

Bayesian Ensembles of Binary-Event Forecasts: When Is It Appropriate to Extremize or Anti-Extremize?

Kenneth C. Lichtendahl Jr.
Victor Richmond R. Jose

Yael Grushka-Cockayne
Robert L. Winkler

Working Paper 19-041



Bayesian Ensembles of Binary-Event Forecasts: When Is It Appropriate to Extremize or Anti-Extremize?

Kenneth C. Lichtendahl Jr.
University of Virginia

Yael Grushka-Cockayne
Harvard Business School

Victor Richmond R. Jose
Georgetown University

Robert L. Winkler
Duke University

Working Paper 19-041

Copyright © 2018 by Kenneth C. Lichtendahl Jr., Yael Grushka-Cockayne, Victor Richmond R. Jose, and Robert L. Winkler

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Bayesian Ensembles of Binary-Event Forecasts: When Is It Appropriate to Extremize or Anti-Extremize?

Kenneth C. Lichtendahl Jr.

Darden School of Business, University of Virginia, Charlottesville, VA 22903, lichtendahlc@darden.virginia.edu

Yael Grushka-Cockayne

Darden School of Business, University of Virginia, Charlottesville, VA 22903, grushkay@darden.virginia.edu

Victor Richmond R. Jose

McDonough School of Business, Georgetown University, Washington, DC 20057, vrj2@georgetown.edu

Robert L. Winkler

The Fuqua School of Business, Duke University, Durham, NC 27708, rwinkler@duke.edu

Many organizations face critical decisions that rely on forecasts of binary events. In these situations, organizations often gather forecasts from multiple experts or models and average those forecasts to produce a single aggregate forecast. Because the average forecast is known to be underconfident, methods have been proposed that create an aggregate forecast more extreme than the average forecast. But is it always appropriate to extremize the average forecast? And if not, when is it appropriate to anti-extremize (i.e., to make the aggregate forecast less extreme)? To answer these questions, we introduce a class of optimal aggregators. These aggregators are Bayesian ensembles because they follow from a Bayesian model of the underlying information experts have. Each ensemble is a generalized additive model of experts' probabilities that first transforms the experts' probabilities into their corresponding information states, then linearly combines these information states, and finally transforms the combined information states back into the probability space. Analytically, we find that these optimal aggregators do not always extremize the average forecast, and when they do, they can run counter to existing methods. On two publicly available datasets, we demonstrate that these new ensembles are easily fit to real forecast data and are more accurate than existing methods.

Key words: Forecast aggregation; linear opinion pool; generalized additive model; generalized linear model; stacking.

Date: March 27, 2018

1. Introduction

Many organizations face forecasting challenges that involve binary events. These forecasts are critical to decisions such as the approval of credit (probability of default), the target of a marketing campaign (probability of a click), the recommendation of a drug (probability of having a disease), and the choice of a national security response (probability of a geopolitical event occurring). Often crowds of experts or models issue probabilities for such events. To aggregate the individual probabilities, we offer a new method based on Bayesian principles. This method generalizes several

proposed aggregators in the literature, exhibits some structural advantages, and may be more accurate in practice.

Since its introduction by Stone (1961), many researchers have found the linear opinion pool to be an attractive way to aggregate forecasts (DeGroot 1974, McConway 1981, Genest and Zidek 1986, DeGroot and Mortera 1991). The most popular way to aggregate expert forecasts is to take the average of the experts' forecasts, which is a linear opinion pool with equal weights (Clemen and Winkler 1986, Larrick and Soll 2006). In aggregating binary-events forecasts, Winkler and Poses (1993, p. 1533) state, "Simple averages of forecasts seem to work as well as or better than fancier combining methods."

Nonetheless, Hora (2004) and Ranjan and Gneiting (2010) show that the linear opinion pool is underconfident. In the context of binary events, the linear opinion pool is, on average, not extreme enough. To address this problem, Ranjan and Gneiting (2010) propose a method that extremizes the linear opinion pool by pushing it closer to its nearer extreme. Many others have employed schemes to extremize the average forecast (Karmarkar 1978, Erev et al. 1994, Ariely et al. 2000, Shlomi and Wallsten 2010, Turner et al. 2014, Mellers et al. 2014, Baron et al. 2014, Satopää et al. 2014). Extremizing methods, such as the logit aggregator, are now used as benchmarks in practice (IARPA Geopolitical Forecasting Challenge 2018). Baron et al. (2014, p. 134) note that "If every forecaster said 0.6, and they were using different information, then someone who knew all of this would have a right to much higher confidence." This updating process is consistent with Bayesian principles under certain assumptions and is the main justification for extremizing the average forecast.

These existing methods, however, are heuristics that do not follow from Bayesian principles. Consequently, they risk being sub-optimal. The idea that a decision maker can use Bayesian reasoning to aggregate experts' forecasts goes back at least to Winkler (1968) and Morris (1974). "To the expert, the probability assessment is a representation of his state of information; to the decision maker, the probability assessment is information." (Morris 1974, p. 1241) Many other researchers have proposed models along these lines (Lindley et al. 1979, French 1980, Clemen 1987). Dawid et al. (1995) make an important contribution to the literature on aggregating binary-event forecasts. They introduce the first set of aggregation methods based on Bayesian principles and a fundamental condition regarding calibration. The condition says that if the decision maker hears from only one calibrated expert, he adopts that expert's forecast as his own. Recently, Satopää et al. (2016), building on the work of Dawid et al. (1995) and generalizing an example in Ranjan and Gneiting

(2010), introduce a probit aggregator that is consistent with Bayesian principles. We call such an aggregator a Bayesian ensemble in the spirit of Chipman et al. (2007).

In this paper, we significantly enlarge the class of Bayesian ensembles. This larger class has three main structural properties, the combination of which offers advantages over previous methods. First, our Bayesian ensembles incorporate the prior-predictive probability into the aggregate forecast. One can think of the prior-predictive probability as the base rate known to the experts and the decision maker alike. For example, when forecasting rain in Phoenix, Arizona, all parties involved in the aggregation may agree on the base rate at which rain occurs daily, say 10%. Second, our Bayesian ensembles are generalized linear models of experts' transformed probabilities where the decision maker can choose a link function beyond those that accompany the logit and probit models. With a customized link function, a generalized linear model can often fit data better. Third, because our Bayesian ensembles are generalized linear models of experts' transformed probabilities, they can be easily and quickly estimated on large datasets using statistical computing software.

With this enlarged class of Bayesian ensembles, we provide new insights into why, when, and how to extremize the average forecast. In this class, it often makes sense to extremize the average forecast, but sometimes it does not. In some cases, it can make sense to anti-extremize, or make the aggregate forecast less extreme than the average forecast. Three different types of information play a crucial role in our understanding of extremizing/anti-extremizing: (i) the prior information known to the decision maker and experts, (ii) each expert's private information, and (iii) the shared information known to all the experts but not the decision maker. When the experts rely on private information only, the decision maker naturally wants to form an aggregate forecast that extremizes the average forecast because each individual forecast contains some weight on the prior information. For example, with two experts, the average forecast double counts the prior information and is pulled too far in the direction of the prior-predictive probability.

When the experts rely on the same amount of information but some of that information is shared information, the decision maker does not have as much new information in the reports to justify the same move away from the prior-predictive probability. Hence, shared information tends to reduce the degree of extremizing. In fact, the presence of shared information can cause a Bayesian ensemble to attenuate the average forecast. In practice, shared information is common, since real-world experts often use similar models or have similar training (Kim et al. 2001, Chen et al. 2004, Marinovic et al. 2012). If models' forecasts are being combined, the models often pick up on similar features from the training set.

In the next two sections, we introduce two types of Bayesian ensembles. The first type, developed in Section 2, involves the use of conjugate pairs of distributions. The conjugate pairs help build intuition for how a Bayesian ensemble emerges from the information experts have. The conjugate-pair Bayesian ensembles also help motivate the assumptions behind our second type of Bayesian ensemble, called the generalized probit ensemble. In Section 3, we provide the form of the generalized probit ensemble and describe how it can be easily fit to data using a generalized linear model with a custom link function. In Section 4, we present two pilot studies where we fit the generalized probit ensemble to large sets of real forecast data. In the first study the goal is to predict loan defaults in Fannie Mae’s single-family loan performance data. In the second study, the task is to predict defective used cars in Kaggle’s *Don’t Get Kicked!* competition. The forecasts our ensemble and benchmark methods aggregate come from several leading statistical and machine learning algorithms. The results of our pilot studies suggest that the generalized probit ensemble may be more accurate than several leading aggregation methods in other applications.

2. Conjugate-Pair Bayesian Ensembles

In this section, we introduce our first type of Bayesian ensemble. Before introducing this ensemble, we provide a formal definition of extremizing—an important definition that will be used throughout the paper.

2.1. Extremizing

Our definition of extremizing is inspired by the definitions of sharpness in Winkler and Jose (2008) and Ranjan and Gneiting (2010). Sharper, or more extreme forecasts, are those farther away from their marginal event frequencies. Winkler and Jose (2008) state, “If climatology c is used as the baseline probability for probability of precipitation forecasts, sharpness should be viewed in terms of shifts from c toward zero or one instead of shifts from 0.5 toward zero or one.” In forecasting rain in Phoenix, Arizona where the historical daily frequency of rain is about 10%, a forecast of 40% would naturally be considered more extreme than a forecast of 30%.

For the following definition, we assume the average forecast is not equal to either the prior-predictive probability (i.e., the forecast based on the prior information only) or the aggregate forecast.

DEFINITION 1 (EXTREMIZING). *The aggregate forecast extremizes the average forecast if it is farther away from the prior-predictive probability in the same direction as the average forecast. Otherwise, the aggregate forecast anti-extremizes the average forecast.*

This definition differs from other definitions in the literature. For example, Baron et al. 2014, Satopää et al. (2014), and Satopää et al. (2016) say the aggregate forecast extremizes the average forecast if it is farther away from one-half in the same direction as the average forecast. Under this definition, a forecast of 40% for rain in Phoenix would be considered less extreme than a forecast of 30%.

2.2. Information Structure

Below we introduce an information structure that describes the information the decision maker and experts have in order to make their forecasts. In this setting, each of $k \geq 2$ experts issues a probability forecast for the binary event y given her private sample information \mathbf{x}_i and the shared sample information \mathbf{x}_s . The event y equals 1 if the event occurs, or it equals 0 if the event does not occur. We denote expert i 's forecast $P(y = 1 | \mathbf{x}_i, \mathbf{x}_s)$ by p_i . After hearing from his $k \geq 2$ experts, the decision maker aggregates the experts' forecasts into a single forecast.

A decision maker's aggregate forecast is a Bayesian ensemble, denoted by \hat{p} , if it is the posterior-predictive probability $P(y = 1 | p_1, \dots, p_k)$ derived from the joint distribution of (p_1, \dots, p_k, y) using Bayes' Theorem. This aggregate forecast is optimal because using any other forecast would yield a worse score when evaluated by a proper scoring rule (Gneiting and Raftery 2007). Later in this section we compare several Bayesian ensembles to the average forecast, $\bar{p} = (1/k) \sum_{i=1}^k p_i$, to see when these ensembles extremize (or anti-extremize) the average forecast.

Suppose there is an exchangeable sequence of data points:

$$\left(\underbrace{x_1, \dots, x_{n_1}}_{\substack{\text{Expert 1's} \\ \text{private sample} \\ \text{of size } n_1}}, \underbrace{x_{n_1+1}, \dots, x_{n_1+n_2}}_{\substack{\text{Expert 2's} \\ \text{private sample} \\ \text{of size } n_2}}, \dots, \underbrace{x_{N_{k-1}+1}, \dots, x_{N_k}}_{\substack{\text{Expert } k\text{'s} \\ \text{private sample} \\ \text{of size } n_k}}, \underbrace{x_{N_k+1}, \dots, x_{N_k+n_s}}_{\substack{\text{Shared} \\ \text{sample} \\ \text{of size } n_s}}, \underbrace{x_{N_k+n_s+1}}_{\substack{\text{Related to} \\ \text{the binary} \\ \text{event } y}} \right). \quad (1)$$

Expert i sees the private sample $\mathbf{x}_i = (x_{N_{i-1}+1}, \dots, x_{N_i})$ of size n_i for $i = 1, \dots, k$, where $N_i = \sum_{l=1}^i n_l$. The shared information—known to all the experts, but not the decision maker—is the sample $\mathbf{x}_s = (x_{N_k+1}, \dots, x_{N_k+n_s})$ of size n_s .

