

Perspectives on Constrained Forecasting

Mike West*

Abstract. This expository paper discusses Bayesian decision analysis perspectives on problems of constrained forecasting. Foundational and pedagogic discussion contrasts decision analytic approaches with the traditional, but typically inappropriate, inferential approach. Illustrative examples include development of novel constrained point forecasting and entropic tilting methodology to explore consistency of a predictive distribution with an imposed or hypothesized constraint. Linear, aggregate constraints define illuminating examples that relate to broadly important problems involving aggregate and hierarchical constraints in commercial and economic forecasting. Discussion explores the impact of different loss functions, questions of how constrained forecasting is impacted by dependencies among outcomes being predicted, and promotes the broader use of decision analysis including routine evaluation of predictive distributions of loss under chosen forecasts/decisions. Extensions to more general constrained forecasting problems, connections with broader interests in forecast reconciliation and other considerations are noted.

Keywords: Bayesian forecasting, Bayesian predictive synthesis, distributions of loss, entropic tilting, forecast reconciliation, hierarchical forecasting, multivariate time series forecasting, optimization.

MSC2020 subject classifications: 62C05, 62C10, 62F15, 62F30.

1 Introduction

There are many areas in which it is of interest to assess a constraint, or set of constraints, against a posterior or predictive distribution. The goals may include conditioning a distribution on constraints as well as understanding compatibility of constraints with a given distribution. This paper discusses this general question using simple examples; the perspective is foundational and pedagogic. The area is open for new research and methodology development, and this is anticipated with discussion of decision theoretic alternatives to traditional probabilistic approaches that are typically inappropriate.

The examples focus on constrained forecasting in which a constraint is to be explored and/or imposed on a given predictive distribution. Motivating settings are in commercial and economic systems where problems of aggregate and constrained forecasting are common and important. Broad issues include consistency of forecast models, distributions and point forecast selection at different levels of aggregation in time and/or dimension, and conditioning forecasts for sets of series on those of others, often higher-level aggregates. The Bayesian forecasting literature has a long history in these areas, but core challenges remain and are increasingly important in large-scale forecasting with many intersecting levels of aggregation. Roles for Bayesian decision analytic approaches have

arXiv: [2007.11037](https://arxiv.org/abs/2007.11037)

*Department of Statistical Science, Duke University, Durham NC 27708, U.S.A., mwest@duke.edu

been relatively unexploited; this paper highlights new opportunities for methodological and applied progress using decision analysis.

Among practitioners of Bayesian forecasting, conditioning forecast distributions on assumed values of totals or aggregates has been routine for decades, with early work going back to the 1970s in formal models. Detailed discussion, with references, can be found in West and Harrison (1997, section 16.3). Such approaches condition predictive distributions on assumed constraints, i.e., take a purely probabilistic– or, more broadly inferential– view that constraints are information to condition upon (e.g. Green and Harrison, 1973; de Alba, 1988, 1992, 1993; West and Harrison, 1997). The increased interest in constrained and hierarchical forecasting with major scaling of data, time series and complex hierarchies has continued to build on these foundations. Beyond fully subjective Bayesian approaches, the field has explored related Bayesian moment-based approaches (de Alba, 2006) as well as related non-Bayesian approaches (Guerrero, 1989; Guerrero and Nieto, 1999; Wickramasuriya et al., 2019). One main theme in such approaches is to exploit variants of constrained least-squares or linear Bayes’ methods (Goldstein and Wooff, 2007). Technically these approaches share features with the inferential results in constrained multivariate normal distributions, while the perspective is again that of adjusting inferences in estimation/prediction settings.

There are both foundational and technical challenges with the solely inferential approach. A first and main point is that probabilistic conditioning is generally not justified in settings where constraints are imposed as a matter of intervention. Implicit in the probabilistic approach is routine Bayesian learning under which realized constraint values– totals or other aggregates– arise as random draws from the underlying model distributions for outcome quantities to be constrained. This is not often the case in applications. There, assumed constraint values are often chosen to explore “what-if” implications, so are imposed externally on the initial forecast distribution. In other contexts they are taken as representative values from a different, external model with a view to understand implications on relevant forecast values for the initial context. Echoing Lindley (1992), this argues, in part, for a decision analytic view, at least as a complement to the traditional inferential view. In addition to providing an alternative view of the problem that may have technical advantages as well as conceptual motivation, decision analysis allows for investigation of full predictive uncertainties of loss. Exploring ranges of potential losses under a constrained predictive distribution– rather than just following the typical “act on the optimal decision”– can highlight potentially important practical considerations. This is well-recognized in areas such as finance (e.g., “value-at-risk” studies) but is under-appreciated in other areas. In the development of a decision-analytic approach to imposing constraints on multivariate forecast distributions, dependencies among variables being forecast can profoundly impact on the implied loss distributions and hence should be routinely explored.

From a technical viewpoint, probabilistic conditioning of joint distributions on totals or other aggregates becomes challenging outside of normal or related least squares approaches. Increasingly prevalent contexts involve time series of non-negative counts (e.g. Chen and Lee, 2017; Chen et al., 2018; Aktekin et al., 2018; Chen et al., 2019; Berry and West, 2020; Berry et al., 2020; West, 2020). In such settings, normal/linear Bayes’

approaches are inappropriate and, if applied, can generate misleading results such as negative point forecasts and always non-integer values. Further challenges arise with increasingly high-dimensional time series for which posited deterministic constraints are increasingly likely to represent outcome regions that are “rare events” (i.e., out in the tails) of joint forecast distributions. This raises foundational and practical questions of how, when and whether to proceed to condition probability forecasts. Finally, predictive distributions are often represented as Monte Carlo samples. Conditioning Monte Carlo representations of joint distributions on constraints can be addressed in various ways, such as with importance sampling or adaptive ABC-style approaches (Bonassi and West, 2015, and references therein). However, such approaches are, inherently limited theoretically (e.g. Li et al., 2013) and very challenged in realistic applied contexts of even modest dimensions.

With this background, the current paper explores conditioning using Bayesian decision analysis to contrast with the probabilistic approach. This begins with discussion of optimal Bayesian point forecast selection subject to constraints. Beyond the basics of exploring the roles of different loss functions and the optimization questions to compute constrained point forecasts, this discussion emphasizes the broader pay-offs of a full Bayesian analysis that can come from exploring the full predictive distributions of outcome losses. The second general decision analysis development involves entropic tilting (ET). This defines a nice methodological bridge between traditional probabilistic conditioning and the developments of Bayesian point forecast and optimization approaches. This development also exemplifies the contrast of direct probabilistic conditioning with the complementary decision analytic approach, showing how the latter can broaden perspective and expand available technical methodology.

One specific motivating applied setting is that of commercial revenue forecasting in large consumer sales companies. The examples draw on this in connection with company revenues that are naturally hierarchically structured. For example, per time period revenues in a large supermarket system can be considered at the company-wide level, but these disaggregate to regional, supermarket store groups within region, individual stores and then categories of items sold within stores. A canonical setting in which constrained forecasting interests arise is that of “what-if?” analyses. One example begins with a joint model of revenue flows across the collection of store groups categorized geographically, and then asks questions about how the joint forecast distribution would be modified conditional on various constraints that represent either potential policy-decisions or “what-if?” values of overall revenues aggregated across the store groups. The former might be, for example, a policy decision to implement a system-wide promotion campaign that has initial costs and a specified target goal in terms of expected revenue increase/impact. The latter might be, for example, a point forecast of system-wide revenue (or revenue inflation) based on a macro-level model incorporating assumptions about longer-term customer growth, inflation, and other factors not already integrated into the finer-scale, store-group model and its resulting predictions.

Section 2 discusses constrained forecasting in the simplest setting of overlaying a sum constraint on an outcome vector to be predicted. With examples, this section discusses the traditional probabilistic approach, Bayesian decision analysis perspectives to define

constraint point forecasts, and introduces the Bayesian entropic tilting approach. Section 3 details relevant loss functions, highlighting some of central use and importance in commercial and economic forecasting with positive outcomes, with additional technical development in the Supplementary Material (West, 2023). Section 4 discusses examples using multivariate lognormal forecast distributions, including a 100-dimensional example related to applications in constrained commercial forecasting. Section 5 concludes with general comments and discussion of extensions. Supporting technical details and additional illustrative examples are given in the Supplementary Material.

2 Total Constrained Forecasting

2.1 Setting and Background

At one time point in a forecasting analysis of a set of n series, the outcome of interest $\mathbf{y} = (y_1, \dots, y_n)'$ has predictive distribution $P(\mathbf{y})$ with margins $P_i(y_i)$, ($i = 1:n$). The corresponding p.d.f.s are $p(\mathbf{y})$ and $p_i(y_i)$ whether the distributions are discrete, continuous or mixed; the density (p.d.f.) terminology is used with the corresponding general Stieltjes notation for expectations with the understanding that this covers all cases. Key interests are in discrete time series including binary and non-negative counts as well as in more traditional contexts where the y_i are continuous and often positive.

Total constrained forecasting of \mathbf{y} is the simplest but most important example context of linear constrained analysis. This conveys the foundational concepts and issues, provides access to some analytically tractable examples that generate insights, and forms the basis of more general cases based on sets of constraints, including hierarchical constraints. In an hierarchical context, conditioning forecasts of the y_i may be defined via a cascade in which the forecasts of Y are themselves based on higher-level totals or other aggregates, with obvious recursive extension to multi-level hierarchies.

With $\mathbf{1} = (1, \dots, 1)'$ as the n -vectors of ones, let $Y = \mathbf{1}'\mathbf{y} = \sum_{i=1:n} y_i$ with implied distribution $P(Y)$ and p.d.f. $p(Y)$. Interest lies in forecasting \mathbf{y} given a specific value F for this total. The value F may be a chosen point forecast, such as $E(Y) = F$ under $P(\cdot)$, a point forecast generated from some external model or source, or a “what-if?” value from a set being explored to understand impact of potential total constraints on the y_i . It may also be just one value generated from an external or alternative model for Y alone, representing one of a set of Monte Carlo draws against which interest lies in understanding implications for \mathbf{y} based on $p(\mathbf{y})$ coupled with the information defined by that external source.

2.2 Traditional Inferential Perspective: Probabilistic Conditioning

The traditional Bayesian/probabilistic view is that conditioning on the value of F is a purely inferential question, and that it implies modifying $P(\mathbf{y})$ to $P(\mathbf{y}|Y = F)$, the conditional distribution of \mathbf{y} given $Y = F$. The implied analysis identifies $P(\mathbf{y}|Y)$ for any total Y , and then plugs-in the value $Y = F$ (West and Harrison, 1997, section 16.3).

