2 \\ & + \sum_{k \in I_{j,t}} w_{t,k} E[(\bar{Y}_t(j) - \bar{p}_t(j))^2 | \mathcal{F}_{t-1}] \\ & - 2E\left[(\bar{Y}_t(j) - \bar{p}_t(j)) \sum_{k \in I_{j,t}} w_{t,k} (Y_{t,k} - p_{t,k}) \right] \end{aligned}$$

by using the fact that conditional on \mathcal{F}_{t-1} the $Y_{t,k}$ are independent Bernoulli. Since $\bar{Y}_t(j) - \bar{p}_t(j) = \sum_{k \in I_{j,t}} (Y_{t,k} - p_{t,k})/n_{j,t}$, we can use this fact again to combine the last two terms of (5.9) into

$$(5.10) \quad - \sum_{k \in I_{j,t}} (w_{t,k}/n_{j,t}) E[(Y_{t,k} - p_{t,k})^2 | \mathcal{F}_{t-1}].$$

Since $w_{t,k} \geq 0$, we can drop the second term in (5.9) to obtain (5.7) from (5.9) and (5.10). Moreover, since this term is actually 0 when the $p_{t,k}$ are all equal for $k \in I_{j,t}$, equality holds in (5.7) in this case.

Let $w_{t,k} = n_{j,t}/(n_{j,t} - 1)$. Then (5.7) reduces to

$$(5.11) \quad E(n_{j,t} \hat{v}_t(j) | \mathcal{F}_{t-1}) \geq \sum_{k \in I_{j,t}} p_{t,k}(1 - p_{t,k}).$$

Under the assumptions of part (iii) of the theorem, we can apply the martingale strong law to obtain

$$(5.12) \quad \sum_{t=1}^T \{n_{j,t} \hat{v}_t(j) - E(n_{j,t} \hat{v}_t(j) | \mathcal{F}_{t-1})\} / n_j \longrightarrow 0 \quad \text{a.s. on } \{n_j \rightarrow \infty\}.$$

Combining (5.11) with (5.12) yields $\hat{v}(j) \geq v(j) + o_p(1)$, with equality when the $p_{t,k}$ are all equal for $k \in I_{j,t}$.

To prove part (i) of the theorem, put $w_{t,k} = \{L(1, \hat{p}_{t,k}) - L(0, \hat{p}_{t,k})\}^2 n_{j,t} / (n_{j,t} - 1)$ in (5.7) and then use the same argument as in the preceding paragraph. The proof of part (ii) is similar. \square

6. Extensions and connections to forecast comparison in econometrics.

Our new approach to evaluating probability forecasts in Section 3 is based on consistent and asymptotically normal estimates of the average loss $n^{-1} \sum_{i=1}^n L(p_i, \hat{p}_i)$, without any assumptions on how the observed indicator variables Y_i and their forecasts \hat{p}_i are generated. The key to this approach is that conditional on \mathcal{F}_{i-1} , Y_i is Bernoulli(p_i), and therefore martingale arguments can be used to derive the results in Section 3. In Section 6.1 we show how this approach can be extended to more general random variables Y_i . As shown in (2.3), when Y_i is an indicator variable, the conditional expectation of the score $L(Y_i, \hat{p}_i)$ given \mathcal{F}_{i-1} is a linear function of p_i , but this does not extend to more general random variables Y_i . In Section 6.2 we review the recent econometrics literature on testing the equality of the expected scores of two forecasts and discuss an alternative approach to statistical inference on the expected difference in average scores of two forecasts.

