BELLMAN CONFORMAL INFERENCE: CALIBRATING PREDICTION INTERVALS FOR TIME SERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

We introduce Bellman Conformal Inference (BCI), a framework that wraps around any time series forecasting models and provides calibrated prediction intervals. Unlike the existing methods, BCI is able to leverage multi-step ahead forecasts and explicitly optimize the average interval lengths by solving a one-dimensional stochastic control problem (SCP) at each time step. In particular, we use the dynamic programming algorithm to find the optimal policy for the SCP. We prove that BCI achieves long-term coverage under arbitrary distribution shifts and temporal dependence, even with poor multi-step ahead forecasts. When the raw forecast intervals are approximately calibrated, we prove that BCI is near optimal for its surrogate model predictive control problem. We find empirically that BCI avoids uninformative intervals with infinite lengths under our main experimental configuration and generates substantially shorter prediction intervals when the raw forecast intervals are poorly calibrated.

1 INTRODUCTION

Uncertainty quantification for time series nowcasting and forecasting is crucial in many areas such as climate science, epidemiology, industrial engineering, and macroeconomics. Ideally, the forecaster would generate a prediction interval at each time period that is calibrated in the sense that the fraction of intervals covering the true outcomes is approximately equal to the target coverage level in the long run. Classical approaches for generating prediction intervals are mostly model-based Box & Jenkins (1976); Engle (1982a); Stock & Watson (2010); Brown (1964); Jorda (2005). However, time series models are often mis-specified due to non-stationarity or changing environments. As a result, the model-based prediction intervals tend to be poorly calibrated (see for instance the gray curves in Figure 1). Moreover, many forecasters have upgraded their workflows by incorporating black-box machine learning algorithms (e.g. Taylor & Letham, 2018; Makridakis et al., 2018; Herzen et al., 2022), for which valid uncertainty quantification proves to be challenging.

Due to complex temporal dependence and distribution shifts, distribution-free uncertainty quantification techniques such as conformal inference (e.g. Saunders et al., 1999; Papadopoulos et al., 2002; Vovk et al., 2005; Lei et al., 2015; 2018; Angelopoulos & Bates, 2021) are not guaranteed to achieve calibration. While many variants of conformal inference have been developed under weaker assumptions, most require restrictive dependence structure Chernozhukov et al. (2018); Tibshirani et al. (2019); Lei & Candès (2021); Oliveira et al. (2022); Candès et al. (2023), or limited distribution shifts Barber et al. (2023), or accurate model estimates Xu & Xie (2021; 2023), or multiple independent copies of the time series Stankeviciute et al. (2021); Dietterich & Hostetler (2022); Sun & Yu (2023).

An important departure from the aforementioned methods is Adaptive Conformal Inference (ACI) proposed by Gibbs & Candès (2021), as well as its variants Gibbs & Candès (2022); Zaffran et al. (2022); Feldman et al. (2023); Angelopoulos et al. (2023), which can produce approximately calibrated prediction intervals without making any assumptions on the time series. ACI generalizes standard conformal prediction for exchangeable data by choosing a time-varying nominal miscoverage rate to achieve calibration. Alternatively, ACI can be formulated as an online gradient update that adjusts the nominal miscoverage rate downwards following each failure to cover the true outcome and upwards otherwise.

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.

While the appeal of assumption-free calibration is evident, ACI lacks a mechanism to explicitly optimize the average interval lengths. For most practical time series forecasting models, multi-step ahead prediction intervals are readily available Box & Jenkins (1976); Fan & Yao (2003); West & Harrison (2006); Politis & Wu (2023). These intervals could be used to trade off between present and future interval lengths. For instance, if a two-day ahead interval accurately represents the one-day ahead interval for the following day, its length (at any given nominal miscoverage rate) could provide valuable guidance for judiciously selecting the nominal miscoverage rate to prevent unnecessarily wide intervals in two days. In this paper, we introduce Bellman Conformal Inference (BCI), which wraps around any multi-step ahead prediction intervals and formulates a stochastic control problem (SCP) to explicitly optimize the average interval lengths. In particular, at each time point, BCI formulates an SCP by taking the nominal miscoverage rate as the action to optimize an objective function that trades off between the average length of multi-step ahead prediction intervals, as proxies for the actual future prediction intervals, and the estimated average future coverage. The SCP problem can be efficiently solved by dynamic programming (DP).