Theoretical Examples

Examples 1: Normal, and Other Elliptically Symmetric Cases. If $P(\mathbf{y})$ is given by $\mathbf{y} \sim N(\mathbf{m}, \mathbf{V})$, then $Y \sim N(M, w)$ where $M = \mathbf{1}'\mathbf{m}$ and $w = \mathbf{1}'\mathbf{c}$ in which \mathbf{c} is the covariance vector $\mathbf{c} \equiv C(\mathbf{y}, Y) = \mathbf{V}\mathbf{1}$. It follows that $P(\mathbf{y}|Y = F)$ is (singular) normal with mean $\mathbf{m}_F = \mathbf{m} + \mathbf{c}(F - M)/w$ and (singular) variance matrix $\mathbf{V}_F = \mathbf{V} - \mathbf{c}\mathbf{c}'/w$.

If $P(\mathbf{y})$ is a multivariate T, or other elliptically symmetric distribution, the location of $P(\mathbf{y}|Y = F)$ is modified as in the normal case, while dispersion depends on F and increases in $|M - F|$. The construction of elliptically symmetric distribution as normal scale mixtures defines the underlying probability calculus and interpretation. For example, $\mathbf{y} \sim T_k(\mathbf{m}, \mathbf{V})$, implies $\mathbf{y}|Y = F \sim T_k(\mathbf{m}_F, v_F\mathbf{V}_F)$ —now singular T_k —with $\mathbf{m}_F, \mathbf{V}_F$ as above and $v_F = \{k + (F - M)^2/q\}/(k + n)$. Uncertainty in $P(\mathbf{y}|Y = F)$ naturally inflates as a function of the lack of concordance of the value of F with $p(Y)$.

Examples 2: Lognormal and Log-T Cases. Major areas of application in business and economic analysis use conditionally normal, linear models—such as dynamic linear models (DLMs: West and Harrison, 1997; Prado et al., 2021)—for log transformed data. Implied predictive distributions are log-T distributions on the original \mathbf{y} data scale. Normal approximations may be used, but are inadequate in contexts of restricted T degrees of freedom such as arise routinely, for example, in multivariate volatility modelling (e.g., Prado et al., 2021, chapter 10, and West, 2020, section 2). The challenge then is that $p(Y)$ and $p(\mathbf{y}|Y)$ are not available analytically, so raising difficult questions of computation. This is a severe constraint generally, but particularly when interest lies in fast and scalable analysis to accurately evaluate aspects of $p(\mathbf{y}|Y = F)$.

Examples 3: Discrete Cases. Similar comments apply to distributions arising in increasingly large-scale models for discrete time series, including binary and non-negative count data (Berry and West, 2020; Berry et al., 2020). Even in relatively simple models based on conditional Poisson forms for univariate series, dependencies across series destroy the ability to evaluate joint and conditional distributions analytically.

Simulation-Based Analysis

In many realistic models $P(\mathbf{y})$ is represented via a Monte Carlo sample, so that evaluating and using $p(\mathbf{y}|Y)$ is a major challenge. Approaches such as adaptive importance sampling (e.g. West, 1993) and sequential Monte Carlo including approximate Bayesian computation (ABC: e.g. Bonassi and West, 2015) can be considered. However, such methods do not reliably deal in any generality with the problem of conditioning on totals, or other aggregates, in problems of practicable dimension. The problem has been considered in other areas too, and shown to be a very major challenge as well as NP-hard (non-deterministic polynomial-time hard) in discrete contexts (Li et al., 2013); it stands as an open challenge to computational statistics. Again, applied contexts increasingly require fast, reliable and scalable analysis, which is simply not (yet) available.

In low-dimensional problems, a vanilla ABC-style accept/reject method can sometimes prove useful. Such an approach defines a subspace S of the sample space of \mathbf{y}

consistent with values of $Y = \mathbf{1}'\mathbf{y}$ “near” F . Simulated values $\mathbf{y} \sim P(\mathbf{y})$ are then accepted if and only if $Y \in S$. The example in Section 4 has positive Y, F and takes $S = \{\mathbf{y} : (1 - \tau)F \leq Y \leq (1 + \tau)F\}$ for positive, small τ so that the maximum percentage error 100τ defines “nearness” of simulated Y values to the target F . Accepted values are then approximately distributed as $P(\mathbf{y}|Y = F)$ with approximation based on τ . The acceptance probability $p_S = Pr(\mathbf{y} \in S)$ under $P(\mathbf{y})$ can be trivially estimated from the Monte Carlo samples and gives a guide to how effective the approach is given any value of τ .

2.3 Decision-Guided Probabilistic Conditioning: Entropic Tilting

Imposing constraints in forecasting is often essentially not an inference problem. Asking questions about how to forecast \mathbf{y} given the total $Y = F$ imposed from an external model or source moves outside the formal probability model; the imposed value of $Y = F$, or a collection of values to consider, did not arise from the $p(Y)$ implied by $p(\mathbf{y})$. The value of F is imposed by intervention, so that asking about how fixing $Y = F$ should impact forecasts for \mathbf{y} is more naturally a decision question. At the least, exploring decision analysis perspectives is an opportunity to broaden the framework and examine approaches complementary to the traditional, probabilistic view. Entropic tilting (ET) is one fully Bayesian decision analysis approach.

Entropic Tilting and Moment Constraints

With respect to the baseline distribution $P(\mathbf{y})$, ET is a maximum entropy related approach that aims to choose a distribution $G(\mathbf{y})$ as a modification of $P(\mathbf{y})$ that satisfies a set of expectation constraints $E_g[\mathbf{q}(\mathbf{y})] = \mathbf{0}$ where $\mathbf{q}(\mathbf{y})$ is a q -vector of functions of \mathbf{y} . An example with $q = 2$ takes $\mathbf{q}(\mathbf{y}) = (Y - F, Y^2 - F^2 - S)'$ with $Y = \mathbf{1}'\mathbf{y}$ and some specified F and $S > 0$; then $G(\mathbf{y})$ implies a distribution for Y that has mean F and variance S (relevant only, of course, in contexts where variances exist). ET is an explicit decision analytic approach that aims to select $G(\mathbf{y})$ “close” to $P(\mathbf{y})$ subject to the expectation constraints; ET defines “close” in terms of the Kullback-Leibler divergence (KLD) of $P(\mathbf{y})$ from $G(\mathbf{y})$. With p.d.f.s $p(\mathbf{y})$ and $g(\mathbf{y})$, the solution is the exponentially-tilted form $g(\mathbf{y}) = c_\gamma \exp\{\boldsymbol{\gamma}'\mathbf{q}(\mathbf{y})\}p(\mathbf{y})$ where the q -vector $\boldsymbol{\gamma}$ explicitly enforces the expectation constraints and $c_\gamma > 0$ is the normalizing constant. Detailed theoretical background, including well-known aspects of the theory of KLD, foundations of ET and illuminating examples, appear in Tallman and West (2022b).

The KLD optimal $\boldsymbol{\gamma}$ can typically be numerically computed using a Newton-Raphson (NR) algorithm to solve the q -vector equation $\int_{\mathbf{y}} \mathbf{q}(\mathbf{y}) \exp\{\boldsymbol{\gamma}'\mathbf{q}(\mathbf{y})\}p(\mathbf{y})d\mathbf{y} = \mathbf{0}$, with NR utilizing the second derivative matrix $\int_{\mathbf{y}} \mathbf{q}(\mathbf{y})\mathbf{q}(\mathbf{y})' \exp\{\boldsymbol{\gamma}'\mathbf{q}(\mathbf{y})\}p(\mathbf{y})d\mathbf{y}$ (with summations replacing integrals in cases of discrete $P(\mathbf{y})$, of course). In introducing ET to the econometrics community, one of the key contributions of Robertson et al. (2005) was to note that Bayesian analysis of $P(\mathbf{y})$ is often/typically based on Monte Carlo simulation; this makes the integrals in such optimization computations accessible via direct Monte Carlo approximation.

In recent years the use of ET in Bayesian econometrics has expanded with applications in policy-oriented macro-economics, in particular (e.g. Krüger et al., 2017; Koop et al., 2019). These and other authors address questions of constraining forecast distributions based on information from external sources, alternative models (e.g., when using models for time series at different frequencies), and related sources such as forecast surveys. The main technical use of ET in these areas has been to impose expectation constraints on some subset of variables being forecast; that is, to modify a predictive $P(\mathbf{y})$ so that the mean of some (possibly vector) function of \mathbf{y} is constrained to an imposed value. The connections with the broad area of Bayesian (and other) forecast combination methods is also evident here; some of these works explicitly take the view that ET is a tool for “reconciling” forecasts from different models by matching specific sets of forecast means across models. This is naturally related to the constrained forecasting perspectives here, while the following development takes the ET approach in new directions including that of exploring “almost exact” constraints.

Novel Development of ET for Conditioning on Constraints

ET can be used to define $G(\mathbf{y})$ to approximately conform with deterministic constraints. Focus now explicitly on the context of a single constraint (whether linear or non-linear), so $q = 1$, $\mathbf{q}(\mathbf{y}) = q(\mathbf{y})$ is a scalar function and $\gamma = \gamma$ is scalar. Take the specific choice $q(\mathbf{y}) = I(\mathbf{y} \in S) - (1 - \epsilon)$ for some subspace S of the sample space of \mathbf{y} , indicator function $I(\cdot)$ and a specified probability ϵ . Applying ET then yields a distribution $G(\mathbf{y})$ under which $Pr(\mathbf{y}_i \in S) = 1 - \epsilon$. This is very general and can be used, for example, to map $P(\mathbf{y})$ to a $G(\mathbf{y})$ having a specified median. Applying this to approximate the constrained problem of the current paper involves taking: (i) the subspace S very small and concentrated around a constrained value of a deterministic function of \mathbf{y} ; and (ii) the tolerance ϵ small, i.e., $1 \gg \epsilon > 0$. The resulting $G(\mathbf{y})$ approximately satisfies the constraint with level of approximation defined by how small S and ϵ are. The example in Section 4 for the context with $Y = \mathbf{1}'\mathbf{y}$ takes $S = \{\mathbf{y} : (1 - \tau)F \leq Y \leq (1 + \tau)F\}$ where $1 \gg \tau > 0$; here τ defines concentration of S based on percentage error of Y from F being no more than 100τ .

In this very specific ET context, the solution $G(\mathbf{y})$ can be directly evaluated without resort to numerical optimization. Note that the p.d.f is $g(\mathbf{y}) \propto \exp\{\gamma I(\mathbf{y} \in S)\}p(\mathbf{y})$, or $g(\mathbf{y}) = c\{\exp(\gamma)I(\mathbf{y} \in S) + 1 - I(\mathbf{y} \in S)\}p(\mathbf{y})$ where c is the normalizing constant; clearly, $c^{-1} = \exp(\gamma)p_s + 1 - p_s$ where $p_s = Pr(\mathbf{y} \in S)$ under $P(\mathbf{y})$. Then, since the expectation of $q(\mathbf{y})$ under $G(\cdot)$ is constrained to be $1 - \epsilon$, the optimal γ is the solution to $1 - \epsilon = c \exp(\gamma)p_s$; this is easily solved to give $\gamma = \log\{(1 - \epsilon)(1 - p_s)/(\epsilon p_s)\}$. Note that this is decreasing in both ϵ and p_s , and will generate $\exp(\gamma) \gg 1$ when ϵ and/or τ are small, consistent with an increasingly binding constraint.