6.1. *Extensions to general predictands.* A characteristic of (\hat{p}_i, Y_i) in probability forecasting is that $E(Y_i | \mathcal{F}_{i-1}) = p_i$ while the \mathcal{G}_{i-1} -measurable forecast \hat{p}_i is an estimate of p_i . The theorems in Section 3 and their martingale proofs in Section 5 can be easily extended to general random variables Y_i when the loss

function is of the form $L(\mu_i, \hat{\mu}_i)$, where $\mu_i = E(Y_i|\mathcal{F}_{i-1})$ and $\hat{\mu}_i$ is a forecast of Y_i given \mathcal{F}_{i-1} . Although $Y_i|\mathcal{F}_{i-1} \sim \text{Bernoulli}(p_i)$ in Section 3, no parametric assumptions are actually needed when we use a loss function of the form $L(\mu_i, \hat{\mu}_i)$. As in (2.2), such loss function is said to have a linear equivalent \tilde{L} if

$$(6.1) \quad \tilde{L}(y, \hat{y}) \text{ is linear in } y \quad \text{and} \quad L(y, \hat{y}) - \tilde{L}(y, \hat{y}) \text{ does not depend on } \hat{y}.$$

The bucket model in Section 3.4 can be extended so that $Y_{t,k}|\mathcal{F}_{t-1}$ have the same mean and variance for all (t, k) belonging to the same bucket. In place of (3.12), we now use

$$(6.2) \quad \hat{v}_t(j) = \sum_{i \in I_{j,t}} (Y_i - \bar{Y}_t(j))^2 / (n_{j,t} - 1)$$

as an unbiased estimate of the common conditional variance of Y_i given \mathcal{F}_{i-1} for $i = (t, k) \in I_{j,t}$, using the same notation as that in the proof of Theorem 4. While the extension of Theorem 3 only needs the first two moments of $Y_{t,k}|\mathcal{F}_{t-1}$ to be equal for all (t, k) belonging to the same bucket, Theorem 4 can also be extended by assuming the first four moments of $Y_{t,k}|\mathcal{F}_{t-1}$ to be equal for all (t, k) belonging to the same bucket, by using Arvesen’s (1969) jackknife estimate of the variance of a U -statistic.

Clearly (5.8), (5.9) and (5.10) also hold with $p_{t,k}$ and $\bar{p}_t(j)$ replaced by $\mu_{t,k}$ and $\bar{\mu}_t(k)$, so Theorem 5 can likewise be extended to quasi-buckets and reliability diagrams for the predicted means $\hat{\mu}_{t,k}$. For sample means in the case of independent observations within each bucket, this extension of Theorem 5 can be viewed as a corollary of the analysis of variance. In fact, the proof of Theorem 5 uses martingale arguments and conditioning to allow dependent observations in each (quasi-)bucket.

6.2. *Inference on expected difference in average scores of two forecasts.*

When the Y_i are indicator variables of events, Theorem 2(ii) establishes asymptotic normality for the difference $\hat{\Delta}_n = n^{-1} \sum_{i=1}^n \{L(Y_i, \hat{p}'_i) - L(Y_i, \hat{p}''_i)\}$ in average scores between two forecasts, from which one can perform inference on

$$(6.3) \quad \begin{aligned} \Delta_n &= n^{-1} \sum_{i=1}^n E\{L(Y_i, \hat{p}'_i) - L(Y_i, \hat{p}''_i) | \mathcal{F}_{i-1}\} \\ &= n^{-1} \sum_{i=1}^n \{\delta_i p_i + L(0, \hat{p}'_i) - L(0, \hat{p}''_i)\}, \end{aligned}$$

where $\delta_i = \{L(1, \hat{p}'_i) - L(0, \hat{p}'_i)\} - \{L(1, \hat{p}''_i) - L(0, \hat{p}''_i)\}$. This simplicity, however, does not extend to general Y_i .