The final data point $x_{N_k+n_s+1}$, abbreviated as x , is related to the binary event y , which is what the decision maker ultimately cares about. If x is in the event occurrence set A , then y equals 1; otherwise, y equals 0. For example, if x is a Bernoulli random variable, the set A might simply be $\{1\}$. Alternatively, if x is a normal random variable, the set A might be the interval $(0, \infty)$.

Data points in the sequence are independent and identically distributed according to a likelihood from a regular, one-parameter exponential family with probability mass or density function $f(x_j | \theta) = a(x_j)b(\theta) \exp(c(\theta)h(x_j))$. The parameter θ is distributed according to a conjugate prior $f(\theta) = [K(\tau_0, \tau_1)]^{-1} [b(\theta)]^{\tau_0} \exp(c(\theta)\tau_1)$, where τ_0 and τ_1 are the prior's hyperparameters

and $K(\tau_0, \tau_1)$ is its normalizing constant. The joint distribution of $(\theta, \mathbf{x}_1, \dots, \mathbf{x}_k, \mathbf{x}_s, x)$ and the event occurrence set A are common knowledge among the decision maker and the experts. The prior/likelihood pair of distributions that describes this joint distribution is called a conjugate pair (Raiffa and Schlaifer 1961, Bernardo and Smith 2000). Hence we call the Bayesian ensemble that a conjugate pair generates a conjugate-pair Bayesian ensemble.

Based on these assumptions, the following function generates a set of predictive distributions—one for the prior-predictive, one for each expert’s posterior-predictive, and one for the decision maker’s posterior-predictive. We call this function the predictive generating function:

$$F_n(t) = \int_{x \in A} a(x) \frac{K(\tau_0 + n + 1, t + h(x))}{K(\tau_0 + n, t)} dx, \quad (2)$$

where n is the relevant sample size and t is the relevant sufficient statistic (Bernardo and Smith 2000). Note that when the random variable x is discrete, the integral in (2), and other integrals like it throughout the paper, naturally become sums. With $n = 0$ and $t = \tau_1$, $F_0(\tau_1)$ is everyone’s prior-predictive probability $P(y = 1)$, denoted by p_0 . With $n = n_i + n_s$ and $t = t_i + t_s$ where $t_i = \sum_{j=N_{i-1}+1}^{N_i} h(x_j)$ and $t_s = \sum_{j=N_k+1}^{N_k+n_s} h(x_j)$, $F_{n_i+n_s}(\tau_1 + t_i + t_s)$ is expert i ’s posterior-predictive probability $P(y = 1 | \mathbf{x}_i, \mathbf{x}_s)$.

With $n = N_k + n_s$ and $t = \sum_{i=1}^k t_i + t_s$, $F_{N_k+n_s}(\tau_1 + \sum_{i=1}^k t_i + t_s)$ is the decision maker’s posterior-predictive probability $P(y = 1 | \mathbf{x}_1, \dots, \mathbf{x}_k, \mathbf{x}_s)$, as if he had access to all the experts’ private and shared information. In reality, the decision maker only hears p_i from expert i , so the best he can do is infer (t_1, \dots, t_k, t_s) , the sufficient statistics for $(\mathbf{x}_1, \dots, \mathbf{x}_k, \mathbf{x}_s)$, from $(t_1 + t_s, \dots, t_k + t_s)$. He can learn each $t_i + t_s$ from p_i if F_n is invertible. If it is invertible, then $t_1 + t_s = F_{n_i+n_s}^{-1}(p_i) - \tau_1$.

2.3. Conjugate Pairs with Private Information Only

In the case of private information only ($n_s = 0$), the conjugate-pair Bayesian ensemble is a generalized additive model in the experts’ probabilities and a generalized linear model in the experts’ sufficient statistics. For this result, we need the following two definitions. A generalized additive model links the conditional expectation of a quantity of interest y to an additive function of some covariates (q_1, \dots, q_k) : $E[y | q_1, \dots, q_k] = g^{-1}(g_0 + g_1(q_1) + \dots + g_k(q_k))$ where g is the link function, g_0 is a constant, and each g_i for $i = 1, \dots, k$ is a smooth function (Hastie and Tibshirani 1986). A generalized linear model is a general additive model where each g_i is a linear function (Nelder and Wedderburn 1972). Proofs of this and other results appear in the Appendix.

PROPOSITION 1 (Private Information Only). *Assume only private information is available to the experts ($n_i > 0$ for $i = 1, \dots, k$ and $n_s = 0$) and the predictive generating function $F_n(t)$ in (2)*

is strictly monotonic in t . Then the conjugate-pair Bayesian ensemble of the experts' probabilities is a generalized additive model:

$$\hat{p} = P(y = 1 | p_1, \dots, p_k) = F_{N_k} \left(- (k-1) F_0^{-1}(p_0) + \sum_{i=1}^k F_{n_i}^{-1}(p_i) \right), \quad (3)$$

where $p_0 = F_0(\tau_1)$, $p_i = F_{n_i}(\tau_1 + t_i)$, and $t_i = \sum_{j=N_{i-1}+1}^{N_i} h(x_j)$. Also, the Bayesian ensemble of the experts' sufficient statistics is a generalized linear model:

$$P(y = 1 | t_1, \dots, t_k) = F_{N_k} \left(\tau_1 + \sum_{i=1}^k t_i \right). \quad (4)$$

This result provides a large class of Bayesian ensembles. The class is as large as the class of regular, one-parameter exponential families. Any ensemble in this class is a generalized additive model of experts' probabilities that first transforms the experts' probabilities into their corresponding information states, then linearly combines these information states, and finally transforms the combined information states back into the probability space.

Below we provide four examples of conjugate-pair Bayesian ensembles. Our first example appears in Dawid et al. (1995), and our third example is a variant of the models studied in Ranjan and Gneiting (2010) and Satopää et al. (2016). The other two examples are new. For details on the derivation of each example's ensemble, e.g., each example's τ_1 , $h(x_j)$, and $F_n(t)$, see the Appendix.

EXAMPLE 1 (BETA/BERNOULLI PAIR). Let x_j given θ be drawn from a Bernoulli distribution with probability θ . The conjugate prior for this likelihood is the beta distribution with shape parameters α and β . Suppose $A = \{1\}$ corresponds to the event that a future borrower defaults on a loan. With the event occurrence set $A = \{1\}$, this conjugate pair leads to the Bayesian ensemble

$$\hat{p} = \left(1 - \sum_{i=1}^k w_{n_i} \right) p_0 + \sum_{i=1}^k w_{n_i} p_i,$$

where $w_{n_i} = (\alpha + \beta + n_i) / (\alpha + \beta + N_k)$. \square

EXAMPLE 2 (GAMMA/POISSON PAIR). Let x_j given θ be drawn from a Poisson distribution with rate θ . The conjugate prior for this likelihood is the gamma distribution with shape α and rate β , denoted by $\theta \sim Ga(\alpha, \beta)$. Suppose $A = \{0\}$ corresponds to the event that a piece of equipment, with exponentially distributed interarrival times of breakdowns, does not break down in the next year. With $A = \{0\}$, this conjugate pair leads to the Bayesian ensemble

$$\hat{p} = \exp \left(- (k-1) \frac{v_{N_k}}{v_0} \log(p_0) + \sum_{i=1}^k \frac{v_{N_k}}{v_{n_i}} \log(p_i) \right),$$

where $v_n = \log((\beta + n) / (\beta + n + 1))$. \square

EXAMPLE 3 (NORMAL/NORMAL PAIR). Let x_j given θ be drawn from a normal distribution with mean θ and variance σ^2 . The conjugate prior for this likelihood is another normal distribution with mean θ_0 and variance σ_0^2 . Suppose $A = (0, \infty)$ corresponds to the event that a new product makes a profit in its first year. With $A = (0, \infty)$, this conjugate pair leads to the Bayesian ensemble

$$\hat{p} = \Phi \left(- (k-1) \sqrt{\frac{v_0}{v_{N_k}}} \Phi^{-1}(p_0) + \sum_{i=1}^k \sqrt{\frac{v_{n_i}}{v_{N_k}}} \Phi^{-1}(p_i) \right),$$

where Φ is the cumulative distribution function (cdf) of the standard normal distribution and $v_n = (\sigma^2/\sigma_0^2 + n)(\sigma^2/\sigma_0^2 + n + 1)\sigma^2$. We call this ensemble a probit ensemble because the inverse link function in this generalized linear model is the standard normal cdf. \square

EXAMPLE 4 (GENERALIZED-GAMMA/GUMBEL PAIR). Let x_j given θ be drawn from a Gumbel distribution with location θ and scale σ . The conjugate prior for this likelihood is the reflection of the generalized gamma distribution in Ahuja and Nash (1967, Equation 2.7): $\exp(\theta/\sigma) \sim Ga(\alpha, \beta)$. Suppose $A = (-\infty, 0)$ corresponds to the event that a hedge-fund manager's best investment makes a loss in some year. With $A = (-\infty, 0)$, this conjugate pair leads to the Bayesian ensemble

$$\hat{p} = \left(\frac{-(k-1)p_0^{v_0}/(1-p_0^{v_0}) + \sum_{i=1}^k p_i^{v_{n_i}}/(1-p_i^{v_{n_i}})}{1 - (k-1)p_0^{v_0}/(1-p_0^{v_0}) + \sum_{i=1}^k p_i^{v_{n_i}}/(1-p_i^{v_{n_i}})} \right)^{1/v_{N_k}},$$

where $v_n = 1/(\alpha + n)$. \square


The ensembles in Examples 1 and 3 are depicted in Figure 1a and 1c, respectively. In both figures, we hold p_1 fixed and vary p_2 . For Figure 1a, we assume $\alpha = \beta = 1$, two experts each privately see two data points, and expert 1 reports $p_1 = 3/4$. For Figure 1c, we assume $\sigma = \sigma_0 = 1$, $\theta_0 = -1.25$, two experts each privately see two data points, and expert 1 reports $p_1 = F_{n_1}(\tau_1) = F_2(-1.25) \approx 0.36$.

With exchangeable experts (i.e., each $n_i = n_1$), the following strict inequality holds for the beta/Bernoulli ensemble in Example 1:


$$\hat{p} = (1 - kw_{n_1})p_0 + \sum_{i=1}^k w_{n_1}p_i = (1 - kw_{n_1})p_0 + kw_{n_1}\bar{p} < (1 - kw_{n_1})\bar{p} + kw_{n_1}\bar{p} = \bar{p}, \quad (5)$$

for $\bar{p} < p_0$. Similarly, for $\bar{p} > p_0$, the inequality in (5) is reversed. Thus, with private information only, the beta/Bernoulli ensemble always extremizes the average forecast. The probit ensemble, however, sometimes does not. Its region of anti-extremizing (in gray) is sizeable.


The intuition for why a Bayesian ensemble tends to extremize the average forecast is that the experts' forecasts each have some weight on the prior information, but when their forecasts are combined, the aggregate forecast only needs that weight once. This intuition explains the coefficient $-(k-1)$ in front of the term $F_0^{-1}(p_0)$ in (3). In other words, extremizing tends to be the result




(a) Beta/Bernoulli (Private Only).



(b) Beta/Bernoulli (Private and Shared).



(c) Normal/Normal (Private Only).



(d) Normal/Normal (Private and Shared).

Figure 1 Conjugate-Pair Bayesian Ensembles from Examples 1 and 5 (Beta/Bernoulli) and Examples 3 and 6 (Normal/Normal).

of removing the redundant weight on the prior information. Because the experts' probabilities are first non-linearly transformed inside the probit ensemble, it does not always extremize the average forecast. In the next subsection, we compare these two ensembles to their shared information counterparts to see what effect this information has on extremizing/anti-extremizing.

2.4. Conjugate Pairs with Private and Shared Information

The following result generalizes Proposition 1 to include the presence of shared information.

PROPOSITION 2 (Private and Shared Information). *Assume private and shared information are available to the experts ($n_i > 0$ for $i = 1, \dots, k$ and $n_s > 0$) and the predictive generating function $F_n(t)$ in (2) is strictly monotonic in t . Then the conjugate-pair Bayesian ensemble of the experts' probabilities is given by*

$$\hat{p} = \int F_{N_k+n_s} \left(- (k-1)F_0^{-1}(p_0) + \sum_{i=1}^k F_{n_i+n_s}^{-1}(p_i) - (k-1)t_s \right) f(t_s|p_1, \dots, p_k) dt_s, \quad (6)$$

where $p_0 = F_0(\tau_1)$, $p_i = F_{n_i+n_s}(\tau_1 + t_i + t_s)$, and $f(t_s|p_1, \dots, p_k)$ is the probability mass or density function of the shared sample's sufficient statistic t_s conditional on the experts' reported probabilities.

