SEQUENTIAL REGULARIZED ENSEMBLE FORECASTING BY ONLINE POLICY LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Aggregating point forecasts from multiple experts often improves accuracy relative to relying on a single expert, but sequential ensemble forecasting becomes difficult for time series data when expert performances change over time. In this paper, we formulate sequential ensemble forecasting as an online policy learning problem and incorporate the recently proposed offline Regularized Ensemble Forecasting (REF) framework. At each forecasting period, the policy objective selects the optimal ensemble weight by balancing the variation of current forecasts, shrinkage toward historical performances, and stabilization toward preceding optimal weights. In the case of a large expert pool, we impose a sparse parameterization that permits exact zero weights and facilitates expert selection over time. Under suitable conditions, we develop an online regret bound that vanishes as the total number of periods goes to infinity. Numerical studies, including applications to the M5 competition and the Survey of Professional Forecasters (SPF), show that the proposed online REF policy performs favorably against standard ensemble benchmarks and offline REF.

1 INTRODUCTION

Ensemble forecasting is a popular and practical forecasting method to improve prediction accuracy where a decision maker properly combines the forecasts from multiple experts instead of relying on the forecast from a single expert. A large body of literature has been dedicated to developing methods to find the optimal linear weights to combine forecasts (see reviews by, e.g., Clemen 1989, Wang et al. 2023). Existing ensemble rules can be broadly categorized as either rules that use only the current set of forecasts, such as the simple mean (Bates & Granger 1969) and trimmed means (Jose & Winkler 2008), or history-based rules that estimate weights only from past forecast errors, such as covariance-based weights (Newbold & Granger 1974) and regression-based or shrinkage combinations (Granger & Ramanathan 1984, Diebold & Pauly 1990).

Recently Su et al. (2026) proposed a Regularized Ensemble Forecasting (REF) framework, which is the first ensemble forecasting approach that leverages both current forecasts and experts' historical performances to obtain the optimal ensemble weights. We call this method *offline REF* since it is developed for cross-sectional ensemble weighting. For a time series forecasting task where forecasts are provided over successive forecasting periods, applying offline REF separately in each period would treat each period as a separate cross-sectional ensemble problem and update the weights using the information available in that period. The resulting sequence of optimal weights for an expert may not fully capture useful continuity across periods.

In this paper, we develop *online REF* for sequential ensemble forecasting of time series where we formulate it as an online policy learning problem. At each forecasting period, we modify offline REF to find an optimal policy that balances current forecast dispersion, shrinkage toward historical prior weights, and penalization on deviations from preceding weights. The weights for the current period are selected before the corresponding outcome is observed. Once the current outcome becomes available and the next set of forecasts is received, the next weighting decision can incorporate the newly available information while also being informed by the preceding weights, which reflect the most recent ensemble decision based on the forecasts and historical information available at that time. Moreover, to make expert selection more direct when the expert pool is large, we incorporate entmax into the policy objective, a sparse parameterization that can assign exact zero weights to

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

certain experts among the currently available experts, while preserving the simplex-weight structure (Martins & Astudillo 2016, Peters et al. 2019).

The contribution of this paper is threefold. First, this paper makes the first attempt to formulate sequential ensemble forecasting as an online policy learning problem. The proposed online REF policy rule not only uses experts' current forecasts and their historical performance but also incorporates preceding optimal weights into the current weighting decision. Second, under suitable conditions, we establish a theoretical guarantee that online REF has vanishing expected average regret relative to the oracle ensemble, which is the best weighted combination selected in hindsight over the forecasting horizon. Third, by numerical studies, including a simulation study and the two benchmark datasets, the M5 competition (Makridakis et al. 2022a,b) and the SPF (Clements et al. 2023, Federal Reserve Bank of Philadelphia 2026), Online REF is shown to achieve competitive forecasting performances relative to offline REF and standard ensemble benchmarks.

2 ONLINE REF

At each forecasting period $t = 1, \ldots, T$, N experts provide their forecasts of a continuous outcome $y_t \in \mathbb{R}$, denoted by $\widehat{y}_t^{(i)}$, $i = 1, \ldots, N$. Write $\widehat{\mathbf{y}}_t = (\widehat{y}_t^{(1)}, \ldots, \widehat{y}_t^{(N)}) \in \mathbb{R}^N$. A decision maker aims to predict y_t by $\widehat{y}_t^{\text{agg}} = \sum_{i=1}^N w_{t,i}\widehat{y}_t^{(i)}$, a linear aggregation of the experts' forecasts based on a weight vector $\boldsymbol{w}_t = (w_{t,1}, \ldots, w_{t,N}) \in \Delta_N$ where $\Delta_N := \{\boldsymbol{w} \in \mathbb{R}^N : w_i \geq 0,\ \sum_{i=1}^N w_i = 1\}$. Our main goal is to find the optimal weights in an online fashion for the best of prediction accuracy. To achieve this goal, we first formulate the sequential forecasting problem and then develop an online policy rule at each period t that generalizes the offline REF proposed by Su et al. (2026), allowing us to sequentially obtain the optimal weights. We then describe how the weights are updated across forecasting periods.

2.1 ONLINE UPDATES FORMULATION

We formulate sequential ensemble forecasting as a contextual decision-making problem. Let $R_{t+1} = -\ell_t(\hat{\boldsymbol{w}}_t)$, where $\ell_t(\boldsymbol{w}) = (y_t - \sum_{i=1}^N w_i\widehat{y}_t^{(i)})^2$ is the squared forecast loss for any $\boldsymbol{w} \in \Delta_N$. Over T forecasting periods, the objective of the decision maker is to maximize the cumulative reward:

$$\sum_{t=1}^T R_{t+1} = -\sum_{t=1}^T \ell_t(\hat{\boldsymbol{w}}_t) = -\sum_{t=1}^T \left(y_t - \sum_{i=1}^N \hat{w}_{t,i}\widehat{y}_t^{(i)} \right)^2 .$$

However, when the decision maker chooses $\hat{\boldsymbol{w}}_t$, y_t has not yet been observed, so the corresponding forecast loss and reward are not yet available. Therefore, each weighting decision has to be made *before* observing y_t.

To address this problem, we adopt the idea of offline REF by Su et al. (2026). At each period t, we propose to find the optimal weight vector $\hat{\boldsymbol{w}}_t$, and obtain the ensemble forecast $\widehat{y}_t^{\text{agg}} = \sum_{i=1}^N \hat{w}_{t,i}\widehat{y}_t^{(i)}$ accordingly, only based on the experts' current forecasts $\widehat{\mathbf{y}}_t$ and historical information summarized by $\boldsymbol{Z}_{t-1}$, including the preceding optimal weights $\hat{\boldsymbol{w}}_{t-1}$. Details about weight optimization at each period t and the specifications of $\boldsymbol{Z}_t$ are given later in Section 2.2.

After y_t is observed, the decision maker can compute R_{t+1} and use it to update $\boldsymbol{Z}_t$, and then find the optimal weight vector $\hat{\boldsymbol{w}}_{t+1}$ based on $\widehat{\mathbf{y}}_{t+1}$ and $\boldsymbol{Z}_t$ at the next period $t + 1$. The updating of $\boldsymbol{Z}_t$ is described in Section 2.3.

2.2 WEIGHT OPTIMIZATION AT EACH PERIOD t

By Su et al. (2026), at each period t, offline REF selects weights that solve

$$\min_{\boldsymbol{w}_t \in \Delta_N} \left\{ f\left(\sum_{i=1}^N w_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right) + \lambda \|\boldsymbol{w}_t - \boldsymbol{s}_t\|_2^2 \right\},$$

where $\bar{\widehat{y}}_t = N^{-1}\sum_{i=1}^N \widehat{y}_t^{(i)}$, $\lambda > 0$ is a tuning parameter, and f is a given monotone increasing function (see further details in Section 3.1). In the first term, $\sum_{i=1}^N w_{t,i}^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2$ is an empirical

approximation of $\mathrm{Var}(\sum_{i=1}^N w_{t,i}\widehat{y}_t^{(i)})$ when $\{\widehat{y}_t^{(i)}\}_{i=1}^N$ are independent with a common mean and finite second moments. See Appendix B for a derivation. In the second term, $\boldsymbol{s}_t = (s_{t,i})_{i=1}^N$ are called "prior weights" that measure the experts' historical performances in forecasting the outcome (the larger the better). Su et al. (2026) adopted the common-correlation (CCR) weights of Soule et al. (2024) as $\boldsymbol{s}_t$, of which details are provided in Section 2.3. Intuitively, offline REF aims to achieve the best balance between the dispersion of experts' current forecasts (first term) and fidelity to their historical performances (second term). The ith expert tends to receive a larger weight if her/his current forecast $\widehat{y}_t^{(i)}$ agrees more with the consensus $\bar{\widehat{y}}_t$ and/or her/his historical performance is better, i.e., $s_{t,i}$ is larger.

To update the weighting decisions sequentially across forecasting periods, we propose the following extension of offline REF by solving, at period t,

$$\hat{\boldsymbol{w}}_t = \arg\min_{\boldsymbol{w}_t \in \Delta_N} \left\{ f\left(\sum_{i=1}^N w_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) + \lambda \|\boldsymbol{w}_t - \boldsymbol{s}_t\|_2^2 + \kappa \|\boldsymbol{w}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 \right\}, \tag{1}$$

where $\kappa > 0$ is a tuning parameter and $\hat{\boldsymbol{w}}_{t-1}$ is the preceding optimal weights $\hat{\boldsymbol{w}}_{t-1}$ obtained at period $t-1$. Compared to offline REF, (1) includes a second regularization term $\|\boldsymbol{w}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2$ that draws the current weights towards the preceding optimal weights. Note that the prior weights $\boldsymbol{s}_t$ and preceding optimal weights $\hat{\boldsymbol{w}}_{t-1}$ are obtained differently and they play different roles in (1). The prior weights $\boldsymbol{s}_t$ summarize experts' overall historical forecast performances up to period $t-1$, whereas the preceding optimal weights $\hat{\boldsymbol{w}}_{t-1}$ record the most recent ensemble decision at period $t-1$. The three terms in (1) indicate that (1) tends to assign a larger weight to the ith expert if her/his current forecast $\widehat{y}_t^{(i)}$ is closer to the consensus $\bar{\widehat{y}}_t$ but also avoids assigning a weight that is substantially different from her/his historical performances or preceding optimal weights.

To sum up, by letting $\boldsymbol{Z}_{t-1} = (\boldsymbol{s}_t, \hat{\boldsymbol{w}}_{t-1})$, we select the optimal weight vector $\hat{\boldsymbol{w}}_t$ by (1) at period t, and accordingly obtain the ensemble forecast $\widehat{y}_t^{\mathrm{agg}} = \sum_{i=1}^N \hat{w}_{t,i}\widehat{y}_t^{(i)}$.

2.3 Online Update via $\boldsymbol{s}_t$

Here we describe how to update $\boldsymbol{Z}_t = (\boldsymbol{s}_{t+1}, \hat{\boldsymbol{w}}_t)$ after y_t is observed, for use in selecting $\hat{\boldsymbol{w}}_{t+1}$. It suffices to update $\boldsymbol{s}_{t+1}$.

Construction of $\boldsymbol{s}_t$. Following Su et al. (2026), we also adopt the CCR weights (Soule et al. 2024) as $\boldsymbol{s}_t$, which assumes a common correlation across experts' historical forecast errors. Explicitly, the prior weight assigned to the ith expert at period t is

$$s_{t,i}(\rho_t) = \frac{\{1 + (N-1)\rho_t\} v_{t,i}^{-2} - \rho_t v_{t,i}^{-1} \sum_{j=1}^N v_{t,j}^{-1}}{\{1 + (N-1)\rho_t\} \sum_{j=1}^N v_{t,j}^{-2} - \rho_t \left(\sum_{j=1}^N v_{t,j}^{-1}\right)^2}, \qquad i = 1, \ldots, N, \tag{2}$$

where $v_{t,i}^2$ denotes the ith expert's forecast-error variance, and ρ_t is the common correlation parameter across experts' historical forecast errors. The prior weights account for differences in recent forecast accuracy through $v_{t,i}^2$ and dependence across experts through ρ_t. When $\rho_t = 0$, CCR weights reduce to variance weights (Bates & Granger 1969). We estimate $v_{t,i}^2$ as described in Section 3.4 and update ρ_t in an online fashion.

Updating ρ_t. At period t, we obtain $\hat{\boldsymbol{w}}_t(\rho_t)$ from (1) based on $\boldsymbol{s}_t(\rho_t)$. After y_t is observed, the decision maker receives $R_{t+1} = -\ell_t(\hat{\boldsymbol{w}}_t(\rho_t))$. We then update ρ_t by projected online gradient descent (OGD) (Zinkevich 2003):

$$\rho_{t+1} = \Pi_{\mathcal{R}_t}\left\{ \rho_t - \eta_t \left. \frac{d}{d\rho} \ell_t\big(\hat{\boldsymbol{w}}_t(\rho)\big) \right|_{\rho=\rho_t} \right\} \tag{3}$$

where $\Pi_{\mathcal{R}_t}$ denotes projection onto the range that yields a valid common-correlation matrix for the current expert pool, and η_t denotes the step size. The update in (3) is applied separately to the three specifications of (1) corresponding to the choices of f described in Section 3.1, since each produces its own weights and aggregate forecast. The derivatives required for these updates are provided in Appendix E. The updated ρ_{t+1} and $v_{t+1,i}^2$ determine $\boldsymbol{s}_{t+1}$ at period $t+1$ to find $\hat{\boldsymbol{w}}_{t+1}$.

3 IMPLEMENTATION DETAILS

This section addresses implementation issues of online REF, including the choices of transformation f, sparse weight parameterization when the expert pool is large, selection of tuning parameters, and nuisance parameter estimations.

3.1 CHOICES OF TRANSFORMATION f

Following Su et al. (2026), we consider three online REF specifications, each defined by one of the following transformations of the current-forecast dispersion term: $f(x) = x$, $f(x) = \log x$, and $f(x) = \log(\sigma_t^2 + x)$, where $\sigma_t^2 = \mathrm{Var}(y_t) > 0$. The estimation of σ_t^2 is described in Section 3.4.

These transformations determine how the dispersion of current forecasts is scaled relative to the regularization terms. The choice of the identity transformation leaves the dispersion term on its original scale, whereas the two logarithmic transformations grow more slowly as dispersion increases. In practice, if there is no clear empirical guidance on the choice of f, one may obtain an ensemble forecast from each of the three specifications of f and then average the three forecasts to form the final online REF forecast.

3.2 ENTMAX: SPARSE WEIGHT PARAMETERIZATION

If the number of experts N is large, selecting an appropriate subset of experts may improve the ensemble forecast accuracy relative to combining the full pool (Kourentzes et al. 2019, Wang et al. 2023). Allowing exact zero weights provides a direct way to identify which experts enter the aggregate forecast at each forecasting period (Diebold & Shin 2019). We use the entmax parameterization (Peters et al. 2019, Martins & Astudillo 2016) to incorporate expert selection into the aggregation policy through sparse weights in Δ_N.

Let $Q_t(\boldsymbol{w})$ denote the objective minimized in (1). We implement the policy by minimizing $Q_t(\mathrm{entmax}_\alpha(\boldsymbol{\theta}))$ over $\boldsymbol{\theta} \in \mathbb{R}^N$ and setting $\hat{\boldsymbol{w}}_t = \mathrm{entmax}_\alpha(\hat{\boldsymbol{\theta}}_t)$, where $\hat{\boldsymbol{\theta}}_t$ is the resulting parameter vector, and α controls the entmax's sparsity. For $\alpha > 1$, entmax is defined by

$$\mathrm{entmax}_\alpha(\boldsymbol{\theta}) = \arg\max_{\boldsymbol{w} \in \Delta_N} \left\{ \boldsymbol{w}^\mathsf{T}\boldsymbol{\theta} + \frac{1 - \sum_{i=1}^N w_i^\alpha}{\alpha(\alpha - 1)} \right\}.$$

The entmax family approaches softmax as α approaches 1 and gives sparsemax at $\alpha = 2$ (Peters et al. 2019, Martins & Astudillo 2016). For $\alpha > 1$, entmax can assign zero weight to some experts. So the experts receiving positive weights form an active subset of the whole available expert pool at each forecasting period. Because the policy is evaluated again as new forecasts and historical information become available, the active expert pool may change over time, and receiving zero weight in one period does not exclude an expert from later periods.

3.3 SELECTION OF TUNING PARAMETERS

To select λ and κ in (1), we extend the rolling-window validation procedure of Su et al. (2026) using validation periods that precede the corresponding evaluation period. Sections 5.2 and 5.3 describe the specific validation and evaluation windows used in the empirical studies. For each candidate (λ, κ), we begin with the first validation period. At each validation period, we solve (1) using the expert forecasts and historical information available at that time, carrying the resulting weights forward as the preceding weights for the next validation period. We select the candidate values with the highest average forecast accuracy across the validation periods as the optimal tuning parameters, denoted by λ^* and κ^*. The tuning parameters are selected separately for the three specifications of f in (1).

For the entmax sparsity parameter α, we also select it by validation. For the averaged online REF forecast, the three specifications of f in (1) share a common entmax sparsity parameter α. For each candidate α, we select λ^* and κ^* for each specification as described above and evaluate the average of the resulting three validation forecasts. We select the candidate value with the highest average validation forecast accuracy as α^*.

3.4 Estimating and Initializing Nuisance Quantities

This section describes the estimation and initialization of nuisance quantities, including σ_t^2 involved in the transformation $f(x) = \log(\sigma_t^2 + x)$, the experts' forecast-error variances $v_{t,i}^2$, the initial correlation parameter ρ_1, and the initial weights $\hat{\boldsymbol{w}}_0$.

Estimating σ_t^2. We estimate σ_t^2, the variance of y_t, in the transformation $f(x) = \log(\sigma_t^2 + x)$ using the simple-mean forecast errors over L preceding periods: $\hat{\sigma}_t^2 = (L-1)^{-1} \sum_{r=t-L}^{t-1} (y_r - \bar{\widehat{y}}_r)^2$, where $\bar{\widehat{y}}_r$ is the simple mean of the expert forecasts available in period r, used to estimate the mean of y_r.