As with ACI, BCI is guaranteed to generate calibrated prediction intervals under Assumption 2.1 that is satisfied by almost all commonly used prediction intervals – in particular, it does not require the nominal multi-step prediction intervals to be well-calibrated. Instead of updating the nominal miscoverage rates directly, BCI applies the online gradient update to adjust a "security parameter" in the SCP that controls the trade-off between average interval lengths and short-term coverage rate. Applying an online gradient update to an alternative parameter is similar to the Rolling RC method introduced by Feldman et al. (2023) or the Conformal Decision Theory from Lekeufack et al. (2023). Both these works extend ACI to online risk control, akin to the extension of conformal inference in the offline setting (Bates et al. (2021); Angelopoulos et al. (2021; 2022)).

Another closely related method is the Conformal PID Control proposed by Angelopoulos et al. (2023). They also extend ACI by leveraging ideas from control theory. This method modifies the ACI update rule by adding two more pieces – error integration and scorecasting. The error integration stabilizes coverage and reduces the risk of the infinite-length intervals that commonly arise with ACI, while scorecasting enables the use of historical observations to take advantage of local stationarity. Both can be achieved by BCI, but through very different mechanisms. In particular, BCI leverages historical data to model dynamics and optimizes interval lengths explicitly.

We apply BCI to forecast stock price volatility, absolute return, and the popularity of Google search trend. Our findings suggest that when the nominal multi-step ahead prediction intervals are poorly-calibrated, BCI can generate substantially shorter intervals compared to ACI and conformal PID, as seen in Figure 1c. Conversely, when the nominal intervals are well-calibrated, BCI generates intervals of comparable lengths to ACI and conformal PID, and, under our main configuration, prevents the occurrence of infinitely long intervals as seen in Figure 1b. We discuss this correspondence in detail in Section 4.3.

2 SETUP

Consider a time series $Y_1, Y_2, \ldots$, where $Y_t \in \mathcal{Y}$ is the outcome of interest the analyst wants to predict. We assume Y_t is not observed until time $t+1$. Further, let $\mathcal{F}_{t-1}$ denote the σ-algebra generated by all information available at time t. In particular, $\mathcal{F}_{t-1}$ includes all past outcomes $Y_{t-1}, Y_{t-2}, \ldots, Y_1$. It can also include other variables that have been observed before t.

2.1 MULTI-STEP AHEAD PREDICTION INTERVALS

The analyst applies the forecasting algorithm that can generate T-step ahead prediction intervals, where T is a positive integer. At each time t and future time s for some $t \le s \le t+T-1$, let $C_{s|t}(1-\beta) \subset \mathcal{Y}$ denote the $(1-\beta)$ prediction interval for Y_s produced by the forecasting algorithm. Here, β stands for the nominal miscoverage rate. We say $C_{s|t}(1-\beta)$ is a nominal prediction interval because $\mathbb{P}(Y_s \in C_{s|t}(1-\beta)) = 1-\beta$ when the model behind the forecasting algorithm is correctly specified. However, when the model is misspecified, $\mathbb{P}(Y_s \in C_{s|t}(1-\beta))$ might deviate from $1-\beta$. Mose time series models generate multi-step prediction intervals. Below are two examples.

GARCH model Engle (1982b); Bollerslev (1986): a GARCH(p, q) model is a generalization of the ARMA model that allows for conditional heteroskedasticitiy. In particular, it assumes $Y_t \sim$