In the presence of shared information, the decision maker can, at most, deduce each expert's sufficient statistic $t_i + t_s$ from p_i : $t_i + t_s = F_{n_i+n_s}^{-1}(p_i) - \tau_1$. So, in the integral (or sum) in (6), $f(t_s|p_1, \dots, p_k)$ is equal to $f(t_s|t_1 + t_s, \dots, t_k + t_s)$, which is not always tractable. The example below illustrates a situation where this conditional distribution is quite simple to evaluate.

EXAMPLE 5 (BETA/BERNOULLI PAIR WITH PRIVATE AND SHARED INFORMATION). Assume $\alpha = \beta = 1$, x_1 and x_2 are private information seen by experts 1 and 2, respectively, x_3 is shared information, and x_4 is related to the binary event y . The binary event y (default or not) is 1 if $x_4 = 1$ and is 0 otherwise. The prior-predictive probability $p_0 = 1/2$. Suppose expert 1 reports $p_1 = 3/4$. Given this report, the decision maker can deduce that $x_1 + x_3 = 2$ (or $x_1 = 1$ and $x_3 = 1$): he knows expert 1 saw two defaults in her private and shared information. Based on expert 1's report, the decision maker also knows that $x_2 + x_3$ can only be either 1 or 2, leading to either $p_2 = 1/4$ or $p_2 = 1/2$. In either case, $f(t_s|p_1, \dots, p_k)$ is a point mass because $P(x_3 = 1|x_1 + x_3 = 2, x_2 + x_3 = 1) = P(x_3 = 1|x_1 + x_3 = 2, x_2 + x_3 = 2) = 1$. Consequently, in this example, the beta/Bernoulli Bayesian ensemble with private and shared information becomes

$$\hat{p} = P(y = 1|p_1 = 3/4, p_2) = \frac{4p_2 + 1}{5}. \quad \square \quad (7)$$

In Figure 1, we compare Example 5's ensemble with private and shared information to Example 1's ensemble with private information only. In both examples, the experts see two data points the decision maker does not see. Turning one piece of private information into shared information, as we move from Figure 1a to Figure 1b, we see a reduction in the degree of extremizing. In Figure 1b, we also see a point of anti-extremizing (in gray) when $p_2 = 0.5$.

The conditional distribution $f(t_s|p_1, \dots, p_k)$ is not always the point mass we see in Example 5. If p_1 were $1/2$ there, then the decision maker could be uncertain about the event $x_3 = 1$ and the

Bayesian ensemble would be a mixture. In this case, the sum in (6) would contain two terms when expert 2 reports 1/2, rather than the single term in (7). In general, this integral (or sum) is difficult to evaluate. For the normal-normal pair though, the Bayesian ensemble is not a mixture and has a tractable form.

PROPOSITION 3 (Private and Shared Information: Normal/Normal Pair). *Assume private and shared information are available to the experts ($n_i > 0$ for $i = 1, \dots, k$ and $n_s > 0$) and the conditions of Example 3's normal/normal conjugate pair hold. Then the normal/normal conjugate-pair leads to the Bayesian ensemble of the experts' probabilities:*

$$\hat{p} = \Phi \left(\frac{-(k-1)\sqrt{v_0}\Phi^{-1}(p_0) + \sum_{i=1}^k \sqrt{v_{n_i+n_s}}\Phi^{-1}(p_i) - (k-1)E[t_s|p_1, \dots, p_k]}{\sqrt{(k-1)^2 \text{Var}[t_s|p_1, \dots, p_k] + v_{N_k+n_s}}} \right), \quad (8)$$

where $v_n = (\sigma^2/\sigma_0^2 + n)(\sigma^2/\sigma_0^2 + n + 1)\sigma^2$. The conditional mean $E[t_s|p_1, \dots, p_k]$ is linear in $(\Phi^{-1}(p_1), \dots, \Phi^{-1}(p_k))$, and the conditional variance $\text{Var}[t_s|p_1, \dots, p_k]$ is constant in (p_1, \dots, p_k) . (The conditional moments are given in the proof of this result.)

Similar to the probit ensemble with private information only, the ensemble in the result above is a linear combination of $(\Phi^{-1}(p_0), \Phi^{-1}(p_1), \dots, \Phi^{-1}(p_k))$ inside the standard normal cdf. Therefore, we call the ensemble in (8) a probit ensemble as well. The two probit ensembles differ only in the weights in their respective linear combinations. These weights determine the degree of extremizing, as the next example illustrates.

EXAMPLE 6 (NORMAL/NORMAL PAIR WITH PRIVATE AND SHARED INFORMATION). Assume $\theta_0 = -1.25$ and $\sigma_0 = \sigma = 1$, x_1 and x_2 are private information seen by experts 1 and 2, respectively, x_3 is shared information, and x_4 is related to the binary event y . The binary event y is 1 if $x_4 > 0$ and is 0 otherwise. Suppose expert 1 reports $p_1 = F_{n_1+n_s}(\tau_1) = F_2(-1.25) \approx 0.36$. Then we have the Bayesian ensemble depicted in Figure 1d for all possible reports from the second expert. Compare this ensemble with private and shared information to Example 3's ensemble with private information only (Figure 1c). The presence of shared information in Figure 1d again reduces the degree of extremizing we see in Figure 1c. The region of anti-extremizing enlarges as one piece of private information becomes shared information. \square

For other settings of the conjugate-pair ensembles above, we get similar results. The ensembles tend to extremize the average forecast, although not always, even when there is no shared information. As more of the information is shared, the degree of extremizing reduces and the region of anti-extremizing enlarges. Anti-extremizing emerges (or increases) because of the reduction in the overall sample size the decision maker uses to aggregate the experts' forecast, as the experts share

more information and see less information privately. The smaller sample size means the decision maker is less confident in his aggregate forecast.

2.5. Comparison with Existing Aggregators

The conjugate-pair ensembles above, because they are optimal aggregators under some reasonable assumptions, suggest an important structural property we would like to see in any aggregation method. An aggregation method should incorporate the prior-predictive probability p_0 . The existing methods in Ranjan and Gneiting (2010), Turner et al. (2014), Baron et al. (2014), Satopää et al. (2014), and Satopää et al. (2016), however, assume p_0 is equal to one-half or do not incorporate p_0 . In many applications though, such as in the two empirical studies of Section 4, p_0 (or the base rate of the y 's in a training set) will be far from one-half.

Ranjan and Gneiting (2010), Turner et al. (2014), and Baron et al. (2014) propose transformations of the linear opinion pool $p_{LOP} = \sum_{i=1}^k w_i p_i$ where $w_i > 0$ for $i = 1, \dots, k$, and $\sum_{i=1}^k w_i = 1$. Ranjan and Gneiting (2010) study the beta-transformed linear opinion pool (BLOP): $p_{BLOP} = F_{Be(\alpha, \beta)}(p_{LOP})$ where $F_{Be(\alpha, \beta)}$ is the cdf of the beta distribution. Turner et al. (2014) and Baron et al. (2014) use the Karmarkar equation $p_{LOP}^a / (p_{LOP}^a + (1 - p_{LOP})^a)$ to extremize the linear opinion pool. We call this aggregator the Karmarkar-transformed linear opinion pool (KLOP), denoted by p_{KLOP} . Both methods involve s-shaped transformations of the linear opinion pool. These two methods, fit to the same data, will be almost identical. For example, it is difficult to distinguish the Karmarkar-transformed linear opinion pool with $a = 2.5$ from the beta-transformed linear opinion pool with $\alpha = \beta = 5$; see Figure 2a where $p_{LOP} = \bar{p}$ and $p_1 = 1/2$.

Satopää et al. (2014) and Turner et al. (2014) propose the logit aggregator, a variant of the logarithmic opinion pool (French 1980, Dawid et al. 1995). This aggregator is a generalized additive model without an intercept: $p_{logit} = F_{Lo(0,1)}(\sum_{i=1}^k (a/k) F_{Lo(0,1)}^{-1}(p_i))$ where $F_{Lo(l,\psi)}(z) = 1/(1 + e^{-(z-l)/\psi})$ is the cdf of the logistic distribution with location l and scale ψ and $F_{Lo(l,\psi)}^{-1}(p) = l + \psi \log(p/(1-p))$ is the inverse of its cdf. Ranjan and Gneiting (2010) and Satopää et al. (2016) study versions of the probit ensemble in Example 3 with $p_0 = 1/2$. When $p_0 = 1/2$, the probit ensemble reduces to a generalized additive model with no intercept. These two ensembles are depicted in Figure 2b. The probit ensemble \hat{p} in Figure 2b has the same settings as the one in Figure 1c (except for $\theta_0 = 0$). This probit ensemble and the logit aggregator with $a = 1.25$ are virtually indistinguishable.

Interestingly, both the logit aggregator and the probit ensemble in Figure 2b involve inverse s-shaped functions. In the next section, we will see that negative correlation between experts' information states results in an s-shaped probit ensemble. The experts' information states according to the probit ensemble of Proposition 3 are always positively correlated (see $v_{i,j}$ in the Proof of




Figure 2 Ensembles With No Informative Prior-Predictive Probability Incorporated.

Proposition 3). Because negative correlation among experts is rare in practice, this result calls into question the use of s-shaped ensembles.

3. Generalized Probit Ensemble

Below we introduce the generalized probit ensemble. For this ensemble, we first assume the experts' information states (e.g., sufficient statistics) are jointly normally distributed, as in the probit ensemble of Example 3, but with a general correlation structure. This general correlation structure can capture more complicated patterns of overlapping information sources, without having to work through detailed combinatorics (Clemen 1987). Second, we assume a generalized linear model of these information states. This second assumption is informed by the Bayesian ensemble in Proposition 1 and, in particular, by its generalized linear model form in (4).

At the outset, the decision maker places beliefs directly on the experts' information states—states that may result from experts seeing some observations in common. He assumes the experts' information states $\mathbf{x} = (x_1, \dots, x_k)'$ are jointly normally distributed with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, denoted $\mathbf{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. The covariance matrix has elements σ_{ij} , and the correlation between x_i and x_j is defined as $\rho_{ij} = \sigma_{ij} / (\sqrt{\sigma_{ii}}\sqrt{\sigma_{jj}})$.

The decision maker also assumes the conditional distribution of y given the information states is given by the generalized linear model of the information states

$$P(y = 1 | x_1, \dots, x_k) = F_{z_0}(\alpha_0 + \alpha_1 x_1 + \dots + \alpha_k x_k), \quad (9)$$

where the inverse link function F_{z_0} is the cdf of a continuous random variable z_0 with the entire real line as its support. In addition, he assumes the inverse link function F_{z_0} is the standard cdf from a location-scale family, i.e., the random variable $z = l + \psi z_0$ has the cdf $F_{z_0}((z-l)/\psi)$ where l is the location parameter and ψ is the scale parameter. Finally, he assumes any expert i reports the probability $p_i = P(y = 1|x_i)$.

In our generalized linear model, we do not restrict ourselves here to the standard normal link function in a generalized linear model of experts' information states. We propose a member of any location-scale family for the inverse link function and later focus on the exponential-power distribution. The exponential-power inverse link function has the flexibility to be either the standard normal near one extreme or the uniform on the other extreme. Therefore, the generalized probit ensemble can take the form of either an inverse s-shape or a linear shape, as needed.

PROPOSITION 4 (Generalized Probit Ensemble). *Given the assumptions in this section, the optimal way to aggregate experts' forecasts is with the generalized probit ensemble*

$$\hat{p} = P(y = 1|p_1, \dots, p_k) = F_{z_0} \left(\beta_0 F_{z_0 + \sqrt{v_0} x_0}^{-1}(p_0) + \sum_{i=1}^k \beta_i F_{z_0 + \sqrt{v_i} x_0}^{-1}(p_i) \right), \quad (10)$$

where $p_0 = F_{z_0 + \sqrt{v_0} x_0}(m_0)$, $p_i = F_{z_0 + \sqrt{v_i} x_0}(m_0 + \beta_i^{-1} \alpha_i (x_i - \mu_i))$, $m_0 = \alpha_0 + \sum_{i=1}^k \alpha_i \mu_i$,

$$\beta_0 = 1 - \sum_{i=1}^k \beta_i, \text{ and } \beta_i = \frac{\alpha_i \sqrt{\sigma_{ii}}}{\sum_{j=1}^k \alpha_j \sqrt{\sigma_{jj} \rho_{ij}}} \text{ for } i = 1, \dots, k.$$

Also, the cdf $F_{z_0 + \sqrt{v_i} x_0}$ is the cdf of the sum of the independent random variables z_0 and $\sqrt{v_i} x_0$ where z_0 is the standard random variable from a location-scale family, x_0 is a standard normal random variable,

$$v_0 = \sum_{i=1}^k \sum_{j=1}^k \alpha_i \alpha_j \sigma_{ij}, \text{ and } v_i = \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \sigma_{jj'} - \frac{(\sum_{j \neq i} \alpha_j \sigma_{ij})^2}{\sigma_{ii}} \text{ for } i = 1, \dots, k.$$

Immediately we see that the generalized probit ensemble in (10) is a generalized additive model of the experts' probabilities. The main benefit of this ensemble is that we do not need to work with a conjugate pair's predictive distribution, which can be difficult to do in all but few cases.