Estimating $v_{t,i}^2$. We estimate expert i's forecast-error variance $v_{t,i}^2$ using its forecast errors over the preceding historical window of length L, with $\hat{v}_{t,i}^2 = L^{-1} \sum_{r=t-L}^{t-1} (y_r - \widehat{y}_r^{(i)})^2$. When the expert pool varies over time, for example in surveys such as the SPF (see Section 5.3), where participation is voluntary and the set of experts varies across periods, the sum includes only the available forecast errors for expert i within this window, and the denominator is the number of these errors. We replace $v_{t,i}$ with $\hat{v}_{t,i}$ in (2) to obtain the prior weights. The historical window lengths L used in the empirical studies are given in Sections 5.2 and 5.3.

Initializing ρ_1. ρ_1 is initialized from the historical window before the first evaluation period following Soule et al. (2024). When the expert pool varies over time, ρ_1 is computed from the experts whose forecast errors are available throughout the initial window. During evaluation, we update ρ_t sequentially using (3), with N equal to the current pool size at period t. Before constructing the current prior weights, we restrict ρ_t to the range allowed by the current number of experts.

Initializing $\hat{\boldsymbol{w}}_0$. For the first period, we initialize $\hat{w}_{0,i} = 1/N$, $i = 1, \ldots, N$, where N is the number of available experts. When the expert pool changes, experts available in both periods use the weights assigned to them in the preceding period, while experts entering or re-entering the pool are assigned a preceding weight of zero.

4 Theoretical Property

We study the cumulative forecast loss of the online REF policy. Algorithm 1 summarizes the general online REF policy learning procedure. Given the forecasts and outcomes observed before period t,

Algorithm 1 Online REF policy learning

Require: Initial historical window, tuning parameters λ, κ.
1: Initialize $\boldsymbol{Z}_0 = (\boldsymbol{s}_1, \hat{\boldsymbol{w}}_0)$.
2: **for** $t = 1, \ldots, T$ **do**
3: Observe experts' forecasts $\widehat{\mathbf{y}}_t$ and set $\boldsymbol{Z}_{t-1} = (\boldsymbol{s}_t, \hat{\boldsymbol{w}}_{t-1})$.
4: Obtain $\hat{\boldsymbol{w}}_t$ by solving (1) based on $\widehat{\mathbf{y}}_t$ and $\boldsymbol{Z}_{t-1}$, and compute $\widehat{y}_t^{\text{agg}} = \sum_{i=1}^{N} \hat{w}_{t,i} \widehat{y}_t^{(i)}$.
5: Observe y_t and update $\boldsymbol{s}_{t+1}$.
6: **end for**

we assume that the current expert forecasts are independent and have a common mean. The full assumptions are given in Appendix A. The following theorem provides an expected average regret bound relative to the best fixed convex combination selected in hindsight.

Theorem 1. *For any $f(z) = z$, $\log z$, or $\log(\sigma_t^2 + z)$, under Assumptions 1–4, as $T \to \infty$, we have*

$$\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} \ell_t(\hat{\boldsymbol{w}}_t) - \min_{\boldsymbol{u} \in \Delta_N} \frac{1}{T}\sum_{t=1}^{T} \ell_t(\boldsymbol{u})\right] = O\left(\frac{1}{\sqrt{T}}\right). \tag{4}$$

The minimum is the average forecast loss of the oracle ensemble, which uses a single fixed weight vector selected in hindsight after all T outcomes are observed.

The result shows that the expected difference in average squared forecast loss between the sequential policy and the best fixed convex combination selected in hindsight vanishes as T increases. Proposition 1 in Appendix D provides an example using the CCR prior weights in (2). Under suitable conditions, the online REF policy using CCR prior weights satisfies the same regret bound. The following remark compares this guarantee with related work by clarifying the decision space and the asymptotic setting considered here.

Remark 1. *Theorem 1 establishes an expected regret guarantee for the online regularized weighting policy in* (1)*. This is consistent with existing regret analyses in online learning, where the number of forecasting periods T goes to infinity while the number of experts N is fixed (e.g., Zinkevich 2003, Shalev-Shwartz 2012). Our asymptotic result differs from Theorem 1 in Su et al. (2026), which develops the rate of convergence for the mean squared prediction error for a forecasting task at a single period as the number of experts increases.*

5 Empirical Evaluations

We evaluate online REF using daily Walmart sales forecasts from the M5 competition and macroeconomic forecasts from the SPF. We first introduce the comparison methods and evaluation metrics, then describe the data and report the results for each study. A simulation study in Appendix F examines forecast accuracy under different forecast-error correlations and the effect of the entmax parameter α on expert selection.

5.1 Methods in comparison and evaluation metrics

In the empirical studies, we average the forecasts from the three specifications of (1) corresponding to the transformations $f(x) = x$, $f(x) = \log x$, and $f(x) = \log(\sigma_t^2 + x)$ described in Section 3.1. The entmax parameter α is selected from $\{1, 1.5, 2, 4, 8, 16\}$ according to validation forecast accuracy as described in Section 3.3. We compare online REF with offline REF and six standard ensemble benchmarks:

- **Offline REF:** The average of the forecasts from the six specifications in Su et al. (2026), combining these three transformations with either L_2 or entropy penalties.
- **Simple Mean:** The equal-weighted average of the current expert forecasts.
- **Trimmed Mean:** The average of the current forecasts after removing the lowest and highest 10% (Jose & Winkler 2008).
- **Winsorized Mean:** The average after replacing the lowest and highest 15% of the current forecasts by the nearest remaining values (Jose & Winkler 2008).
- **Variance Weights:** A weighted average of the current expert forecasts, with weights proportional to the inverse of each expert's estimated historical forecast-error variance (Bates & Granger 1969).
- **Contribution-Weighted Method (CWM):** The simple average of the current forecasts from experts screened by their historical contribution to the simple-mean ensemble (Budescu & Chen 2015).
- **Common Correlation Weights (CCR):** A weighted average based on the experts' historical forecast-error variances and a common correlation parameter across their forecast errors (Soule et al. 2024).

To compare forecast accuracy across time series measured on different scales, we use the root mean squared scaled error (RMSSE), as used in the M5 competition (Makridakis et al. 2022a). By scaling the mean squared forecast error using the in-sample error of a one-step naive forecast, RMSSE facilitates comparisons across series with different magnitudes and historical variation. For each series, let $t_1, \ldots, t_2$ denote the validation or evaluation periods, and let $y_1^{\mathrm{pre}}, \ldots, y_{T_{\mathrm{pre}}}^{\mathrm{pre}}$ denote the pre-validation series used to compute the scale. The RMSSE of the aggregate forecast is

$$\mathrm{RMSSE} = \left\{ \frac{\sum_{t=t_1}^{t_2} (y_t - \widehat{y}_t^{\mathrm{agg}})^2 / (t_2 - t_1 + 1)}{\sum_{r=2}^{T_{\mathrm{pre}}} (y_r^{\mathrm{pre}} - y_{r-1}^{\mathrm{pre}})^2 / (T_{\mathrm{pre}} - 1)} \right\}^{1/2}. \tag{5}$$

The denominator is the mean squared error of a one-step naive forecast on the pre-validation observations. We compute RMSSE separately for each series before averaging across series or horizons.

5.2 M5 COMPETITION

The M5 competition, hosted on Kaggle in 2020, focused on forecasting daily unit sales of Walmart products over a 28-day horizon (Makridakis et al. 2022a,b). The data cover 3,049 products, with sales aggregated across products and stores into 12 hierarchical levels defined by state, store, product category, and department. Following Su et al. (2026), we focus on the 154 sales series in the first nine levels. These include total sales and aggregates by state, store, category, department, state–category, state–department, store–category, and store–department. We use all 50 publicly released team forecasts from the M5 accuracy competition, allowing online REF to select experts from a large expert pool using the entmax parameterization described in Section 3.2.

We use Days 1–21 as the initial historical window, with Days 15–21 serving as the validation period, and Days 22–28 as the online evaluation period. Each forecasting period uses a rolling historical window of length $L = 14$. Following the rolling-window validation described in Section 3.3, we select λ^* and κ^* separately for each series based on forecast accuracy over Days 15–21. We select α^* by minimizing validation RMSSE averaged across the nine aggregation levels over the same validation period. The details are reported in Appendix G.

During Days 22–28, the forecasts are aggregated sequentially. We report RMSSE averaged across the series within each of the nine aggregation levels (avgRMSSE). Overall performance is summarized by the mean of the nine level-specific avgRMSSE values and the ranking of methods, with rank 1 indicating the best performance. Following the M5 definition of RMSSE (Makridakis et al. 2022a), we compute the denominator in (5) using the 1,941-day sales history before the 28-day forecasting period, starting with the first nonzero sale.

Table 1 shows that online REF produces the most accurate forecasts at all nine aggregation levels compared with the other benchmarks. Offline REF ranks second overall, followed by CCR. Although offline REF performs better overall than methods using only current forecasts, such as Simple Mean and Trimmed Mean, or historical performance, such as Variance Weights and CCR, online REF achieves further improvements by incorporating preceding optimal weights and expert selection through the entmax parameterization. The full results, including standard deviation (Std.) values and percentage improvements relative to Simple Mean, are provided in Table A.2.

Table 1: Performance on M5 by aggregation level (avgRMSSE). Bold numbers indicate the lowest avgRMSSE among the methods shown at each level.

Level	**Online REF**	Offline REF	Simple Mean	Trimmed Mean	Winsorized Mean	Variance Weights	CWM	CCR
1	**.115**	.126	.147	.147	.147	.139	.129	.121
2	**.249**	.258	.292	.296	.295	.289	.288	.285
3	**.322**	.331	.347	.348	.349	.346	.354	.341
4	**.226**	.233	.243	.247	.248	.240	.256	.235
5	**.331**	.343	.374	.381	.383	.356	.356	.341
6	**.318**	.325	.350	.352	.353	.345	.337	.333
7	**.453**	.457	.490	.493	.494	.478	.466	.468
8	**.444**	.455	.464	.464	.464	.463	.455	.464
9	**.578**	.582	.600	.602	.602	.594	.584	.582
Avg.	**.337**	.346	.368	.370	.371	.361	.359	.352
Rank[a]	1	2	6	7	8	5	4	3

[a] Ranked by avgRMSSE across aggregation levels.

5.3 SURVEY OF PROFESSIONAL FORECASTERS

The SPF, maintained by the Federal Reserve Bank of Philadelphia, is a quarterly survey in which professional forecasters participate voluntarily and provide forecasts of major economic indicators (Clements et al. 2023, Federal Reserve Bank of Philadelphia 2026).

We consider three commonly used economic indicators, including annualized quarter-over-quarter growth in nominal gross domestic product (GDP), annualized quarter-over-quarter Consumer Price Index (CPI) inflation, and the quarterly change in the unemployment rate. For each economic indicator, we study the current-quarter forecast, denoted by $h = 0$, and the one-quarter-ahead forecast, denoted by $h = 1$. The realizations are constructed from the initial-release vintages of the Real-Time Data Set for Macroeconomists (Croushore & Stark 2001).

The evaluation period runs from 2000Q1 through 2025Q2 and contains 102 quarters. Following the rolling-window validation procedure in Section 3.3, for each evaluation quarter, we use the preceding 8 quarters for validation to select λ^* and κ^*. Each forecasting quarter uses a rolling historical window of length $L = 8$. We select α^* using validation RMSSE based on the validation window for each economic indicator. The details are reported in Appendix G.

Unlike the M5 competition, the SPF has a varying expert pool, as forecasters may enter, leave, and later return to the survey. In each evaluation quarter, we include forecasters who provide the relevant forecast and have at least one historical forecast over the preceding eight-quarter historical window. Missing forecasts are not imputed during evaluation. During validation, missing forecasts are replaced by the mean forecast across the available experts in that quarter. When a forecaster included in the current evaluation quarter has no forecast errors available in the historical window preceding a given validation quarter, $\widehat{v}^{-2}$ is replaced by the mean of the corresponding estimates across the other experts.

We report results separately for GDP growth, CPI inflation, and the change in the unemployment rate, and for horizons $h = 0$ and $h = 1$. For SPF, the RMSSE denominator in (5) uses the historical values of each economic indicator before 1996Q1, the beginning of the initial historical window.

Table 2 shows that online REF performs best across both horizons for all three indicators. The improvement of online REF over offline REF is more apparent for CPI inflation and is largest for the forecast of the change in the unemployment rate at $h = 0$. Compared with the M5 study, CCR is less competitive overall, whereas Winsorized Mean is more competitive. This pattern suggests that historical performance may be less informative for weighting in the SPF, making current forecasts relatively more useful. Offline REF combines these two sources of information and outperforms both methods overall. Online REF additionally incorporates preceding optimal weights into the weighting decision and allows expert selection through the entmax parameterization, achieving further improvements in forecast accuracy. Full results, including Std. values and percentage improvements relative to Simple Mean, are reported in Table A.3 of Appendix G.

Table 2: Performance on SPF by economic indicator and forecast horizon (RMSSE). Bold numbers indicate the lowest RMSSE among the methods shown for each indicator and horizon.

Indicator	Horizon	**Online** REF	Offline REF	Simple Mean	Trimmed Mean	Winsorized Mean	Variance Weights	CWM	CCR
Nominal GDP growth	$h = 0$	**.561**	.562	.600	.602	.600	.612	.621	.623
	$h = 1$	**1.506**	1.517	1.597	1.593	1.594	1.591	1.600	1.584
	Avg.	**1.033**	1.040	1.098	1.097	1.097	1.102	1.111	1.104
	Rank[a]	1	2	5	4	3	6	8	7
CPI inflation rate	$h = 0$	**.656**	.675	.758	.754	.755	.812	.799	.794
	$h = 1$	**1.129**	1.192	1.195	1.195	1.193	1.215	1.253	1.203
	Avg.	**.892**	.933	.977	.974	.974	1.013	1.026	.998
	Rank[a]	1	2	5	4	3	7	8	6
Change in unemployment rate	$h = 0$	**.454**	.609	.814	.829	.825	.786	2.620	.963
	$h = 1$	**2.481**	2.495	2.713	2.665	2.672	2.594	2.614	2.578
	Avg.	**1.468**	1.552	1.764	1.747	1.749	1.690	2.617	1.771
	Rank[a]	1	2	6	4	5	3	8	7

[a] Ranked within each economic indicator by average RMSSE across forecast horizons.

6 DISCUSSION

In this paper, we develop online REF for sequential ensemble forecasting by extending offline REF to an online policy learning framework. The policy uses experts' current forecasts and historical performance and additionally incorporates preceding optimal weights into the current weighting decision. The prior weights are updated using newly observed forecast errors and the aggregate forecast loss, while the entmax parameterization allows exact zero weights and expert selection. Under suitable conditions, we establish a vanishing expected average regret bound relative to the best fixed convex combination of expert forecasts selected in hindsight. Numerical studies, including a simulation study and two empirical applications to the M5 competition and the SPF, show that online REF performs favorably against offline REF and standard ensemble benchmarks. A future direction is to extend online REF to combine predictive distributions (e.g., Billio et al. 2013, Bassetti et al. 2018).

AI USE STATEMENT

The research ideas and reasoning in this paper were developed by the authors. We used generative AI tools to assist with mathematical proofs, identify relevant literature, and improve the language and presentation of text written by the authors. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We provide implementation code for online REF at `https://anonymous.4open.science/r/online-regularized-ensemble-forecasting-0301/`. The repository also includes simulation code for data generation, parameter selection, and sequential weight updates. Details on implementation and tuning parameters are given in Section 3. The simulation design and empirical settings are described in Appendix F and Section 5, respectively. The theoretical assumptions and proofs are provided in Appendices A, C, and D.

REFERENCES

Federico Bassetti, Roberto Casarin, and Francesco Ravazzolo. Bayesian nonparametric calibration and combination of predictive distributions. *Journal of the American Statistical Association*, 113(522):675–685, 2018.

John M. Bates and Clive W. J. Granger. The combination of forecasts. *Journal of the Operational Research Society*, 20(4):451–468, 1969.

Monica Billio, Roberto Casarin, Francesco Ravazzolo, and Herman K. van Dijk. Time-varying combinations of predictive densities using nonlinear filtering. *Journal of Econometrics*, 177(2):213–232, 2013.

Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

David V. Budescu and Eva Chen. Identifying expertise to extract the wisdom of crowds. *Management Science*, 61(2):267–280, 2015.

Robert T. Clemen. Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5(4):559–583, 1989.

Michael P. Clements, Robert W. Rich, and Joseph S. Tracy. Surveys of professionals. In *Handbook of Economic Expectations*, pp. 71–106. Elsevier, 2023.

Dean Croushore and Tom Stark. A real-time data set for macroeconomists. *Journal of Econometrics*, 105(1): 111–130, 2001.

Marie Devaine, Pierre Gaillard, Yannig Goude, and Gilles Stoltz. Forecasting electricity consumption by aggregating specialized experts: A review of the sequential aggregation of specialized experts, with an application to Slovakian and French country-wide one-day-ahead (half-)hourly predictions. *Machine Learning*, 90(2):231–260, 2013.

Francis X. Diebold and Peter Pauly. The use of prior information in forecast combination. *International Journal of Forecasting*, 6(4):503–508, 1990.

Francis X. Diebold and Minchul Shin. Machine learning for regularized survey forecast combination: Partially-egalitarian LASSO and its derivatives. *International Journal of Forecasting*, 35(4):1679–1691, 2019.

Federal Reserve Bank of Philadelphia. Survey of Professional Forecasters documentation. Online documentation, philadelphiafed.org/. . . /survey-of-professional-forecasters, 2026. Last updated: August 14, 2026.

Clive W. J. Granger and Ramu Ramanathan. Improved methods of combining forecasts. *Journal of Forecasting*, 3(2):197–204, 1984.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

Roger A. Horn and Charles R. Johnson. *Matrix analysis*. Cambridge University Press, 2nd edition, 2012.

Victor Richmond R. Jose and Robert L. Winkler. Simple robust averages of forecasts: Some empirical results. *International Journal of Forecasting*, 24(1):163–169, 2008.

Nikolaos Kourentzes, Devon Barrow, and Fotios Petropoulos. Another look at forecast selection and combination: Evidence from forecast pooling. *International Journal of Production Economics*, 209:226–235, 2019.

Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4):1346–1364, 2022a.

Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. The M5 competition: Background, organization, and implementation. *International Journal of Forecasting*, 38(4):1325–1336, 2022b.

André F. T. Martins and Ramón F. Astudillo. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1614–1623. PMLR, 2016.

Paul Newbold and Clive W. J. Granger. Experience with forecasting univariate time series and the combination of forecasts. *Journal of the Royal Statistical Society: Series A (General)*, 137(2):131–146, 1974.

Ben Peters, Vlad Niculae, and André F. T. Martins. Sparse sequence-to-sequence models. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1504–1519. Association for Computational Linguistics, 2019.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.

David Soule, Yael Grushka-Cockayne, and Jason Merrick. A heuristic for combining correlated experts when there are few data. *Management Science*, 70(10):6637–6668, 2024.

Han Su, Xiaojia Guo, and Xiaoke Zhang. Regularized ensemble forecasting for learning weights from historical and current forecasts. arXiv preprint arXiv:2602.11379, 2026.

Xiaoqian Wang, Rob J. Hyndman, Feng Li, and Yanfei Kang. Forecast combinations: An over 50-year review. *International Journal of Forecasting*, 39(4):1518–1547, 2023.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning*, pp. 928–936. AAAI Press, 2003.

APPENDIX

Appendix A states the assumptions for Theorem 1. Appendix B derives the variance term in (1), providing the intuition behind it. Appendix C provides the proof of Theorem 1. Appendix D establishes Proposition 1, which provides an example of CCR prior weights satisfying Assumption 4 under suitable conditions. Under the remaining assumptions of Theorem 1, the online REF policy using these CCR prior weights satisfies the same regret bound. Appendix E provides the derivatives used to update ρ_t in (3). Appendix F examines forecast accuracy under different forecast-error correlations and the effect of the entmax parameter α on expert selection. Appendix G reports additional results for the M5 and SPF studies.

CONTENTS

A Technical Assumptions **13**

B Motivation for the Variance Term **14**

C Proof of Theorem 1 . . . . . **14**

D An Example: CCR Prior Weights **25**

E Derivatives used in the online update **31**

F Simulation Study **32**

G Additional Empirical Results **34**

A TECHNICAL ASSUMPTIONS

The number $N \geq 2$ of experts is fixed. Let $\Delta_N = \{\boldsymbol{w} \in \mathbb{R}^N : w_i \geq 0,\ \mathbf{1}^\top \boldsymbol{w} = 1\}$, $\widehat{\boldsymbol{y}}_t = (\widehat{y}_t^{(1)}, \ldots, \widehat{y}_t^{(N)})^\top$, and $\bar{\widehat{y}}_t = N^{-1}\sum_{i=1}^N \widehat{y}_t^{(i)}$. In the analysis below, σ_t^2 denotes the variance parameter used in $f(x) = \log(\sigma_t^2 + x)$, whether known or estimated from history.

Before period t, the prior weights $\boldsymbol{s}_t$, the preceding optimal weights $\hat{\boldsymbol{w}}_{t-1}$, and any historically estimated variance parameter σ_t^2 have been constructed without using y_t. The current expert forecasts may then be used to select $\hat{\boldsymbol{w}}_t$, but the current outcome is observed only after the aggregate forecast has been formed. Let $\mathcal{F}_t$ denote the σ-field generated by the initial history and the expert forecasts and outcomes observed through period t. Thus $\boldsymbol{s}_t$, $\hat{\boldsymbol{w}}_{t-1}$, and σ_t^2 are $\mathcal{F}_{t-1}$-measurable. The weights $\hat{\boldsymbol{w}}_t$ are selected using this historical information and the current expert forecasts before y_t is observed.

The prior weights satisfy $\mathbf{1}^\top \boldsymbol{s}_t = 1$, and $\hat{\boldsymbol{w}}_0 \in \Delta_N$. For each horizon T, write $\lambda = \lambda_T > 0$ and $\kappa = \kappa_T \geq 0$, both deterministic. The policy is run separately for each T, so the resulting weights could be denoted by $\hat{\boldsymbol{w}}_{t,T}$. For notational simplicity, we write $\hat{\boldsymbol{w}}_t$ throughout. The online REF policy rule is given in (1). Recall that, for $\boldsymbol{w} \in \Delta_N$, the squared forecast loss is $\ell_t(\boldsymbol{w}) = (y_t - \sum_{i=1}^N w_i \widehat{y}_t^{(i)})^2 = \{\sum_{i=1}^N w_i(\widehat{y}_t^{(i)} - y_t)\}^2$.

Assumption 1. *For every period t, conditional on $\mathcal{F}_{t-1}$, the current expert forecasts are independent, have a common conditional mean with $\mathbb{E}[\widehat{y}_t^{(i)} \mid \mathcal{F}_{t-1}] = \mu_t$ and have finite conditional second moments with $\mathbb{E}[(\widehat{y}_t^{(i)})^2 \mid \mathcal{F}_{t-1}] < \infty$, for $i = 1, \ldots, N$. The transformation $f(x)$ in* (1) *satisfies one of the following three sets of conditions.*

(i) For $f(z) = z$, no additional condition is imposed.

(ii) For $f(z) = \log(\sigma_t^2 + z)$, there is a constant $\underline{\sigma}^2 > 0$, independent of t and T, such that $\sigma_t^2 \geq \underline{\sigma}^2$ almost surely.

(iii) For $f(z) = \log z$, the conditional density of each $\widehat{y}_t^{(i)}$ given $\mathcal{F}_{t-1}$ is bounded above on $\mathbb{R}$ by a finite constant $\bar{g} > 0$, independent of i, t, and T, almost surely.

Assumption 2. *The forecast errors satisfy the following conditions:*

(i) There exists a finite constant $H > 0$, independent of t and T, such that $\max_{1 \leq i \leq N} |\widehat{y}_t^{(i)} - y_t| \leq H$, almost surely.

(ii) There exists a deterministic fixed matrix $\boldsymbol{M} = (M_{ij})_{i,j=1}^N$ such that, for every t and all i, j,

$$\mathbb{E}\left[(\widehat{y}_t^{(i)} - y_t)(\widehat{y}_t^{(j)} - y_t) \mid \mathcal{F}_{t-1}\right] = M_{ij}.$$

Assumption 3. *The tuning parameters satisfy $\sup_T \kappa_T < \infty$. Moreover, $\lambda_T \asymp \sqrt{T}$, for $f(z) = z$ or $f(z) = \log(\sigma_t^2 + z)$, and $\lambda_T \asymp T$ for $f(z) = \log z$.*

Assumption 4. *Define the population prior weights by $\boldsymbol{s}^\star \in \arg\min_{\boldsymbol{w} \in \Delta_N} \boldsymbol{w}^\top \boldsymbol{M} \boldsymbol{w}$. There is a finite constant C_1, independent of t and T, such that the projected prior weights $\widetilde{\boldsymbol{s}}_t = \Pi_{\Delta_N}(\boldsymbol{s}_t)$ satisfy $\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 \leq C_1/t$, $t = 1, \ldots, T$. Here $\Pi_{\Delta_N}(\boldsymbol{s})$ is the vector in Δ_N closest to $\boldsymbol{s}$ in Euclidean distance.*

For $f(x) = \log(\sigma_t^2 + x)$, the positive lower bound on σ_t^2 in Assumption 1(ii) keeps the transformation bounded below, while for $f(x) = \log x$, the density condition in Assumption 1(iii) controls behavior near zero for $f(x) = \log x$. Bounded forecast errors in Assumption 2(i), a standard condition in sequential forecast aggregation (Devaine et al. 2013), provide the boundedness needed for the regret analysis. Assumption 2(ii) keeps the conditional forecast-error second-moment matrix fixed across periods, so the population prior weights $\boldsymbol{s}^\star$ minimize the same conditional expected squared loss in every period. The conditions on λ_T and κ_T in Assumption 3 ensure that the online REF weights remain sufficiently close to the projected prior weights for the regret bound. Assumption 4 requires the projected historical prior weights to approach $\boldsymbol{s}^\star$ at mean squared rate $1/t$. Proposition 1 provides an example in which CCR prior weights constructed from a rolling window satisfy this assumption.

B MOTIVATION FOR THE VARIANCE TERM

At period t, given $\mathcal{F}_{t-1}$, consider a candidate weight vector $\boldsymbol{w}_t \in \Delta_N$. Under Assumption 1, the current expert forecasts $\{\widehat{y}_t^{(i)}\}_{i=1}^N$ are independent with a common mean and finite second moments. Hence,

$$\begin{aligned}\mathrm{Var}\left(\sum_{i=1}^N w_{t,i}\widehat{y}_t^{(i)}\right) &= \sum_{i=1}^N w_{t,i}^2\,\mathrm{Var}(\widehat{y}_t^{(i)}) \\ &= \sum_{i=1}^N w_{t,i}^2\mathbb{E}[(\widehat{y}_t^{(i)} - \mathbb{E}[\widehat{y}_t^{(i)}])^2].\end{aligned}$$

We estimate the common mean by $\bar{\widehat{y}}_t = N^{-1}\sum_{i=1}^N \widehat{y}_t^{(i)}$. Replacing each expected squared deviation with the corresponding observed squared deviation gives the empirical approximation $\sum_{i=1}^N w_{t,i}^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2$, which is the variance term used in (1).

C PROOF OF THEOREM 1

The following three lemmas provide the auxiliary results used in the proof of Theorem 1.

Lemma 1. *Under Assumption 1(iii), and Assumption 2(i), there exists a constant $C_2 < \infty$, independent of t and T, such that, for every deterministic $\lambda > 0$,*

$$\mathbb{E}\left[\sqrt{\frac{1}{\lambda}\log\left(\frac{4NH^2}{m_t^2}\right)}\right] \le \frac{C_2}{\sqrt{\lambda}} < \infty,$$

where $m_t := \min_{1\le i\le N}|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t|$. Moreover, $m_t > 0$ almost surely.

Proof. Fix t, i, and $x > 0$. Since $\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t = ((N-1)/N)\widehat{y}_t^{(i)} - (1/N)\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} = ((N-1)/N)\{\widehat{y}_t^{(i)} - (1/(N-1))\sum_{\ell\neq i}\widehat{y}_t^{(\ell)}\}$, we have

$$\begin{aligned}|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le x &\iff \left|\widehat{y}_t^{(i)} - \frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)}\right| \le \frac{Nx}{N-1} \\ &\iff \frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} - \frac{Nx}{N-1} \le \widehat{y}_t^{(i)} \le \frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} + \frac{Nx}{N-1}.\end{aligned}$$

For fixed values of the other forecasts, the admissible interval for $\widehat{y}_t^{(i)}$ therefore has length $2Nx/(N-1)$.

Let $g_{t,i}(a \mid \mathcal{F}_{t-1})$ be the conditional density of $\widehat{y}_t^{(i)}$ at $a \in \mathbb{R}$. Conditional independence in Assumption 1 implies that this remains its density after conditioning on all the other current forecasts. Using condition (iii) and the interval above, we have

$$\begin{aligned}&\Pr\left(|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le x \;\middle|\; \mathcal{F}_{t-1}, (\widehat{y}_t^{(\ell)})_{\ell\neq i}\right) \\ &= \Pr\left(\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} - \frac{Nx}{N-1} \le \widehat{y}_t^{(i)} \le \frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} + \frac{Nx}{N-1} \;\middle|\; \mathcal{F}_{t-1}, (\widehat{y}_t^{(\ell)})_{\ell\neq i}\right) \\ &= \int_{\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} - \frac{Nx}{N-1}}^{\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} + \frac{Nx}{N-1}} g_{t,i}(a \mid \mathcal{F}_{t-1})\,da \le \bar{g}\int_{\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} - \frac{Nx}{N-1}}^{\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} + \frac{Nx}{N-1}} 1\,da \\ &= \bar{g}\left\{\left(\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} + \frac{Nx}{N-1}\right) - \left(\frac{1}{N-1}\sum_{\ell\neq i}\widehat{y}_t^{(\ell)} - \frac{Nx}{N-1}\right)\right\} = \frac{2N\bar{g}}{N-1}x.\end{aligned}$$

Taking conditional expectations given $\mathcal{F}_{t-1}$ gives

$$\Pr\Big(|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le x \mid \mathcal{F}_{t-1}\Big) = \mathbb{E}\Big[\Pr\Big(|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le x \;\Big|\; \mathcal{F}_{t-1}, (\widehat{y}_t^{(\ell)})_{\ell \neq i}\Big)\Big|\mathcal{F}_{t-1}\Big]$$
$$\le \mathbb{E}\left[\frac{2N\bar{g}}{N-1}x\middle|\mathcal{F}_{t-1}\right] = \frac{2N\bar{g}}{N-1}x.$$

For $m_t = \min_{1\le i\le N} |\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t|$, the conditional union bound yields

$$\Pr(m_t \le x \mid \mathcal{F}_{t-1}) = \Pr\left(\bigcup_{i=1}^{N}\{|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le x\}\middle|\mathcal{F}_{t-1}\right)$$
$$\le \sum_{i=1}^{N}\Pr\Big(|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le x \mid \mathcal{F}_{t-1}\Big) \le \sum_{i=1}^{N}\frac{2N\bar{g}}{N-1}x = \frac{2N^2\bar{g}}{N-1}x.$$

Taking expectations gives, for every $x > 0$,

$$\Pr(m_t \le x) = \mathbb{E}[\Pr(m_t \le x \mid \mathcal{F}_{t-1})] \le \frac{2N^2\bar{g}}{N-1}x. \tag{6}$$

In particular, $0 \le \Pr(m_t = 0) \le \Pr(m_t \le x) \le (2N^2\bar{g})/(N-1)x \to 0$ as $x \downarrow 0$. Hence $m_t > 0$ almost surely. By Assumption 2(i),

$$|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| = \left|(\widehat{y}_t^{(i)} - y_t) - \frac{1}{N}\sum_{j=1}^{N}(\widehat{y}_t^{(j)} - y_t)\right|$$
$$\le |\widehat{y}_t^{(i)} - y_t| + \frac{1}{N}\sum_{j=1}^{N}|\widehat{y}_t^{(j)} - y_t| \le H + \frac{1}{N}\sum_{j=1}^{N}H = 2H.$$

Thus $0 < m_t \le 2H$ almost surely.

Define $X_t = \sqrt{\log(4NH^2/m_t^2)}$. Since $m_t = \min_{1\le i\le N}|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le 2H$, then $4NH^2/m_t^2 \ge N \ge 2$, so X_t is finite and nonnegative almost surely.

The bound in (6) holds for every $x > 0$. For every $z \ge 0$, since $m_t > 0$ almost surely,

$$\Pr(X_t > z) = \Pr\Big(\sqrt{\log(4NH^2/m_t^2)} > z\Big) = \Pr\big(\log(4NH^2/m_t^2) > z^2\big)$$
$$= \Pr\Big(m_t^2 < 4NH^2e^{-z^2}\Big) = \Pr\Big(m_t < \sqrt{4NH^2}\,e^{-z^2/2}\Big)$$
$$\le \Pr\Big(m_t \le \sqrt{4NH^2}\,e^{-z^2/2}\Big) \le \frac{2N^2\bar{g}}{N-1}\sqrt{4NH^2}\,e^{-z^2/2}.$$

Since X_t is nonnegative almost surely, the tail-integral formula gives

$$\mathbb{E}X_t = \int_0^\infty \Pr(X_t > z)\,dz \le \int_0^\infty \frac{2N^2\bar{g}}{N-1}\sqrt{4NH^2}\,e^{-z^2/2}\,dz$$
$$= \frac{2N^2\bar{g}}{N-1}\sqrt{4NH^2}\int_0^\infty e^{-z^2/2}\,dz = \frac{2N^2\bar{g}}{N-1}\sqrt{4NH^2}\sqrt{\frac{\pi}{2}} =: C_2 < \infty.$$

Here we used $\int_0^\infty e^{-z^2/2}\,dz = \sqrt{\pi/2}$. Since N is fixed and $\bar{g}$ and H are independent of t and T, C_2 is also independent of t and T. Therefore, for every deterministic $\lambda > 0$,

$$\mathbb{E}\left[\sqrt{\frac{1}{\lambda}\log\Big(\frac{4NH^2}{m_t^2}\Big)}\right] = \frac{1}{\sqrt{\lambda}}\mathbb{E}\left[\sqrt{\log\Big(\frac{4NH^2}{m_t^2}\Big)}\right] = \frac{1}{\sqrt{\lambda}}\mathbb{E}X_t \le \frac{C_2}{\sqrt{\lambda}} < \infty.$$

Moreover, $m_t > 0$ almost surely and, for every $\boldsymbol{w} \in \Delta_N$,

$$\sum_{i=1}^{N} w_i^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \ge m_t^2\sum_{i=1}^{N}w_i^2 \ge m_t^2/N > 0.$$

Therefore, the objective in (1) is finite and continuous on Δ_N almost surely. □

Lemma 2. *Under Assumptions 1–2, for every t, the projected prior weights $\widetilde{\boldsymbol{s}}_t$ satisfy*

$$\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \le \begin{cases} \dfrac{8H^2 + 2\sqrt{2}\kappa}{\lambda}, & f(x) = x, \\ \dfrac{8H^2/\underline{\sigma}^2 + 2\sqrt{2}\kappa}{\lambda}, & f(x) = \log(\sigma_t^2 + x), \\ \dfrac{2\sqrt{2}\kappa}{\lambda} + \sqrt{\dfrac{1}{\lambda}\log\left(\dfrac{4NH^2}{m_t^2}\right)}, & f(x) = \log x. \end{cases}$$

Proof. For prior weights $\boldsymbol{s}_t$ satisfying $\mathbf{1}^\top \boldsymbol{s}_t = 1$, define the projected prior weights $\widetilde{\boldsymbol{s}}_t = \Pi_{\Delta_N}(\boldsymbol{s}_t)$. Since $\hat{\boldsymbol{w}}_t$ are optimal weights, we have:

$$\begin{aligned} & f\left(\sum_{i=1}^N (\hat{w}_{t,i})^2 \big(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t\big)^2\right) + \lambda\|\hat{\boldsymbol{w}}_t - \boldsymbol{s}_t\|_2^2 + \kappa\|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 \\ & \quad \le f\left(\sum_{i=1}^N \widetilde{s}_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) + \lambda\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t\|_2^2 + \kappa\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 \\ \Rightarrow & \lambda\big\{\|\hat{\boldsymbol{w}}_t - \boldsymbol{s}_t\|_2^2 - \|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t\|_2^2\big\} \le f\left(\sum_{i=1}^N \widetilde{s}_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) - f\left(\sum_{i=1}^N (\hat{w}_{t,i})^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) \\ & \quad + \kappa\big\{\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 - \|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2\big\}. \end{aligned}$$