Proper scoring rules for probability forecasts of categorical and continuous variables Y_i have been an active area of research; see the review by Gneiting and Raftery (2007). Another active area of research is related to the extension

of $\hat{\Delta}_n$ to general Y_i in the econometrics literature, beginning with the seminal paper of [Diebold and Mariano \(1995\)](#). They consider the usual forecast errors $e_t := Y_t - \hat{Y}_{t|t-1}$ in time series analysis, where $\hat{Y}_{t|t-1}$ is the one-step ahead forecast of Y_t based on observations up to time $t - 1$. Unlike a probability forecast that gives a predictive distribution of Y_t as in [Gneiting and Raftery \(2007\)](#), $\hat{Y}_{t|t-1}$ is a nonprobabilistic forecast that predicts the value of Y_t [see [Wilks \(2005\)](#), Section 7.3]. The score used by [Diebold and Mariano \(1995\)](#) is of the form $L(Y_t, \hat{Y}_{t|t-1}) = g(e_t)$, and they consider the *average loss differential*

$$(6.4) \quad \hat{\Delta}_n = n^{-1} \sum_{t=1}^n \{L(Y_t, \hat{Y}'_{t|t-1}) - L(Y_t, \hat{Y}''_{t|t-1})\}$$

between two forecasts $Y'_{t|t-1}$ and $Y''_{t|t-1}$, $1 \leq t \leq n$. Assuming a probability measure Q under which $d_t := g(e'_t) - g(e''_t)$ is covariance stationary with absolutely summable autocovariances γ_k so that $f(0) := \sum_{k=-\infty}^{\infty} \gamma_k / (2\pi)$ is the spectral density at frequency 0, they use the asymptotic normality of $\hat{\Delta}_n$ under the null hypothesis $H_0: E_Q(d_t) = 0$ and a window estimate $\hat{f}(0)$ of $f(0)$ so that the test statistic $\sqrt{n} \hat{\Delta}_n / (2\pi \hat{f}(0))^{1/2}$ has a limiting standard normal distribution under H_0 as $n \rightarrow \infty$. This asymptotic normality result, however, requires additional assumptions, such as stationary mixing, which they do not mention explicitly. Their work has attracted immediate attention and spawned many subsequent developments in the econometrics literature on this topic.

[Giacomini and White \(2006\)](#), hereafter abbreviated as G&W, review some of the developments and propose a refinement of H_0 for which the asymptotic normality of $\hat{\Delta}_n$ can be established under precisely stated conditions that can also allow nonstationarity. They formulate the null hypothesis of equal predictive ability of two forecasting models or methods as “a problem of inference about conditional expectations of forecasts and forecast errors that nests the unconditional expectations that are the sole focus of the existing literature.” The econometrics literature they refer to is primarily concerned with “forecast models;” thus Q in the previous paragraph is the probability measure associated with the forecast model being evaluated, or with a more general model than the two competing forecast models whose predictive abilities are compared.

G&W evaluate not only the forecasting model but also the forecasting method, which includes “the forecasting model along with a number of choices,” such as the estimation procedure and the window of past data, used to produce the forecast. They consider k -step ahead forecasts, for which $\hat{Y}_{t|t-1}$ is replaced by $\hat{Y}_{t|t-k}$, and assume that the forecasts are based on finite-memory models involving unknown parameters, that is,

$$(6.5) \quad \hat{Y}_{t+k|t} = h(Y_t, \dots, Y_{t-m+1}, \mathbf{x}_t, \dots, \mathbf{x}_{t-m+1}; \boldsymbol{\beta}),$$

where h is a known function, m is the order of the model, \mathbf{x}_t is a covariate vector at time t and $\boldsymbol{\beta}$ is a parameter vector to be estimated from some specified window of past data. Their formulation generalizes that of [West \(1996\)](#) who considers

regression models. Whereas West assumes that the data are actually generated by the regression model with true parameter β^* , G&W allow model misspecification, and therefore their assumptions do not involve β^* . They consider two such nominal models, resulting in the forecasts $\hat{Y}'_{t|t-k}$ and $\hat{Y}''_{t|t-k}$ that use the same covariates but different estimates $\hat{\beta}'_t$ and $\hat{\beta}''_t$. Their null hypothesis

$$(6.6) \quad H_0 : E\{L(Y_t, \hat{Y}'_{t|t-k}) - L(Y_t, \hat{Y}''_{t|t-k}) | \mathcal{G}_{t-k}\} = 0 \quad \text{a.s. } \forall t \geq 1$$