Another immediate consequence of the result is that each expert is calibrated. An expert is calibrated if $P(y = 1|p_i) = p_i$ (French 1986, Murphy and Winkler 1987). As was the case in the previous section, this calibration means that the decision maker, if he heard from only one expert, would use the same probability the expert reported. Note that if the expert is not calibrated, then

the decision maker could re-calibrate the expert's reported probability using the re-calibration function $R_i(p_i)$ so that $P(y = 1|x_i) = R(p_i)$. All results related to the generalized probit ensemble hold with $R(p_i)$ in place of p_i , if an expert is not calibrated.

To interpret the coefficients in the ensemble, it is helpful to consider the case of exchangeable experts. If the experts are exchangeable, then the weights are given by

$$\beta_0 = 1 - \frac{k}{(k-1)\rho+1} < 0 \text{ and } \beta_i = \frac{1}{(k-1)\rho+1} > 0 \text{ for } i = 1, \dots, k,$$

Note that for the exchangeable distribution of information states to be a proper distribution, we must have $\rho > -1/(k-1)$, which ensures that the matrix Σ is positive definite.

3.1. Exponential-Power Inverse Link Function

Below we propose a family of inverse link functions to use when fitting the generalized probit ensemble to data. The family is based on the exponential-power distribution, also known as a generalized error or generalized Gaussian distribution (Subbotin 1923, Box and Tiao 1973, Mineo and Ruggieri 2005, Zhang et al. 2012). This family contains the normal and Laplace distributions as special cases. Before we give the family's general form, we look at the example of the normal distribution.

EXAMPLE 7 (NORMAL DISTRIBUTION). Let the inverse link function F_{z_0} in Proposition 4 be standard normal cdf. In this case, $F_{z_0+\sqrt{v_i}x_0}(u) = \Phi(u/\sqrt{1+v_i})$ so that $F_{z_0+\sqrt{v_i}x_0}^{-1}(p_i) = \sqrt{1+v_i}\Phi^{-1}(p_i)$. The probit ensemble, according to the assumptions in this section, is given by

$$\hat{p} = \Phi\left(\beta_0\sqrt{1+v_0}\Phi^{-1}(p_0) + \sum_{i=1}^k \beta_i\sqrt{1+v_i}\Phi^{-1}(p_i)\right).$$

□

To estimate this probit ensemble's coefficients $\beta'_i = \beta_i\sqrt{1+v_i}$ from data, one can fit a generalized linear model of y on $(\Phi^{-1}(p_1), \dots, \Phi^{-1}(p_k))$ using a standard normal inverse link function.

EXAMPLE 8 (EXPONENTIAL-POWER DISTRIBUTION). Let the inverse link function be the cdf F_{z_0} with density

$$f_{z_0}(z) = \frac{1}{2\eta^{1/\eta}\Gamma(1+1/\eta)\psi} \exp\left(-\frac{1}{\eta}\left|\frac{z-l}{\psi}\right|^\eta\right).$$

This is the density of an exponential-power distribution. We denote this distribution by $z_0 \sim EP(l, \psi, \eta)$ where l is the mean, $\eta^{2/\eta}\psi^2\Gamma(3/\eta)/\Gamma(1/\eta)$ is the variance, and $\eta > 0$ is the power parameter. For a fixed power parameter, this distribution is a member of a location-scale family. □

For the ensemble based on the exponential-power distribution, the inside function is $F_{z_0 + \sqrt{v_i}x_0}$, which is not tractable (Soury and Alouini 2015). Nonetheless, we can approximate this inside function closely by matching the first two moments. We choose v'_i in $\sqrt{1 + v'_i z_{EP(0,1,\eta)}}$ so that the variance of this random variable equals the variance of $z_{EP(0,1,\eta)} + \sqrt{v_i}x_0$. Their means are both zero by construction. Consequently, the approximate generalized probit ensemble with an exponential-power inverse link function, denoted by \tilde{p} , is given by

$$\tilde{p} = F_{EP(0,1,\eta)} \left(\beta_0 \sqrt{1 + v'_0} F_{EP(0,1,\eta)}^{-1}(p_0) + \sum_{i=1}^k \beta_i \sqrt{1 + v'_i} F_{EP(0,1,\eta)}^{-1}(p_i) \right). \quad (11)$$

This ensemble is useful in applications because one can estimate the coefficients $\beta'_i = \beta_i \sqrt{1 + v'_i}$ by fitting a generalized linear model of y on $(F_{EP(0,1,\eta)}^{-1}(p_1), \dots, F_{EP(0,1,\eta)}^{-1}(p_k))$ using the exponential-power inverse link function $F_{EP(0,1,\eta)}$. For $\eta = 2$, the exponential-power distribution becomes the normal distribution, and the resulting ensemble is a probit ensemble. For $\eta = 1$, the exponential-power distribution becomes the Laplace distribution. For η strictly less (greater) than 2, the distribution has fat (thin) tails. As $\eta \rightarrow \infty$, the distribution of $z_{EP(l,\psi,\eta)}$ goes to a uniform distribution on $(l - \psi, l + \psi)$ and the resulting ensemble approaches a linear ensemble, like the one in the beta/Bernoulli ensemble of Example 1.

In Figure 3, we show several approximate generalized probit ensembles of two exchangeable experts' forecasts. We plot these ensembles as a function of p_2 with $p_1 = 1/2$, $\alpha = 1$, $\sigma = 1/20$, and $\rho = 3/4$. The weights on these positively correlated experts are $\beta_1 = \beta_2 = 0.57$. Figure 3a depicts three approximate generalized probit ensembles: the ensemble with $\eta = 1$, the probit ensemble ($\eta = 2$), and the ensemble with $\eta = 4$. As the power parameter increases, the ensemble becomes more linear. Also, we see in Figure 3b that the ensemble of positively correlated experts is inverse s-shaped, while the ensemble of negatively correlated ($\rho = -1/2$) experts is s-shaped.

4. Empirical Studies

In this section, we present two empirical studies where we fit the generalized probit ensemble and compare its out-of-sample forecasting performance to several leading aggregation methods. The challenge in the first study is to predict defaults on loans acquired by Fannie Mae in 2007, just before the Great Recession. These data are available at <https://loanperformancedata.fanniemae.com/lppub/index.html>. In the Fannie Mae data, there are 1,056,724 records on acquired loans and 20 independent variables, such as the borrower's credit score, the home's loan-to-value, and borrower's debt-to-income ratio. This year of acquisitions had the highest rate of defaults at 8.5% in the period 2000-2015. The second study's challenge is to predict bad used-car buys by Carvana




Figure 3 Approximate Generalized Probit Ensembles As the Power Parameter (η) and the Correlation Coefficient (ρ) Change.

(a large used-car retailer) during the period 2009-2010. These data are available at <https://www.kaggle.com/c/DontGetKicked/data>. In the Carvana data, there are 72,983 records on used-car purchases and 34 independent variables, such as the vehicle's age, the vehicle's odometer, the vehicle's model, the buyer's id number, and the auction's location. This prediction challenge was part of a Kaggle competition called *Don't Get Kicked!*. The base rate of defective cars in the training set from Kaggle is 12.3%.

4.1. Training and Stacking the Models

In the machine learning literature, forecast aggregation is known as stacking (Wolpert 1992, Breiman 1996, Smyth and Wolpert 1999, Džeroski and Ženko 2004). The idea is that the predictions from several base models become features in a second-stage stacker model. Breiman (1996) and Smyth and Wolpert (1999), for example, both consider stacker models that are linear opinion pools of base models' probabilities. They choose optimal weights in a linear opinion pool that maximize the likelihood on a training set.

For each study, we trained three base models using the covariates available in the datasets. Each base model is a leading statistical or machine learning algorithm or part of a competition-winning ensemble. The first model is a regularized logistic regression (RLR), the lasso proposed by Tibshirani (1996). The second model is the random forest (RF) introduced by Breiman (2001).

The third model is the extreme gradient boosted trees model called xgboost (XGB) (Chen and Guestrin 2016). This model is an extension of the gradient boosted trees model introduced by Friedman (2001). We could add more more base models to our ensemble, but our goal here is to demonstrate the plausibility of our approach. The xgboost model, by itself, represents a difficult benchmark to beat. It was a part of 17 winning solutions published on Kaggle in 2015, and it was used by every team in the top 10 at KDD Cup 2015 (Chen and Guestrin 2016).

To ensure that an ensemble was trained on out-of-sample probabilities from the base models, we employed a two-step process for building a stacker model. First, we randomly split the data into 10 equal folds and used these folds for both steps. In Step 1, the base models were fit to the first nine folds of data using the available covariates (e.g., credit score and loan-to-value) and the outcomes of the binary event of interest (e.g., loan defaults). Then the base models were used to predict the binary event of interest in the tenth fold. Next the base models were trained on Folds 1-8 and 10 and used to predict the binary events of interest in the Fold 9. This process continued until each fold was held out once and out-of-sample predictions were made for it. In Step 2, an ensemble, or stacker model, was trained on the out-of-sample predictions made by the base models in Step 1. For example, a stacker model was trained on Folds 1-9 and tested (or evaluated) on Fold 10. This training and testing of a stacker model was done 10 times, with each fold serving as the hold-out sample once. For more details on how these models were trained and tested, see the supplemental materials. All data and code are available from the authors.

4.2. Scoring the Models' Forecasts

To evaluate out-of-sample forecasts, we use three different scoring rules: (i) the log score (LS), (ii) the asymmetric log score (ALS), and (iii) the area under the curve (AUC). The first scoring rule is negatively oriented (lower score is better), and the second and third scoring rules are positively oriented (higher score is better). The log score of a probability forecast p for a binary event y is given by $LS(p, y) = -(y \log(p) + (1 - y) \log(1 - p))$ for $0 < p < 1$. This score is consistent with maximum likelihood estimation; the model (or ensemble) with the lowest average log score maximizes the log-likelihood. The asymmetric log score of a probability forecast p for a binary event y is given by $ALS(y, p) = (LS(c, y) - LS(p, y)) / LS(c, I_{p>c})$ where $I_{p>c}$ equals 1 if $p > c$ and equals 0 otherwise (Winkler 1994). This score is “adjusted for the difficulty of the forecast task . . . with the value of $c \in (0, 1)$ adapted to reflect a baseline probability” (Gneiting and Raftery 2007, p. 365). In the results we report below, c is taken to be the base rate of occurrence of y (denoted \bar{y}) in the training set. The area under the curve is a popular score in the machine-learning community (Hand and Till 2001).

4.3. Results

Table 1 reports the average scores of out-of-sample predictions from the three base models (p_{RLR} , p_{RF} , and p_{XGB}), four existing aggregation models (the linear opinion pool with equal weights \bar{p} , the linear opinion pool with optimal weights p_{OLOP} , the beta-transformation of the linear opinion pool with optimal weights p_{BLOP} , the logit aggregator p_{logit}), and the best approximate generalized probit ensemble with an exponential-power inverse link function, \tilde{p} in (11). The best \tilde{p} has the power parameter η^* that minimizes the average log loss out-of-sample. The estimates of η^* are 40 and 9 in the two studies, respectively. In both studies, the best approximate generalized probit ensemble outperforms all other models on average over the 10 cross-validation folds.

Model	Fannie Mae			Carvana		
	<i>LS</i>	<i>ALS</i>	<i>AUC</i>	<i>LS</i>	<i>ALS</i>	<i>AUC</i>
p_{RLR}	0.2362	0.2921	0.8162	0.3109	0.1414	0.7552
p_{RF}	0.2450	0.2496	0.7940	0.2993	0.1568	0.7612
p_{XGB}	0.2337	0.3058	0.8206	0.2946	0.1687	0.7708
\bar{p}	0.2347	0.2951	0.8205	0.2966	0.1719	0.7742
p_{OLOP}	0.2330	0.3083	0.8229	0.2932	0.1765	0.7753
p_{BLOP}	0.2359	0.2924	0.8229	0.2953	0.1663	0.7753
p_{logit}	0.2329	0.3090	0.8229	0.2928	0.1775	0.7751
\tilde{p}	0.2327	0.3116	0.8230	0.2925	0.1777	0.7753

Table 1 Average Scores of Out-of-Sample Predictions in the Two Studies.