1) For the left-hand side $\|\hat{\boldsymbol{w}}_t - \boldsymbol{s}_t\|_2^2 - \|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t\|_2^2$:

By the definition of Euclidean projection, $\widetilde{\boldsymbol{s}}_t = \Pi_{\Delta_N}(\boldsymbol{s}_t)$ minimizes $\|\boldsymbol{s}_t - \boldsymbol{z}\|_2^2$ over $\boldsymbol{z} \in \Delta_N$. For every $\boldsymbol{w} \in \Delta_N$ and $\alpha \in [0,1]$, the convexity of Δ_N gives $\widetilde{\boldsymbol{s}}_t + \alpha(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t) \in \Delta_N$. Therefore, for $\alpha \in (0,1]$,

$$\begin{aligned} \|\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 &\le \|\boldsymbol{s}_t - \{\widetilde{\boldsymbol{s}}_t + \alpha(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t)\}\|_2^2 = \|\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t - \alpha(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t)\|_2^2 \\ &= \|\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 - 2\alpha(\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t)^\top(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t) + \alpha^2\|\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t\|_2^2, \\ 0 &\le -2(\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t)^\top(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t) + \alpha\|\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t\|_2^2. \end{aligned}$$

Letting $\alpha \downarrow 0$ yields $(\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t)^\top(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t) \le 0$. Then,

$$\begin{aligned} \|\hat{\boldsymbol{w}}_t - \boldsymbol{s}_t\|_2^2 - \|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t\|_2^2 &= \|(\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t) + (\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t)\|_2^2 - \|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t\|_2^2 \\ &= \|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 + 2(\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t)^\top(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}_t) \ge \|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 \cdot ((\boldsymbol{s}_t - \widetilde{\boldsymbol{s}}_t)^\top(\boldsymbol{w} - \widetilde{\boldsymbol{s}}_t) \le 0) \end{aligned}$$

2) For $\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 - \|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2$:

$$\begin{aligned} & \big|\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 - \|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2\big| \\ &= \big|(\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t)^\top(\widetilde{\boldsymbol{s}}_t + \hat{\boldsymbol{w}}_t - 2\hat{\boldsymbol{w}}_{t-1})\big| && (\|\boldsymbol{a}\|_2^2 - \|\boldsymbol{b}\|_2^2 = (\boldsymbol{a}-\boldsymbol{b})^\top(\boldsymbol{a}+\boldsymbol{b})) \\ &\le \|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2 \|\widetilde{\boldsymbol{s}}_t + \hat{\boldsymbol{w}}_t - 2\hat{\boldsymbol{w}}_{t-1}\|_2 && (\text{Cauchy–Schwarz: } |\boldsymbol{x}^\top\boldsymbol{y}| \le \|\boldsymbol{x}\|_2\|\boldsymbol{y}\|_2) \\ &\le \|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2 \{\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2 + \|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2\} && (\text{triangle inequality: } \|\boldsymbol{x}+\boldsymbol{y}\|_2 \le \|\boldsymbol{x}\|_2 + \|\boldsymbol{y}\|_2) \\ &\le 2\sqrt{2}\,\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2 && (\|\boldsymbol{x}-\boldsymbol{y}\|_2 \le \sqrt{2} \text{ for } \boldsymbol{x}, \boldsymbol{y} \in \Delta_N). \end{aligned}$$

Then,

$$\begin{aligned} \lambda\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 &\le f\left(\sum_{i=1}^N \widetilde{s}_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) - f\left(\sum_{i=1}^N (\hat{w}_{t,i})^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) \\ &\quad + \kappa\big\{\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 - \|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2\big\} \\ &\le \left|f\left(\sum_{i=1}^N \widetilde{s}_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) - f\left(\sum_{i=1}^N (\hat{w}_{t,i})^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right)\right| \\ &\quad + \kappa\big|\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2 - \|\hat{\boldsymbol{w}}_t - \hat{\boldsymbol{w}}_{t-1}\|_2^2\big| \\ &\quad (\text{triangle inequality: } x + \kappa y \le |x| + \kappa|y| \text{ for } \kappa \ge 0) \\ &\le \left|f\left(\sum_{i=1}^N \widetilde{s}_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right) - f\left(\sum_{i=1}^N (\hat{w}_{t,i})^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right)\right| + 2\sqrt{2}\,\kappa\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2. \end{aligned}$$

3) For $|f(\sum_{i=1}^N \widetilde{s}_{t,i}^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2) - f(\sum_{i=1}^N (\hat{w}_{t,i})^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2)|$:

By definition of $\bar{\widehat{y}}_t$, $\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t = (\widehat{y}_t^{(i)} - y_t) - N^{-1}\sum_{j=1}^N (\widehat{y}_t^{(j)} - y_t)$, and Assumption 2(i) gives $\max_i |\widehat{y}_t^{(i)} - y_t| \le H$. Then, by triangle inequality: $|x+y| \le |x| + |y|$, we have $|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le |\widehat{y}_t^{(i)} - y_t| + N^{-1}\sum_{j=1}^N |\widehat{y}_t^{(j)} - y_t| \le 2H$. Direct differentiation gives

$$
\begin{aligned}
&\left\| \nabla_{\boldsymbol{w}} \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right\|_2 \\
=&2 \left\| \left((\widehat{y}_t^{(1)} - \bar{\widehat{y}}_t)^2 w_1, \ldots, (\widehat{y}_t^{(N)} - \bar{\widehat{y}}_t)^2 w_N \right)^\top \right\|_2 && (\partial(w_i^2)/\partial w_i = 2w_i) \\
\le&2 \max_i (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \|\boldsymbol{w}\|_2 && (\|(a_i w_i)_{i=1}^N\|_2 \le (\max_i |a_i|)\|\boldsymbol{w}\|_2) \\
\le&2(2H)^2 \|\boldsymbol{w}\|_2 && (|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le 2H) \\
\le&8H^2 && (\|\boldsymbol{w}\|_2 \le \|\boldsymbol{w}\|_1 = 1).
\end{aligned}
$$

a) If $f(x) = x$:

By mean-value inequality: $|g(\boldsymbol{w}) - g(\boldsymbol{u})| \le \sup_{\boldsymbol{v} \in [\boldsymbol{w}, \boldsymbol{u}]} \|\nabla g(\boldsymbol{v})\|_2 \|\boldsymbol{w} - \boldsymbol{u}\|_2$; here $g(\boldsymbol{w}) = \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2$, then

$$
\begin{aligned}
&\left| \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 - \sum_{i=1}^N u_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right| \le \sup_{\boldsymbol{v} \in \Delta_N} \left\| \nabla_{\boldsymbol{v}} \sum_{i=1}^N v_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right\|_2 \|\boldsymbol{w} - \boldsymbol{u}\|_2 \\
&\le 8H^2 \|\boldsymbol{w} - \boldsymbol{u}\|_2.
\end{aligned}
$$

Applying the mean-value bound with $(\boldsymbol{w}, \boldsymbol{u}) = (\widetilde{\boldsymbol{s}}_t, \hat{\boldsymbol{w}}_t)$ gives

$$
\begin{aligned}
\lambda \|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 &\le \left| \sum_{i=1}^N \widetilde{s}_{t,i}^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 - \sum_{i=1}^N (\hat{w}_{t,i})^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right| + 2\sqrt{2}\kappa \|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2 \\
&\le (8H^2 + 2\sqrt{2}\kappa) \|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2.
\end{aligned}
$$

Thus,

$$
\begin{aligned}
&\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 \le (8H^2 + 2\sqrt{2}\,\kappa) \|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2 / \lambda \\
\Rightarrow &\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \le (8H^2 + 2\sqrt{2}\,\kappa)/\lambda. \qquad (7)
\end{aligned}
$$

If $\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 > 0$, dividing both sides by this norm gives the desired bound. The bound holds trivially when the norm is zero.

b) If $f(x) = \log(\sigma_t^2 + x)$ under Assumption 1(ii):

Assumption 1(ii) gives $\sigma_t^2 + \sum_i w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \ge \underline{\sigma}^2 > 0$. By the chain rule, $\nabla \log g(\boldsymbol{w}) = \nabla g(\boldsymbol{w})/g(\boldsymbol{w})$.

$$
\begin{aligned}
&\left\| \nabla_{\boldsymbol{w}} \log \left\{ \sigma_t^2 + \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right\} \right\|_2 = \frac{\left\| \nabla_{\boldsymbol{w}} \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right\|_2}{\sigma_t^2 + \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2} \\
&\le \frac{8H^2}{\sigma_t^2 + \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2} \le \frac{8H^2}{\underline{\sigma}^2}.
\end{aligned}
$$

Then, by mean-value inequality, for every $\boldsymbol{w}, \boldsymbol{u} \in \Delta_N$,

$$
\left| \log \left\{ \sigma_t^2 + \sum_{i=1}^N w_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right\} - \log \left\{ \sigma_t^2 + \sum_{i=1}^N u_i^2 (\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \right\} \right| \le \frac{8H^2}{\underline{\sigma}^2} \|\boldsymbol{w} - \boldsymbol{u}\|_2.
$$

Applying the preceding mean-value bound with $(\boldsymbol{w}, \boldsymbol{u}) = (\widetilde{\boldsymbol{s}}_t, \hat{\boldsymbol{w}}_t)$ gives

$$\begin{aligned}\lambda\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 &\le \left|\log\left\{\sigma_t^2 + \sum_{i=1}^N \widetilde{s}_{t,i}^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right\} - \log\left\{\sigma_t^2 + \sum_{i=1}^N (\hat{w}_{t,i})^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right\}\right| \\ &\quad + 2\sqrt{2}\kappa\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2 \\ &\le \left(\frac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\kappa\right)\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2.\end{aligned}$$

Thus,

$$\begin{aligned}\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 &\le (\frac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa)\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2/\lambda \\ \Rightarrow\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 &\le (\frac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa)/\lambda \qquad (8)\end{aligned}$$

If $\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 > 0$, dividing both sides by this norm gives the desired bound. The bound holds trivially when the norm is zero.

c) If $f(x) = \log(x)$ under Assumption 1(iii):

Define $m_t = \min_{1\le i\le N} |\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t|$. By Assumption 2(i), $|\bar{\widehat{y}}_t - y_t| = |N^{-1}\sum_{j=1}^N(\widehat{y}_t^{(j)} - y_t)| \le N^{-1}\sum_{j=1}^N |\widehat{y}_t^{(j)} - y_t| \le H$. Consequently, for every i, $|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| = |(\widehat{y}_t^{(i)} - y_t) - (\bar{\widehat{y}}_t - y_t)| \le |\widehat{y}_t^{(i)} - y_t| + |\bar{\widehat{y}}_t - y_t| \le 2H$. For every $\boldsymbol{w} \in \Delta_N$,

$$\begin{aligned}\sum_{i=1}^N w_i^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 &\ge m_t^2\sum_{i=1}^N w_i^2 && (|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \ge m_t \text{ for every } i) \\ &\ge \frac{m_t^2}{N} && (\text{Cauchy–Schwarz:} (\textstyle\sum_i w_i)^2 \le N\sum_i w_i^2, \text{ with } \sum_i w_i = 1), \\ \sum_{i=1}^N w_i^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 &\le 4H^2\sum_{i=1}^N w_i^2 && (|\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t| \le 2H \text{ for every } i) \\ &\le 4H^2 && (\textstyle\sum_i w_i^2 \le (\sum_i w_i)^2 = 1 \text{ for } w_i \ge 0).\end{aligned}$$

Thus, $m_t^2/N \le \sum_{i=1}^N w_i^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2 \le 4H^2$. By Lemma 1, $m_t > 0$ almost surely. For $\boldsymbol{w}, \boldsymbol{u} \in \Delta_N$, the above bounds give

$$\left|\log\left\{\sum_{i=1}^N w_i^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right\} - \log\left\{\sum_{i=1}^N u_i^2(\widehat{y}_t^{(i)} - \bar{\widehat{y}}_t)^2\right\}\right| \le \log\left(\frac{4H^2}{m_t^2/N}\right) = \log\left(\frac{4NH^2}{m_t^2}\right).$$

Applying the preceding derived bound with $(\boldsymbol{w}, \boldsymbol{u}) = (\widetilde{\boldsymbol{s}}_t, \hat{\boldsymbol{w}}_t)$ gives

$$\lambda\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2^2 \le \log\left(\frac{4NH^2}{m_t^2}\right) + 2\sqrt{2}\,\kappa\|\widetilde{\boldsymbol{s}}_t - \hat{\boldsymbol{w}}_t\|_2.$$

Define $r = \|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \ge 0$. Then $\lambda r^2 - 2\sqrt{2}\,\kappa r - \log(4NH^2/m_t^2) \le 0$. Solving this quadratic inequality gives

$$\begin{aligned}r &\le \frac{2\sqrt{2}\,\kappa + \sqrt{8\kappa^2 + 4\lambda\log(4NH^2/m_t^2)}}{2\lambda} \\ &\le \frac{2\sqrt{2}\,\kappa + 2\sqrt{2}\,\kappa + 2\sqrt{\lambda\log(4NH^2/m_t^2)}}{2\lambda} && (\sqrt{8\kappa^2 + 4\lambda L} \le 2\sqrt{2}\,\kappa + 2\sqrt{\lambda L},\ L = \log(4NH^2/m_t^2)) \\ &= \frac{2\sqrt{2}\,\kappa}{\lambda} + \sqrt{\frac{1}{\lambda}\log\left(\frac{4NH^2}{m_t^2}\right)}.\end{aligned}$$

Thus, we get

$$\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \le \frac{2\sqrt{2}\,\kappa}{\lambda} + \sqrt{\frac{1}{\lambda}\log\left(\frac{4NH^2}{m_t^2}\right)}. \qquad (9)$$

□

Lemma 3. *Under Assumption 2, let $\boldsymbol{s}^\star \in \arg\min_{\boldsymbol{u}\in\Delta_N} \boldsymbol{u}^\top \boldsymbol{M}\boldsymbol{u}$ be a fixed vector of population prior weights. Then*

$$0 \le \mathbb{E}\left[\sum_{t=1}^{T} \ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^{T} \ell_t(\boldsymbol{u})\right] \le 4H^2\sqrt{2T\log(2N^2)}.$$

Consequently, for fixed N, the expected difference between the cumulative loss of the population prior weights and that of the best fixed convex weights selected in hindsight is $O(\sqrt{T})$.

Proof. Let $\widehat{\boldsymbol{M}}_T = (\widehat{M}_{T,ij})_{i,j=1}^N$, where $\widehat{M}_{T,ij} = T^{-1}\sum_{t=1}^T (\widehat{y}_t^{(i)} - y_t)(\widehat{y}_t^{(j)} - y_t)$. For any $\boldsymbol{u} \in \Delta_N$, we have

$$\begin{aligned}
\sum_{t=1}^{T} \ell_t(\boldsymbol{u}) &= \sum_{t=1}^{T} \left(\sum_{i=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)\right)^2 \\
&= \sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{j=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)(\widehat{y}_t^{(j)} - y_t)u_j \\
&= T\sum_{i=1}^{N}\sum_{j=1}^{N} u_i \widehat{M}_{T,ij} u_j = T\boldsymbol{u}^\top \widehat{\boldsymbol{M}}_T \boldsymbol{u}.
\end{aligned}$$

Then,

$$\mathbb{E}\left[\sum_{t=1}^{T} \ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^{T} \ell_t(\boldsymbol{u})\right] = T\,\mathbb{E}\left[\boldsymbol{s}^{\star\top}\widehat{\boldsymbol{M}}_T\boldsymbol{s}^\star - \min_{\boldsymbol{u}\in\Delta_N} \boldsymbol{u}^\top \widehat{\boldsymbol{M}}_T \boldsymbol{u}\right].$$

Define $\delta_T = \|\widehat{\boldsymbol{M}}_T - \boldsymbol{M}\|_{\max} = \max_{1\le i,j\le N} |\widehat{M}_{T,ij} - M_{ij}|$. For any $\boldsymbol{u} \in \Delta_N$, we have:

$$\begin{aligned}
&\left|\boldsymbol{u}^\top\left(\widehat{\boldsymbol{M}}_T - \boldsymbol{M}\right)\boldsymbol{u}\right| = \left|\sum_{i=1}^{N}\sum_{j=1}^{N} u_i\left(\widehat{M}_{T,ij} - M_{ij}\right)u_j\right| \\
&\le \sum_{i=1}^{N}\sum_{j=1}^{N} |u_i||u_j|\left|\widehat{M}_{T,ij} - M_{ij}\right| \quad (\,|\textstyle\sum_{i,j} a_{ij}| \le \sum_{i,j}|a_{ij}|, \text{ and } |abc| = |a||b||c|) \\
&\le \delta_T \sum_{i=1}^{N}\sum_{j=1}^{N} |u_i||u_j| = \delta_T \|\boldsymbol{u}\|_1^2 = \delta_T.
\end{aligned}$$

So, $\boldsymbol{u}^\top\boldsymbol{M}\boldsymbol{u} - \delta_T \le \boldsymbol{u}^\top\widehat{\boldsymbol{M}}_T\boldsymbol{u} \le \boldsymbol{u}^\top\boldsymbol{M}\boldsymbol{u} + \delta_T$ for any $\boldsymbol{u} \in \Delta_N$. Consequently, we have

$$\boldsymbol{s}^{\star\top}\widehat{\boldsymbol{M}}_T\boldsymbol{s}^\star \le \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star + \delta_T, \qquad \min_{\boldsymbol{u}\in\Delta_N} \boldsymbol{u}^\top\widehat{\boldsymbol{M}}_T\boldsymbol{u} \ge \min_{\boldsymbol{u}\in\Delta_N} \boldsymbol{u}^\top\boldsymbol{M}\boldsymbol{u} - \delta_T = \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star - \delta_T.$$

Subtracting the two inequalities gives

$$\boldsymbol{s}^{\star\top}\widehat{\boldsymbol{M}}_T\boldsymbol{s}^\star - \min_{\boldsymbol{u}\in\Delta_N} \boldsymbol{u}^\top\widehat{\boldsymbol{M}}_T\boldsymbol{u} \le \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star + \delta_T - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star + \delta_T = 2\delta_T.$$

Then,

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{T} \ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^{T} \ell_t(\boldsymbol{u})\right] = T\,\mathbb{E}\left[\boldsymbol{s}^{\star\top}\widehat{\boldsymbol{M}}_T\boldsymbol{s}^\star - \min_{\boldsymbol{u}\in\Delta_N} \boldsymbol{u}^\top\widehat{\boldsymbol{M}}_T\boldsymbol{u}\right] \\
&\le 2T\,\mathbb{E}[\delta_T] = 2T\,\mathbb{E}\left\|\widehat{\boldsymbol{M}}_T - \boldsymbol{M}\right\|_{\max}.
\end{aligned}$$

For each pair (i,j), we define $Z_{t,ij} = (\widehat{y}_t^{(i)} - y_t)(\widehat{y}_t^{(j)} - y_t) - M_{ij}$. By the information structure above, $Z_{t,ij}$ is $\mathcal{F}_t$-measurable.