seems to be stronger than $EL(Y_t, Y'_{t|t-k}) = EL(Y_t, Y''_{t|t-k}) \quad \forall t$ considered by Diebold and Mariano (1995) for the case $k = 1$. On the other hand, (6.6) in the case $k = 1$ just says that $L(Y_t, Y'_{t|t-1}) - L(Y_t, Y''_{t|t-1})$ is a martingale difference sequence under H_0 so that the martingale central limit theorem can be applied to derive the limiting χ^2 -distribution of G&W's test statistics under H_0 . Unlike Diebold and Mariano (1995) for the case $k = 1$, G&W do not use test statistics of the form (6.4) and their test statistics involve more complicated weighted sums of $L(Y_t, \hat{Y}'_{t|t-k}) - L(Y_t, \hat{Y}''_{t|t-k})$. These weights are chosen to improve the power of the test and require additional mixing and moment assumptions on (\mathbf{x}_t, Y_t) given in their Theorems 1–3.

The methodology developed in Section 3 and its extension outlined in Section 6.1 suggest an alternative approach to comparing econometric forecasts. As in (6.4), we consider the average score difference

$$(6.7) \quad \hat{\Delta}_n = n^{-1} \sum_{i=1}^n \{L(Y_t, \hat{Y}'_t) - L(Y_t, \hat{Y}''_t)\},$$

in which \hat{Y}'_t and \hat{Y}''_t are forecasts that are \mathcal{G}_{t-1} -measurable. Since k -step ahead forecasts of Y_t are \mathcal{G}_{t-1} -measurable for any $k \geq 1$, the theory applies to all k -step ahead forecasts of Y_t , as illustrated in Table 1. Instead of hypothesis testing, our approach is targeted toward estimating

$$(6.8) \quad \Delta_n = n^{-1} \sum_{i=1}^n E\{L(Y_t, \hat{Y}'_t) - L(Y_t, \hat{Y}''_t) | \mathcal{G}_{t-1}\},$$

in which E is with respect to the actual but unknown probability measure P . Analogously to Theorem 2(i) for the case of binary Y_i , we can apply the martingale central limit theorem to establish the asymptotic normality of $\hat{\Delta}_n - \Delta_n$. In many applications, one can make use of the bucket structure of the type in Section 3.4 to estimate the asymptotic variance of $\hat{\Delta}_n$. In particular, this structure is inherent in dynamic panel data in econometrics and longitudinal data in epidemiology, which is beyond the scope of this paper on forecasting probabilities of events and will be treated elsewhere. Note that the bucket structure is only used in estimating the asymptotic variance of $\hat{\Delta}_n$ by (6.2), and that Theorem 5 and its extension outlined in Section 6.1 imply that the variance estimate tends to be conservative if the assumed bucket structure actually fails to hold.

7. Discussion. The average score $n^{-1} \sum_{i=1}^n L(Y_i, \hat{p}_i)$ measures the divergence of the predicted probabilities \hat{p}_i , which lie between 0 and 1, from the indicator variables Y_i that can only have values 0 or 1. As noted by [Lichtendahl and Winkler \(2007\)](#), this tends to encourage more aggressive bets on the binary outcomes, rather than the forecaster's estimates of the event probabilities. For example, an estimate of 95% probability may lead to a probability forecast of 100% for a higher reward associated with the indicator variable Y_i ; see also [Mason \(2008\)](#), who gives an example in which a forecaster is encouraged to give such "dishonest" forecasts. This difficulty would disappear if one uses L to compare \hat{p}_i with the actual p_i , rather than with the Bernoulli(p_i) random variable Y_i . Because the p_i are unknown, this is not feasible and the importance of using a proper score $L(Y_i, \hat{p}_i)$ to evaluate a probability forecast has been emphasized to address the issue of dishonest forecasts. In Section 3.2 we have shown that it is possible to use $L(p_i, \hat{p}_i) - L(p_i, \hat{p}'_i)$ for comparing two forecasters and to construct confidence intervals of the average loss difference. A key idea underlying this development is the linear equivalent of a loss function introduced in Section 2. [Schervish \(1989\)](#), Section 3, has used a framework of two-decision problems involving these loss functions to develop a method for comparing forecasters. Our approach that considers $L(p_i, \hat{p}'_i) - L(p_i, \hat{p}''_i)$ can be regarded as a further step in this direction.