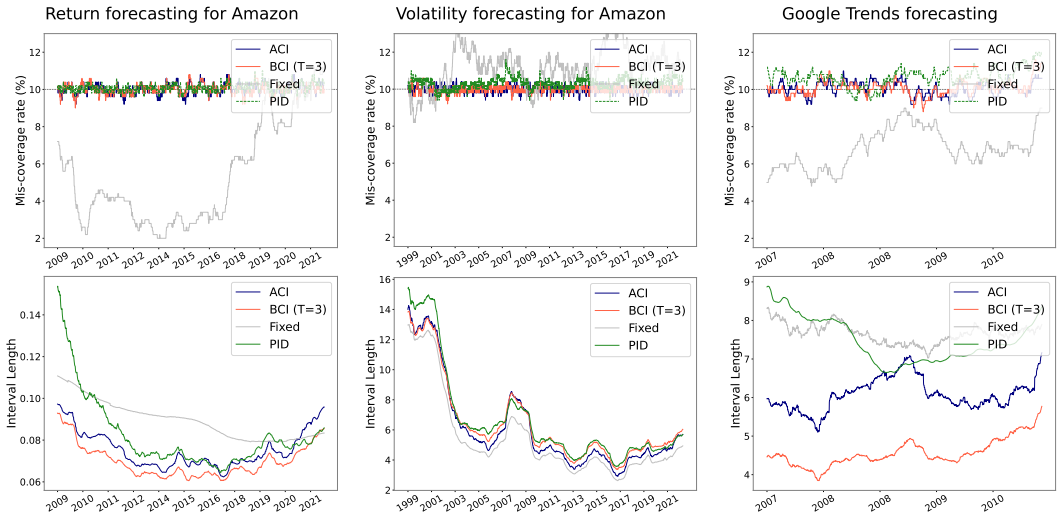

(a) Return forecasting (b) Volatility forecasting (c) Google trend

Figure 1: Online time series forecasting for three different tasks: return forecasting on Amazon stock data, volatility forecasting on Amazon stock data, and Google search popularity data for the keyword "deep learning". **Top panel:** moving averages of miscoverage rates over 500 data points. **Bottom panel:** moving averages of prediction interval lengths. In all figures, the red and green curves correspond to our proposed BCI algorithm and the conformal PID method; the blue curves correspond to the ACI algorithm with stepsize 0.1; the gray curves correspond to setting $\alpha_t = \overline{\alpha}$ for all t.

$\mathcal{N}(0, \sigma_t^2)$ where

$$\sigma_t^2 = \omega + b_1\sigma_{t-1}^2 + \ldots + b_p\sigma_{t-p}^2 + a_1\epsilon_{t-1}^2 + \ldots + a_q\epsilon_{t-q}^2,$$

and $\epsilon_1, \epsilon_2, \ldots \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma^2)$. When the GARCH model is applied to forecast Y_t^2, $C_{s|t}(1-\beta)$ can be formed by the $(\beta/2)$-th and $(1-\beta/2)$-th quantile of $(\hat{\sigma}_{t+s-1}Z)^2$ where $Z \sim \mathcal{N}(0,1)$ and $\hat{\sigma}_{t+s-1}^2$ is the estimate of σ_{t+s-1}^2.

Sequence-to-sequence neural networks Hochreiter & Schmidhuber (1997); Salinas et al. (2020); Vaswani et al. (2017): this approach first models the marginal distribution of each Y_t (e.g., $Y_t \sim \mathcal{N}(\mu_t, \sigma_t^2)$). Then it uses a sequence-to-sequence neural network, such as a transformer, to predict the likelihood parameters of a future observation Y_t (e.g. $\hat{\mu}_t, \hat{\sigma}_t$) from the lagged observations $\{Y_{t-1}, Y_{t-2}, \ldots\}$. The prediction intervals $C_{s|t}(1-\beta)$ can be formed using $\beta/2$ and $1-\beta/2$-th quantiles of Y_t with likelihood parameters imputed by the model forecast.

Throughout the paper we make the following mild assumptions on the nominal prediction intervals, akin to Gupta et al. (2022) and Feldman et al. (2023).

Assumption 2.1. For any $t, s \geq 1$, the prediction interval $C_{s|t}(1-\beta) \subset \mathcal{Y}$ satisfies the following three conditions:

- **Continuity:** Almost surely, $\beta \mapsto 1\{Y_s \in C_{s|t}(1-\beta)\}$ is right-continuous on $[0, 1)$; We set $C_{s|t}(1-\beta) = C_{s|t}(1)$ for all $\beta < 0$ and $C_{s|t}(1-\beta) = C_{s|t}(0)$ for all $\beta > 1$.
- **Monotonicity:** $C_{s|t}(1-\beta_1) \subset C_{s|t}(1-\beta_2)$ if $\beta_1 > \beta_2$;
- **Safeguard:** $C_{s|t}(1-0) = \mathcal{Y}, C_{s|t}(1-1) = \emptyset$, i.e., $\mathbb{P}(Y \in C_{s|t}(1)) = 1, \mathbb{P}(Y \in C_{s|t}(0)) = 0$.