Since $|\widehat{y}_t^{(i)} - y_t| \le H$ for all i, then $|Z_{t,ij}| = |(\widehat{y}_t^{(i)} - y_t)(\widehat{y}_t^{(j)} - y_t) - M_{ij}| \le |(\widehat{y}_t^{(i)} - y_t)(\widehat{y}_t^{(j)} - y_t)| + |M_{ij}| \le H^2 + H^2 = 2H^2$.

By Assumption 2(ii), $\{Z_{t,ij}\}_{t=1}^{T}$ is a martingale-difference sequence, and $|Z_{t,ij}| \le 2H^2$. Applying Hoeffding's lemma (Hoeffding 1963, Eq. (4.16)) conditionally gives, for every $\eta \in \mathbb{R}$,

$$\mathbb{E}[\exp(\eta Z_{t,ij}) \mid \mathcal{F}_{t-1}] \le \exp(2\eta^2 H^4).$$

Iterating conditional expectations over t gives

$$\mathbb{E}\exp\left(\eta \sum_{t=1}^{T} Z_{t,ij}\right) \le \exp(2\eta^2 H^4 T).$$

Similarly, for $-Z_{t,ij}$, we have

$$\mathbb{E}\exp\left(-\eta \sum_{t=1}^{T} Z_{t,ij}\right) \le \exp(2\eta^2 H^4 T).$$

Applying the log-sum-exp bound $\max_{1\le r\le m} x_r \le \eta^{-1}\log(\sum_{r=1}^{m} e^{\eta x_r})$ (Boyd & Vandenberghe 2004, p. 72), for any $\eta > 0$, to the $2N^2$ signed sums gives

$$\begin{aligned}
\mathbb{E}\max_{i,j}\left|\sum_{t=1}^{T} Z_{t,ij}\right| &\le \frac{1}{\eta}\,\mathbb{E}\log\left[\sum_{i,j=1}^{N}\left\{\exp\left(\eta\sum_{t=1}^{T} Z_{t,ij}\right) + \exp\left(-\eta\sum_{t=1}^{T} Z_{t,ij}\right)\right\}\right] \\
&\le \frac{1}{\eta}\,\log\left[\sum_{i,j=1}^{N}\left\{\mathbb{E}\left[\exp\left(\eta\sum_{t=1}^{T} Z_{t,ij}\right)\right] + \mathbb{E}\left[\exp\left(-\eta\sum_{t=1}^{T} Z_{t,ij}\right)\right]\right\}\right] \\
&\qquad \text{(Jensen's inequality: } \mathbb{E}\log X \le \log \mathbb{E}X \text{ for } X > 0\text{)} \\
&\le \frac{1}{\eta}\log\left[2N^2\exp(2\eta^2 H^4 T)\right] = \frac{\log(2N^2)}{\eta} + 2\eta H^4 T.
\end{aligned}$$

Let $\eta = \sqrt{\log(2N^2)/(2H^4T)}$. We have $\mathbb{E}\max_{i,j}|\sum_{t=1}^{T} Z_{t,ij}| \le 2H^2\sqrt{2T\log(2N^2)}$.

Since $\widehat{M}_{T,ij} - M_{ij} = T^{-1}\sum_t Z_{t,ij}$, we have

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right] &\le 2T\,\mathbb{E}\left\|\widehat{\boldsymbol{M}}_T - \boldsymbol{M}\right\|_{\max} \\
&\le 2\mathbb{E}\max_{i,j}\left|\sum_{t=1}^{T} Z_{t,ij}\right| \le 4H^2\sqrt{2T\log(2N^2)}.
\end{aligned}$$

□

Proof of Theorem 1. By Assumption 2(i), $\max_i|\widehat{y}_t^{(i)} - y_t| \le H$. Hence, for $\boldsymbol{w}, \boldsymbol{u} \in \Delta_N$:

$$\begin{aligned}
|\ell_t(\boldsymbol{w}) - \ell_t(\boldsymbol{u})| &= \left|\left(\sum_{i=1}^{N} w_i(\widehat{y}_t^{(i)} - y_t)\right)^2 - \left(\sum_{i=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)\right)^2\right| \\
&= \left|\sum_{i=1}^{N} w_i(\widehat{y}_t^{(i)} - y_t) - \sum_{i=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)\right| \times \left|\sum_{i=1}^{N} w_i(\widehat{y}_t^{(i)} - y_t) + \sum_{i=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)\right| \\
&\le \left|\sum_{i=1}^{N} w_i(\widehat{y}_t^{(i)} - y_t) - \sum_{i=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)\right| \times \left(\left|\sum_{i=1}^{N} w_i(\widehat{y}_t^{(i)} - y_t)\right| + \left|\sum_{i=1}^{N} u_i(\widehat{y}_t^{(i)} - y_t)\right|\right) \\
&\qquad \text{(triangle inequality: } |a+b| \le |a| + |b|\text{)} \\
&\le 2H\left|\sum_{i=1}^{N}\left[w_i(\widehat{y}_t^{(i)} - y_t) - u_i(\widehat{y}_t^{(i)} - y_t)\right]\right| \qquad (\max_i|\widehat{y}_t^{(i)} - y_t| \le H) \\
&\le 2H\left|\sum_{i=1}^{N}(\widehat{y}_t^{(i)} - y_t)(w_i - u_i)\right| \\
&\le 2H^2\|\boldsymbol{w} - \boldsymbol{u}\|_1 \qquad (\text{Hölder: } |\boldsymbol{x}^\top\boldsymbol{y}| \le \|\boldsymbol{x}\|_\infty\|\boldsymbol{y}\|_1) \\
&\le 2H^2\sqrt{N}\|\boldsymbol{w} - \boldsymbol{u}\|_2 \qquad (\|\boldsymbol{x}\|_1 \le \sqrt{N}\|\boldsymbol{x}\|_2).
\end{aligned}$$

Applying this inequality with $(\boldsymbol{w}, \boldsymbol{u}) = (\hat{\boldsymbol{w}}_t, \widetilde{\boldsymbol{s}}_t)$ gives

$$|\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)| \le 2H^2\sqrt{N}\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2. \tag{10}$$

a) For $f(x) = x$:

$$\begin{aligned}
\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u}) &= \{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)\} + \{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u})\} \\
&\le \ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u}) + 2H^2\sqrt{N}\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 && (|\ell_t(\boldsymbol{w}) - \ell_t(\boldsymbol{u})| \le 2H^2\sqrt{N}\|\boldsymbol{w} - \boldsymbol{u}\|_2) \\
&\le \ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u}) + \frac{2H^2\sqrt{N}}{\lambda}(8H^2 + 2\sqrt{2}\,\kappa) && \text{(by (7)).}
\end{aligned}$$

Summing this inequality gives

$$\sum_{t=1}^{T}\{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u})\} \le \sum_{t=1}^{T}\{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u})\} + \frac{2H^2\sqrt{N}\,T}{\lambda}(8H^2 + 2\sqrt{2}\,\kappa).$$

b) For $f(x) = \log(\sigma_t^2 + x)$ under Assumption 1(ii):

Lemma 2 gives

$$\begin{aligned}
\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u}) &\le \ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u}) + 2H^2\sqrt{N}\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \\
&\le \ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u}) + \frac{2H^2\sqrt{N}}{\lambda}\left(\frac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa\right) && \text{(by (8)).}
\end{aligned}$$

Summing gives

$$\sum_{t=1}^{T}\{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u})\} \le \sum_{t=1}^{T}\{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u})\} + \frac{2H^2\sqrt{N}\,T}{\lambda}\left(\frac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa\right).$$

c) For $f(x) = \log x$ under Assumption 1(iii):

By Lemma 2, we have

$$\begin{aligned}
\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u}) &\le \ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u}) + 2H^2\sqrt{N}\|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \\
&\le \ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u}) + 2H^2\sqrt{N}\left\{\frac{2\sqrt{2}\,\kappa}{\lambda} + \sqrt{\frac{1}{\lambda}\log\left(\frac{4NH^2}{m_t^2}\right)}\right\} && \text{(by (9)).}
\end{aligned}$$

Summing gives

$$\sum_{t=1}^{T}\{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u})\} \le \sum_{t=1}^{T}\{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u})\} + 2H^2\sqrt{N}\sum_{t=1}^{T}\left\{\frac{2\sqrt{2}\,\kappa}{\lambda} + \sqrt{\frac{1}{\lambda}\log\left(\frac{4NH^2}{m_t^2}\right)}\right\}.$$

Combining the three cases, for any $\boldsymbol{u} \in \Delta_N$, we have

$$\begin{aligned}
&\sum_{t=1}^{T}\{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\boldsymbol{u})\} \le \sum_{t=1}^{T}\{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{u})\} + B_T, \\
\Rightarrow &\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t) - \sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star) \le \left\{\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t) - \sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star)\right\} + B_T; \\
\text{and, } &\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u}) \le \left\{\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right\} + B_T.
\end{aligned}$$

where

$$B_T := \begin{cases}
\dfrac{2H^2\sqrt{N}\,T}{\lambda}(8H^2 + 2\sqrt{2}\,\kappa), & f(x) = x, \\
\dfrac{2H^2\sqrt{N}\,T}{\lambda}\left(\dfrac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa\right), & f(x) = \log(\sigma_t^2 + x), \\
2H^2\sqrt{N}\displaystyle\sum_{t=1}^{T}\left\{\dfrac{2\sqrt{2}\,\kappa}{\lambda} + \sqrt{\dfrac{1}{\lambda}\log\left(\dfrac{4NH^2}{m_t^2}\right)}\right\}, & f(x) = \log x.
\end{cases}$$

Taking expectations on both sides and applying Lemma 1 gives

$$\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t)-\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star)\right]\le\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t)-\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star)\right]+\mathbb{E}[B_T],$$
$$\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t)-\min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right]\le\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t)-\min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right]+\mathbb{E}[B_T].$$

where

$$\mathbb{E}[B_T]\le A_T:=\begin{cases}\dfrac{2H^2\sqrt{N}\,T}{\lambda}(8H^2+2\sqrt{2}\,\kappa), & f(x)=x,\\[2ex] \dfrac{2H^2\sqrt{N}\,T}{\lambda}\left(\dfrac{8H^2}{\underline{\sigma}^2}+2\sqrt{2}\,\kappa\right), & f(x)=\log(\sigma_t^2+x),\ .\\[2ex] 2H^2\sqrt{N}\,T\left\{\dfrac{2\sqrt{2}\,\kappa}{\lambda}+\dfrac{C_2}{\sqrt{\lambda}}\right\}, & f(x)=\log x.\end{cases}$$

We next bound the expected loss of the projected prior weights relative to $\boldsymbol{s}^\star$. For any $\mathcal{F}_{t-1}$-measurable vector $\boldsymbol{x}\in\Delta_N$, by Assumption 2(ii),

$$\mathbb{E}\left[\left\{\sum_{i=1}^{N}x_i(\widehat{y}_t^{(i)}-y_t)\right\}^2\middle|\mathcal{F}_{t-1}\right]=\sum_{i=1}^{N}\sum_{j=1}^{N}x_ix_j\mathbb{E}\left[(\widehat{y}_t^{(i)}-y_t)(\widehat{y}_t^{(j)}-y_t)\middle|\mathcal{F}_{t-1}\right]$$
$$=\sum_{i=1}^{N}\sum_{j=1}^{N}x_iM_{ij}x_j=\boldsymbol{x}^\top\boldsymbol{M}\boldsymbol{x}.$$

Since the prior weights $\boldsymbol{s}_t$ are $\mathcal{F}_{t-1}$-measurable and Π_{Δ_N} is a deterministic continuous function, the projected prior weights $\widetilde{\boldsymbol{s}}_t=\Pi_{\Delta_N}(\boldsymbol{s}_t)$ are also $\mathcal{F}_{t-1}$-measurable. The population prior weights $\boldsymbol{s}^\star$ are deterministic. Therefore, set $\boldsymbol{x}=\widetilde{\boldsymbol{s}}_t$ and $\boldsymbol{x}=\boldsymbol{s}^\star$, respectively. Recalling that

$$\ell_t(\boldsymbol{w})=\left(y_t-\sum_{i=1}^{N}w_i\widehat{y}_t^{(i)}\right)^2=\left\{\sum_{i=1}^{N}w_i(\widehat{y}_t^{(i)}-y_t)\right\}^2,$$

we obtain

$$\mathbb{E}\left[\ell_t(\widetilde{\boldsymbol{s}}_t)-\ell_t(\boldsymbol{s}^\star)|\mathcal{F}_{t-1}\right]=\widetilde{\boldsymbol{s}}_t^\top\boldsymbol{M}\widetilde{\boldsymbol{s}}_t-\boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star.\tag{11}$$

The right side of (11) is nonnegative because $\widetilde{\boldsymbol{s}}_t\in\Delta_N$ and $\boldsymbol{s}^\star\in\arg\min_{\boldsymbol{w}\in\Delta_N}\boldsymbol{w}^\top\boldsymbol{M}\boldsymbol{w}$. Taking expectations and summing gives

$$\mathbb{E}\sum_{t=1}^{T}\{\ell_t(\widetilde{\boldsymbol{s}}_t)-\ell_t(\boldsymbol{s}^\star)\}=\sum_{t=1}^{T}\mathbb{E}\Big[\mathbb{E}\{\ell_t(\widetilde{\boldsymbol{s}}_t)-\ell_t(\boldsymbol{s}^\star)\mid\mathcal{F}_{t-1}\}\Big]=\sum_{t=1}^{T}\mathbb{E}[\widetilde{\boldsymbol{s}}_t^\top\boldsymbol{M}\widetilde{\boldsymbol{s}}_t-\boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star].$$

When comparing the expected actual squared-loss regret relative to the population prior weights $\boldsymbol{s}^\star$, we have

$$\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t)-\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star)\right]\le\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t)-\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star)\right]+\mathbb{E}[B_T]$$
$$\le\sum_{t=1}^{T}\mathbb{E}[\widetilde{\boldsymbol{s}}_t^\top\boldsymbol{M}\widetilde{\boldsymbol{s}}_t-\boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star]+A_T.$$

We next compare the cumulative forecast loss of online REF with that of the oracle ensemble, defined as the best fixed convex combination of expert forecasts selected in hindsight after all T

outcomes are observed. By Lemma 3, we have

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right] \le \mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right] + \mathbb{E}[B_T] \\
&= \mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\widetilde{\boldsymbol{s}}_t) - \sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star)\right] + \mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right] + \mathbb{E}[B_T] \\
&\le \sum_{t=1}^{T}\mathbb{E}[\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star] + A_T + \mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right] \\
&\le \sum_{t=1}^{T}\mathbb{E}[\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star] + 4H^2\sqrt{2T\log(2N^2)} + A_T.
\end{aligned}$$

It remains to bound $\sum_{t=1}^{T}\mathbb{E}[\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star]$. We have:

$$\begin{aligned}
&\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star = \left(\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\boldsymbol{s}^\star\right) + \left(\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\boldsymbol{s}^\star - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star\right) \\
=&\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right) + \left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\boldsymbol{s}^\star \\
=&\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t + \left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\boldsymbol{s}^\star \quad \text{(Since } \boldsymbol{M} \text{ is symmetric)} \\
=&\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\left\{\boldsymbol{s}^\star + \left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)\right\} + \left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\boldsymbol{s}^\star \\
=&2\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\boldsymbol{s}^\star + \left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right).
\end{aligned}$$

By Cauchy-Schwarz inequality and Rayleigh quotient bound (Horn & Johnson 2012, Section 4.2):

$$\begin{aligned}
2\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\boldsymbol{s}^\star &\le 2\|\boldsymbol{M}\boldsymbol{s}^\star\|_2\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2, \\
\text{and, } \left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right)^\top \boldsymbol{M}\left(\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\right) &\le \mu_{\max}(\boldsymbol{M})\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2,
\end{aligned}$$

where $\mu_{\max}$ is the largest eigenvalue of $\boldsymbol{M}$.

Because the population prior weights satisfy $\boldsymbol{s}^\star \in \Delta_N$, for $t = 1, \ldots, T$, $\Pi_{\Delta_N}(\boldsymbol{s}^\star) = \boldsymbol{s}^\star$. Projection nonexpansiveness gives $\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2 = \|\Pi_{\Delta_N}(\boldsymbol{s}_t) - \Pi_{\Delta_N}(\boldsymbol{s}^\star)\|_2 \le \|\boldsymbol{s}_t - \boldsymbol{s}^\star\|_2$.

By Assumption 4, $\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 \le C_1/t$, and Jensen's inequality gives $\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2 \le \sqrt{\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2} \le \sqrt{C_1/t}$.

Because $|M_{ij}| \le H^2$ and $\boldsymbol{M}$ is positive semidefinite, then we have

$$\mu_{\max}(\boldsymbol{M}) = \|\boldsymbol{M}\|_2 \le \|\boldsymbol{M}\|_F = \sqrt{\sum_{i=1}^{N}\sum_{j=1}^{N}M_{ij}^2} \le NH^2.$$

Also, since $\|\boldsymbol{s}^\star\|_2 \le 1$, $\|\boldsymbol{M}\boldsymbol{s}^\star\|_2 \le \|\boldsymbol{M}\|_2\|\boldsymbol{s}^\star\|_2 \le NH^2$.