As noted in Section 3.5, an important assumption underlying statistical inference in the verification of probability forecasts in meteorology is that the forecast-observation pairs are independent realizations from the joint distribution of forecasts and observations. Although [Mason \[\(2008\), page 32\]](#) has pointed out that this assumption cannot hold "if the verification score is calculated using forecasts for different locations, or if both the forecasts and observations are not independent temporally," not much has been done to address this problem other than using moving-blocks bootstrap [[Mason \(2008\)](#), [Wilks \(2005\)](#)] because traditional statistical inference does not seem to provide much help in tackling more general forecast-observation pairs. The new approach in Section 3 can be used to resolve this difficulty. It uses martingale theory to allow the forecast-observation pairs to be generated by general stochastic systems, without the need to model the underlying system in carrying out the inference. The treatment of spatial dependence is also covered in Section 3.5, in which dependence of the events at K_t locations at time t is encapsulated in the highly complex joint distribution of their generating probabilities $p_{t,1}, \dots, p_{t,K_t}$, which our approach does not need to model in performing inference on forecast validation. Our viewpoint in forecast evaluation is that one should try not to make unnecessary or arbitrary assumptions on the underlying data-generating mechanism, especially in regulatory settings such as regulatory supervision of a bank's internal ratings models of loan default probabilities; see Section 3.4 and [Lai and Wong \(2008\)](#). A convenient but incorrect data-generating model that is assumed can unduly bias the comparison.

Acknowledgment. The author thanks the Epidemiology unit at INSERM, France, and the Statistics Department at Tel Aviv University for their generous hospitality while working on revisions of this paper.