Considering cases when $P(\mathbf{y})$ is represented as a Monte Carlo sample, this ties intimately with the ABC-style accept/reject approach noted above. Suppose, with no loss of generality, that the Monte Carlo sample is equally weighted. The ABC analysis leads to a weighted sample in which a sampled vector $\mathbf{y} \sim P(\mathbf{y})$ is given a weight of 1 if $\mathbf{y} \in S$, 0 otherwise. The ET analysis defines weights proportional to $\exp(\gamma)$ for samples $\mathbf{y} \in S$, and proportional to 1 otherwise. As noted, γ will tend to be large in

practice, so that this ET reweighting is close to accept/reject. It can further be shown that, with a large Monte Carlo sample size, the effective sample size of the normalized ET weights converges around $c^{-2}/\{\exp(2\gamma)p_s + 1 - p_s\}$; in practical cases with $\gamma \gg 1$, this is approximately p_s , the acceptance probability of the ABC-style analysis.

2.4 Point Forecast Decision Analysis Perspective

More traditional Bayesian decision analysis approaches focus on optimal point forecast selection under $P(\mathbf{y})$ but now subject to the imposed constraints.

Let $L(\mathbf{y}, \mathbf{f})$ be a loss function chosen to score a point forecast vector \mathbf{f} of outcome \mathbf{y} . Standard Bayesian decision analysis chooses that point forecast vector \mathbf{f}^* that minimizes the expected loss subject to the constraints. That is,

$$\mathbf{f}^* = \operatorname{argmin}_{\mathbf{f}} R(\mathbf{f}) \quad \text{subject to} \quad \mathbf{1}'\mathbf{f} = F, \quad \text{where} \quad R(\mathbf{f}) = \int_{\mathbf{y}} L(\mathbf{y}, \mathbf{f}) dP(\mathbf{y}). \quad (2.1)$$

The following section concerns technical developments using specific loss functions. Some general comments on loss function structure and the broader applied perspective on Bayesian decision analysis are first noted.

Additive Losses. In many applied problems, loss functions will be additive over outcomes, i.e., $L(\mathbf{f}, \mathbf{y}) = \sum_{i=1:n} L_i(y_i, f_i)$ for individual loss functions $L_i(\cdot, \cdot)$ in each dimension. In forecasting sales or demand for sets of items i , for example, the translation to revenue (gained or lost) per item is a primary consideration. Other contexts might extend to loss functions reflecting cross-item scores. Development below focuses on additive loss functions, leaving extensions to the reader and future, customized applications.

Generalizations with Multiple Constraints. The broader class of problems for hierarchical and other sets of constraints simply extends the above formulation to involve the constraint $\mathbf{A}'\mathbf{f} = \mathbf{F}$ where \mathbf{A} is a specified $n \times k$ matrix of full rank $k < n$, and \mathbf{F} is a specified k -vector. Analysis then targets minimization of $R(\mathbf{f})$ subject to these k constraints. For example, constraints on sets of intersecting subtotals can be defined by a matrix \mathbf{A} of zeros and ones, while other, more general weighted averages are obvious extensions.

Broader View: Distributions of Loss. Section 1 has already raised the central question and potential importance of exploring predicted loss distributions, i.e., considering aspects of $P(L(\mathbf{y}, \mathbf{f}))$ implied under $P(\mathbf{y})$ for $\mathbf{f} = \mathbf{f}^*$ (and possibly other values, perhaps “close to” \mathbf{f}^*). This perspective is one of evaluation and presentation of uncertainties in loss outcomes in decision analysis akin to the usual “uncertainty quantification” view in inference. Again, this simply argues for the broader view of decision analysis in exploring loss distributions, and this is a point of emphasis throughout this paper in the specific settings of constrained forecasting.

3 Decision Analysis and Classes of Loss Functions

3.1 Lagrangian Formulation

The Lagrangian formulation for optimization in (2.1) is to choose (\mathbf{f}, λ) to minimize

$$R(\mathbf{f}) + \lambda(F - \mathbf{1}'\mathbf{f}) \quad (3.1)$$

with a real-valued Lagrange multiplier λ . Assuming a minimizing solution $\mathbf{f}^*(\lambda)$ given any allowable value of λ , solving $F = \mathbf{1}'\mathbf{f}^*(\lambda)$ for λ^* defines the optimal forecast vector $\mathbf{f}^* \equiv \mathbf{f}^*(\lambda^*) = (f_1^*, \dots, f_n^*)'$.

Details for specific loss functions commonly used in forecasting applications are noted below (with additional technical details in Supplementary Material). The standard use in unconstrained Bayesian decision analysis— in forecasting and parameter estimation— is background (e.g. French and Insua, 2010; Smith, 2010). A main interest is to present examples of optimal constrained forecasts for such loss functions, and highlight differences and implications. As noted above, the development uses an additive loss function $L(\mathbf{f}, \mathbf{y}) = \sum_{i=1:n} L_i(y_i, f_i)$ so that

$$R(\mathbf{f}) = \sum_{i=1:n} R_i(f_i) \quad \text{with} \quad R_i(f_i) = \int_{y_i} L_i(y_i, f_i) dP_i(y_i) dy_i, \quad i = 1 : n, \quad (3.2)$$

where $P_i(y_i)$ is the marginal predictive distribution of y_i . The resulting \mathbf{f}^* does not involve dependencies among the y_i . However, it is critical for practical application to be aware that the resulting distributions of loss at the optimum (or at any other value of \mathbf{f}) are of course very much impacted by the joint structure of $P(\mathbf{y})$, as examples in Section 4 below illustrate.

In most practical contexts, there is no direct analytic solution to the implied optimization problem; numerical methods are needed. Assuming $\mathbf{f}^*(\lambda)$ is available for any λ , the optimal λ^* is solution to $q(\lambda) = 0$ where $q(\lambda) = \mathbf{1}'\mathbf{f}^*(\lambda) - F$. A direct Newton-Raphson (NR) algorithm is typically most efficient and effective in solving this, relying on the derivative function $\dot{q}(\cdot)$. The basis of NR iterations is as follows.

- *Initialize:* Set iterate count $t = 0$, and Lagrange multiplier value $\lambda = \lambda^0$, a chosen initial value; set $F^0 = \mathbf{1}'\mathbf{f}^*(\lambda^0)$.
- *Iterate:* For steps $t \geq 1$, compute $\lambda^t = \lambda^{t-1} - q(\lambda^{t-1})/\dot{q}(\lambda^{t-1})$, then update the implied $\mathbf{f}^*(\lambda^t)$ and the sum $F^t = \mathbf{1}'\mathbf{f}^*(\lambda^t)$.
- *Stop:* When changes in the sequence of scalars λ^t and/or $|F - F^t|$ become “small enough”, set $\lambda^* = \lambda^t$ and $\mathbf{f}^* = \mathbf{f}^*(\lambda^*)$, and stop.

Examples in Section 4 utilize this, with NR iterations typically converging very fast. Depending on the chosen loss function, the range of values of λ is restricted. This allows the analysis to be self-monitoring in that NR iterates moving λ to a lower or upper bound would indicate incompatibility of the conditioning value F with the predictive

distribution $P(Y)$. Such contexts are those in which enforcing the constraint might be questioned, and the algorithm will signal that. Finally, in contexts where predictions are based on Monte Carlo samples from $P(\mathbf{y})$, implied Monte Carlo estimates will be used to evaluate $\mathbf{f}^*(\lambda)$ via direct, weighted or importance sampling.

3.2 Squared Error Loss

Squared error (SE) loss is not of main applied interest in commercial forecasting applications compared to other choices noted below. However, details are tractable and illuminating. SE loss is, of course, restricted to models in which $P(\mathbf{y})$ has finite second-order moments. Supposing this, let m_i be the mean of $P_i(y_i)$ and $\mathbf{m} = (m_1, \dots, m_n)'$ with sum $M = \mathbf{1}'\mathbf{m}$. The m_i are optimal unconstrained point forecasts under SE.

Take $L_i(y_i, f_i) = (y_i - f_i)^2/c_i$ where the $c_i > 0$ can represent different scales or simply different weightings of forecast errors across the n outcomes. Write $\mathbf{c} = (c_1, \dots, c_n)'$ and $C = \mathbf{1}'\mathbf{c}$. Then simple quadratic optimization yields $f_i^*(\lambda) = m_i + \lambda c_i/2$ for each $i = 1:n$. Imposing the total constraint yields $\lambda^* = 2(F - M)/C$ and thus $f_i^* = m_i + (F - M)c_i/C$. These are the marginally optimal means m_i corrected by the term $(F - M)c_i/C$; this naturally represents an upward (downward) correction if F exceeds (falls short of) the forecast mean of the total $E(Y) = M$. While natural, it is clear that the scope for relevant application is proscribed; in addition to earlier comments on constraints for relevant applications, many applied interests concern integer, count, non-negative or bounded outcomes, and the inherent “constrained least squares” results lead to theoretically optimal forecasts that violate such inherent requirements.

3.3 Absolute Deviation Loss

Absolute Deviation (AD) loss is perhaps the most important and widely used loss function in commercial forecasting as in other areas. Take $L_i(y_i, f_i) = |y_i - f_i|/c_i$ where again $c_i > 0$ are known weights. Assuming finite first moments of the $P_i(y_i)$, it follows that, for any given λ , (3.1) is minimized over the f_i at the values satisfying $2P_i(f_i) - 1 = c_i\lambda$ (see details in Supplementary Material). Thus $f_i^*(\lambda) = P_i^-((1 + \lambda c_i)/2)$ where $P_i^-(\cdot)$ is the inverse c.d.f. (quantile function) for each i , whether discrete or continuous. Note that the usual unconstrained forecast is the median of $P_i(y_i)$ in the case $\lambda = 0$. Otherwise, $f_i^*(\lambda)$ is the $100(1 + \lambda c_i)/2$ percentile of $P_i(y_i)$. Note further that λ must lie in $[-r, r]$ where $r = 1/\max_{i=1:n} c_i$.

Example: Exponential Models. A purely illustrative, analytically tractable example highlights the analysis. Suppose the y_i are marginally exponential, $y_i \sim \text{Exp}(1/m_i)$ with $m_i = E(y_i)$. The marginal medians are $f_i = m_i \log(2)$. Take $c_i = 1$ so $C = n$ and $\lambda \in [-1, 1)$. It follows that $f_i^*(\lambda) = m_i \log(2/(1 - \lambda))$ for each i , and imposing $F = \mathbf{1}'\mathbf{f}^*(\lambda)$ yields $\lambda^* = 1 - 2 \exp(-F/M)$. As a result, the optimal forecast is $\mathbf{f}^* = \mathbf{f}^*(\lambda^*) = \mathbf{m}F/M$. That is, each marginal mean m_i is simply- and very naturally-scaled by the positive constant F/M so that the resulting $f_i^* = m_i F/M$ sum to F .