The monotonicity condition assumes that a smaller nominal miscoverage rate gives a wider interval, which is a natural condition for any reasonable prediction interval. The safeguard condition assumes that it is always safe to set $\beta = 0$. It is important for the coverage guarantee in Theorem 3.1. Otherwise, nature can always choose the time series adversarially to escape the prediction interval. The monotonicity condition can be enforced by redefining $C_{s|t}(1-\beta) = \bigcup_{\beta \leq \beta'} C_{s|t}(1-\beta')$ and $C_{s|t}(1) = \mathcal{Y}$.

2.2 CALIBRATING PREDICTION INTERVALS

Given the multi-step ahead prediction intervals $C_{s|t}(\cdot)$, BCI selects a nominal miscoverage index α_t and outputs $C_{t|t}(1-\alpha_t) \subset \mathcal{Y}$ as the prediction interval for Y_t. We write $C_t(\cdot) := C_{t|t}(\cdot)$ for notational convenience. Following Gibbs & Candès (2021), our goal is to generate a sequence of nominal miscoverage indices $\{\alpha_t : t \in [K]\}$ such that α_t only depends on $\mathcal{F}_{t-1}$ and

$$\lim_{K\to\infty} \frac{1}{K}\sum_{t=1}^{K} 1(Y_t \notin C_t(1-\alpha_t)) \le \overline{\alpha} \quad \text{almost surely} \tag{1}$$

for some pre-specified target miscoverage level $\overline{\alpha}$ (e.g. $\overline{\alpha} = 0.1$). In particular, the inequality must hold for every joint distribution of $\{Y_t : t \ge 1\}$ satisfying Assumption 2.1, including deterministic sequences satisfying these conditions. The ACI algorithm chooses an initial nominal coverage α_0 and updates α_t by

$$\alpha_t = \alpha_{t-1} + \gamma(\overline{\alpha} - \mathrm{err}_{t-1}), \tag{2}$$

where $\mathrm{err}_{t-1} = 1(Y_{t-1} \notin C_{t-1}(1-\alpha_{t-1}))$. Above, γ is a fixed stepsize, which can be made data-adaptive using more advanced techniques in online learning Gibbs & Candès (2022). Gibbs & Candès (2021) prove that, the average miscoverage over the first T time periods is at most $\overline{\alpha}+2/T\gamma$.

Let β_t be the supremum of the nominal miscoverage rates at which Y_t is covered,

$$\beta_t = \sup_{Y_t \in C_t(1-\beta)} \beta. \tag{3}$$

We shall refer to β_t as the uncalibrated probability inverse transform (PIT). Note that at time t, β_{t-1} is observed as Y_{t-1} is. Under Assumption 2.1, right-continuity and the safeguard at $a = 0$ give $\beta_t > 0$. Monotonicity gives coverage for $0 \le a < \beta_t$; right-continuity excludes coverage at $a = \beta_t < 1$, and the safeguard excludes it at $a = \beta_t = 1$. Thus $\{a \in [0,1] : Y_t \in C_t(1-a)\} = [0, \beta_t)$ with $0 < \beta_t \le 1$, and $\mathrm{err}_t = 1\{Y_t \notin C_t(1-\alpha_t)\} = 1\{\alpha_t \ge \beta_t\}$, almost surely. Thus, for ACI, α_t depends exclusively on $(\alpha_{t-1}, \beta_{t-1})$. While ACI indeed provides a valid calibration guarantee, it disregards two sources of information: the first is the length of $C_t(1-\beta)$ and multi-step ahead intervals, which can inform a better choice of α_t. The second is the historical sequence of uncalibrated PITs $\{\beta_j\}_{j<t}$ that allows a better estimate of β_{t+1} than just using β_t alone. As we shall see, the BCI method we introduce in this paper chooses α_t as a function of all available information including the past actions $\alpha_{t-1}, \alpha_{t-2}, \ldots$, the past uncalibrated PITs $\beta_{t-1}, \beta_{t-2}, \ldots$, and the multi-step ahead intervals $C_{t|t}(\cdot), \ldots, C_{t+T-1|t}(\cdot)$.

3 BELLMAN CONFORMAL INFERENCE

3.1 BCI AS MODEL PREDICTIVE CONTROL

At a high level, BCI is analogous to Model Predictive Control (MPC) (Borrelli et al., 2017). At each time point t, BCI models the "dynamics" of the process $(Y_1, \beta_1), (Y_2, \beta_2), \ldots$ from past observations, simulates the "system" over the next T steps, and plans the "control" α_t by minimizing the "cost" driven by both the average interval length and miscoverage. Unlike traditional reactive control approaches, the proactive MPC approach is more suitable for forecasting problems under substantial distribution shifts.