REFERENCES

- ARVESEN, J. N. (1969). Jackknifing U -statistics. *Ann. Math. Statist.* **40** 2076–2100. [MR0264805](#)
- BASEL COMMITTEE ON BANKING SUPERVISION (2006). Basel II: International convergence of capital measurement and capital standards: A revised framework. Available at <http://www.bis.org/publ/bcbs128.htm>.
- BRIER, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review* **78** 1–3.
- BRÖCKER, J. and SMITH, L. A. (2007). Increasing the reliability of reliability diagrams. *Weather and Forecasting* **22** 651–661.
- COX, D. R. (1958). Two further applications of a model for binary regression. *Biometrika* **45** 562–565.
- DAWID, A. P. (1982). The well-calibrated Bayesian. *J. Amer. Statist. Assoc.* **77** 605–613. [MR0675887](#)
- DE FINETTI, B. (1975). *Theory of Probability: A Critical Introductory Treatment*. Vol. 2. Wiley, London. Translated from the Italian by Antonio Machi and Adrian Smith. [MR0440641](#)
- DEGROOT, M. H. and FIENBERG, S. E. (1983). The comparison and evaluation of forecasters. *Statistician* **32** 12–22.
- DIEBOLD, F. X. and MARIANO, R. S. (1995). Comparing predictive accuracy. *J. Bus. Econom. Statist.* **13** 253–263.
- FOX, C. R. and BIRKE, R. (2002). Forecasting trial outcomes: Lawyers assign higher probability to possibilities that are described in greater detail. *Law Hum. Behav.* **26** 159–173.
- GIACOMINI, R. and WHITE, H. (2006). Tests of conditional predictive ability. *Econometrica* **74** 1545–1578. [MR2268409](#)
- GNEITING, T., BALABDAOUI, F. and RAFTERY, A. E. (2007). Probabilistic forecasts, calibration and sharpness. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **69** 243–268. [MR2325275](#)
- GNEITING, T. and RAFTERY, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* **102** 359–378. [MR2345548](#)
- GOOD, I. J. (1952). Rational decisions. *J. Roy. Statist. Soc. Ser. B* **14** 107–114. [MR0077033](#)
- GRÜNWARD, P. D. and DAWID, A. P. (2004). Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Ann. Statist.* **32** 1367–1433. [MR2089128](#)
- HARI, P. N., ZHANG, M.-J., ROY, V., PÉREZ, W. S., BASHEY, A., TO, L. B., ELFENBEIN, G., FREYTES, C. O., GALE, R. P., GIBSON, J., KYLE, R. A., LAZARUS, H. M., MCCARTHY, P. L., MILONE, G. A., PAVLOVSKY, S., REECE, D. E., SCHILLER, G., VELA-OJEDA, J., WEISDORF, D. and VESOLE, D. (2009). Is the international staging system superior to the Durie–Salmon staging system? A comparison in multiple myeloma patients undergoing autologous transplant. *Leukemia* **23** 1528–1534.
- LAI, T. L. and WONG, S. P.-S. (2008). Statistical models for the Basel II internal ratings-based approach to measuring credit risk of retail products. *Stat. Interface* **1** 229–241. [MR2476741](#)
- LICHTENDAHL, K. C. JR. and WINKLER, R. L. (2007). Probability elicitation, scoring rules, and competition among forecasters. *Management Sci.* **53** 1745–1755.
- MASON, S. J. (2008). Understanding forecast verification statistics. *Meteorol. Appl.* **15** 31–40.
- MURPHY, A. H. and WINKLER, R. L. (1984). Probability forecasting in meteorology. *J. Amer. Statist. Assoc.* **79** 489–500.
- RANJAN, R. and GNEITING, T. (2010). Combining probability forecasts. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **72** 71–91. [MR2751244](#)

- REDELMEIER, D. A., BLOCH, D. A. and HICKAM, D. H. (1991). Assessing predictive accuracy: How to compare Brier scores. *J. Clin. Epidemiol.* **44** 1141–1146.
- SCHERVISH, M. J. (1989). A general method for comparing probability assessors. *Ann. Statist.* **17** 1856–1879. [MR1026316](#)
- SEILLIER-MOISEWITSCH, F. and DAWID, A. P. (1993). On testing the validity of sequential probability forecasts. *J. Amer. Statist. Assoc.* **88** 355–359. [MR1212496](#)
- SPIEGELHALTER, D. J. (1986). Probabilistic prediction in patient management and clinical trials. *Stat. Med.* **5** 421–433.
- WEST, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica* **64** 1067–1084. [MR1403232](#)
- WILKS, D. (2005). *Statistical Methods in the Atmospheric Sciences*, 2nd ed. *International Geophysics* **91**. Academic Press, New York.
- WILLIAMS, D. (1991). *Probability with Martingales*. Cambridge Univ. Press, Cambridge. [MR1155402](#)
- WINKLER, R. L. (1994). Evaluating probabilities: Asymmetric scoring rules. *Management Sci.* **40** 1395–1405.

T. L. LAI
DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
SEQUOIA HALL, 390 SERRA MALL
STANFORD, CALIFORNIA 94305-4065
USA
E-MAIL: lait@stanford.edu

S. T. GROSS
ZICKLIN SCHOOL OF BUSINESS
BARUCH COLLEGE
CITY UNIVERSITY OF NEW YORK
BOX 11-220, 1 BERNARD M. BARUCH WAY
NEW YORK, NEW YORK 10010
USA
E-MAIL: Shulamith.Gross@baruch.cuny.edu

D. B. SHEN
UBS
677 WASHINGTON BOULEVARD
STAMFORD, CONNECTICUT 06901
USA
E-MAIL: happybo7@gmail.com