More generally, the direct Newton-Raphson (NR) algorithm for optimization solves $q(\lambda) = 0$ where $q(\cdot)$ and its derivative $\dot{q}(\cdot)$ are now given by

$$q(\lambda) = \sum_{i=1:n} f_i^*(\lambda) - F \quad \text{and} \quad \dot{q}(\lambda) = \sum_{i=1:n} c_i \{2p_i(f_i^*(\lambda))\}^{-1}.$$

These are easily calculated when the marginal forecast distributions are of parametric forms, and via Monte Carlo approximations using forecast samples in other cases.

3.4 Absolute Percent Error Loss and Variants

APE Loss

For strictly positive outcomes $y_i > 0$, the modification of AD loss to a percent scale defines the absolute percent error (APE) loss that is simply key in commercial applications. APE puts forecast errors on a common scale (percent revenue, percent sales of numbers of items, etc.) so as to enable easy comparisons across outcomes and contexts (e.g. Berry and West, 2020). Assuming $P_i(\cdot)$ has support $y_i > 0$ (perhaps bounded above), take $L_i(y_i, f_i) = |y_i - f_i|/(y_i c_i)$ where again $c_i > 0$ are known weights. Then the risk function component $R_i(f_i)$ for outcome i has the form of the expected value of AD loss $|y_i - f_i|/c_i$ with respect to the modified distribution with density function $g_i(y_i) \propto p_i(y_i)/y_i$. If this defines a p.d.f., then $g_i(y_i) = k_i p_i(y_i)/y_i$ for some normalizing constant $k_i > 0$ and

$$R_i(f_i) = c_i^{-1} \int_{y_i > 0} |y_i - f_i| y_i^{-1} dP_i(y_i) = (c_i k_i)^{-1} \int_{y_i} |y_i - f_i| dG_i(y_i),$$

where $G_i(\cdot)$ is the c.d.f. implied by p.d.f. $g_i(\cdot)$. Hence the AD analysis above applies with each $P_i(\cdot)$ replaced by $G_i(\cdot)$ and the weights c_i replaced by $c_i k_i$. That is, theoretically and in the numerical evaluation using NR, each $f_i^*(\lambda)$ is the $100(1 + \lambda c_i k_i)/2$ percentile of $G_i(y_i)$. The following details and examples are to be noted.

- When $\lambda = 0$ so that the constraint does not apply, the optimal forecasts f_i^* are the medians of the $G_i(\cdot)$, also known as the (-1) -medians of $P_i(\cdot)$. With typical positively skewed distributions on $y_i > 0$, these lie below the medians due to the greater mass at lower values under $G_i(\cdot)$ than under $P_i(\cdot)$. This feature is inherited in the constrained decision analysis as the relevant percentiles of $G_i(\cdot)$ for any given λ will be similarly lower than those of $P_i(\cdot)$.
- Practical models include cases when the predictive distributions have forms related to those of compound shifted Poisson, compound gamma, lognormal and others. As one theoretically tractable example revisited in Section 4 below, suppose that $P_i(\cdot)$ is lognormal, $y_i \sim LN(m_i, v_i)$ with mode, median and mean of y_i given by $\hat{f}_i = \exp(m_i - v_i)$, $\tilde{f}_i = \exp(m_i)$ and $\bar{f}_i = \exp(m_i + v_i/2)$, respectively. It easily follows that $G_i(\cdot)$ is $LN(m_i - v_i, v_i)$ and $k_i = \exp(m_i - v_i/2)$. Note that the (-1) -median of $P_i(y_i)$ is exactly its mode in this case. Percentiles of $G_i(\cdot)$ relevant in the constrained decision analysis solutions can be very substantially smaller than those of $P_i(\cdot)$ when predictions are uncertain.

- In contexts where predictions are based on Monte Carlo samples from $P(\mathbf{y})$, implied Monte Carlo estimates of the percentiles of $G_i(\cdot)$ are easily evaluated using weighted or importance sampling.
- In some cases, this analysis is infeasible as $p_i(y_i)/y_i$ is not integrable, whether available analytically or via simulation. Key cases with real practical importance again include models generating log-T predictive distributions; truncating the distributions to finite ranges is one modification enabling the analysis.

ZAPE Loss

In discrete cases when forecast distributions have non-zero probabilities on $y_i = 0$ for some $i = 1:n$, APE loss is not applicable. Extension to zero-adjusted absolute percent error (ZAPE) loss functions is then of interest (Berry and West, 2020; Berry et al., 2020). Suppose $y_i \geq 0$ and that the predictive distribution has a non-zero point mass $\pi_{i0} = P_i(0)$ at $y_i = 0$. A ZAPE loss function is $L_i(y_i, f_i) = w_i(f_i)I(y_i = 0) + |y_i - f_i|/(y_i c_i)I(y_i > 0)$ where $w_i(f_i) > 0$ penalizes point forecast f_i when $y_i = 0$. Berry and West (2020) show the relevance of ZAPE in forecasting sales of large numbers of consumer items when there are appreciable probabilities of “no sales”. In the current context, the constrained APE analysis is easily extended; the emerging $f_i^*(\lambda)$ may now include exact zero values for some outcomes i across ranges of values of λ .

For example, take $w_i(f_i) = f_i/c_i$ so that a point forecast $f_i = 1$ when $y_i = 0$ is penalized exactly as a point forecast $f_i = 0$ when $y_i = 1$ (Berry and West, 2020). Define the c.d.f. $P_i^+(\cdot)$ for the c.d.f. $P_i(\cdot)$ constrained and renormalized on $y_i > 0$, with corresponding p.d.f. $p_i^+(y_i)$. Then

$$P_i(y_i) = \pi_{i0}I(y_i = 0) + (1 - \pi_{i0})P_i^+(\cdot)I(y_i > 0).$$

Then, define $G_i(\cdot)$ as the c.d.f. with p.d.f. $g_i(y_i) = k_i p_i^+(y_i)/y_i$ on $y_i > 0$ where k_i is the with appropriate normalizing constant. With $w_i(f_i) = f_i/c_i$, the risk component $R_i(f_i)$ satisfies

$$c_i R_i(f_i) = \pi_{i0} f_i + (1 - \pi_{i0}) k_i^{-1} \int_{y_i > 0} |y_i - f_i| dG_i(y_i).$$

It follows that, for any given λ , (3.1) is minimized over the f_i at values given by

$$f_i^*(\lambda) = \begin{cases} 0, & \text{if } u_i(\lambda) \leq 0, \\ G_i^-(u_i(\lambda)), & \text{if } u_i(\lambda) > 0, \end{cases} \quad \text{with } u_i(\lambda) = \frac{1}{2} \left\{ 1 + k_i \frac{(\lambda c_i - \pi_{i0})}{(1 - \pi_{i0})} \right\},$$

and where $G_i^-(\cdot)$ is the inverse of the c.d.f. $G_i(\cdot)$ (see additional details in Supplementary Material). Here λ must lie in $[s, r)$ with bounds given by $s = \max_{i=1:n} \{((k_i + 1)\pi_{i0} - 1)/(c_i k_i)\}$ and $r = \min_{i=1:n} \{((k_i - 1)\pi_{i0} + 1)/(c_i k_i)\}$. The results for APE are confirmed when $\pi_{i0} = 0$ for all i . Otherwise, higher probabilities π_{i0} will lead to optimal point forecasts at zero. Extending to constrained forecasting is particularly interesting in such contexts. Technically, only minor modifications to the NR algorithm arise, with $q(\cdot)$ and its derivative now given by

$$q(\lambda) = \sum_{i=1:n} f_i^*(\lambda) - F \quad \text{and} \quad \dot{q}(\lambda) = \sum_{i=1:n} I(u_i(\lambda) > 0) c_i k_i \{2g_i(f_i^*(\lambda))\}^{-1}.$$

Practically important modifications of the above example of ZAPE include choices of the $w_i(f_i)$ penalty at $y_i = 0$ that less heavily penalize larger values of the point forecasts f_i . In some contexts, the linear in f_i penalty is too dominant for larger values of f_i , pushing the optimal f_i^* to zero more aggressively than desired. In such settings, bounded weight functions such as $w_i(f_i) = f_i/(1 + f_i)$ or $\min\{1, f_i\}$ are more relevant. Examples using these forms can be easily implemented using extensions of the above optimization method, and bear out the effectiveness in reducing the overly aggressive shrinkage to zero of optimal forecasts while adding only modestly to computational load.

4 Illustrative Examples

4.1 General Comments

As discussed above, some motivating applications involve non-negative outcomes in commercial and allied areas. Two illustrative examples reflect this, with multivariate lognormal distributions that allow ranges of dependencies among the y_i . This setting provides access to some analytic tractability that aids in generating insights. Related examples (not shown) using count data in which conditional Poisson models linked via latent factors share similar general features, though lack analytic tractability. The examples touch on differences in constrained point forecasts based on choice of loss function, and on how these vary with dependencies among the y_i . They also focus on aspects of predictive distributions of losses as well as optimal point forecasts, a point stressed earlier that should always be part of the broader Bayesian decision analysis.

4.2 Bivariate Lognormal Example

Setting and Optimal Forecasts

A first set of examples has $n = 2$ so that $\mathbf{y}' = (y_1, y_2)$. The contours in Figure 1 are those of three bivariate lognormal distributions $\mathbf{y} \sim LN(\mathbf{m}, \mathbf{V})$ whose parameters are the mean and variance matrix of the underlying bivariate normal for $(\log(y_1), \log(y_2))'$. The examples have $\mathbf{m}' = (\log(7), \log(14))$, $\text{diag}(\mathbf{V}) = (v_1, v_2) = (0.04, 0.09)$, and the off-diagonal entry of \mathbf{V} is 0.06ρ for dependence parameter $\rho \in (-1, 1)$. The univariate lognormal margins $P_i(y_i)$ have modes— that are also the (-1) -medians— at $\{6.73, 12.80\}$, medians at $\{7, 14\}$ and means at $\{7.14, 14.64\}$. The contours are those of the highest predictive density regions under $P(\mathbf{y})$ with $\{0.01, 0.25, 0.5, 0.75, 0.9, 0.95\}$ probability content. The three examples show contours for the cases of $\rho \in \{-0.7, 0, 0.7\}$.