To set up the cost minimization problem at time t, we denote by $L_{s|t}(\beta) = |C_{s|t}(1-\beta)|$ the function that maps the miscoverage rate to the length of the nominal multi-step ahead prediction interval for Y_s and $F_{s|t}$ the marginal distribution of β_s estimated using the past observations. We do not make any assumption on how $L_{s|t}(\cdot)$ and $F_{s|t}$ are generated. For all experiments in this paper, we simply set $F_{s|t} \equiv F_t$ where F_t is the empirical CDF of $\{\beta_{t-1}, \ldots, \beta_{t-B}\}$ for some large B. In our simulation, we set $B = 100$.

At time t, we solve the following optimization problem:

$$\min_{\alpha_{t|t},\ldots,\alpha_{t+T-1|t}} \mathbb{E}\left[\underbrace{\sum_{s=t}^{t+T-1} L_{s|t}(\alpha_{s|t})}_{\text{Efficiency: interval length}} + \lambda_t \underbrace{\max\left\{ \frac{1}{T}\sum_{s=t}^{t+T-1} \mathrm{err}_{s|t} - \bar{\alpha}, 0 \right\}}_{\text{Validity: miscoverage rate}} \right] + \tau K(\alpha_{t|t}, \bar{\alpha}), \tag{4}$$

where $(\beta_{t|t}, \ldots, \beta_{t+T-1|t}) \sim F_{t|t} \otimes \cdots \otimes F_{t+T-1|t}$, $\alpha_{s|t}$ denotes the planned action for time s, $\mathrm{err}_{s|t} = 1(\alpha_{s|t} \geq \beta_{s|t})$ denotes the error indicator for the realized $\beta_{s|t}$, and λ_t denotes the relative weight on the miscoverage that BCI uses to achieve the coverage guarantee, and $\tau K(\alpha_{t|t}, \bar{\alpha})$ is a regularization term. We will set $\tau = 0$ for the applications and consider $\tau > 0$ and a KL regularizer in Appendix F. When the minimum is not attained, we select a solution whose objective value is within ϵ of the infimum in (4), for a sufficiently small $\epsilon > 0$. Throughout the paper, we assume that the selected solution at time t is $\mathcal{F}_{t-1}$-measurable. The first term measures the average interval length and the second term measures the rescaled average miscoverage rate within the receding horizon. Here we assume $\beta_{t|t}, \ldots, \beta_{t+T-1|t}$ are independent so that the cost function only depends on the marginal distributions. This is not required for the coverage guarantee (see Theorem 3.1) but makes the cost-minimization problem more tractable.

Remark 3.1. For BCI, the above argument leads to a simple i.i.d. model for the dynamics of $\beta_{t+s|t}$ as well as a simple estimator – the empirical distribution of past β_ts. Thus, BCI does not rely on another machine learning algorithm to learn the dynamics. By contrast, even when the underlying prediction intervals are perfectly calibrated, the dynamics of the conformity scores remain unspecified. We present a theoretical result (Theorem 3.3) that justifies this choice.

Note that

$$\lambda_t \leq 0 \Longrightarrow \alpha^*_{s|t} = 1, \ s = t, \ldots, t+T-1, \tag{5}$$

since for a negative miscoverage penalty λ_t, the interval length and the negative miscoverage rate are both decreasing in $\alpha_{s|t}$, and hence the optimum is achieved at $\alpha_{s|t} = 1$. The analyst can apply any algorithm to minimize (4) as long as it satisfies (5). In particular, we discuss a DP-based algorithm in Section 3.3 that provides the exact solution and is computationally efficient for moderate values of T. Let $\alpha^*_{s|t}$ be the solution, which does not have to be exact. Clearly, $\alpha^*_{s|t}$ only depends on the model forecasts $(L_{s|t}(\cdot), F_{s|t})$ and the relative weight λ_t. The standard MPC would set $\alpha_t = \alpha^*_{t|t}$ and discard all other planned actions. To achieve the coverage guarantee in Theorem 3.1, BCI would modify α_t by considering both $\alpha^*_{t|t}$ and λ_t.