After performing 10-fold cross validation for each study, we estimate the coefficients in the generalized linear model for the best approximate generalized probit ensemble on the entire dataset. Table 2 lists these estimates. Not surprisingly, the ensemble puts the most weight on the best base model: xgboost. We also report in Table 2 the base rate at which the binary events occur in each dataset. In addition, we provide the percentage of times in out-of-sample forecasting that the best approximate generalized probit ensemble extremizes the average forecast.

Figure 4 depicts the best approximate generalized probit ensemble as a function of the best base model's forecasts, with the other two base models' forecasts set to twice the base rate in their respective dataset. For these settings, the ensembles lie somewhere in between all weight on the best base model (xgboost) and equal weights on the base models. In each plot, we also highlight the anti-extremizing region in gray. In addition, because the estimated power parameters are high in these studies, we can see that the ensembles are nearly linear over much of the domain of p_{XGB} .

	Fannie Mae	Carvana
Constant	0.0496	0.0653
Coefficient β'_{RLR}	0.3710	-0.0281
Coefficient β'_{RF}	0.0810	0.3864
Coefficient β'_{XGB}	0.6070	0.7176
Power parameter η^*	40	9
\tilde{p} extremizes \bar{p}	79.6%	81.3%
Base rate \bar{y}	0.0849	0.1230
No. of observations	1,056,724	72,983

Table 2 Final Estimation of Best Approximate Generalized Probit Ensembles.

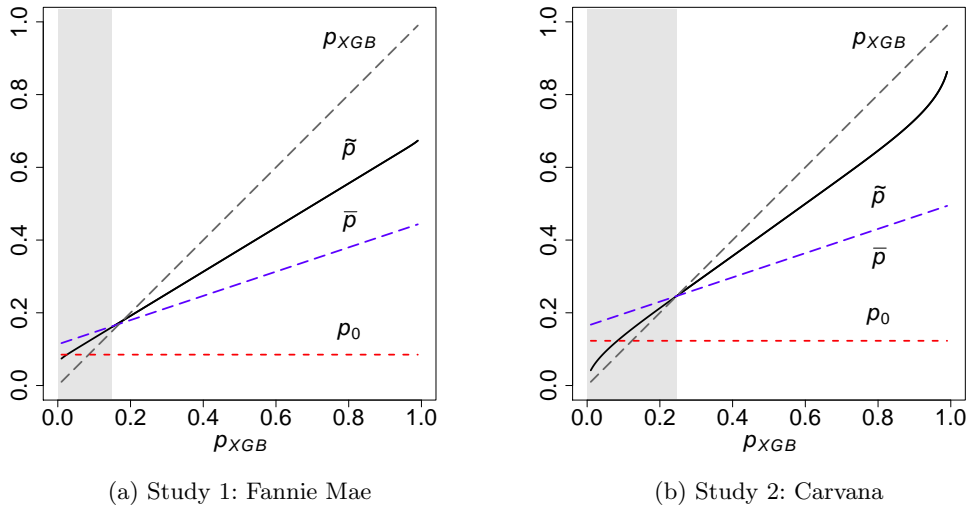


Figure 4 Best Approximate Generalized Probit Ensemble as a Function of the Best Base Model's Forecast.

5. Summary and Conclusions

In this paper, we introduce a large class of Bayesian ensembles. Because the ensembles in this class are based on Bayesian reasoning, they can help us understand why, when, and how much extremizing is appropriate. It is appropriate to extremize the average forecast in order to remove the redundant reports of the prior-predictive probability in the experts' forecasts. When some information is shared by the experts, however, the decision maker may sometimes want to anti-extremize. Due to the shared information, the decision maker does not have as much total information upon which to base his aggregate forecast, and he may in turn issue a less confident aggregate forecast. This theoretical result matches the empirical findings of Tetlock and Gardner (2015). They also

find that when teams are good at sharing information, like teams of superforecasters, extremizing does not help them much.

The two types of Bayesian ensembles we introduce here—conjugate-pair Bayesian ensembles and the generalized probit ensemble—take the form of a generalized additive model. In the first type, the link function is determined by a regular, one-parameter exponential family and its conjugate prior. This type is easy to interpret because of the natural sampling process that underlies the ensemble. In the second type, the inverse link function is any distribution from a location-scale family. This type is more flexible; and therefore more suitable for applications. In practice, we find that the exponential-power distribution is a good choice for the inverse link function in the generalized probit ensemble. Within the family of generalized probit ensembles are the probit ensemble near one extreme and a linear ensemble at the other extreme, depending on where the power parameter is set. Using a generalized linear modeling framework, this power parameter can be tuned to fit real world training data. Importantly, our ensembles include a constant term inside their generalized additive model forms. This constant term incorporates in the aggregate forecast the prior-predictive probability, e.g., the base rate of occurrences of the binary event in a training set.

In our two empirical studies, the generalized probit ensemble performed well in making out-of-sample predictions. It outperformed three base models and four leading aggregation methods. The base models in these studies were leading statistical and machine learning algorithms: the lasso, random forest, and xgboost. The latter two models are ensembles themselves and can be difficult to beat in practice, especially xgboost. In outperforming xgboost, it is important to stress that the Bayesian ensembles here are not replacements for an ensemble like xgboost. The Bayesian ensemble we fit used xgboost as one of its inputs. More importantly, the empirical results presented here demonstrate the plausibility of our generalized probit ensemble, not its superiority. We see several possible avenues for future work, including applications of the generalized probit ensemble to (possibly poorly calibrated) human judgments.

Appendix

In this appendix, we provide proofs of Propositions 1-4 and derivations of Examples 1-4. We also state a lemma and its proof. This lemma is useful in the proof of Proposition 4.

Proof of Proposition 1.

By Bernardo and Smith (2000, Prop. 5.5), expert i 's forecast of x given \mathbf{x}_i is given by

$$f(x|\mathbf{x}_i) = a(x) \frac{K(\tau_0 + n_i + 1, \tau_1 + t_i + h(x))}{K(\tau_0 + n_i, \tau_1 + t_i)},$$

where $t_i = \sum_{j=N_{i-1}+1}^{N_i} h(x_j)$ is the sufficient statistic for \mathbf{x}_i . Expert i 's forecast of $(y|\mathbf{x}_i)$ is given by $P(y=1|\mathbf{x}_i) = \int_{x \in A} f(x|\mathbf{x}_i) dx$. Hence, $p_i = P(y=1|\mathbf{x}_i) = F_{n_i}(\tau_1 + t_i)$. The decision maker, if he had direct access to the experts' samples, would form the following forecast of $(y|\mathbf{x}_1, \dots, \mathbf{x}_k)$:

$$\begin{aligned} P(y=1|\mathbf{x}_1, \dots, \mathbf{x}_k) &= \int_{x \in A} a(x) \frac{K(\tau_0 + \sum_{i=1}^k n_i + 1, \tau_1 + \sum_{i=1}^k t_i + h(x))}{K(\tau_0 + \sum_{i=1}^k n_i, \tau_1 + \sum_{i=1}^k t_i)} dx \\ &= F_{N_k} \left(\tau_1 + \sum_{i=1}^k t_i \right). \end{aligned} \quad (12)$$

The decision maker knows all the sample sizes because of the common knowledge assumption. The prior-predictive probability is given by

$$p_0 = P(y=1) = \int_{x \in A} a(x) \frac{K(\tau_0 + 1, \tau_1 + h(x))}{K(\tau_0, \tau_1)} dx = F_0(\tau_1).$$

By the assumption that $F_n(t)$ is strictly monotonic in t , we can invert $F_n(t)$. By substitution according to $\tau_1 = F_0^{-1}(p_0)$ and $t_i = F_{n_i}^{-1}(p_i) - \tau_1$ into (12), we have the result.

Derivation of Example 1's Ensemble.

The Bernoulli distribution is a regular, one-parameter exponential family with pmf

$$f(x_j|\theta) = \theta^{x_j} (1-\theta)^{1-x_j} = (1-\theta) \exp \left(\log \left(\frac{\theta}{1-\theta} \right) x_j \right) \text{ for } x_j \in \{0, 1\},$$

where $a(x_j) = 1$, $b(\theta) = 1 - \theta$, $c(\theta) = \log(\frac{\theta}{1-\theta})$, and $h(x_j) = x_j$. Its conjugate prior has the pdf

$$f(\theta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1} (1-\theta)^{\beta-1} = \frac{1}{K(\tau_0, \tau_1)} (1-\theta)^{\tau_0} \exp \left(\log \left(\frac{\theta}{1-\theta} \right) \tau_1 \right) \text{ for } \theta \in (0, 1),$$

where $\tau_0 = \alpha + \beta - 2$, $\tau_1 = \alpha - 1$, and $K(\tau_0, \tau_1) = \Gamma(\tau_1 + 1)\Gamma(\tau_0 - \tau_1 + 1)/\Gamma(\tau_0 + 2)$. With $A = \{1\}$,

$$\begin{aligned} F_n(t) &= \int_{x \in A} a(x) \frac{K(\tau_0 + n + 1, t + h(x))}{K(\tau_0 + n, t)} dx \\ &= \frac{\Gamma(t + 1 + 1)\Gamma(\tau_0 + n + 1 - (t + 1) + 1)}{\Gamma(\tau_0 + n + 1 + 2)} \frac{\Gamma(\tau_0 + n + 2)}{\Gamma(t + 1)\Gamma(\tau_0 + n - t + 1)} = \frac{t + 1}{\alpha + \beta + n}, \end{aligned}$$

and $F_{n_i}^{-1}(p_i) = (\alpha + \beta + n_i)p_i - 1$.

Derivation of Example 2's Ensemble.

The Poisson distribution is a regular, one-parameter exponential family with pmf

$$f(x_j|\theta) = \frac{\theta^{x_j} \exp(-\theta)}{x_j!} = \frac{1}{x_j!} \exp(-\theta) \exp(\log(\theta)x_j) \text{ for } x_j \in \{0, 1, 2, \dots\},$$

where $a(x_j) = 1/x_j!$, $b(\theta) = \exp(-\theta)$, $c(\theta) = \log(\theta)$, and $h(x_j) = x_j$. Its conjugate prior has the pdf

$$f(\theta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \theta^{\alpha-1} \exp(-\beta\theta) = \frac{1}{K(\tau_0, \tau_1)} \theta^{\tau_1} \exp(-\tau_0\theta) \text{ for } \theta \in (0, \infty),$$

where $\tau_0 = \beta$, $\tau_1 = \alpha - 1$, and $K(\tau_0, \tau_1) = \Gamma(\tau_1 + 1)/\tau_0^{\tau_1 + 1}$. With $A = \{0\}$,

$$F_n(t) = \int_{x \in A} a(x) \frac{K(\tau_0 + n + 1, t + h(x))}{K(\tau_0 + n, t)} dx = \frac{\Gamma(t + 1)}{(\tau_0 + n + 1)^{t+1}} \frac{(\tau_0 + n)^{t+1}}{\Gamma(t + 1)} = \exp(v_n(t + 1)),$$

where $v_n = \log((\beta + n)/(\beta + n + 1))$, and $F_{n_i}^{-1}(p_i) = \log(p_i)/v_{n_i} - 1$.

Derivation of Example 3's Ensemble.

This example's normal distribution is a regular, one-parameter exponential family with pdf

$$\begin{aligned} f(x_j|\theta) &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(x_j - \theta)^2\right) \\ &= \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{x_j^2}{2\sigma^2}\right) \exp\left(-\frac{\theta^2}{2\sigma^2}\right) \exp\left(\frac{\theta}{\sigma^2}x_j\right) \text{ for } x_j \in (-\infty, \infty), \end{aligned}$$

where $a(x) = (1/\sqrt{2\pi\sigma^2}) \exp(-x_j^2/(2\sigma^2))$, $b(\theta) = \exp(-\frac{\theta^2}{2\sigma^2})$, $c(\theta) = \theta$, and $h(x_j) = x_j$. Its conjugate prior has the pdf

$$\begin{aligned} f(\theta) &= \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left(-\frac{1}{2\sigma_0^2}(\theta - \theta_0)^2\right) \\ &= \frac{1}{\sqrt{2\pi\sigma_0^2}} \exp\left(-\frac{\theta_0^2}{2\sigma_0^2}\right) \left(\exp\left(-\frac{\theta^2}{2\sigma^2}\right)\right)^{\tau_0} \exp\left(\frac{\theta}{\sigma^2}\tau_1\right) \text{ for } \theta \in (-\infty, \infty), \end{aligned}$$

where $\tau_0 = \sigma^2/\sigma_0^2$, $\tau_1 = \sigma^2\theta_0/\sigma_0^2$, and $K(\tau_0, \tau_1) = \sqrt{(2\pi\sigma^2)/\tau_0} \exp(\tau_1^2/(2\sigma^2\tau_0))$. With $A = (0, \infty)$,

$$\begin{aligned} F_n(t) &= \int_{x \in A} a(x) \frac{K(\tau_0 + n + 1, t + h(x))}{K(\tau_0 + n, t)} dx \\ &= \int_0^\infty \frac{1}{\sqrt{2\pi \frac{\tau_0+n+1}{\tau_0+n} \sigma^2}} \exp\left(-\frac{\tau_0+n}{2\sigma^2(\tau_0+n+1)}\left(x - \frac{t}{\tau_0+n}\right)^2\right) dx \\ &= 1 - \Phi\left(\frac{-\frac{t}{\tau_0+n}}{\sqrt{\frac{\tau_0+n+1}{\tau_0+n}\sigma^2}}\right) = \Phi\left(\frac{\frac{t}{\tau_0+n}}{\sqrt{\frac{\tau_0+n+1}{\tau_0+n}\sigma^2}}\right) = \Phi\left(\frac{t}{\sqrt{v_n}}\right), \end{aligned}$$

where $v_n = (\sigma^2/\sigma_0^2 + n)(\sigma^2/\sigma_0^2 + n + 1)\sigma^2$, and $F_{n_i}^{-1}(p_i) = \sqrt{v_{n_i}}\Phi^{-1}(p_i)$. The second integrand in (13) is the posterior-predictive distribution for the normal/normal model with known precision in Bernardo and Smith (2000, p. 439). For example, with $t = \tau_1 + t_i$ and $n = n_i$, this integrand is expert i 's posterior-predictive distribution of x .