We also use $\sum_{t=1}^{T} t^{-1/2} \le 2\sqrt{T}$ and $\sum_{t=1}^{T} t^{-1} \le 1 + \log T$.

Then,

$$\begin{aligned}
&\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star \le 2\|\boldsymbol{M}\boldsymbol{s}^\star\|_2\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2 + \mu_{\max}(\boldsymbol{M})\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 \\
&\le 2NH^2\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2 + NH^2\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 \\
&\Rightarrow \mathbb{E}\left[\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star\right] \le 2NH^2\,\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2 + NH^2\,\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 \\
&\le 2NH^2\sqrt{\frac{C_1}{t}} + NH^2\frac{C_1}{t} \\
&\Rightarrow \sum_{t=1}^{T}\mathbb{E}\left[\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star\right] \le 4NH^2\sqrt{C_1 T} + NH^2C_1(1 + \log T).
\end{aligned}$$

Therefore

$$\sum_{t=1}^{T}\mathbb{E}\left[\widetilde{\boldsymbol{s}}_t^\top \boldsymbol{M}\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^{\star\top}\boldsymbol{M}\boldsymbol{s}^\star\right] = \mathcal{O}(\sqrt{T}). \tag{12}$$

By Assumption 3, $\sup_T \kappa < \infty$, take $\lambda \asymp \sqrt{T}$ when $f(x) = x$ and $f(x) = \log(\sigma_t^2 + x)$, and $\lambda \asymp T$ when $f(x) = \log(x)$. We have

$$A_T = \begin{cases} \dfrac{2H^2\sqrt{N}\,T}{\lambda}(8H^2 + 2\sqrt{2}\,\kappa) = \mathcal{O}(\sqrt{T}), & f(x) = x, \\ \dfrac{2H^2\sqrt{N}\,T}{\lambda}\left(\dfrac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa\right) = \mathcal{O}(\sqrt{T}), & f(x) = \log(\sigma_t^2 + x), \\ 2H^2\sqrt{N}\,T\left\{\dfrac{2\sqrt{2}\,\kappa}{\lambda} + \dfrac{C_2}{\sqrt{\lambda}}\right\} = \mathcal{O}(\sqrt{T}), & f(x) = \log x. \end{cases} \tag{13}$$

Based on (12) and (13), we have $\mathbb{E}[\sum_{t=1}^T \ell_t(\hat{\boldsymbol{w}}_t) - \sum_{t=1}^T \ell_t(\boldsymbol{s}^\star)] \le C_3\sqrt{T}$. Combining this bound with Lemma 3 gives

$$\mathbb{E}\left[\sum_{t=1}^T \ell_t(\hat{\boldsymbol{w}}_t) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^T \ell_t(\boldsymbol{u})\right] \le C_4\sqrt{T}.$$

It remains to obtain the corresponding lower bound. Based on (10) and Lemma 2,

$$\sum_{t=1}^T |\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)| \le 2H^2\sqrt{N}\sum_{t=1}^T \|\hat{\boldsymbol{w}}_t - \widetilde{\boldsymbol{s}}_t\|_2 \le B_T,$$

where

$$B_T = \begin{cases} \dfrac{2H^2\sqrt{N}\,T}{\lambda}(8H^2 + 2\sqrt{2}\,\kappa), & f(x) = x, \\ \dfrac{2H^2\sqrt{N}\,T}{\lambda}\left(\dfrac{8H^2}{\underline{\sigma}^2} + 2\sqrt{2}\,\kappa\right), & f(x) = \log(\sigma_t^2 + x), \\ 2H^2\sqrt{N}\displaystyle\sum_{t=1}^T \left\{\dfrac{2\sqrt{2}\,\kappa}{\lambda} + \sqrt{\dfrac{1}{\lambda}\log\left(\dfrac{4NH^2}{m_t^2}\right)}\right\}, & f(x) = \log x. \end{cases}.$$

Consequently,

$$\mathbb{E}\sum_{t=1}^T \{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)\} \ge -\mathbb{E}\sum_{t=1}^T |\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)| \ge -\mathbb{E}[B_T] \ge -A_T.$$

Moreover, $\mathbb{E}\sum_{t=1}^T \{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{s}^\star)\} \ge 0$ by the optimality of the population prior weights $\boldsymbol{s}^\star$, and

$$\mathbb{E}\left[\sum_{t=1}^T \ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^T \ell_t(\boldsymbol{u})\right] \ge 0$$

by definition. For the fixed population prior weights $\boldsymbol{s}^\star$,

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^T \ell_t(\hat{\boldsymbol{w}}_t) - \sum_{t=1}^T \ell_t(\boldsymbol{s}^\star)\right] = \mathbb{E}\sum_{t=1}^T \{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)\} + \mathbb{E}\sum_{t=1}^T \{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{s}^\star)\} \\ &\ge -A_T \ge -C_5\sqrt{T}. \end{aligned}$$

Similarly, for the oracle ensemble, we have

$$\begin{aligned} &\mathbb{E}\left[\sum_{t=1}^T \ell_t(\hat{\boldsymbol{w}}_t) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^T \ell_t(\boldsymbol{u})\right] \\ &= \mathbb{E}\sum_{t=1}^T \{\ell_t(\hat{\boldsymbol{w}}_t) - \ell_t(\widetilde{\boldsymbol{s}}_t)\} + \mathbb{E}\sum_{t=1}^T \{\ell_t(\widetilde{\boldsymbol{s}}_t) - \ell_t(\boldsymbol{s}^\star)\} + \mathbb{E}\left[\sum_{t=1}^T \ell_t(\boldsymbol{s}^\star) - \min_{\boldsymbol{u}\in\Delta_N} \sum_{t=1}^T \ell_t(\boldsymbol{u})\right] \\ &\ge -A_T \ge -C_6\sqrt{T}. \end{aligned}$$

Here C_5 and C_6 are finite constants independent of T. Thus the expected hindsight regret lies between $-C_6\sqrt{T}$ and $C_4\sqrt{T}$. Therefore, relative to the oracle ensemble, which uses the best fixed convex combination of expert forecasts selected in hindsight after all T outcomes are observed,

$$\mathbb{E}\left[\sum_{t=1}^{T}\ell_t(\hat{\boldsymbol{w}}_t) - \min_{\boldsymbol{u}\in\Delta_N}\sum_{t=1}^{T}\ell_t(\boldsymbol{u})\right] = \mathcal{O}(\sqrt{T}).$$

Dividing by T gives (4) and completes the proof. □

D An Example: CCR Prior Weights

This section provides an example in which CCR prior weights constructed from a rolling window satisfy Assumption 4 under suitable conditions. Under the remaining assumptions of Theorem 1, the online REF policy using these CCR prior weights satisfies the same regret bound.

For an integer window length $L_T \geq 2$, the estimates for period t use the forecast errors from periods $t - L_T, \ldots, t-1$. We write the CCR prior weights in (2) in matrix form. For positive scales $\boldsymbol{v} = (v_1, \ldots, v_N)^\top$ and $-1/(N-1) < \rho < 1$, define

$$\boldsymbol{M}(\boldsymbol{v}, \rho) = \boldsymbol{D}(\boldsymbol{v})\big[(1-\rho)\boldsymbol{I}_N + \rho\mathbf{1}\mathbf{1}^\top\big]\boldsymbol{D}(\boldsymbol{v}), \qquad \boldsymbol{D}(\boldsymbol{v}) = \mathrm{diag}(v_1, \ldots, v_N).$$

This matrix has diagonal entries v_i^2 and off-diagonal entries $\rho v_i v_j$, and is positive definite under the stated conditions. The CCR prior weights are

$$\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}, \rho) = \frac{\boldsymbol{M}(\boldsymbol{v}, \rho)^{-1}\mathbf{1}}{\mathbf{1}^\top\boldsymbol{M}(\boldsymbol{v}, \rho)^{-1}\mathbf{1}}. \tag{14}$$

To verify the equivalence with (2), note that

$$\big[(1-\rho)\boldsymbol{I}_N + \rho\mathbf{1}\mathbf{1}^\top\big]^{-1} = \frac{1}{1-\rho}\left[\boldsymbol{I}_N - \frac{\rho}{1+(N-1)\rho}\mathbf{1}\mathbf{1}^\top\right].$$

Substituting this inverse into (14) and setting $(\boldsymbol{v}, \rho) = (\widehat{\boldsymbol{v}}_t, \widehat{\rho}_t)$ gives (2) with $\rho_t = \widehat{\rho}_t$. These prior weights sum to one, although individual components may be negative. We use the MAP estimator with the uninformative prior in Soule et al. (2024, Proposition 1). Their historical sample size n is L_T here, and their scale $\widehat{\sigma}_i$ is denoted by $\widehat{v}_{t,i}$. Define the historical second moments by

$$\overline{M}_{t,ij} = \frac{1}{L_T}\sum_{r=t-L_T}^{t-1}(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r).$$

The scale estimates and the MAP dependence estimate are

$$\widehat{v}_{t,i}^2 = \frac{1}{L_T - 1}\sum_{r=t-L_T}^{t-1}(\widehat{y}_r^{(i)} - y_r)^2 = \frac{L_T}{L_T-1}\overline{M}_{t,ii}, \qquad \widehat{v}_{t,i} = \sqrt{\widehat{v}_{t,i}^2}, \tag{15}$$

$$\bar{\rho}_t = \frac{2}{N(N-1)}\sum_{i<j}\frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}}, \qquad \widehat{\rho}_t = \frac{L_T - 1}{L_T + 2N + 3}\bar{\rho}_t. \tag{16}$$

Thus $\bar{\rho}_t$ averages the normalized error cross-products; no sample-mean centering is applied. The common factor $L_T/(L_T-1)$ in the sample second moments cancels from each normalized ratio. We analyze these explicit formulas under Assumption 2.

Whenever the displayed empirical quantities are well defined, define the CCR prior weights and their projections by

$$\boldsymbol{s}_t = \boldsymbol{s}^{\mathrm{CCR}}(\widehat{\boldsymbol{v}}_t, \widehat{\rho}_t), \qquad \widetilde{\boldsymbol{s}}_t = \Pi_{\Delta_N}(\boldsymbol{s}_t). \tag{17}$$

The components of $\boldsymbol{s}_t$ sum to one and may be negative. The projected prior weights $\widetilde{\boldsymbol{s}}_t$ are used as a comparison point in the proof.

Proposition 1. *Suppose that $\boldsymbol{M} = \boldsymbol{M}(\boldsymbol{v}^\star, \rho^\star)$ for some $\boldsymbol{v}^\star \in [\underline{v}, \overline{v}]^N$ and $\rho^\star \in \mathcal{R}$, where $0 < \underline{v} < \overline{v} < \infty$ and $\mathcal{R} = [\underline{\rho}, \overline{\rho}] \subset (-1/(N-1), 1)$, with all bounds independent of T. For each horizon T, let the rolling-window length $L_T \geq 2$ be a deterministic integer satisfying $L_T \asymp T$. The same window length is used for $t = 1, \ldots, T$. The forecast errors from periods $1 - L_T, \ldots, 0$ are available before period 1. For the historical periods, $\mathcal{F}_r$ denotes the information available through period r. For every $r = 1 - L_T, \ldots, T$, suppose Assumption 2 holds with the same H and $\boldsymbol{M}$, and each $\widehat{y}_r^{(i)} - y_r$ has a continuous distribution, $i = 1, \ldots, N$. Assume also that the population CCR prior weights are nonnegative, $[\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star, \rho^\star)]_i \geq 0$, $i = 1, \ldots, N$. Then, for the CCR prior weights in* (17) *and fixed N, there is a finite constant C_1, independent of t and T, such that*

$$\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 \leq \frac{C_1}{t}, \qquad 1 \leq t \leq T. \tag{18}$$

Under the common-correlation structure in Proposition 1, $M_{ii} = v_i^{\star 2}$ and $M_{ij} = \rho^\star v_i^\star v_j^\star$ for $i \neq j$. Thus, $v_i^\star = \sqrt{M_{ii}}$ measures the size of expert i's forecast errors, while $\rho^\star$ describes the common dependence among the forecast errors of different experts. When the forecast errors have zero conditional means, $\rho^\star$ is their common conditional correlation. The population CCR prior weights $\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star, \rho^\star)$ sum to one. The nonnegativity condition therefore ensures that $\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star, \rho^\star) \in \Delta_N$. As shown below, these weights equal the population prior weights $\boldsymbol{s}^\star$. The estimated prior weights $s_{t,i}$, computed from a finite rolling window, may still be negative.

Proof. We first bound the mean squared error of the projected CCR prior weights relative to the projected population CCR prior weights. We then show that the population CCR prior weights equal $\boldsymbol{s}^\star$.

We verify that $\widehat{v}_{t,i} > 0$ for every i and $-1/(N-1) < \widehat{\rho}_t < 1$ almost surely, so that the CCR prior weights are well defined. The continuity condition ensures that $\overline{M}_{t,ii} > 0$ almost surely. Indeed, for each i,

$$\begin{aligned}\Pr(\overline{M}_{t,ii} = 0) &= \Pr\Big(\widehat{y}_r^{(i)} - y_r = 0 \text{ for every } r = t - L_T, \ldots, t-1\Big) \\ &\leq \Pr\Big(\widehat{y}_{t-L_T}^{(i)} - y_{t-L_T} = 0\Big) = 0.\end{aligned}$$

Hence $\widehat{v}_{t,i}$ and $\widehat{\rho}_t$ are well defined almost surely. Moreover, the Cauchy–Schwarz inequality and a nonnegative squared sum give

$$\begin{aligned}&|\overline{M}_{t,ij}| \leq \sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}, \\ &0 \leq \frac{1}{L_T}\sum_{r=t-L_T}^{t-1}\left(\sum_{i=1}^{N}\frac{\widehat{y}_r^{(i)} - y_r}{\sqrt{\overline{M}_{t,ii}}}\right)^2 = N + N(N-1)\bar{\rho}_t.\end{aligned}$$

Consequently,

$$-\frac{1}{N-1} \leq \bar{\rho}_t \leq 1, \qquad -\frac{1}{N-1} < \widehat{\rho}_t < 1,$$

where the strict bounds follow from $0 < (L_T - 1)/(L_T + 2N + 3) < 1$. Thus $\boldsymbol{M}(\widehat{\boldsymbol{v}}_t, \widehat{\rho}_t)$ is positive definite, and the CCR prior weights $\boldsymbol{s}_t$ are well defined almost surely. Because these quantities use observations only through period $t-1$, they are $\mathcal{F}_{t-1}$-measurable.

For each pair (i, j), by Assumption 2(ii), we have $\mathbb{E}[(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij} \mid \mathcal{F}_{r-1}] = 0$. By Assumption 2(i),

$$\begin{aligned}|M_{ij}| &= \Big|\mathbb{E}\Big[(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) \mid \mathcal{F}_{r-1}\Big]\Big| \leq \mathbb{E}\Big[\Big|(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r)\Big| \mid \mathcal{F}_{r-1}\Big] \\ &= \mathbb{E}\Big[\Big|\widehat{y}_r^{(i)} - y_r\Big|\Big|\widehat{y}_r^{(j)} - y_r\Big| \mid \mathcal{F}_{r-1}\Big] \leq \mathbb{E}[H \cdot H \mid \mathcal{F}_{r-1}] = H^2.\end{aligned}$$

Then, $|(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}| \leq |(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r)| + |M_{ij}| = |\widehat{y}_r^{(i)} - y_r||\widehat{y}_r^{(j)} - y_r| + |M_{ij}| \leq H \cdot H + H^2 = 2H^2$. Thus, $(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}$ is a bounded martingale difference.