Under $P(\mathbf{y})$, the sum of medians of the y_i is 21, and the mean is $E(Y) = 21.9$. Figure 1 is based on $F = 14.7$, well into the lower tail of the forecast distribution $P(Y)$ so will lead to larger adjustments to constrained point forecasts. The dashed lines define the total constraint, so optimal point forecasts lie on these lines. The NR algorithm converges in two or three steps to a high degree of precision. Searching for the AD optimal is initialized at marginal medians, and for the APE optimal at marginal (-1) -medians. The figures show the values of the constrained AD, APE and SE point

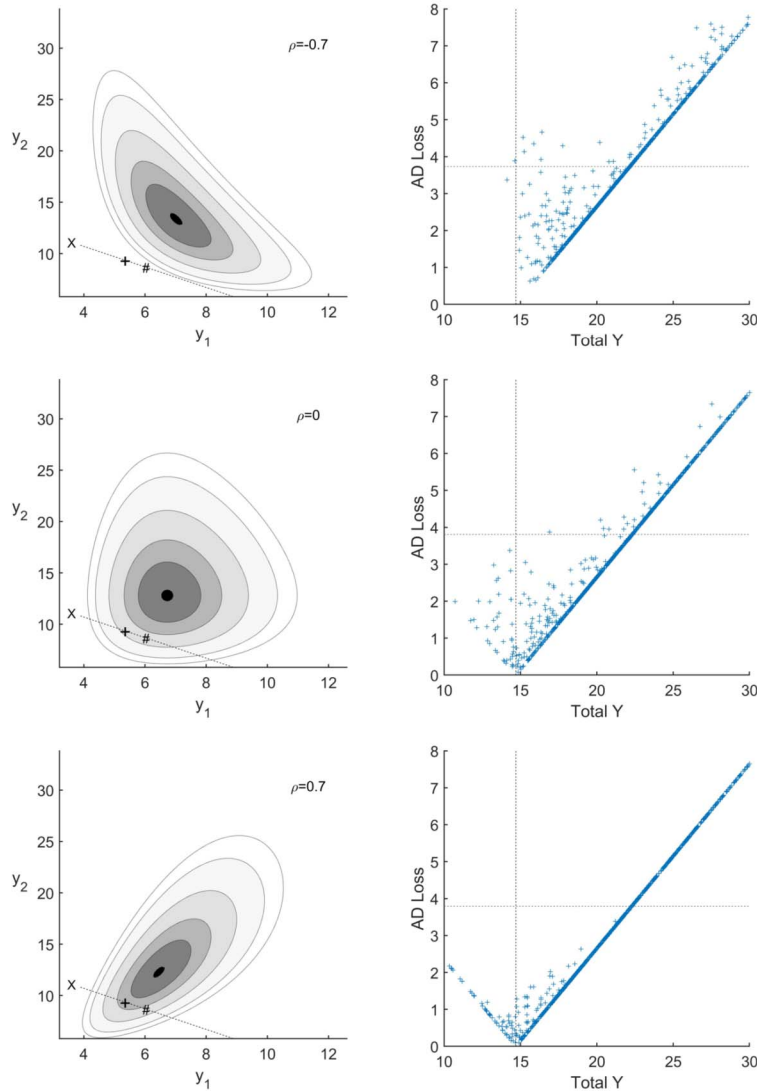


Figure 1: Bivariate lognormal example with $F = 14.7$, lying quite far into lower tail of $P(Y)$. *Left column:* Contours of $p(\mathbf{y})$ for three different levels of dependence $\rho \in \{-0.7, 0, 0.7\}$ and with the constraint $\mathbf{1}'\mathbf{y} = F$ indicated as the dashed line. The symbols indicate the optimal point forecast vector \mathbf{f}^* under AD loss ($+$), APE loss ($\#$) and SE loss (X). While marginal APE optimal forecasts are always lower than those under AD loss, the joint constrained APE optimal forecast can be higher than AD optimal in some dimensions, simply due to the total constraint. *Right column:* Scatter plots of the corresponding joint predictive distributions of the outcome total and the *per dimension* loss at the value of \mathbf{f}^* i.e., a Monte Carlo sample from $P(Y, L(\mathbf{y}, \mathbf{f}^*)/2)$. The vertical dashed line marks the value of the constraint, $Y = F$; the horizontal dashed line marks the value of the expected loss at the minimum, i.e., the optimized risk $R(\mathbf{f}^*)$.

forecasts, showing differences due to the choice of loss function. While marginal (-1) -medians are always lower than marginal medians, the imposition of the constraint will often change the ordering in some dimensions. Note also that the use of SE loss would be questioned in this context. The Supplementary Material provides additional illustration with the same model but now using two other values of F — one in the center of the predictive distribution $P(Y)$ and one in the upper tail— with corresponding summaries.

Loss Distributions

Figure 1 also explores distributions of AD loss. For each value of ρ , Monte Carlo samples of \mathbf{y} give samples from the joint distribution of $\{Y, L(\mathbf{y}, \mathbf{f})\}$ at any chosen \mathbf{f} . The figure shows resulting scatter plots at the AD-optimal $\mathbf{f} = \mathbf{f}^*$, with the simulated loss values scaled by $1/n = 1/2$ so that the vertical axis is on a *per dimension* loss scale. While \mathbf{f}^* is the same in all three cases, the distributions of optimized losses depend on the full joint $P(\mathbf{y})$. The predictive distribution $p(Y)$ is more diffuse for positive values of ρ than for zero or negative values, and this naturally translates into greater dispersion of the resulting distribution of losses. As ρ varies in $\{-0.7, 0, 0.7\}$, the medians of the loss distributions are approximately $\{2.9, 2.7, 2.5\}$, and the means are approximately $\{3.38, 3.18, 3.10\}$; in each case, the minimum value of the risk function reduces as ρ increases. However, the loss uncertainty increases as ρ increases; for example, the upper 95% points of the loss distributions are approximately $\{6.86, 7.45, 7.90\}$ at these three values of ρ . Thus, while average or median risks define one order, the tail behaviour of loss distributions raises additional considerations of possible “downside” losses. Note also that there is appreciable probability on loss outcomes that are lower than the optimized risk, i.e., corresponding to the potential “upside” outcomes. This argues for the broader view of decision analysis to understand aspects of loss distributions at optimal— or other— chosen point forecasts. It should be stressed that this is a general point— not specific to constrained forecasting, but highlighted in this context. While the general concept has been well-recognized in areas such as finance (with “value-at-risk” studies resulting) it is not generally appreciated in other areas of decision analysis.

Probabilistic Conditioning

This example is a case in which the conditioning value of F lies is well into the tails of $p(Y)$ which, as discussed earlier, represents challenges to the purely probabilistic approach of summarizing aspects of $p(\mathbf{y}|Y = F)$. That said, in this simple illustrative example in only 2-dimensions, it is easy to generate very large Monte Carlo samples from $p(\mathbf{y})$ and apply vanilla ABC-style methods. Figure 2 gives examples in the case of $\rho = 0.7$. A large sample from $p(\mathbf{y})$ was conditioned to simulated values of \mathbf{y} such that $Y = \mathbf{1}'\mathbf{y}$ was “close” to the conditioning F value, illustrating results for both the low and high values of F . Closeness was specified by $|Y - F|/F < \tau$ where 100τ is the percent tolerance on this natural “closeness to constraint” metric. Examples displayed use $\tau = 0.005$; resulting scatter plots (not shown) of the constrained samples appear visually indistinguishable from the line $\mathbf{1}'\mathbf{y} = F$. ABC acceptance rates at this tolerance are around 0.5–1.5%, and smaller for negative values of ρ , indicative of the




Figure 2: Bivariate lognormal example with $\rho = 0.7$, showing marginal and approximate constrained p.d.f.s. Analysis generates a large Monte Carlo sample and accepts \mathbf{y} if, and only if, the sum Y satisfies $|Y - F|/F < \tau$ with percent tolerance $100\tau = 0.5$. The histograms represent ABC-approximate conditionals $p_i(y_i|Y = F)$ based on the joint prior $p(\mathbf{y})$ whose lognormal margins $p_i(y_i)$ are displayed as curves. Analyses are based on the very low value of $F = 14.7$ (left) and the high value of $F = 24.15$ (right).

challenges of using probabilistic conditioning. In realistic, higher-dimensional settings, this ABC-style analysis is simply not an option as (i) it becomes really challenging to define relevant tolerance ranges, and even with that in place (ii) the acceptance rates decay exponentially with dimension. In contrast, the decision analysis approach has different goals, and generates useful and informative results of direct applied value in such contexts, and is at least complementary to the purely probabilistic approach.

Entropic Tilting

In the setting of Section 2.3, take ET function $q(\mathbf{y}) = I(\mathbf{y} \in S) - (1 - \epsilon)$ where the probability ϵ is close to zero and $S = \{\mathbf{y} : (1 - \tau)F \leq Y \leq (1 + \tau)F\}$. This leads to ET optimal $G(\mathbf{y}) \approx P(\mathbf{y}|Y = F)$ with p.d.f. $g(\mathbf{y}) \propto \exp\{\gamma I(\mathbf{y} \in S)\}p(\mathbf{y})$ with $\gamma = \log\{(1 - \epsilon)(1 - p_s)/(\epsilon p_s)\}$. In the current example, $P(\mathbf{y})$ is defined in terms of a direct Monte Carlo sample \mathbf{y}^j for $j = 1 : J$ where J is the Monte Carlo sample size. These are reweighted according to weights $w_j = \exp(\gamma)$ for $\mathbf{y} \in S$, and $w_j \propto 1$ for $\mathbf{y} \notin S$, subject to summing to 1. As noted earlier, the optimal solution has $\gamma = \log\{(1 - \epsilon)(1 - p_s)/(\epsilon p_s)\}$. The optimal set of weights w_j represent an importance sampling approximation to the optimal $G(\mathbf{y})$ subject to the assessment of importance sampling accuracy using the usual metrics (e.g. West, 1993). Inference on \mathbf{y} conditional on the constraints then follows; one easy and oft-used step in importance sampling is to simply resample \mathbf{y} from the set \mathbf{y}^j according to the weights w_j , and then proceed based on that resample (which, of course, will generally include replicates).

In the running bivariate lognormal example with $\rho = 0.7$, the ET analysis has been explored for a range of choices of small values of the tolerance parameters τ, ϵ . There is a high level of robustness with respect to these values. Take $\tau = 0.005$ as in the accept/reject analysis, and, for example, $\epsilon = 0.001$. Monte Carlo samples of $J = 1 \times 10^6$ generate empirical distributions of resampled \mathbf{y}^j values that are visually indistinguishable from the direct accept/reject results shown in Figure 2. The resulting constrained medians of each element of \mathbf{y} are equal to those from the traditional accept/reject analysis up to two decimal places (on the practically relevant scale of 0–20). This, and multiple other examples, supports ET as a novel approach to decision-guided conditioning on almost-exact deterministic constraints.