Derivation of Example 4's Ensemble.

This example's Gumbel distribution is a regular, one-parameter exponential family with pdf

$$\begin{aligned} f(x_j|\theta) &= \frac{1}{\sigma} \exp\left(-\frac{x_j - \theta}{\sigma}\right) \exp\left(-\exp\left(-\frac{x_j - \theta}{\sigma}\right)\right) \\ &= \frac{1}{\sigma} \exp\left(-\frac{x_j}{\sigma}\right) \exp\left(\frac{\theta}{\sigma}\right) \exp\left(-\exp\left(-\frac{x_j - \theta}{\sigma}\right)\right) \text{ for } x_j \in (-\infty, \infty), \end{aligned}$$

where $a(x_j) = \exp(-x_j/\sigma)/\sigma$, $b(\theta) = \exp(\theta/\sigma)$, $c(\theta) = -\exp(\theta/\sigma)$, and $h(x_j) = \exp(-x_j/\sigma)$. Its conjugate prior has the pdf

$$f(\theta) = \frac{\beta^\alpha}{\sigma\Gamma(\alpha)} \left(\exp\left(\frac{\theta}{\sigma}\right) \right)^\alpha \exp\left(-\beta \exp\left(\frac{\theta}{\sigma}\right)\right) \text{ for } \theta \in (-\infty, \infty),$$

where $\tau_0 = \alpha$, $\tau_1 = \beta$, and $K(\tau_0, \tau_1) = \Gamma(\tau_0)\tau_1^{-\tau_0}\sigma$. With $A = (-\infty, 0)$,

$$\begin{aligned} F_n(t) &= \int_{x \in A} a(x) \frac{K(\tau_0 + n + 1, t + h(x))}{K(\tau_0 + n, t)} dx \\ &= \int_{-\infty}^0 \frac{1}{\sigma} \exp\left(-\frac{x}{\sigma}\right) \frac{\Gamma(\alpha + n + 1)(t + \exp(-\frac{x}{\sigma}))^{-(\alpha+n+1)}}{\Gamma(\alpha + n)t^{-(\alpha+n)}} dx \\ &= \frac{1}{\sigma} t^{\alpha+n} \int_{-\infty}^0 \frac{(\alpha + n) \exp(-\frac{x}{\sigma})}{(t + \exp(-\frac{x}{\sigma}))^{\alpha+n+1}} dx = t^{\alpha+n} \left[\frac{1}{(t + \exp(-\frac{x}{\sigma}))^{\alpha+n}} \right]_{-\infty}^0 = \left(\frac{t}{1+t} \right)^{\alpha+n}, \end{aligned}$$

and $F_{n_i}^{-1}(p_i) = p_i^{v_{n_i}} / (1 - p_i^{v_{n_i}})$, where $v_n = 1/(\alpha + n)$.

Proof of Proposition 2.

By Bernardo and Smith (2000, Prop. 5.5), expert i 's forecast of x given the private sample \mathbf{x}_i and the shared sample \mathbf{x}_s is given by

$$f(x|\mathbf{x}_i, \mathbf{x}_s) = a(x) \frac{K(\tau_0 + n_i + n_s + 1, \tau_1 + t_i + t_s + h(x))}{K(\tau_0 + n_i + n_s, \tau_1 + t_i + t_s)},$$

where t_i and t_s are the sufficient statistics for the samples \mathbf{x}_i and \mathbf{x}_s , respectively. Expert i 's forecast of $(y|\mathbf{x}_i, \mathbf{x}_s)$ is given by $P(y = 1|\mathbf{x}_i, \mathbf{x}_s) = \int_{x \in A} f(x|\mathbf{x}_i, \mathbf{x}_s) dx$. Hence, $p_i = P(y = 1|\mathbf{x}_i, \mathbf{x}_s) = F_{n_i+n_s}(\tau_1 + t_i + t_s)$. The decision maker, if he had direct access to the experts' private samples and the shared sample, would form the following forecast of $(y|\mathbf{x}_1, \dots, \mathbf{x}_k, \mathbf{x}_s)$:

$$\begin{aligned} P(y = 1|\mathbf{x}_1, \dots, \mathbf{x}_k, \mathbf{x}_s) &= \int_{x \in A} a(x) \frac{K(\tau_0 + \sum_{i=1}^k n_i + n_s + 1, \tau_1 + \sum_{i=1}^k t_i + t_s + h(x))}{K(\tau_0 + \sum_{i=1}^k n_i + n_s, \tau_1 + \sum_{i=1}^k t_i + t_s)} dx \\ &= F_{N_k+n_s} \left(\tau_1 + \sum_{i=1}^k t_i + t_s \right). \end{aligned} \tag{13}$$

The decision maker knows all the sample sizes because of the common knowledge assumption. From the proof of Proposition 1, the prior-predictive probability is given by $p_0 = P(y = 1) = F_0(\tau_1)$. By

the assumption that $F_n(t)$ is strictly monotonic in t , we can invert $F_n(t)$. By substitution according to $\tau_1 = F_0^{-1}(p_0)$ and $t_i = F_{n_i+n_s}^{-1}(p_i) - \tau_1 - t_s$ into (13), we have

$$\begin{aligned} P(y = 1 | t_1, \dots, t_k, t_s) &= F_{N_k+n_s} \left(\tau_1 + \sum_{i=1}^k t_i + t_s \right) \\ &= F_{N_k+n_s} \left(F_0^{-1}(p_0) + \sum_{i=1}^k (F_{n_i+n_s}^{-1}(p_i) - F_0^{-1}(p_0) - t_s) + t_s \right) \\ &= F_{N_k+n_s} \left(-(k-1)(F_0^{-1}(p_0) + t_s) + \sum_{i=1}^k F_{n_i+n_s}^{-1}(p_i) \right) \\ &= P(y = 1 | p_1, \dots, p_k, t_s). \end{aligned}$$

Consequently, because $P(y = 1 | p_1, \dots, p_k) = E[E(y | p_1, \dots, p_k, t_s)]$, we have the result.

Proof of Proposition 3

The first step in the proof is to derive $f(t_s | p_1, \dots, p_k)$, which is equivalent to $f(t_s | t_1 + t_s, \dots, t_k + t_s)$ because $t_i + t_s = \sqrt{v_{n_i+n_s}} \Phi^{-1}(p_i) - \tau_1$ is a known function of p_i . The fact that $(t_s, t_1 + t_s, \dots, t_k + t_s)$ is jointly normally distributed follows from two facts: (a) $(x_1, \dots, x_{N_k+n_s})$ is jointly normally distributed according to Bernardo and Smith's (2000, Prop. 5.5(ii)) and (b) $(t_s, t_1 + t_s, \dots, t_k + t_s)$ is a linear combination of $(x_1, \dots, x_{N_k+n_s})$. The requisite means are

$$\begin{aligned} m_s &= E[t_s] = E \left[E \left[\sum_{j=N_k+1}^{N_k+n_s} x_j \middle| \theta \right] \right] = E[n_s \theta] = n_s \theta_0 \\ m_i &= E[t_i + t_s] = E \left[E \left[\sum_{j=N_{i-1}+1}^{N_i} x_j + \sum_{j=N_k+1}^{N_k+n_s} x_j \middle| \theta \right] \right] = E[(n_i + n_s) \theta] = (n_i + n_s) \theta_0. \end{aligned}$$

The variance of $t_i + t_s$ is given by

$$\begin{aligned} v_{i,i} &= \text{Var}[t_i + t_s] = E[\text{Var}[t_i + t_s | \theta]] + \text{Var}[E[t_i + t_s | \theta]] \\ &= (n_i + n_s) \sigma^2 + \text{Var}[(n_i + n_s) \theta] = (n_i + n_s) \sigma^2 + (n_i + n_s)^2 \sigma_0^2. \end{aligned}$$

Similarly, the variance of t_s is given by $v_{s,s} = \text{Var}[t_s] = n_s \sigma^2 + n_s^2 \sigma_0^2$. The covariance between $t_i + t_s$ and $t_j + t_s$ (for $i \neq j$) is given by

$$\begin{aligned} v_{i,j} &= \text{Cov}[t_i + t_s, t_j + t_s] = E[\text{Cov}[t_i + t_s, t_j + t_s | \theta]] + \text{Cov}[E[t_i + t_s | \theta], E[t_j + t_s | \theta]] \\ &= E[\text{Cov}[t_i, t_j | \theta] + \text{Cov}[t_i, t_s | \theta] + \text{Cov}[t_s, t_j | \theta] + \text{Cov}[t_s, t_s | \theta]] \\ &\quad + \text{Cov}[E[t_i | \theta] + E[t_s | \theta], E[t_j | \theta] + E[t_s | \theta]] \\ &= n_s \sigma^2 + \text{Cov}[n_i \theta, n_j \theta] + \text{Cov}[n_i \theta, n_s \theta] + \text{Cov}[n_s \theta, n_j \theta] + \text{Cov}[n_s \theta, n_s \theta] \\ &= n_s \sigma^2 + (n_i n_j + n_i n_s + n_s n_j + n_s^2) \sigma_0^2, \end{aligned}$$

which follows by the law of total covariance and the fact that t_i and t_j are conditionally independent given θ . Similarly, the covariance between t_s and $t_j + t_s$ is given by $v_{s,j} = \text{Cov}[t_s, t_j + t_s] = n_s \sigma^2 + (n_s n_j + n_s^2) \sigma_0^2$.

Thus, $(t_s, t_1 + t_s, \dots, t_k + t_s)' \sim N(\mathbf{m}, \mathbf{V})$ (West and Harrison 1997, p. 637). The mean vector is $\mathbf{m} = (\mathbf{m}_1, \mathbf{m}_2)'$ where $\mathbf{m}_1 = m_s$ and $\mathbf{m}_2 = (m_1, \dots, m_k)'$. The covariance matrix is

$$\mathbf{V} = \begin{pmatrix} \mathbf{V}_{11} & \mathbf{V}_{12} \\ \mathbf{V}_{21} & \mathbf{V}_{22} \end{pmatrix},$$

where $\mathbf{V}_{11} = v_{s,s}$, $\mathbf{V}_{12} = (v_{s,1}, \dots, v_{s,k})$, and \mathbf{V}_{22} has elements $v_{i,j}$ for $i, j \in \{1, \dots, k\}$.

According to West and Harrison (1997, p. 637), $(t_s | t_1 + t_s, \dots, t_k + t_s)$ is normally distributed with mean and variance:

$$E[t_s | t_1 + t_s, \dots, t_k + t_s] = \mathbf{m}_1 + \mathbf{V}_{12} \mathbf{V}_{22}^{-1} ((t_1 + t_s, \dots, t_k + t_s)' - \mathbf{m}_2)$$

$$\text{Var}[t_s | t_1 + t_s, \dots, t_k + t_s] = \mathbf{V}_{11} - \mathbf{V}_{12} \mathbf{V}_{22}^{-1} \mathbf{V}_{21}.$$

Hence, $(t_s | p_1, \dots, p_k)$ is normally distributed with mean and variance

$$E[t_s | p_1, \dots, p_k] = \mathbf{m}_1 + \mathbf{V}_{12} \mathbf{V}_{22}^{-1} ((\sqrt{v_{n_1+n_s}} \Phi^{-1}(p_1) - \tau_1, \dots, \sqrt{v_{n_k+n_s}} \Phi^{-1}(p_k) - \tau_1)' - \mathbf{m}_2)$$

$$\text{Var}[t_s | p_1, \dots, p_k] = \mathbf{V}_{11} - \mathbf{V}_{12} \mathbf{V}_{22}^{-1} \mathbf{V}_{21}.$$

Note that $\tau_1 = \sigma^2 \theta_0 / \sigma_0^2$ and $p_0 = F_0(\tau_1) = \Phi(\tau_1 / \sqrt{v_0})$, which implies that $\tau_1 = \sqrt{v_0} \Phi^{-1}(p_0)$ and $\theta_0 = (\sigma_0^2 / \sigma^2) \sqrt{v_0} \Phi^{-1}(p_0)$.