Expanding the square and including the cross terms gives

$$\mathbb{E}|\overline{M}_{t,ij} - M_{ij}|^2 = \frac{1}{L_T^2}\mathbb{E}\left[\left\{\sum_{r=t-L_T}^{t-1}\left[(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\right]\right\}^2\right]$$

$$= \frac{1}{L_T^2}\sum_{r=t-L_T}^{t-1}\mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\}^2\Big]$$

$$+ \frac{2}{L_T^2}\sum_{t-L_T\le r<s\le t-1}\mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\} \times \big\{(\widehat{y}_s^{(i)} - y_s)(\widehat{y}_s^{(j)} - y_s) - M_{ij}\big\}\Big].$$

For $r < s$, the centered error product from period r is $\mathcal{F}_{s-1}$-measurable. Hence the law of iterated expectations and the martingale-difference property give

$$\begin{aligned}
&\mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\}\big\{(\widehat{y}_s^{(i)} - y_s)(\widehat{y}_s^{(j)} - y_s) - M_{ij}\big\}\Big] \\
&\quad= \mathbb{E}\Big[\mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\}\big\{(\widehat{y}_s^{(i)} - y_s)(\widehat{y}_s^{(j)} - y_s) - M_{ij}\big\} \mid \mathcal{F}_{s-1}\Big]\Big] \\
&\quad= \mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\}\mathbb{E}\Big[(\widehat{y}_s^{(i)} - y_s)(\widehat{y}_s^{(j)} - y_s) - M_{ij} \mid \mathcal{F}_{s-1}\Big]\Big] \\
&\quad= \mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\}\cdot 0\Big] = 0.
\end{aligned}$$

Thus all cross terms have zero expectation. Since each centered error product is bounded in absolute value by $2H^2$ and the historical window contains L_T observations, it follows that

$$\begin{aligned}
\mathbb{E}|\overline{M}_{t,ij} - M_{ij}|^2 &= \frac{1}{L_T^2}\sum_{r=t-L_T}^{t-1}\mathbb{E}\Big[\big\{(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r) - M_{ij}\big\}^2\Big] \\
&\le \frac{1}{L_T^2}\sum_{r=t-L_T}^{t-1}(2H^2)^2 = \frac{L_T\cdot 4H^4}{L_T^2} = \frac{4H^4}{L_T}. \qquad (19)
\end{aligned}$$

Here and below, C denotes a finite positive constant independent of t and T whose value may change from one bound to another. Since $(v_i^\star)^2 = M_{ii}$ and $\widehat{v}_{t,i} + v_i^\star \ge \underline{v} > 0$, then

$$|\widehat{v}_{t,i} - v_i^\star|^2 = \frac{|\widehat{v}_{t,i}^2 - (v_i^\star)^2|^2}{(\widehat{v}_{t,i} + v_i^\star)^2} = \frac{|\widehat{v}_{t,i}^2 - M_{ii}|^2}{(\widehat{v}_{t,i} + v_i^\star)^2} \le \frac{|\widehat{v}_{t,i}^2 - M_{ii}|^2}{\underline{v}^2}.$$

By (15), the squared-scale estimation error can be decomposed as

$$\begin{aligned}
\widehat{v}_{t,i}^2 - M_{ii} &= \frac{L_T}{L_T-1}\overline{M}_{t,ii} - M_{ii} \\
&= \frac{L_T}{L_T-1}(\overline{M}_{t,ii} - M_{ii}) + \left(\frac{L_T}{L_T-1} - 1\right)M_{ii} \\
&= \frac{L_T}{L_T-1}(\overline{M}_{t,ii} - M_{ii}) + \frac{M_{ii}}{L_T-1}.
\end{aligned}$$

Applying $(a+b)^2 \le 2a^2 + 2b^2$ and taking expectations therefore yields

$$\begin{aligned}
\mathbb{E}|\widehat{v}_{t,i} - v_i^\star|^2 &\le \frac{1}{\underline{v}^2}\mathbb{E}\left[\left\{\frac{L_T}{L_T-1}(\overline{M}_{t,ii} - M_{ii}) + \frac{M_{ii}}{L_T-1}\right\}^2\right] \\
&\le \frac{2}{\underline{v}^2}\left[\left(\frac{L_T}{L_T-1}\right)^2\mathbb{E}|\overline{M}_{t,ii} - M_{ii}|^2 + \frac{M_{ii}^2}{(L_T-1)^2}\right] \\
&\le \frac{2}{\underline{v}^2}\left[\left(\frac{L_T}{L_T-1}\right)^2\frac{4H^4}{L_T} + \frac{H^4}{(L_T-1)^2}\right].
\end{aligned}$$

Here the last inequality follows from (19) and $0 \le M_{ii} \le H^2$. Since $L_T \ge 2$, we have

$$\frac{L_T}{L_T-1} \le 2, \qquad \frac{1}{(L_T-1)^2} \le \frac{4}{L_T^2} \le \frac{2}{L_T}.$$

Consequently,

$$\mathbb{E}|\widehat{v}_{t,i} - v_i^\star|^2 \le \frac{2}{\underline{v}^2}\left[\frac{16H^4}{L_T} + \frac{2H^4}{L_T}\right] = \frac{36H^4}{\underline{v}^2 L_T} = \frac{C}{L_T},$$

where $C = 36H^4/\underline{v}^2$ is independent of t, i, and L_T.

The bounded-error condition in Assumption 2 implies

$$0 \le \overline{M}_{t,ii} = \frac{1}{L_T}\sum_{r=t-L_T}^{t-1}(\widehat{y}_r^{(i)} - y_r)^2 \le H^2, \qquad v_i^\star = \sqrt{M_{ii}} \le H.$$

Since $(v_i^\star)^2 = M_{ii}$ and $v_i^\star \ge \underline{v}$, the difference-of-squares identity gives

$$\left|\sqrt{\overline{M}_{t,ii}} - v_i^\star\right| = \frac{|\overline{M}_{t,ii} - (v_i^\star)^2|}{\sqrt{\overline{M}_{t,ii}} + v_i^\star} = \frac{|\overline{M}_{t,ii} - M_{ii}|}{\sqrt{\overline{M}_{t,ii}} + v_i^\star} \le \frac{|\overline{M}_{t,ii} - M_{ii}|}{\underline{v}}.$$

Adding and subtracting $\sqrt{\overline{M}_{t,ii}}v_j^\star$, and then applying the triangle inequality, yields

$$\begin{aligned}
\left|\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}} - v_i^\star v_j^\star\right| &= \left|\sqrt{\overline{M}_{t,ii}}\left(\sqrt{\overline{M}_{t,jj}} - v_j^\star\right) + v_j^\star\left(\sqrt{\overline{M}_{t,ii}} - v_i^\star\right)\right| \\
&\le \sqrt{\overline{M}_{t,ii}}\left|\sqrt{\overline{M}_{t,jj}} - v_j^\star\right| + v_j^\star\left|\sqrt{\overline{M}_{t,ii}} - v_i^\star\right| \\
&\le H\frac{|\overline{M}_{t,jj} - M_{jj}|}{\underline{v}} + H\frac{|\overline{M}_{t,ii} - M_{ii}|}{\underline{v}} \\
&= \frac{H}{\underline{v}}\left(|\overline{M}_{t,ii} - M_{ii}| + |\overline{M}_{t,jj} - M_{jj}|\right).
\end{aligned}$$

Moreover, the Cauchy–Schwarz inequality applied to the historical forecast errors gives

$$\begin{aligned}
|\overline{M}_{t,ij}| &= \left|\frac{1}{L_T}\sum_{r=t-L_T}^{t-1}(\widehat{y}_r^{(i)} - y_r)(\widehat{y}_r^{(j)} - y_r)\right| \\
&\le \left\{\frac{1}{L_T}\sum_{r=t-L_T}^{t-1}(\widehat{y}_r^{(i)} - y_r)^2\right\}^{1/2}\left\{\frac{1}{L_T}\sum_{r=t-L_T}^{t-1}(\widehat{y}_r^{(j)} - y_r)^2\right\}^{1/2} \\
&= \sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}.
\end{aligned}$$

Since $M_{ij}/(v_i^\star v_j^\star) = \rho^\star$ for $i < j$, adding and subtracting $\overline{M}_{t,ij}/(v_i^\star v_j^\star)$ gives

$$\begin{aligned}
\frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}} - \rho^\star &= \frac{\overline{M}_{t,ij} - M_{ij}}{v_i^\star v_j^\star} + \overline{M}_{t,ij}\left(\frac{1}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}} - \frac{1}{v_i^\star v_j^\star}\right) \\
&= \frac{\overline{M}_{t,ij} - M_{ij}}{v_i^\star v_j^\star} + \frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}}\frac{v_i^\star v_j^\star - \sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}}{v_i^\star v_j^\star}.
\end{aligned}$$

Taking absolute values and using the preceding bounds together with $v_i^\star v_j^\star \ge \underline{v}^2$, we obtain

$$\begin{aligned}
\left|\frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}} - \rho^\star\right| &\le \frac{|\overline{M}_{t,ij} - M_{ij}|}{v_i^\star v_j^\star} + \frac{|\overline{M}_{t,ij}|}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}}\frac{|v_i^\star v_j^\star - \sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}|}{v_i^\star v_j^\star} \\
&\le \frac{1}{\underline{v}^2}\left[|\overline{M}_{t,ij} - M_{ij}| + \frac{H}{\underline{v}}|\overline{M}_{t,ii} - M_{ii}| + \frac{H}{\underline{v}}|\overline{M}_{t,jj} - M_{jj}|\right].
\end{aligned}$$

Squaring and applying $(a+b+c)^2 \le 3(a^2+b^2+c^2)$ yields

$$\begin{aligned}
&\mathbb{E}\left|\frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}} - \rho^\star\right|^2 \\
&\le \frac{3}{\underline{v}^4}\left[\mathbb{E}|\overline{M}_{t,ij} - M_{ij}|^2 + \frac{H^2}{\underline{v}^2}\mathbb{E}|\overline{M}_{t,ii} - M_{ii}|^2 + \frac{H^2}{\underline{v}^2}\mathbb{E}|\overline{M}_{t,jj} - M_{jj}|^2\right] \\
&\le \frac{3}{\underline{v}^4}\left[\frac{4H^4}{L_T} + \frac{H^2}{\underline{v}^2}\frac{4H^4}{L_T} + \frac{H^2}{\underline{v}^2}\frac{4H^4}{L_T}\right] = \frac{12H^4}{\underline{v}^4 L_T}\left(1 + \frac{2H^2}{\underline{v}^2}\right),
\end{aligned}$$

where the second inequality uses (19) for each of the three entries. Since there are $N(N-1)/2$ distinct pairs, the definition in (16) implies

$$\bar{\rho}_t - \rho^\star = \frac{2}{N(N-1)}\sum_{i<j}\left(\frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}} - \rho^\star\right).$$

Jensen's inequality for the squared pairwise average therefore gives

$$\begin{aligned}
\mathbb{E}|\bar{\rho}_t - \rho^\star|^2 &\le \frac{2}{N(N-1)}\sum_{i<j}\mathbb{E}\left|\frac{\overline{M}_{t,ij}}{\sqrt{\overline{M}_{t,ii}\overline{M}_{t,jj}}} - \rho^\star\right|^2 \\
&\le \frac{2}{N(N-1)}\frac{N(N-1)}{2}\frac{12H^4}{\underline{v}^4 L_T}\left(1 + \frac{2H^2}{\underline{v}^2}\right) = \frac{C}{L_T},
\end{aligned}$$

where $C = (12H^4/\underline{v}^4)(1 + 2H^2/\underline{v}^2)$. By the definition of the MAP estimator, adding and subtracting $((L_T-1)/(L_T+2N+3))\rho^\star$ gives

$$\begin{aligned}
\widehat{\rho}_t - \rho^\star &= \frac{L_T-1}{L_T+2N+3}\bar{\rho}_t - \rho^\star \\
&= \frac{L_T-1}{L_T+2N+3}(\bar{\rho}_t - \rho^\star) + \left(\frac{L_T-1}{L_T+2N+3} - 1\right)\rho^\star \\
&= \frac{L_T-1}{L_T+2N+3}(\bar{\rho}_t - \rho^\star) + \frac{(L_T-1)-(L_T+2N+3)}{L_T+2N+3}\rho^\star \\
&= \frac{L_T-1}{L_T+2N+3}(\bar{\rho}_t - \rho^\star) - \frac{2N+4}{L_T+2N+3}\rho^\star.
\end{aligned}$$

For fixed N, the second term is deterministic and of order L_T^{-1}, since $|\rho^\star| \le 1$ implies

$$\left|\frac{2N+4}{L_T+2N+3}\rho^\star\right| \le \frac{2N+4}{L_T+2N+3} \le \frac{2N+4}{L_T}.$$

Applying $(a-b)^2 \le 2a^2 + 2b^2$ and taking expectations yields

$$\begin{aligned}
\mathbb{E}|\widehat{\rho}_t - \rho^\star|^2 &\le 2\left(\frac{L_T-1}{L_T+2N+3}\right)^2\mathbb{E}|\bar{\rho}_t - \rho^\star|^2 + 2\left(\frac{2N+4}{L_T+2N+3}\right)^2(\rho^\star)^2 \\
&\le 2\mathbb{E}|\bar{\rho}_t - \rho^\star|^2 + \frac{2(2N+4)^2}{L_T^2} \\
&\le \frac{24H^4}{\underline{v}^4 L_T}\left(1 + \frac{2H^2}{\underline{v}^2}\right) + \frac{2(2N+4)^2}{L_T^2} \\
&\le \frac{1}{L_T}\left[\frac{24H^4}{\underline{v}^4}\left(1 + \frac{2H^2}{\underline{v}^2}\right) + (2N+4)^2\right] = \frac{C}{L_T}.
\end{aligned}$$

The second inequality uses $0 \le (L_T-1)/(L_T+2N+3) \le 1$ and $|\rho^\star| \le 1$, the third uses the preceding bound for $\bar{\rho}_t$, and the fourth uses $L_T \ge 2$. The constant C depends only on H, $\underline{v}$, and the fixed number of forecasts N.

Choose $\delta := (1/2)\min\{\underline{v},\ \underline{\rho} + 1/(N-1),\ 1-\overline{\rho}\} > 0$. On the parameter range $[\underline{v}/2, \overline{v}+\delta]^N \times [\underline{\rho}-\delta, \overline{\rho}+\delta]$, the matrix $\boldsymbol{M}(\boldsymbol{v},\rho)$ is positive definite, and

$$\mathbf{1}^\top \boldsymbol{M}(\boldsymbol{v},\rho)^{-1}\mathbf{1} \geq \frac{N}{\lambda_{\max}(\boldsymbol{M}(\boldsymbol{v},\rho))} \geq \frac{N}{\mathrm{tr}(\boldsymbol{M}(\boldsymbol{v},\rho))} \geq \frac{1}{(\overline{v}+\delta)^2} > 0.$$

Thus the derivative of $\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v},\rho)$ is uniformly bounded on this set. Whenever $\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2 \leq \delta^2$, both the estimated and population parameters lie within the bounds above.

We consider two cases according to the size of the parameter estimation error.

1) If

$$\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2 \leq \delta^2,$$

the parameter estimates lie in the range specified above. The mean-value inequality and nonexpansiveness of Euclidean projection therefore imply

$$\left\|\widetilde{\boldsymbol{s}}_t - \Pi_{\Delta_N}\left(\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)\right\|_2^2 \leq \left\|\boldsymbol{s}_t - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right\|_2^2 \leq C\left(\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2\right).$$

2) If $\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2 > \delta^2$, then multiplying this inequality by $2/\delta^2$ gives

$$2 < \frac{2}{\delta^2}\left\{\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2\right\}.$$

Since both projections belong to Δ_N, their squared distance is at most 2. Consequently, in this second case,

$$\left\|\widetilde{\boldsymbol{s}}_t - \Pi_{\Delta_N}\left(\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)\right\|_2^2 \leq 2 \leq \frac{2}{\delta^2}\left(\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2\right).$$

The coefficient is C in the first case and $2/\delta^2$ in the second. Taking their maximum therefore gives a bound that holds in both cases:

$$\left\|\widetilde{\boldsymbol{s}}_t - \Pi_{\Delta_N}\left(\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)\right\|_2^2 \leq \max\left\{C, \frac{2}{\delta^2}\right\}\left(\sum_{i=1}^N |\widehat{v}_{t,i} - v_i^\star|^2 + |\widehat{\rho}_t - \rho^\star|^2\right).$$

Taking expectations and applying the preceding parameter estimation bounds yields

$$\mathbb{E}\left\|\widetilde{\boldsymbol{s}}_t - \Pi_{\Delta_N}\left(\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)\right\|_2^2 \leq \max\left\{C, \frac{2}{\delta^2}\right\}\left(\sum_{i=1}^N \mathbb{E}|\widehat{v}_{t,i} - v_i^\star|^2 + \mathbb{E}|\widehat{\rho}_t - \rho^\star|^2\right) \leq \frac{C_7}{L_T}. \tag{20}$$

Under the nonnegativity condition in Proposition 1, the population CCR prior weights are nonnegative and sum to one, so $\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star) \in \Delta_N$. By (14),

$$\boldsymbol{M}\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star) = \boldsymbol{M}\frac{\boldsymbol{M}^{-1}\mathbf{1}}{\mathbf{1}^\top\boldsymbol{M}^{-1}\mathbf{1}} = \frac{\mathbf{1}}{\mathbf{1}^\top\boldsymbol{M}^{-1}\mathbf{1}}.$$

For any $\boldsymbol{w}$ satisfying $\mathbf{1}^\top\boldsymbol{w} = 1$, the difference between $\boldsymbol{w}$ and the population CCR prior-weight vector has components summing to zero, $\mathbf{1}^\top(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)) = 1 - 1 = 0$. Consequently, the cross term in the quadratic expansion vanishes:

$$\left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)^\top \boldsymbol{M}\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star) = \frac{\left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)^\top \mathbf{1}}{\mathbf{1}^\top\boldsymbol{M}^{-1}\mathbf{1}} = 0.$$

Expanding $\boldsymbol{w} = \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star) + (\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star))$ therefore gives

$$\begin{aligned}\boldsymbol{w}^\top\boldsymbol{M}\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)^\top\boldsymbol{M}\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star) &= \left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)^\top \boldsymbol{M}\left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right) \\ &\quad + 2\left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)^\top \boldsymbol{M}\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star) \\ &= \left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right)^\top \boldsymbol{M}\left(\boldsymbol{w} - \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star,\rho^\star)\right) \geq 0.\end{aligned}$$

Positive definiteness implies that equality holds only when $\boldsymbol{w} = \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star, \rho^\star)$. Since this vector belongs to Δ_N, it is the unique minimizer over Δ_N. Thus

$$\boldsymbol{s}^\star = \boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star, \rho^\star), \qquad \boldsymbol{M}\boldsymbol{s}^\star = \frac{\mathbf{1}}{\mathbf{1}^\top \boldsymbol{M}^{-1}\mathbf{1}}.$$

Since $L_T \asymp T$ and every L_T is positive, we may choose $C_1 := C_7 \sup_{T\geq 1} T/L_T < \infty$. Because the population prior weights satisfy $\boldsymbol{s}^\star \in \Delta_N$, their projection leaves them unchanged. Applying (20), for any $1 \leq t \leq T$, gives

$$\mathbb{E}\|\widetilde{\boldsymbol{s}}_t - \boldsymbol{s}^\star\|_2^2 = \mathbb{E}\big\|\widetilde{\boldsymbol{s}}_t - \Pi_{\Delta_N}\big(\boldsymbol{s}^{\mathrm{CCR}}(\boldsymbol{v}^\star, \rho^\star)\big)\big\|_2^2 \leq \frac{C_7}{L_T} = \frac{C_7}{T}\frac{T}{L_T} \leq \frac{C_1}{T} \leq \frac{C_1}{t}.$$

This proves (18) and verifies Assumption 4. □

E Derivatives used in the online update

This section gives the derivatives used to update ρ_t in (3). For notational simplicity, we omit the period and specification indices below. Let $\mathcal{A}$ be the current expert set and $\mathcal{I} = \{i \in \mathcal{A} : w_i > 0\}$. Write $\widetilde{w}_i^*$ for the preceding optimal weight assigned to expert i, and $\dot{w}_i = dw_i/d\rho$.