Concordance of ET with probabilistic conditioning is also clear in questions about constrained values that are extreme under $P(\mathbf{y})$. One gauge of this in the ABC accept/reject analysis is the empirical estimate of acceptance rate $100p_s\%$. In the ET analysis, the percentage effective sample size (ESS) of the importance sampling weights, $ESS = 100/\sum_j(Iw_j^2)$, is comparable. In the example with $\tau = 0.005$ and $\epsilon = 0.001$ at the optimized values of γ , the summaries are as follows. When $F = 14.7$, the low value, the effective % sample size of the importance sampling weights, these two measures are each 0.56% to two decimal places; when $F = 24.15$, the high value, they are each approximately 1.36%. These very low values are again indicative of the challenges of conditioning. Constraints that are unlikely under $P(\mathbf{y})$ will generate low ESS and showing the instability of the results. In such contexts, taking ϵ very small and m closer to 1– as required for theoretical reliance on the approach– is increasingly fragile without access to very large samples from $P(\mathbf{y})$. That said, it is very worthwhile to explore both probabilistic conditioning using ABC-style analysis and the ET approach together in a given context. The formal Bayesian decision analysis approach– with its different goals and outputs– generates additional results and insights, as is now further exemplified.

Constraint Sensitivity Analysis

Exploring loss outcomes under perturbations of the chosen conditioning value F defines a local sensitivity analysis: perturb a “nominal” constrained value F and reevaluate \mathbf{f}^* across perturbed values $F + \delta$ for some δ in a specified discrete range of values. In this example context with positive outcomes, perturbations of a chosen sum F are best couched as percentage changes, i.e., taking $\delta = \pm\epsilon F$ for small ϵ on a discrete range of specified values. To illustrate this here, Figure 3 summarizes results in this bivariate lognormal example for two cases of the nominal F , low and high values with respect to $p(Y)$, and taking $\epsilon = 0.1$. The resulting ranges of \mathbf{f}^* values as F varies within 10% of the chosen nominal values indicate quite tight ranges, in this example context. Note how the optimal point forecast values track the levels of the joint p.d.f. as F varies, with trajectories that appear to move very naturally along a “ridge” in the p.d.f. Also, while these regions of optimal point forecasts are of course very different conceptually to probability intervals under the purely probabilistic framework, note the concordance with the ABC-approximate conditional p.d.f.s in Figure 2.

The NR algorithm is fast; running it for multiple values $F + \delta$ is computationally easy. A first-order approximation is available to define a computational short-cut, perhaps at

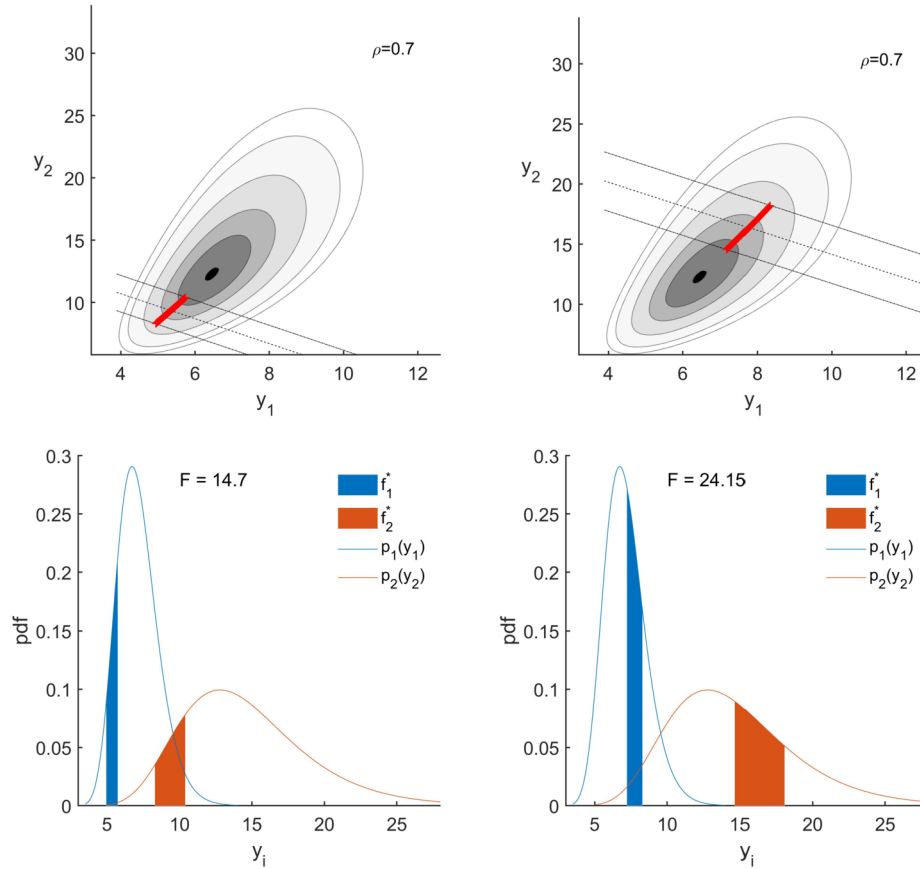


Figure 3: Bivariate lognormal example with $\rho = 0.7$, showing the joint p.d.f. and its margins with constrained ranges of each f_i^* indicated. This is based on the sensitivity analysis under AD loss and varying the constrained total in $\{0.9F, 1.1F\}$ — i.e., within $\pm 10\%$ of the nominal value F . Implications of the earlier used low ($F = 14.7$, left figures) and high ($F = 24.15$, right figures) values for the— now nominal— constrained total are shown. The upper frames show three straight lines representing the nominal constraint with now lower and upper bounds also indicated; the central of the three lines is that defining the nominal constraint. Superimposed is a shaded region (in red) that maps how \mathbf{f}^* changes as F varies across its range. The lower frames show the marginal p.d.f.s with ranges of each f_i^* indicated by the shaded regions.

least for initial exploratory analysis. Based on the NR update equation, a perturbation of F to $F + \delta$ for some small δ yields the update to $\lambda = \lambda^* + \delta/\dot{q}(\lambda^*)$. In the current example context this translates to $\lambda = \lambda^* + \epsilon F/\dot{q}(\lambda^*)$ where ϵ takes value in a small range of % changes. This provides a trivial short-cut approximation to evaluating λ over the range, and then computing the implied ranges $\mathbf{f}^*(\lambda)$ in the sensitivity analysis.

Finally, in some applied settings it may be of interest to explore sensitivity analyses with different ranges of values of F , such as defined by predictive intervals for Y under some external model for the outcome total that is imposed on the predictive model $p(\mathbf{y})$ in the spirit of information aggregation, or predictive synthesis (West and Crosse, 1992; West, 1992; West and Harrison, 1997, section 16.3; McAlinn and West, 2019).

4.3 A 100-Dimensional Example

Some of the main motivating areas of application are in constrained forecasting in commercial or economic settings. In commercial sales forecasting in large companies, key example contexts including those of projecting point forecasts throughout hierarchies of sales or revenues, and of ensuring consistency of forecasts for sales at increasingly fine levels of disaggregation. Forecasts for “high-level” sales or revenue must be consistent with sets of forecasts at “lower levels” (e.g. Green and Harrison, 1973; West and Harrison, 1997, section 16.3). Here the general framework concerns conditioning sets of forecasts on information about totals and aggregates, all of which can be represented via sets of linear constraints on the uncertain outcomes being predicted. Such questions arise commonly in “what-if” evaluation in policy decision contexts such as above and in other areas including macro-economic forecasting over multiple time periods (McAlinn et al., 2020). Related questions arise in ensuring compatibility of forecasts at different resolutions in time, often with different forecast models generating predictions at different time scales (e.g. Ferreira and Lee, 2007; Molina et al., 2010; Berry et al., 2020).

An example in $n = 100$ dimensions is summarized in Figure 4, linked to applied studies in supermarket sales modelling and forecasting. The data come from $n = 100$ stores with monthly revenue $\$y_i$ in the same consumer goods sector in each store, recorded each month over several years. The data are scale transformed for confidentiality. The snapshot here concerns a forecast distribution $P(\mathbf{y})$ where \mathbf{y} represents the one-month ahead revenue vector at a chosen time point. Here $P(\mathbf{y})$ is a 100-dimensional lognormal $\mathbf{y} \sim LN(\mathbf{m}, \mathbf{V})$ with (\mathbf{m}, \mathbf{V}) set at values based on the historical record and model analysis. Conditioning on the sales total Y is then of relevance from the above noted perspectives. A point forecast of Y may be generated from an “external” source and the goal is to explore consistency with $P(\mathbf{y})$ and/or to proceed to constrained forecasting of \mathbf{y} . The external source may be an alternative model of aggregate sales, or one of a number of values generated as “what-if” or scenario values for consideration. A main point for illustration is to complement the above examples in highlighting the role of dependencies in a total-constrained decision analysis and the impact of the decision perspective. This is also a higher-dimensional example and it should be noted that the computational load in evaluation of decision analytic constrained forecasts remains almost trivial using the NR algorithm.

Figure 4 displays summaries of correlations in \mathbf{V} . There are dependencies across stores, with both negative and positive dependencies exhibited in the displays of correlations defined by \mathbf{V} . The decision analysis is summarized through evaluation of optimal point forecasts using absolute deviation (AD) loss; revenue outcomes are all on the \$ scale so are directly comparable and APE loss is less relevant, while the context is such that results under AD, APE and SE are in any case similar. Marginal point forecasts