The second step in the proof is to apply Proposition 2. Let $L = -(k-1) \sqrt{v_0} \Phi^{-1}(p_0) + \sum_{i=1}^k \sqrt{v_{n_i+n_s}} \Phi^{-1}(p_i)$. According to Proposition 2 and the derivation of Example 3 where $F_n(t) = \Phi(t / \sqrt{v_n})$, we have

$$\begin{aligned} P(y=1 | p_1, \dots, p_k) &= \int_{-\infty}^{\infty} P(y=1 | p_1, \dots, p_k, t_s) f(t_s | p_1, \dots, p_k) dt_s \\ &= \int_{-\infty}^{\infty} \Phi\left(\frac{L - (k-1)t_s}{\sqrt{v_{N_k+n_s}}}\right) f(t_s | p_1, \dots, p_k) dt_s \\ &= \int_{-\infty}^{\infty} \left(1 - \Phi\left(\frac{t_s - L/(k-1)}{\sqrt{v_{N_k+n_s}}/(k-1)}\right)\right) f(t_s | p_1, \dots, p_k) dt_s \\ &= 1 - \int_{-\infty}^{\infty} \left(\int_{-\infty}^{t_s} \frac{k-1}{\sqrt{v_{N_k+n_s}}} \phi\left(\frac{z - L/(k-1)}{\sqrt{v_{N_k+n_s}}/(k-1)}\right) dz\right) f(t_s | p_1, \dots, p_k) dt_s \\ &= 1 - \int \int_{z \leq t_s} \frac{k-1}{\sqrt{v_{N_k+n_s}}} \phi\left(\frac{z - L/(k-1)}{\sqrt{v_{N_k+n_s}}/(k-1)}\right) f(t_s | p_1, \dots, p_k) dz dt_s \\ &= 1 - P(z \leq t_s | p_1, \dots, p_k) = P(t_s \leq z | p_1, \dots, p_k) = P(t_s - z \leq 0 | p_1, \dots, p_k), \end{aligned}$$

where $(z | p_1, \dots, p_k) \sim N(L/(k-1), v_{N_k+n_s}/(k-1)^2)$ and z and t_s are independent given (p_1, \dots, p_k) . Consequently, $t_s - z \sim N(E[t_s | p_1, \dots, p_k] - L/(k-1), \text{Var}[t_s | p_1, \dots, p_k] + v_{N_k+n_s}/(k-1)^2)$, which gives us the result.

Lemma 1 and its Proof.

The result below provides key properties of the linear combination of information states. We use these properties in the proof of Proposition 4 below.

LEMMA 1 (Linear Combination of Information States). *Assume $\mathbf{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Then the following statements hold.*

(i) *The distribution of $\alpha_0 + \sum_{i=1}^k \alpha_i x_i$ is normal with mean $m_0 = \alpha_0 + \sum_{i=1}^k \alpha_i \mu_i$ and variance $v_0 = \boldsymbol{\alpha} \boldsymbol{\Sigma} \boldsymbol{\alpha}' = \sum_{i=1}^k \sum_{j=1}^k \alpha_i \alpha_j \sigma_{ij}$.*

(ii) *The joint distribution of $\alpha_0 + \sum_{j \neq i} \alpha_j x_j$ and $\alpha_i x_i$ is normal with mean $\mathbf{a} + \mathbf{A}_i \boldsymbol{\mu}$ and variance $\mathbf{A}_i \boldsymbol{\Sigma} \mathbf{A}_i'$, where*

$$\mathbf{a} = \begin{pmatrix} \alpha_0 \\ 0 \end{pmatrix} \quad \text{and} \quad \mathbf{A}_i = \begin{pmatrix} \alpha_1 & \cdots & \alpha_{i-1} & 0 & \alpha_{i+1} & \cdots & \alpha_k \\ 0 & \cdots & 0 & \alpha_i & 0 & \cdots & 0 \end{pmatrix}.$$

(iii) *The conditional distribution of $\alpha_0 + \sum_{j \neq i} \alpha_j x_j$ given $\alpha_i x_i$ is normal with mean $m_i = \theta_{-i} + b_{-i,i} b_{i,i}^{-1} (\alpha_i x_i - \theta_i)$ and variance $v_i = b_{-i,-i} - b_{-i,i} b_{i,i}^{-1} b_{i,-i}$, where*

$$\begin{pmatrix} \theta_{-i} \\ \theta_i \end{pmatrix} = \mathbf{a} + \mathbf{A}_i \boldsymbol{\mu}, \quad \mathbf{B}_i = \begin{pmatrix} b_{-i,-i} & b_{-i,i} \\ b_{i,-i} & b_{i,i} \end{pmatrix} = \mathbf{A}_i \boldsymbol{\Sigma} \mathbf{A}_i',$$

$$b_{-i,-i} = \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \sigma_{jj'}, \quad b_{-i,i} = \sum_{j \neq i} \alpha_j \alpha_i \sigma_{ij}, \quad \text{and} \quad b_{i,i} = \alpha_i^2 \sigma_{ii}.$$

A second expression for the variance v_i is $v_i = \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \sigma_{jj'} - \sigma_{ii}^{-1} (\sum_{j \neq i} \alpha_j \sigma_{ij})^2$.

Proof. Statements (i) and (ii) follow directly from results in West and Harrison (1997, p. 637) on linear transformations of jointly distributed normal random variables. Statement (iii) follows from Statement (ii) and West and Harrison's (1997, p. 637) result on conditional distributions of jointly distributed normal random variables. The variance of $\alpha_0 + \sum_{i=1}^k \alpha_i x_i$ follows from the formula for the covariance of two linear transformations:

$$v_0 = \text{Cov} \left[\alpha_0 + \sum_{i=1}^k \alpha_i x_i, \alpha_0 + \sum_{i=1}^k \alpha_i x_i \right] = \sum_{i=1}^k \sum_{j=1}^k \alpha_i \alpha_j \text{Cov}[x_i, x_j] = \sum_{i=1}^k \sum_{j=1}^k \alpha_i \alpha_j \sigma_{ij}.$$

The elements of \mathbf{B}_i follow from the same formula:

$$b_{-i,-i} = \text{Cov} \left[\alpha_0 + \sum_{j \neq i} \alpha_j x_j, \alpha_0 + \sum_{j \neq i} \alpha_j x_j \right] = \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \text{Cov}[x_j, x_{j'}] = \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \sigma_{jj'},$$

$$b_{-i,i} = \text{Cov} \left[\alpha_0 + \sum_{j \neq i} \alpha_j x_j, \alpha_i x_i \right] = \sum_{j \neq i} \alpha_j \alpha_i \text{Cov}[x_j, x_i] = \sum_{j \neq i} \alpha_j \alpha_i \sigma_{ij},$$

and $b_{i,i} = \text{Cov}[\alpha_i x_i, \alpha_i x_i] = \alpha_i^2 \sigma_{ii}$, so that

$$\begin{aligned} v_i &= b_{-i,-i} - b_{-i,i} b_{i,i}^{-1} b_{i,-i} = b_{-i,-i} - \frac{b_{-i,i}^2}{b_{i,i}} = \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \sigma_{jj'} - \frac{\left(\sum_{j \neq i} \alpha_j \alpha_i \sigma_{ij} \right)^2}{\alpha_i^2 \sigma_{ii}} \\ &= \sum_{j \neq i} \sum_{j' \neq i} \alpha_j \alpha_{j'} \sigma_{jj'} - \frac{\left(\sum_{j \neq i} \alpha_j \sigma_{ij} \right)^2}{\sigma_{ii}}. \end{aligned}$$

Proof of Proposition 4.

The conditional probability of $(y|\alpha_i x_i)$ is given by

$$P(y=1|\alpha_i x_i) = \int_{-\infty}^{\infty} f(x|\alpha_i x_i) F_{z_0}(x + \alpha_i x_i) dx,$$

where $x = \alpha_0 + \sum_{j \neq i} \alpha_j x_j$, $f(x|\alpha_i x_i)$ is the normal density according to the joint normality assumption and Lemma 1(iii), and $P(y=1|\mathbf{x}') = F_{z_0}(x + \alpha_i x_i)$ comes from the generalized linear model assumption. We can rewrite this expression as

$$\begin{aligned} P(y=1|\alpha_i x_i) &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{v_i}} \phi\left(\frac{x - m_i}{\sqrt{v_i}}\right) \left(\int_{-\infty}^x f_{z_0}(z + \alpha_i x_i) dz \right) dx \\ &= \int \int_{z \leq x} \frac{1}{\sqrt{v_i}} \phi\left(\frac{x - m_i}{\sqrt{v_i}}\right) f_{z_0}(z + \alpha_i x_i) dz dx = P(z \leq x) = P(z - x \leq 0), \end{aligned} \quad (14)$$

where on the right-hand side of the third equality, z and $x \sim N(m_i, v_i)$ become random variables. By the assumption about the link function, the random variable $z = \psi z_0 + l_i$ is from a location-scale family with the standard cdf F_{z_0} , location parameter $l_i = -\alpha_i x_i$, and scale parameter $\psi = 1$. Since the last integrand involves a product of these random variable's densities, z and x are independent, as are z and $-x \sim N(-m_i, v_i)$. The random variable $-x$, being normally distributed, is also from a location-scale family where $-x = \sqrt{v_i} x_0 - m_i$ and $x_0 \sim N(0, 1)$. We can express the last probability in (14) as

$$P(z - x \leq 0) = P(z_0 + l_i + \sqrt{v_i} x_0 - m_i \leq 0) = P(z_0 + \sqrt{v_i} x_0 \leq m_i - l_i).$$

Let $F_{z_0 + \sqrt{v_i} x_0}(u) = P(z_0 + \sqrt{v_i} x_0 \leq u)$, the cdf of the sum of the independent random variables z_0 and $\sqrt{v_i} x_0$. Because $m_i - l_i = \theta_{-i} - b_{-i,i} b_{i,i}^{-1} \theta_i + (b_{-i,i} b_{i,i}^{-1} + 1) \alpha_i x_i$, we have

$$P(y=1|\alpha_i x_i) = F_{z_0 + \sqrt{v_i} x_0}(\theta_{-i} - b_{-i,i} b_{i,i}^{-1} \theta_i + (b_{-i,i} b_{i,i}^{-1} + 1) \alpha_i x_i),$$

which is strictly monotonic in x_i .