Derivative of the prior weights. Write $N = |\mathcal{A}|$. With the period subscripts suppressed, write the CCR prior weights in (2) as $s_i(\rho)$, $i \in \mathcal{A}$. Let $N_i(\rho)$ and $D(\rho)$ denote, respectively, the numerator and denominator in (2), with sums taken over $\mathcal{A}$. Their derivatives are

$$N_i'(\rho) = (N-1)\widehat{v}_i^{-2} - \widehat{v}_i^{-1}\sum_{j\in\mathcal{A}}\widehat{v}_j^{-1}, \qquad D'(\rho) = (N-1)\sum_{j\in\mathcal{A}}\widehat{v}_j^{-2} - \left(\sum_{j\in\mathcal{A}}\widehat{v}_j^{-1}\right)^2.$$

Provided $D(\rho) \neq 0$, the quotient rule gives

$$\frac{ds_i(\rho)}{d\rho} = \frac{N_i'(\rho)D(\rho) - N_i(\rho)D'(\rho)}{D(\rho)^2}.$$

$f(x) = x$. For brevity, write $c_i = (\widehat{y}^{(i)} - \bar{\widehat{y}})^2$. On $\mathcal{I}$, the stationarity equations are

$$2(c_i + \lambda + \kappa)w_i - 2\lambda s_i(\rho) - 2\kappa\widetilde{w}_i^* + \gamma = 0, \qquad i \in \mathcal{I}, \qquad \sum_{i\in\mathcal{I}} w_i = 1.$$

Here, γ is the Lagrange multiplier for the simplex equality constraint. On an interval where $\mathcal{I}$ is unchanged and $c_i + \lambda + \kappa \neq 0$ for $i \in \mathcal{I}$, differentiating these equations with respect to ρ and solving together with $\sum_{i\in\mathcal{I}} \dot{w}_i = 0$ gives

$$\dot{w}_i = \begin{cases} \dfrac{\lambda}{c_i + \lambda + \kappa}\left\{\dfrac{ds_i(\rho)}{d\rho} - \dfrac{\displaystyle\sum_{j\in\mathcal{I}} \frac{ds_j(\rho)/d\rho}{c_j + \lambda + \kappa}}{\displaystyle\sum_{j\in\mathcal{I}} \frac{1}{c_j + \lambda + \kappa}}\right\}, & i \in \mathcal{I}, \\ 0, & i \notin \mathcal{I}. \end{cases} \tag{21}$$

The set $\mathcal{I}$ is determined by the weights obtained using entmax, and (21) gives $dw_i/d\rho$ while $\mathcal{I}$ remains unchanged.

$f(x) = \log x$ and $f(x) = \log(\sigma_t^2 + x)$. Set $U = \sum_{j\in\mathcal{A}} c_j w_j^2$ for $f(x) = \log x$ and $U = \hat{\sigma}_t^2 + \sum_{j\in\mathcal{A}} c_j w_j^2$ for $f(x) = \log(\sigma_t^2 + x)$. On $\mathcal{I}$, write the stationarity and simplex equations as

$$F_i(\boldsymbol{w}_\mathcal{I}, \gamma; \rho) = \frac{2c_i w_i}{U} + 2\lambda\{w_i - s_i(\rho)\} + 2\kappa\{w_i - \widetilde{w}_i^*\} + \gamma = 0, \qquad i \in \mathcal{I},$$
$$F_0(\boldsymbol{w}_\mathcal{I}) = \sum_{i\in\mathcal{I}} w_i - 1 = 0.$$

For $i, m \in \mathcal{I}$, write $H_{im} = \partial F_i / \partial w_m$. Then

$$H_{im} = 2c_i \left\{ \frac{\mathbf{1}\{i = m\}}{U} - \frac{2w_i w_m c_m}{U^2} \right\} + 2(\lambda + \kappa)\mathbf{1}\{i = m\}.$$

The remaining Jacobian entries are $\partial F_i / \partial \gamma = 1$, $\partial F_0 / \partial w_m = 1$, and $\partial F_0 / \partial \gamma = 0$. On an interval where $\mathcal{I}$ is unchanged, provided $U > 0$ and the bordered Jacobian is nonsingular, ρ enters these equations only through the prior weights, and implicit differentiation gives

$$\begin{pmatrix} H & \mathbf{1} \\ \mathbf{1}^\top & 0 \end{pmatrix} \begin{pmatrix} d\boldsymbol{w}_{\mathcal{I}}/d\rho \\ d\gamma/d\rho \end{pmatrix} = \begin{pmatrix} 2\lambda \, d\boldsymbol{s}_{\mathcal{I}}/d\rho \\ 0 \end{pmatrix}.$$

The derivative is set to zero outside $\mathcal{I}$.

Finally, for all three specifications, the derivative of the squared forecast loss is

$$\frac{d}{d\rho} \left(y - \sum_{i \in \mathcal{A}} w_i(\rho) \widehat{y}^{(i)} \right)^2 = -2 \left(y - \sum_{i \in \mathcal{A}} w_i(\rho) \widehat{y}^{(i)} \right) \sum_{i \in \mathcal{A}} \widehat{y}^{(i)} \dot{w}_i.$$

This is the derivative used in (3).

F Simulation Study

We consider a simulation study in which forecast accuracy differs across experts and changes over time. Specifically, some experts experience periods of greater forecast dispersion and bias, while the remaining experts follow the baseline forecast process. Expert accuracy changes gradually over time, with different experts following different patterns. This design examines whether online REF can improve accuracy as expert performance changes over time under different forecast-error correlations. We also examine the effect of the entmax parameter α on expert selection. In period t, the quantity of interest y_t and expert forecasts $\widehat{y}_{ti}$ are generated by

$$y_t = m + u_t, \qquad \widehat{y}_{ti} = m + h_{ti}(a_{ti} + r_{ti}) + b_{ti}, \qquad a_{ti} = \phi a_{t-1,i} + \epsilon_{ti},$$

where $u_t \sim N(0, \sigma_y^2)$, $\epsilon_{ti} \sim N\{0, (1-\phi^2)\sigma_a^2\}$, and $r_{ti} \sim N(0, \sigma_\mu^2 - \sigma_a^2)$. The AR(1) component is initialized from $N(0, \sigma_a^2)$. All random draws in $\{u_t\}$, $\{\epsilon_{ti}\}$, and $\{r_{ti}\}$ are mutually independent and are independent of the initial values $\{a_{0i}\}$. The term a_{ti} carries over to the next period through the AR(1) process, while r_{ti} varies independently across periods. The multiplier h_{ti} controls the magnitude of these deviations, and b_{ti} adds bias. When $h_{ti} = 1$ and $b_{ti} = 0$, the forecast equation reduces to the baseline process $\widehat{y}_{ti} = m + a_{ti} + r_{ti}$.

We consider an expert pool containing 50 experts. In each replication, we randomly select 40 experts whose dispersion and bias vary over time; the remaining 10 experts have $h_{ti} = 1$ and $b_{ti} = 0$. The selected experts remain the same over time. For each selected expert, we set

$$h_{ti} = 1 + A \left\{ 1 + \sin\left(\frac{2\pi(t-1)}{q} + \theta_i \right) \right\}, \qquad b_{ti} = \frac{B}{2} \left\{ 1 + \sin\left(\frac{2\pi(t-1)}{q} + \theta_i \right) \right\},$$

where $\theta_i \sim \mathrm{Uniform}(0, 2\pi)$ is drawn once for each selected expert, independently across experts and independently of the other random draws. The 10 remaining experts are relatively reliable, while the other 40 become less accurate when their dispersion and bias increase.

We use 100 Monte Carlo replications for each design point. We set $m = 100$, $\sigma_y = 10$, $\sigma_a = 4$, and $\phi = 0.5$, with $A = 0.25$, $B = 30$, and $q = 40$. Each simulated dataset contains 50 experts and a total of 80 periods: the first 40 periods form the initial historical window, the next 20 periods are used for validation and tuning parameter selection, and the final 20 periods form the evaluation period. We use a rolling historical window of 40 periods. We vary $\rho_0 \in \{0.2, 0.4, 0.6, 0.8\}$, where ρ_0 is the correlation between the forecast errors of any two of the 10 experts with $h_{ti} = 1$ and $b_{ti} = 0$, with $\sigma_\mu^2 = \sigma_y^2(1 - \rho_0)/\rho_0$. A larger ρ_0 corresponds to a smaller variance of r_{ti}, while A and B remain fixed. Correlations involving the other 40 experts depend on their values of h_{ti} and can therefore differ across expert pairs and vary over time. Since $-1 \le \sin(x) \le 1$ and $A = 0.25$, we have $1 \le h_{ti} \le 1 + 2A = 1.5$. For given values of the experts' h_{ti} and b_{ti} at period t, the forecast errors of two different experts have covariance σ_y^2, and expert i's forecast-error

variance is $\sigma_y^2\{1 + h_{ti}^2(1-\rho_0)/\rho_0\}$. Dividing this covariance by the product of the two forecast-error Std. values gives their correlation. As either expert's h_{ti} increases, the error variance increases and the correlation decreases. The correlation is therefore bounded above by ρ_0, obtained when $h_{ti} = h_{tj} = 1$, and below by $\rho_0/\{\rho_0 + 2.25(1-\rho_0)\}$, obtained by setting $h_{ti} = h_{tj} = 1.5$. For $\rho_0 = 0.2, 0.4, 0.6$, and 0.8, these bounds are $[0.1, 0.2]$, $[0.23, 0.4]$, $[0.4, 0.6]$, and $[0.64, 0.8]$, respectively.

We compare online REF with offline REF and the six benchmarks described in Section 5.1. Tuning parameters are selected following Section 3.3. The RMSSE denominator is computed from the initial historical window.

Table A.1 reports forecasting performance across the 100 replications. Online REF ranks first overall, followed by offline REF and CCR, with a 29.54% reduction in average RMSSE relative to Simple Mean. Online REF produces the most accurate forecasts at $\rho_0 = 0.2, 0.6$, and 0.8, while CCR performs best at $\rho_0 = 0.4$. Figure A.1 shows the number of experts receiving positive weights in the final online REF forecast at $\rho_0 = 0.2$. As α increases, the number of active experts receiving positive weights generally decreases from all 50 experts at $\alpha = 1$ to smaller subsets of the expert pool.

Table A.1: Forecast accuracy under different baseline correlations ρ_0 (avgRMSSE). Std. across 100 replications are shown in parentheses. Bold numbers indicate the best performance in each row.

ρ_0	**Online REF**	Offline REF	Simple Mean	Trimmed Mean	Winsorized Mean	Variance Weights	CWM	CCR
0.2	**.935**	.941	1.183	1.137	1.160	1.090	1.120	.947
	(.200)	(.187)	(.211)	(.206)	(.209)	(.203)	(.215)	(.198)
0.4	.818	.822	1.167	1.124	1.142	1.034	1.068	**.813**
	(.170)	(.172)	(.205)	(.201)	(.202)	(.190)	(.212)	(.170)
0.6	**.778**	.779	1.162	1.123	1.139	.993	1.022	.784
	(.160)	(.162)	(.203)	(.199)	(.200)	(.182)	(.224)	(.165)
0.8	**.760**	.762	1.159	1.125	1.141	.961	.984	.794
	(.153)	(.158)	(.201)	(.199)	(.200)	(.176)	(.235)	(.171)
Avg.[a]	**.823**	.826	1.168	1.127	1.146	1.020	1.048	.835
Impr. (%)[b]	29.54%	29.25%	0.00%	3.47%	1.88%	12.69%	10.22%	28.52%
Rank[c]	1	2	8	6	7	4	5	3

[a] Averaged across the four values of ρ_0.
[b] Percentage reduction relative to Simple Mean.
[c] Ranked by overall avgRMSSE.

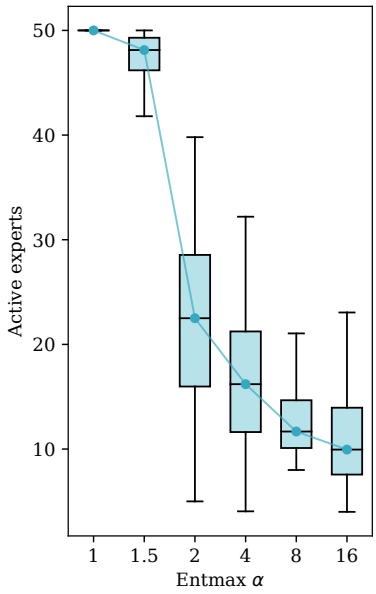

Figure A.1: Number of active experts receiving positive weights in the final online REF forecast for different values of α at $\rho_0 = 0.2$, averaged over the 20 evaluation periods. Boxplots summarize 100 replications, and connected points show the medians.

G Additional Empirical Results

Tables A.2 and A.3 report the full results for the M5 and SPF studies, including Std. values and percentage improvements relative to Simple Mean. The evaluation periods and RMSSE definitions are the same as in the main analysis.

Because α controls the sparsity of the online REF weights, we report the validation-selected values together with the resulting numbers of experts receiving positive weights. The final online REF forecast averages the forecasts from the three specifications. We count an expert as receiving a positive weight when its weight averaged across the three specifications is positive. For M5, validation selects $\alpha^* = 4$. At this value, the final online REF forecast assigns positive weights to about 24 of the 50 available experts (48%) on average across series and evaluation days. For SPF, validation selects $\alpha^* = 8$ for GDP growth, $\alpha^* = 4$ for CPI inflation, and $\alpha^* = 2$ for the change in the unemployment rate. Averaged across both horizons and evaluation quarters, the final online REF forecast assigns positive weights to about 11 experts from an average pool of 37 for GDP growth, 8 from an average pool of 37 for CPI inflation, and 16 from an average pool of 38 for the change in the unemployment rate. These values correspond to 30%, 21%, and 43% of the available experts, respectively.

Table A.2: Performance on M5 by aggregation level (avgRMSSE). Std. values across series within each level are shown in parentheses, except for Level 1, which contains only one series. Bold numbers indicate the lowest avgRMSSE among the methods shown at each level.

Level	**Online REF**	Offline REF	Simple Mean	Trimmed Mean	Winsorized Mean	Variance Weights	CWM	CCR
1	**.115**	.126	.147	.147	.147	.139	.129	.121
	(-)	(-)	(-)	(-)	(-)	(-)	(-)	(-)
2	**.249**	.258	.292	.296	.295	.289	.288	.285
	(.104)	(.097)	(.070)	(.067)	(.066)	(.074)	(.060)	(.087)
3	**.322**	.331	.347	.348	.349	.346	.354	.341
	(.102)	(.087)	(.096)	(.092)	(.092)	(.088)	(.107)	(.096)
4	**.226**	.233	.243	.247	.248	.240	.256	.235
	(.138)	(.139)	(.123)	(.128)	(.127)	(.126)	(.137)	(.129)
5	**.331**	.343	.374	.381	.383	.356	.356	.341
	(.177)	(.173)	(.189)	(.197)	(.201)	(.175)	(.173)	(.177)
6	**.318**	.325	.350	.352	.353	.345	.337	.333
	(.094)	(.090)	(.075)	(.075)	(.074)	(.079)	(.087)	(.090)
7	**.453**	.457	.490	.493	.494	.478	.466	.468
	(.165)	(.167)	(.169)	(.172)	(.172)	(.165)	(.171)	(.173)
8	**.444**	.455	.464	.464	.464	.463	.455	.464
	(.172)	(.170)	(.178)	(.178)	(.177)	(.174)	(.172)	(.168)
9	**.578**	.582	.600	.602	.602	.594	.584	.582
	(.217)	(.218)	(.221)	(.222)	(.222)	(.217)	(.219)	(.214)
Avg.[a]	**.337**	.346	.368	.370	.371	.361	.359	.352
Impr. (%)[b]	8.21%	5.98%	0.00%	-0.64%	-0.81%	1.74%	2.47%	4.18%
Rank[c]	1	2	6	7	8	5	4	3

[a] Averaged across the nine levels.
[b] Percentage reduction relative to Simple Mean.
[c] Ranked by overall avgRMSSE.

Table A.3: Performance on SPF by economic indicator and forecast horizon (RMSSE). Sample Std. values of quarter-level scaled absolute errors are shown in parentheses. Bold numbers indicate the lowest RMSSE among the methods shown for each indicator and horizon.

Indicator	Horizon	**Online REF**	Offline REF	Simple Mean	Trimmed Mean	Winsorized Mean	Variance Weights	CWM	CCR
Nominal GDP growth	$h=0$	**.561** (.442)	.562 (.449)	.600 (.495)	.602 (.501)	.600 (.499)	.612 (.487)	.621 (.505)	.623 (.501)
	$h=1$	**1.506** (1.271)	1.517 (1.269)	1.597 (1.355)	1.593 (1.351)	1.594 (1.351)	1.591 (1.316)	1.600 (1.330)	1.584 (1.311)
	Avg.[a]	**1.033**	1.040	1.098	1.097	1.097	1.102	1.111	1.104
	Impr. (%)[b]	5.93%	5.32%	0.00%	0.10%	0.16%	-0.29%	-1.13%	-0.48%
	Rank[c]	1	2	5	4	3	6	8	7
CPI inflation rate	$h=0$	**.656** (.450)	.675 (.458)	.758 (.519)	.754 (.508)	.755 (.511)	.812 (.574)	.799 (.578)	.794 (.571)
	$h=1$	**1.129** (.814)	1.192 (.877)	1.195 (.885)	1.195 (.884)	1.193 (.883)	1.215 (.898)	1.253 (.916)	1.203 (.882)
	Avg.[a]	**.892**	.933	.977	.974	.974	1.013	1.026	.998
	Impr. (%)[b]	8.64%	4.44%	0.00%	0.23%	0.24%	-3.76%	-5.07%	-2.21%
	Rank[c]	1	2	5	4	3	7	8	6
Change in unemployment rate	$h=0$	**.454** (.360)	.609 (.528)	.814 (.748)	.829 (.768)	.825 (.763)	.786 (.708)	2.620 (2.478)	.963 (.894)
	$h=1$	**2.481** (2.365)	2.495 (2.382)	2.713 (2.576)	2.665 (2.533)	2.672 (2.539)	2.594 (2.448)	2.614 (2.462)	2.578 (2.437)
	Avg.[a]	**1.468**	1.552	1.764	1.747	1.749	1.690	2.617	1.771
	Impr. (%)[b]	16.78%	11.99%	0.00%	0.92%	0.85%	4.18%	-48.40%	-0.40%
	Rank[c]	1	2	6	4	5	3	8	7

[a] Averaged across $h=0$ and $h=1$ within each indicator.
[b] Percentage reduction relative to Simple Mean.
[c] Ranked by average RMSSE within each indicator.