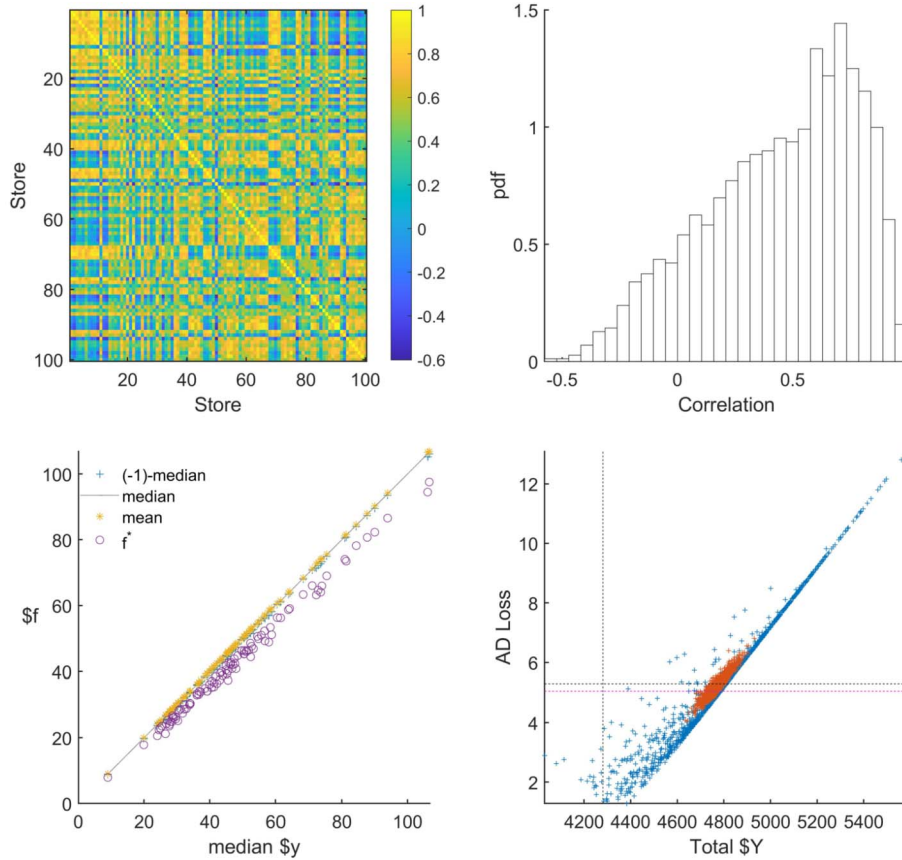


Figure 4: Supermarket sector sales revenue example with $n = 100$ and $\mathbf{y} \sim LN(\mathbf{m}, \mathbf{V})$. The top frames exhibit the correlations underlying \mathbf{V} in heat-map and histogram forms. The lower left frame plots the marginal (-1) -medians, medians and means of the y_i against their medians, and overlays a scatter plot of the AD-optimal f_i^* point forecasts in the case of a total constraint $F = 4,281$. This value of $Y = F$ lies somewhat in the lower tail of $P(Y)$ as exhibited in the lower right frame that scatter plots a Monte Carlo sample from $P(Y, L(\mathbf{y}, \mathbf{f}^*)/n)$ (blue +). The vertical dashed line marks the value of the constraint, $Y = F$; the horizontal dashed lines mark the value of the expected loss at the minimum, i.e., the optimized risk $R(\mathbf{f}^*)$ (black dashed line) and the median of the loss distribution (red dashed line). Overlaid is a scatter plot (+) of a corresponding sample from a modified $P(\mathbf{y})$ that has the same location and scale parameters but sets dependencies to zero, i.e., $\mathbf{y} \sim LN(\mathbf{m}, \mathbf{V}_0)$ where \mathbf{V}_0 has the same diagonal elements as \mathbf{V} but zero off-diagonal entries.

for each y_i are closely similar across stores i , as illustrated in the figure. Using AD loss and constraining to the total $F = \mathbf{1}'\mathbf{f}$ that is fixed at a value somewhat (though not extreme) in the lower tail of $p(Y)$ shows optimal constrained forecasts that are clearly

downward adjustments to the marginal values. The figure then displays Monte Carlo samples from $P(Y, L(\mathbf{y}, \mathbf{f}^*)/n)$ as in the earlier examples of Section 4.2. That is, a scatter plot of samples from the joint distribution of the total revenue over stores together with the realized loss *per store* under the AD loss function. The imposed conditional total value here is $F = 4,281$, lying in the tail— though not really substantially extreme— of the forecast distribution $P(Y)$ that is close to symmetric with median 4,775.

The distribution of loss at the AD-optimal value \mathbf{f}^* is spread over 2–12 on the scale defined in this analysis, while the mean and median of the loss distribution are around 5.1–5.2. There is appreciable probability of loss values much less than this, as well as reasonable chances of higher losses up to the 9–12 range (akin to “value at risk”). These are key and potentially critical presentations of realistic outcomes from the decision analysis.

A further point speaks to the impact of dependencies in $P(\mathbf{y})$ on the implied loss distributions. The figure highlights this in contrasting summaries of the predictive distribution $P(Y, L(\mathbf{y}, \mathbf{f}^*)/n)$ under the dependent model with a model in which \mathbf{V} is replaced by a diagonal matrix \mathbf{V}_0 having the same diagonal elements. This is not a strange choice for comparison; applied analyses of such problems will often analyze data independently across stores, so this is a relevant benchmark. The resulting $P(Y)$ and hence the joint $P(Y, L(\mathbf{y}, \mathbf{f}^*)/n)$ are very concentrated relative to the original model analysis. In the dependent model, there are ranges of negative and positive dependencies among the y_i , but the preponderance and magnitudes of positive values lead to overall increased uncertainty about Y and hence about the potential loss outcomes. This is typical in such commercial applications, where dependencies often arise through common factors such a seasonality and management policies that are comparable across stores (Berry and West, 2020; Berry et al., 2020). The comparison analysis that fixes all correlations in \mathbf{V}_0 to zero defines, in contrast, a rather concentrated predictive distribution for Y and hence losses. While the optimized expected losses are the same under the two models, the independence model massively understates the levels of realistic uncertainty in the outcome total Y and hence in the loss distribution; this leads to the potential to generate substantial over-confidence in the selection of the optimal constrained point forecasts. Other comparisons could be made, but this practically-grounded example serves to again highlight the main point of examining loss distributions along with optimal point forecasts.

5 Summary Comments

Integrating decision analysis into problems of constrained forecasting is complementary to the traditional inferential view, and examples here illustrate the benefits of ensuring attention to the decision theoretic “Yang” of Bayesian analysis as well as to the inferential “Yin”. This view is, of course, not specific to the contexts of Bayesian constrained forecasting, but decision perspectives are under-regarded and under-represented in applied forecasting as in other areas (e.g. Lindley, 1992; Lavine et al., 2021; West, 2020, section 2.3). The potential for methodological advance as well as more comprehensive analysis is highlighted in the motivating context of constrained forecasting. Examples

with features faithful to applied settings demonstrate the relevance of the broader view of decision analysis that expands from a main focus on optimal decisions to always explore implied loss distributions. Constrained forecasting examples where multivariate dependencies can have major impact on ranges of likely losses highlight the practical import. It is also noted that, in many common, practical contexts, expected loss functions are undefined while distributions of losses are valid and accessible. Key examples involve log-T distributions that routinely arise in linear models of many kinds, including dynamic linear models (DLMs) in time series forecasting. Here blind adoption of optimal point forecasts based on minimizing expected losses is unfounded implied predictive distributions of losses are perfectly well defined.

Methodologically, connections of decision analysis and inferential approaches are nicely enhanced by new ideas and methodology exploiting entropic tilting. As an approach to understanding compatibility of constraints with a forecast distribution, and to minimally modifying the distribution to approximately accord with deterministic constraints, ET is a fully Bayesian decision analysis framework that can be exploited to extend methodology in practical ways. Its connections to Monte Carlo methods for probabilistic conditioning nicely highlight its use and complementarities relative to constrained point forecasting alone. Technically, use of more elaborate, non-linear constraints will raise challenges of optimization. While many, general algorithms exist for numerical optimization problems, linkages with Bayesian emulation approaches such as have been used in portfolio optimization subject to various constraints (e.g. Irie and West, 2019) may provide relevant and novel opportunities for progress.

Detailed application in live forecasting require extensions to contexts of multiple constraints such as arise in hierarchies. Extension of the basic decision problem of (2.1) to a set of constraints is theoretically immediate. That is, the optimization is generalized to condition on k constraints $\mathbf{A}'\mathbf{f} = \mathbf{F}$ where \mathbf{A} is a given $n \times k$ matrix of full rank $k < n$ and \mathbf{F} a given constraint k -vector. Developments with intersecting sets of subtotal constraints— in which case \mathbf{A} is a matrix with zero/one entries— are one main class of interest. Evaluation of optimal constrained forecasts based on multivariate Newton-Raphson is immediate.

Problems in which the loss functions are not additive in outcomes $i = 1 : n$ are also of interest. In commercial forecasting with positive, or non-negative, outcomes (such as with consumer sales of items or batches of items, revenues in multiple sectors or markets) it can be relevant to consider loss functions that involve cross-talk between outcomes. Sales of one product may be inversely related to those of another due to substitution effects, and overall sales across categories might be of main interest; similar comments apply to revenue forecasting over multiple sectors. Customized losses $L(\mathbf{y}, \mathbf{f})$ that are not additive will be of interest.

Important extensions concern problems in which the predictive distribution $P(\mathbf{y})$ is itself impacted by the future actions based on chosen optimal point forecasts. As one example, consider supermarket sales where y_i is the sales outcome, and f_i^* defines the store manager's decision to stock f_i^* (or, perhaps, $f_i^* +$ a few more) items of a specific consumer item. Then, necessarily $y_i \leq f_i^*$ since no more than that are available for sale in the next time period. Extension of this applies to cases of other constraints on \mathbf{y} ,

such as are relevant in macro-economic forecasting where one or more of the y_i are putatively controllable as policy instruments. A macro-economic forecasting model is then applied across a range of “what-if?” values of one variable, corresponding to one (linear) constraint on the outcome vector rather than a total constraint. Such conditional forecasting across multiple time periods lies at the heart of applied Bayesian forecasting in macro-economics (e.g. Del Negro and Otrok, 2008; Nakajima and West, 2013; McAlinn et al., 2020). This argues for extensions in which $P(\mathbf{y})$ is modified by multiple constraints in the optimization setting. This is not a new concept (e.g. Harrison and Smith, 1980) but is certainly under-regarded in both forecasting and Bayesian analysis literatures. It is an extension of significant potential practical importance.

Connections with forecast reconciliation and combination in Bayesian econometrics go beyond the use of entropic tilting (as discussed in Section 2.3). In particular, there is much interest in questions of constraining forecasts (based on point forecasts from alternative models, or information from surveys of forecasters, etc) in settings where there are multiple, intersecting hierarchies (e.g. Koop et al., 2020). A key such area is that of mixed-frequency time series, where some variables are observed at finer time scales than others, so naturally inducing temporal aggregation questions (e.g. Koop et al., 2019). Coupled with this are questions of consistency of forecasts with respect to aggregation over time, as well as contexts where there are inherent cross-sectional hierarchies (e.g., monthly consumption is an aggregate over several sectors). The decision analysis approach coupled with ET methodology offers promise for expanding Bayesian methodology in such areas.

Finally, the discussion has touched on connections to the broader area of integration of forecast information, and so-called forecast reconciliation. In the Bayesian literature, this area has evolved from basic forecast combination to more fully subjective Bayesian approaches to correcting for biases and more general calibration, and combining forecasts from multiple, potentially related sources. This line of literature (e.g., West and Crosse, 1992; West, 1992; West and Harrison, 1997, section 16.3; McAlinn and West, 2019) intersects intimately with the apparently narrower and more specific goals of constrained forecasting, but as detailed in examples in West and Harrison (1997, section 16.3), clearly forms part of a broader context; decision analytic perspective of this paper can be expected to be applicable in these broader settings. The use of entropic tilting opens up new directions for exploiting constraints generally, as well as in broadening the Bayesian thinking about comparison, calibration and combination of forecasts from multiple models. Some recent relevant developments in Tallman and West (2022a), that integrate decision goals into forecast model combination and explicitly rely on some new uses of entropic tilting, indicate some such opportunities. Further, the constructive use of ET also yields computational benefits that will be of use in more complicated settings. For example, in settings with non-linear constraints and/or non-additive loss functions the ET methodology based on Monte Carlo samples from forecast distributions is accessible and no more computationally demanding than in the simpler example settings of this paper.