3.2 BCI UPDATE RULE FOR λ_t AND α_t

The key lever of BCI is the relative weight λ_t in (4). Intuitively, we shall increase λ_t to penalize the miscoverage term more when errors occur too often and decrease λ_t otherwise. Motivated by ACI Gibbs & Candès (2021) and the Rolling RC method Feldman et al. (2023), we update λ_t as follows:

$$\lambda_{t+1} = \lambda_t - \gamma[\overline{\alpha} - \mathrm{err}_t], \tag{6}$$

where γ is a user-defined step size that controls how fast the procedure adapts to past errors and $\mathrm{err}_t = 1(\alpha_t \geq \beta_t)$ as before.

As mentioned in Section 2, the safeguard property in Assumption 2.1 is crucial to guarantee long-term coverage even in the adversarial case. Here, we force α_t to be zero, so that $C_{t|t}(1-\alpha_t) = \mathcal{Y}$ under Assumption 2.1, when λ_t passes some pre-specified threshold $\lambda_{\max} > 0$:

$$\alpha_t = \begin{cases} 0, & \text{if } \lambda_t \geq \lambda_{\max}, \\ \alpha^\star_{t|t}, & \text{otherwise}, \end{cases} \tag{7}$$

where $\alpha^*_{t|t}$ is the minimizer of (4).

To match the scale, we can set $\gamma = c\lambda_{\max}$ for some $c \in (0, 1)$. Unlike ACI that updates α_t directly through a step of online gradient descent, BCI updates α_t indirectly through an online gradient descent-type update on λ_t. This allows BCI to take input other than past (α_t, β_t)'s.

3.3 SOLVING $\alpha^*_{t|t}$ VIA DYNAMIC PROGRAMMING

When the value of T is moderate (e.g. T=3), we can minimize (4) exactly and efficiently by DP. In fact, we do not recommend choosing a large T because the multi-step ahead prediction intervals tend to be less accurate or informative for events way into the future. For all our experiments, we choose $T = 3$.

The reader may refer to Appendix B.1 for some background on the usage of DP in solving stochastic control problems. To set up the DP algorithm for our problem, we define the state variable $\rho_{s|t}$ as

$$\rho_{t|t} = 0,\ \rho_{s|t} = \sum_{k=t}^{s-1} 1(\alpha_{k|t} \geq \beta_{k|t}),\ s = t+1, \ldots, t+T.$$

The dynamics for the state variable can be written as

$$\rho_{s+1|t} = \rho_{s|t} + 1(\alpha_{s|t} \geq \beta_{s|t}).$$

Clearly, $\rho_{s|t}$ takes values in $\{0, 1, \ldots, s-t\}$. DP then minimizes (4) in a backward fashion. We relegate the details of the algorithm to Appendix B.1. In a nutshell, the computational cost is $O(T^2)$, which is negligible for small T as in applications (e.g., $T = 3$).

We summarize BCI with the DP algorithm described in this subsection in Algorithm 1.

Algorithm 1 Bellman conformal inference at time t

1: **Hyperparameters:** target miscoverage level $\overline{\alpha}$; length of receding horizon T; maximum weight $\lambda_{\max}$; initial value $\lambda_1 \in [0, \lambda_{\max}]$; step size $\gamma = c\lambda_{\max}$ for $c \in (0, 1)$.
2: **Input:** Estimated marginal CDF of future uncalibrated PITs $F_{t|t}, \ldots, F_{t+T-1|t}$;
Length functions for multi-step ahead prediction intervals $L_{t|t}(\cdot), \ldots, L_{t+T-1|t}(\cdot)$;
Weight λ_{t-1} and miscoverage indicator err_{t-1} from the previous iteration.
3: **Step 1:** Update the security parameter as $\lambda_t = \lambda_{t-1} - \gamma[\overline{\alpha} - \mathrm{err}_{t-1}]$.
4: **Step 2:** Using $L_{s|t}$ and F_t, instantiate the stochastic control problem defined in (4) from Section 3.1.
5: **Step 3:** Apply the DP algorithm described in Section 3.3 to get $\tilde{\alpha}_{t+T-1|t}(\cdot), \ldots, \tilde{\alpha}_{t|t}(\cdot)$
6: **Output:** $\alpha_t = \tilde{\alpha}_{t|t}(0)$ if $\lambda_t \leq \lambda_{\max}$ and $\alpha_t = 0$ otherwise.