By Bayes' Theorem and a change of variable, we have that $P(y=1|q) = P(y=1|x_i)$ where $q = \alpha_i x_i$:

$$\begin{aligned} P(y=1|q) &= \frac{P(y=1) f_{q|y=1}(q|y=1)}{f_q(q)} = \frac{P(y=1) f_{q|y=1}(\alpha_i x_i|y=1) \alpha_i}{f_q(\alpha_i x_i) \alpha_i} \\ &= \frac{P(y=1) f_{x_i|y=1}(x_i|y=1)}{f_{x_i}(x_i)} = P(y=1|x_i). \end{aligned}$$

Then by the assumption that $p_i = P(y = 1|x_i)$, we have that $p_i = P(y = 1|\alpha_i x_i)$. Consequently,

$$\begin{aligned} p_i &= F_{z_0 + \sqrt{v_i}x_0}(\theta_{-i} - b_{-i,i}b_{i,i}^{-1}\theta_i + (b_{-i,i}b_{i,i}^{-1} + 1)\alpha_i x_i) \\ \iff \alpha_i x_i &= \frac{F_{z_0 + \sqrt{v_i}x_0}^{-1}(p_i) + b_{-i,i}b_{i,i}^{-1}\theta_i - \theta_{-i}}{(b_{-i,i}b_{i,i}^{-1} + 1)}. \end{aligned}$$

Substitution of this last expression into the assumed generalized linear model gives us

$$\alpha_0 + \sum_{i=1}^k \alpha_i x_i = \alpha_0 + \sum_{i=1}^k \frac{b_{-i,i}b_{i,i}^{-1}\theta_i - \theta_{-i}}{(b_{-i,i}b_{i,i}^{-1} + 1)} + \sum_{i=1}^k \beta_i F_{z_0 + \sqrt{v_i}x_0}^{-1}(p_i),$$

where

$$\beta_i = \frac{1}{b_{-i,i}b_{i,i}^{-1} + 1} = \frac{b_{i,i}}{b_{-i,i} + b_{i,i}} = \frac{\alpha_i^2 \sigma_{ii}}{\sum_{j=1}^k \alpha_j \alpha_i \sigma_{ij}} = \frac{\alpha_i \sqrt{\sigma_{ii}}}{\sum_{j=1}^k \alpha_j \sqrt{\sigma_{jj} \rho_{ij}}}.$$

Because $\theta_i = \alpha_i \mu_i$ and $\theta_{-i} = \alpha_0 + \sum_{j \neq i} \alpha_j \mu_j$, which implies that $\theta_i + \theta_{-i} = m_0$ from Lemma 1, we have that

$$\begin{aligned} \alpha_0 + \sum_{i=1}^k \frac{b_{-i,i}b_{i,i}^{-1}\theta_i - \theta_{-i}}{(b_{-i,i}b_{i,i}^{-1} + 1)} &= \alpha_0 + \sum_{i=1}^k (1 - \beta_i)\theta_i - \sum_{i=1}^k \beta_i \theta_{-i} = \alpha_0 + \sum_{i=1}^k \theta_i - \sum_{i=1}^k \beta_i (\theta_i + \theta_{-i}) \\ &= \alpha_0 + \sum_{i=1}^k \alpha_i \mu_i - \sum_{i=1}^k \beta_i m_0 = \left(1 - \sum_{i=1}^k \beta_i\right) m_0. \end{aligned}$$

From the expressions above, we can rewrite expert i 's reported probability as follows:

$$\begin{aligned} p_i &= F_{z_0 + \sqrt{v_i}x_0}(\theta_{-i} - b_{-i,i}b_{i,i}^{-1}\theta_i + (b_{-i,i}b_{i,i}^{-1} + 1)\alpha_i x_i) \\ &= F_{z_0 + \sqrt{v_i}x_0}(\theta_i + \theta_{-i} - (b_{-i,i}b_{i,i}^{-1} + 1)\theta_i + (b_{-i,i}b_{i,i}^{-1} + 1)\alpha_i x_i) \\ &= F_{z_0 + \sqrt{v_i}x_0}(\theta_i + \theta_{-i} + (b_{-i,i}b_{i,i}^{-1} + 1)\alpha_i (x_i - \mu_i)) \\ &= F_{z_0 + \sqrt{v_i}x_0} \left(\alpha_0 + \sum_{i=1}^k \alpha_i \mu_i + \beta_i^{-1} \alpha_i (x_i - \mu_i) \right). \end{aligned}$$

Finally, $P(y = 1) = \int_{-\infty}^{\infty} f(w) F_{z_0}(w) dw$ where $w = \alpha_0 + \boldsymbol{\alpha} \mathbf{x}$ and $f(w)$ is the normal density according to the joint normality assumption, the generalized linear model assumption, and Lemma 1(i). This expression can be simplified:

$$\begin{aligned} P(y = 1) &= \int_{-\infty}^{\infty} \frac{1}{\sqrt{v_0}} \phi\left(\frac{w - m_0}{\sqrt{v_0}}\right) \left(\int_{-\infty}^w f_{z_0}(z) dz \right) dw = \int \int_{z \leq w} \frac{1}{\sqrt{v_0}} \phi\left(\frac{w - m_0}{\sqrt{v_0}}\right) f_{z_0}(z) dz dw \\ &= P(z_0 \leq w) = P(z_0 - w \leq 0), \end{aligned}$$

where on the right-hand side of the third equality, z_0 and w become independent random variables.

Because $-w = \sqrt{v_0}x_0 - m_0$ and $x_0 \sim N(0, 1)$, we can express the last probability as $P(z_0 - w \leq 0) = P(z_0 + \sqrt{v_0}x_0 \leq m_0)$. Let $F_{z_0 + \sqrt{v_0}x_0}(u) = P(z_0 + \sqrt{v_0}x_0 \leq u)$. Then $m_0 = F_{z_0 + \sqrt{v_0}x_0}^{-1}(p_0)$ and $\beta_0 = 1 - \sum_{i=1}^k \beta_i$.

References

- Ahuja, J C, S W Nash. 1967. The generalized Gompertz-Verhulst family of distributions. *Sankhyā: The Indian Journal of Statistics, Series A* **29**(2) 141–156.
- Ariely, D, W T Au, R H Bender, D V Budescu, C B Dietz, H Gu, T S Wallsten, G Zauberan. 2000. The effects of averaging subjective probability estimates between and within judges. *Journal of Experimental Psychology Applied* **6**(2) 130–146.
- Baron, J, B A Mellers, P E Tetlock, E Stone, L H Ungar. 2014. Two reasons to make aggregated probability forecasts more extreme. *Decision Analysis* **11**(2) 133–145.
- Bernardo, J M, A F M Smith. 2000. *Bayesian Theory*. John Wiley & Sons, Ltd, Chichester, England.
- Box, G E P, G C Tiao. 1973. *Bayesian Inference in Statistical Inference*. Adison-Wesley, Reading, MA.
- Breiman, L. 1996. Bagging predictors. *Machine Learning* **24**(2) 123–140.
- Breiman, L. 2001. Random forests. *Machine Learning* **45**(1) 5–32.
- Chen, K, L R Fine, B A Huberman. 2004. Eliminating public knowledge biases in information-aggregation mechanisms. *Management Science* **50** 983–994.
- Chen, T, C Guestrin. 2016. Xgboost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 785–794.
- Chipman, Hugh A., Edward I. George, Robert E. McCulloch. 2007. Bayesian ensemble learning. P B Schölkopf, J C Platt, T Hoffman, eds., *Advances in Neural Information Processing Systems 19*. MIT Press, Cambridge, MA, 265–272.
- Clemen, R T. 1987. Combining overlapping information. *Management Science* **33**(3) 373–380.
- Clemen, R T, R L Winkler. 1986. Combining economic forecasts. *Journal of Business & Economic Statistics* **4**(1) 39–46.
- Dawid, A P, M H DeGroot, J Mortera. 1995. Coherent combination of experts' opinions. *Test* **4**(2) 263–313.
- DeGroot, M H. 1974. Reaching a consensus. *Journal of the American Statistical Association* **69**(345) 118–121.
- DeGroot, M H, J Mortera. 1991. Optimal linear opinion pools. *Management Science* **37**(5) 546–558.
- Džeroski, S, B Ženko. 2004. Is combining classifiers with stacking better than selecting the best one? *Machine Learning* **54**(3) 255–273.
- Erev, I, T S Wallsten, D V Budescu. 1994. Simultaneous over- and underconfidence: The role of error in judgment processes. *Psychological Review* **101**(3) 519–527.
- French, S. 1980. Updating of belief in the light of someone else's opinion. *Journal of the Royal Statistical Society, Series A* **143**(1) 43–48.
- French, S. 1986. Calibration and the expert problem. *Management Science* **32**(3) 315–321.

- Friedman, J H. 2001. Greedy function approximation: A gradient boosting machine. *Annals of Statistics* **29**(5) 1189–1232.
- Genest, C, J V Zidek. 1986. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science* **1**(1) 114–135.
- Gneiting, T, A E Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* **102**(477) 359–378.
- Hand, D J, R J Till. 2001. A simple generalisation of the area under the ROC curve for multiple class classification problems. *Machine Learning* **45**(2) 171–186.
- Hastie, T, R Tibshirani. 1986. Generalized additive models. *Statistical Science* **1**(3) 297–318.
- Hora, S C. 2004. Probability judgments for continuous quantities: Linear combinations and calibration. *Management Science* **50**(5) 597–604.
- IARPA. 2018. Geopolitical forecasting challenge. <https://www.iarpa.gov/challenges/gfchallenge.html>.
- Karmarkar, U S. 1978. Subjectively weighted utility: A descriptive extension of the expected utility model. *Organizational Behavior and Human Performance* **21**(1) 61–72.
- Kim, O, S C Lim, K W Shaw. 2001. The inefficiency of the mean analyst forecast as summary of forecast of earnings. *Journal of Accounting Research* **39** 329–335.
- Larrick, R P, J B Soll. 2006. Intuitions about combining opinions: Misappreciation of the averaging principle. *Management Science* **52**(1) 111–127.
- Lindley, D V, A Tversky, R V Brown. 1979. On the reconciliation of probability assessments. *Journal of the Royal Statistical Society, Series A* **142**(2) 146–180.
- Marinovic, I, M Ottaviani, P N Sorensen. 2013. Forecasters’ objectives and strategies. G Elliott, A Timmermann, eds., *Handbook of Economic Forecasting*, vol. 2. Elsevier, Amsterdam, Netherlands, 690–720.
- McConway, K J. 1981. Marginalization and linear opinion pools. *Journal of the American Statistical Association* **76**(374) 410–414.
- Mellers, B, L Ungar, J Baron, J Ramos, B Gurcay, K Fincher, S E Scott, D Moore, P Atanasov, S A Swift, T Murray, E Stone, P E Tetlock. 2014. Psychological strategies for winning a geopolitical forecasting tournament. *Psychological Science* **25**(5) 1106–1115.
- Mineo, A M, M Ruggieri. 2005. A software tool for the exponential power distribution: The normalp package. *Journal of Statistical Software* **12**(4) 1–24.
- Morris, P A. 1974. Decision analysis expert use. *Management Science* **20**(9) 1233–1241.
- Murphy, A H, R L Winkler. 1987. A general framework for forecast verification. *Monthly Weather Review* **115**(7) 1330–1338.
- Nelder, J A, R W M Wedderburn. 1972. Generalized linear models. *Journal of the Royal Statistical Society, Series A* **135**(3) 370–385.

- Raiffa, H, R Schlaifer. 1961. *Applied Statistical Decision Theory*. Harvard Business School, Boston, MA.
- Ranjan, R, T Gneiting. 2010. Combining probability forecasts. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **72**(1) 71–91.
- Satopää, V A, J Baron, D P Foster, B A Mellers, P E Tetlock, L H Ungar. 2014. Combining multiple probability predictions using a simple logit model. *International Journal of Forecasting* **30** 344–356.
- Satopää, V A, R Pemantle, L H Ungar. 2016. Modeling probability forecasts via information diversity. *Journal of the American Statistical Association* **111**(516) 1623–1633.
- Shlomi, Y, T S Wallsten. 2010. Subjective recalibration of advisors’ probability estimates. *Psychonomic Bulletin & Review* **17**(4) 492–498.
- Smyth, P, D Wolpert. 1999. Linearly combining density estimators via stacking. *Machine Learning* **36**(1-2) 59–83.
- Soury, H, M S Alouini. 2015. New results on the sum of two generalized Gaussian random variables. *2015 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*. IEEE, 1017–1021.
- Stone, M. 1961. The opinion pool. *The Annals of Mathematical Statistics* **32**(4) 1339–1342.
- Subbotin, M T. 1923. On the law of frequency of error. *Matematicheskii Sbornik* **31**(2) 296–301.
- Tetlock, P, D Gardner. 2016. *Superforecasting: The Art and Science of Prediction*. Broadway Books, New York, NY.
- Tibshirani, R. 1996. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B* **58**(1) 267–288.
- Turner, B M, M Steyvers, E C Merkle, D V Budescu, T S Wallsten. 2014. Forecast aggregation via recalibration. *Machine Learning* **95**(3) 261–289.
- West, M, J Harrison. 1997. *Bayesian Forecasting and Dynamic Models, 2nd ed.*. Springer, New York, NY.
- Winkler, R L. 1968. The consensus of subjective probability distributions. *Management Science* **15**(2) B61–B75.
- Winkler, R L. 1994. Evaluating probabilities: Asymmetric scoring rules. *Management Science* **40**(11) 1395–1405.
- Winkler, R L, V R R Jose. 2008. Comments on: Assessing probabilistic forecasts of multivariate quantities, with an application to ensemble predictions of surface winds. *Test* **17** 251–255.
- Winkler, R L, R M Poses. 1993. Evaluating and combining physicians’ probabilities of survival in an intensive care unit. *Management Science* **39**(12) 1526–1543.
- Wolpert, D H. 1992. Stacked generalization. *Neural Networks* **5**(2) 241–259.
- Zhang, Z, S Wang, D Liu, M I Jordan. 2012. EP-GIG priors and applications in Bayesian sparse learning. *Journal of Machine Learning Research* **13** 2031–2061.