Supplementary Material

Perspectives on Constrained Forecasting. Supplementary Material: Further Details and Examples (DOI: [10.1214/23-BA1379SUPP](https://doi.org/10.1214/23-BA1379SUPP); .pdf). Supplementary details and examples are provided in West (2023).

References

- Aktekin, T., Polson, N. G., and Soyer, R. (2018). “Sequential Bayesian analysis of multivariate count data.” *Bayesian Analysis*, 13: 385–409. MR3780428. doi: <https://doi.org/10.1214/17-BA1054>. 1014
- Berry, L. R., Helman, P., and West, M. (2020). “Probabilistic forecasting of heterogeneous consumer transaction-sales time series.” *International Journal of Forecasting*, 36: 552–569. 1014, 1017, 1024, 1031, 1033
- Berry, L. R. and West, M. (2020). “Bayesian forecasting of many count-valued time series.” *Journal of Business and Economic Statistics*, 38: 872–887. MR4154894. doi: <https://doi.org/10.1080/07350015.2019.1604372>. 1014, 1017, 1023, 1024, 1033
- Bonassi, F. V. and West, M. (2015). “Sequential Monte Carlo with adaptive weights for approximate Bayesian computation.” *Bayesian Analysis*, 10: 171–187. MR3420901. doi: <https://doi.org/10.1214/14-BA891>. 1015, 1017
- Chen, C. W. S. and Lee, S. (2017). “Bayesian causality test for integer-valued time series models with applications to climate and crime data.” *Journal of the Royal Statistical Society (Series C: Applied Statistics)*, 66: 797–814. MR3670418. doi: <https://doi.org/10.1111/rssc.12200>. 1014
- Chen, X., Banks, D., and West, M. (2019). “Bayesian dynamic modeling and monitoring of network flows.” *Network Science*, 7: 292–318. 1014
- Chen, X., Irie, K., Banks, D., Haslinger, R., Thomas, J., and West, M. (2018). “Scalable Bayesian modeling, monitoring and analysis of dynamic network flow data.” *Journal of the American Statistical Association*, 113: 519–533. MR3832205. doi: <https://doi.org/10.1080/01621459.2017.1345742>. 1014
- de Alba, E. (1988). “Disaggregation and forecasting: A Bayesian analysis.” *Journal of Business and Economic Statistics*, 6: 197–206. 1014
- de Alba, E. (1992). “Constrained forecasting in regression: Bayesian analysis.” *Estadística*, 44: 27–72. MR1257038. 1014
- de Alba, E. (1993). “Constrained forecasting in autoregressive time series models: A Bayesian analysis.” *International Journal of Forecasting*, 9: 95–108. 1014
- de Alba, E. (2006). “The Bayesian method of moments (BMOM) in some aggregation problems in econometrics.” *Applied Stochastic Models in Business and Industry*, 22: 95–112. MR2277012. doi: <https://doi.org/10.1002/asmb.621>. 1014

- Del Negro, M. and Otrok, C. M. (2008). “Dynamic factor models with time-varying parameters: Measuring changes in international business cycles.” Staff Report 326, New York Federal Reserve. 1035
- Ferreira, M. A. R. and Lee, H. K. (2007). *Multiscale Modeling: A Bayesian Perspective*. Springer. MR2348732. 1031
- French, S. and Insua, D. R. (2010). *Statistical Decision Theory: Kendall’s Library of Statistics 9*. Wiley. 1021
- Goldstein, M. and Wooff, D. (2007). *Bayes Linear Statistics: Theory and Methods*. Wiley. MR2335584. doi: <https://doi.org/10.1002/9780470065662>. 1014
- Green, M. and Harrison, P. J. (1973). “Fashion forecasting for a mail order company using a Bayesian approach.” *Journal of Operational Research Society*, 24: 193–205. 1014, 1031
- Guerrero, V. M. (1989). “Optimal conditional ARIMA forecasts.” *Journal of Forecasting*, 8: 215–229. 1014
- Guerrero, V. M. and Nieto, F. H. (1999). “Temporal and contemporaneous disaggregation of multiple economic time series.” *Test*, 8: 459–489. MR1748282. doi: <https://doi.org/10.1007/BF02595880>. 1014
- Harrison, P. J. and Smith, J. Q. (1980). “Discontinuity, decision and conflict (with discussion).” In Bernardo, J. M., Groot, M. H. D., Lindley, D. V., and Smith, A. F. M. (eds.), *Bayesian Statistics I*, 99–142. Valencia University Press. Proceedings of the First Valencia International Meeting 1979. MR0638875. 1035
- Irie, K. and West, M. (2019). “Bayesian emulation for multi-step optimization in decision problems.” *Bayesian Analysis*, 14: 137–160. MR3910041. doi: <https://doi.org/10.1214/18-BA1105>. 1034
- Koop, G., McIntyre, S., and Mitchell, J. (2019). “UK regional nowcasting using a mixed frequency vector auto-regressive model with entropic tilting.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 183: 91–119. MR4049656. 1019, 1035
- Koop, G., McIntyre, S., Mitchell, J., and Poon, A. (2020). “Regional output growth in the United Kingdom: More timely and higher frequency estimates from 1970.” *Journal of Applied Econometrics*, 35(2): 176–197. MR4098531. doi: <https://doi.org/10.1002/jae.2748>. 1035
- Krüger, F., Clark, T. E., and Ravazzolo, F. (2017). “Using entropic tilting to combine BVAR forecasts with external nowcasts.” *Journal of Business and Economic Statistics*, 35(3): 470–485. MR3663786. doi: <https://doi.org/10.1080/07350015.2015.1087856>. 1019
- Lavine, I., Lindon, M., and West, M. (2021). “Adaptive variable selection for sequential prediction in multivariate dynamic models.” *Bayesian Analysis*. MR4381127. doi: <https://doi.org/10.1214/20-BA1245>. 1033

- Li, L., Ramsundar, B., and Russell, S. (2013). “Dynamic scaled sampling for deterministic constraints.” In Carvalho, C. M. and Ravikumar, P. (eds.), *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, 397–405. Proceedings of Machine Learning Research 13 (PMLR, <http://proceedings.mlr.press/>). 1015, 1017
- Lindley, D. V. (1992). “Is our view of Bayesian statistics too narrow?” In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M. (eds.), *Bayesian Statistics 4*, Valencia IV Conference President’s Address. Oxford University Press. Proceedings of the Fourth Valencia International Meeting: Dedicated to the Memory of Morris H. DeGroot, 1931–1989: April 15–20, 1991. MR1380267. 1014, 1033
- McAlinn, K., Aastveit, K. A., Nakajima, J., and West, M. (2020). “Multivariate Bayesian predictive synthesis in macroeconomic forecasting.” *Journal of the American Statistical Association*, 115: 1092–1110. MR4143452. doi: <https://doi.org/10.1080/01621459.2019.1660171>. 1031, 1035
- McAlinn, K. and West, M. (2019). “Dynamic Bayesian predictive synthesis in time series forecasting.” *Journal of Econometrics*, 210: 155–169. MR3944768. doi: <https://doi.org/10.1016/j.jeconom.2018.11.010>. 1031, 1035
- Molina, G., Han, C.-H., and Fouque, J.-P. (2010). “MCMC estimation of multiscale stochastic volatility models.” *Mathematics and Computers in Simulation*, 103: 199–212. MR3213768. doi: <https://doi.org/10.1016/j.matcom.2013.07.005>. 1031
- Nakajima, J. and West, M. (2013). “Bayesian analysis of latent threshold dynamic models.” *Journal of Business and Economic Statistics*, 31: 151–164. MR3055329. doi: <https://doi.org/10.1080/07350015.2012.747847>. 1035
- Prado, R., Ferreira, M. A. R., and West, M. (2021). *Time Series: Modeling, Computation & Inference*. Chapman & Hall/CRC Press, 2nd edition. 1017
- Robertson, J. C., Tallman, E. W., and Whitemann, C. H. (2005). “Forecasting using relative entropy.” *Journal of Money, Credit, and Banking*, 37: 383–401. 1018
- Smith, J. Q. (2010). *Decision Analysis: A Bayesian Approach*. Statistics Texts. Chapman and Hall. MR0941143. 1021
- Tallman, E. and West, M. (2022a). “Bayesian predictive decision synthesis.” *Invited revision under journal review*. ArXiv:2206.03815. 1035
- Tallman, E. and West, M. (2022b). “On entropic tilting and predictive conditioning.” *Technical Report, Department of Statistical Science, Duke University*. ArXiv:2207.10013. 1018
- West, M. (1992). “Modelling agent forecast distributions.” *Journal of the Royal Statistical Society (Ser. B)*, 54: 553–567. MR1160482. 1031, 1035
- West, M. (1993). “Approximating posterior distributions by mixtures.” *Journal of the Royal Statistical Society (Ser. B)*, 54: 553–568. 1017, 1028

- West, M. (2020). “Bayesian forecasting of multivariate time series: Scalability, structure uncertainty and decisions (with discussion).” *Annals of the Institute of Statistical Mathematics*, 72: 1–44. [1014](#), [1017](#), [1033](#)
- West, M. (2023). “Perspectives on Constrained Forecasting. Supplementary Material: Further Details and Examples.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/23-BA1379SUPP>. [1016](#), [1036](#)
- West, M. and Crosse, J. (1992). “Modelling of probabilistic agent opinion.” *Journal of the Royal Statistical Society (Ser. B)*, 54: 285–299. [MR1157726](#). [1031](#), [1035](#)
- West, M. and Harrison, P. J. (1997). *Bayesian Forecasting & Dynamic Models*. Springer, 2nd edition. [MR1482232](#). [1014](#), [1016](#), [1017](#), [1031](#), [1035](#)
- Wickramasuriya, S. L., Athanasopoulos, G., and Hyndman, R. J. (2019). “Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization.” *Journal of the American Statistical Association*, 114: 804–819. [MR3963182](#). doi: <https://doi.org/10.1080/01621459.2018.1448825>. [1014](#)

Acknowledgments

Research reported here was motivated by applied forecasting problems with researchers at 84.51° Labs., 100 West 5th Street, Cincinnati, OH 45202, and benefited in particular from discussions with Andrew Cron, Paul Helman and Christoph Hellmayr. Discussions with Domenico Giannoni of Amazon, and with past graduate students in Statistical Science at Duke University including Lindsay Berry, Daniel Deng, Isaac Lavine and Anna Yanchenko, contributed further perspectives. The author also acknowledges the most constructive comments and suggestions from the Editor-in-Chief, the Associate Editor, and two anonymous referees.