3.4 Validity Theory of BCI

We start by proving that BCI achieves approximate average coverage control for any large consecutive blocks. The proof for Theorem 3.1 can be found in Appendix C.

Theorem 3.1. *Consider Algorithm 1. Under Assumption 2.1, for any $m \geq 0$, almost surely,*

$$\left| \frac{1}{K} \sum_{t=m+1}^{m+K} \mathrm{err}_t - \overline{\alpha} \right| \leq \frac{c+1}{cK}. \tag{8}$$

In particular, (1) *holds by letting $m = 0$ and $K \to \infty$.*

Next, we consider the case where $(Y_t)_{t\geq 1}$ is a time-homogeneous Markov process of order p, i.e., $Y_{t+1} \perp (Y_{t-p}, Y_{t-p-1}, \ldots) \mid (Y_t, \ldots, Y_{t-p+1})$. In this case, we prove that the average coverage guarantee can be strengthened to a guarantee on marginal coverage $\mathbb{E}[\mathrm{err}_t]$ when $t \to \infty$, insofar as $\bar{\alpha} = 1/q$ for some integer $q \geq 2$ (e.g., $q = 10$ when $\bar{\alpha} = 0.1$). We present an informal version below and defer the formal statement and proofs to Appendix D.

Theorem 3.2 (Informal)**.** *Assume $(Y_t)_{t\geq 1}$ is a time-homogeneous Markov process of order p, $\bar{\alpha} = 1/q$ for some integer $q \geq 2$, and B is a constant. Under the technical assumptions in Appendix D.2, for any $\lambda_1 \in (0, \lambda_{\max})$, as $t \to \infty$*

$$\lim_{t\to\infty} \frac{1}{q} \sum_{j=0}^{q-1} \mathbb{E}[\mathrm{err}_{t+j}] = \bar{\alpha}.$$

Furthermore, the joint process of the p most recent observations, the B most recent uncalibrated PITs, and the weight admits an invariant distribution $\overline{\mu}$. If

$$\big((Y_0, Y_{-1}, \ldots, Y_{-(p-1)}), (\beta_{1-B}, \ldots, \beta_0), \lambda_1\big) \sim \overline{\mu}, \tag{9}$$

then $\mathbb{E}[\mathrm{err}_t] = \bar{\alpha}$ for every $t \geq 1$.

3.5 EFFICIENCY THEORY OF BCI

Next, we prove that modeling $(\beta_{t|t}, \ldots, \beta_{t+T-1|t})$ as i.i.d. draws from the empirical distribution $\hat{\mu}_t^{(B)}$ of $\beta_{t-1}, \ldots, \beta_{t-B}$ in the stochastic control problem yields an approximately optimal planned actions $\{\hat{\alpha}_{s|t} : s \in [t, t+T-1]\}$ when the nominal prediction intervals are approximately calibrated and B is large. The proof is presented in Appendix E.

Theorem 3.3. *Let $\mathcal{F}_{t-1}$ denote the information available before the time-t action is chosen. Define the true and estimated conditional distribution of future βs as*

$$P_t^\star := \mathrm{Law}\big((\beta_t, \beta_{t+1}, \ldots, \beta_{t+T-1}) \mid \mathcal{F}_{t-1}\big), \quad P_t^{(B)} := \big(\widehat{\mu}_t^{(B)}\big)^{\otimes T},$$

where $\widehat{\mu}_t^{(B)}$ denotes the empirical distribution of $(\beta_{t-1}, \ldots, \beta_{t-B})$ for some $B \geq 4$. Define $G_s^\circ(u) := \mathbb{P}(\beta_s \leq u \mid \mathcal{F}_{s-1})$ for $u \in [0,1]$. Assume

$$\sup_{u \in [0,1]} \big|G_s^\circ(u) - u\big| \leq \kappa_s, \tag{10}$$

for some deterministic, non-increasing, and non-negative sequence κ_s, and, for some constant $M > 0$,

$$\big|G_s^\circ(u) - G_s^\circ(v)\big| \leq M|u - v|, \quad \forall s \geq 1,\ \forall u, v \in [0,1]. \tag{11}$$

Further, let $J_t^P(a)$ denote the objective function (4) with $(\beta_{t|t}, \ldots, \beta_{t+T-1|t}) \sim P$ and $a = (\alpha_{t|t}, \ldots, \alpha_{t+T-1|t})$. Let $\hat{a}_t$ denote the BCI estimates $\mathrm{argmin}_a J_t^{\hat{P}_t^{(B)}}(a)$ and $a_t^\star$ the oracle $\mathrm{argmin}_a J_t^{\hat{P}_t^\star}(a)$. Then, with probability at least $1 - \delta$,

$$0 \leq J_t^{P_t^\star}(\widehat{a}_t) - J_t^{P_t^\star}(a_t^\star) \leq 4\lambda_t \left[\frac{\big(\sqrt{\log(B/\delta)} + M\big)}{\sqrt{B}} + 2\kappa_{t-B}. \right]. \tag{12}$$

In Appendix F, we consider a setting in which the expected interval lengths of ACI are provably larger than those of BCI. In this setting, we assume that both ACI and BCI have stationary trajectories. Such trajectories can be obtained by initializing the respective joint processes from invariant distributions, whose existence is established under the assumptions of Gibbs & Candès (2021, Appendix A.5) for ACI and Theorem 3.2 for BCI.

4 EMPIRICAL RESULTS

In this section, we present empirical experiments on real time series forecasting problems to demonstrate the effectiveness of BCI. We consider auto-regressive time series forecasting task for three datasets: daily Google trend popularity for keyword `deep learning`, daily stock return for companies `AMD, Amazon, Nvidia`, and stock volatility for the same companies. We use the ACI procedure (2) as a baseline for comparison. Making fair comparisons among ACI, BCI, and conformal PID is not straightforward due to the different structures of tuning parameters. To address this issue, we propose a matching method for fair comparison in Section 4.2.

4.1 DATASET AND MODEL FITTING

Return forecasting. Our first example studies the relative return of stock prices from various companies. We download the daily stock price of companies `Amazon, AMD, Nvidia` from the *Wall Street Journal Market Data*.[1] The length of history varies from company to company. That said, for all companies we have approximately 15 years of data with roughly 250 trading days each year. For each company, we compute the one-day log return Y_t of the stock as

$$Y_t = \log(P_t / P_{t-1}).$$

On each day t, we use the lagged sequence $\{Y_{t-1}, \ldots, Y_{t-100}\}$ as our predictors. We use a decoder-only transformer Vaswani et al. (2017) of embedding size 16 with 8 heads and 2 layers. This choice leads to a head size of $16/8 = 2$, which makes this a low-dimensional self-attention operation that

[1] `https://www.wsj.com/market-data`

is suitable for univariate time series. For the output, we use a linear layer to map the embedding to a $2 \times T$ dimensional vector, representing the predicted mean and standard deviation of the for the next T days. Mathematically, we can think of our transformer as a mapping from $\mathbb{R}^{100}$ to $\mathbb{R}^{2\times T}$, denoted by $[\mu_{s|t}(\mathbf{X}_t;\boldsymbol{\theta}), \sigma_{s|t}(\mathbf{X}_t;\boldsymbol{\theta})]$ for $s = t, \ldots, t+T-1$. Here, $\mathbf{X}_t = \{Y_{t-1}, \ldots, Y_{t-100}\}$ are the lagged returns and $\boldsymbol{\theta}$ denotes the weights of the neural network. At time t, we estimate $\boldsymbol{\theta}$ by solving the following optimization problem:

$$\min_{\boldsymbol{\theta}} \sum_{\tau<T} \sum_{s=\tau-T}^{\tau-1} -\log \mathcal{N}\left(Y_{t+s}; \mu_{s|t}(\mathbf{X}_t;\boldsymbol{\theta}), \sigma_{s|t}(\mathbf{X}_t;\boldsymbol{\theta})\right).$$

This formulation assumes the distribution of Y_{t+s} follows a normal distribution $\mathcal{N}\left(\mu_{s|t}(\mathbf{X}_t;\boldsymbol{\theta}), \sigma_{s|t}(\mathbf{X}_t;\boldsymbol{\theta})\right)$. It is a straightforward adaptation from the standard techniques for training variational auto-encoders Kingma & Welling (2014). Then we can construct $C_{s|t}(1-\beta)$ as $\left(\mu_{s|t}(\mathbf{X}_t;\boldsymbol{\theta}) - z_{1-\beta/2}\sigma_{s|t}(\mathbf{X}_t;\boldsymbol{\theta}),\ \mu_{s|t}(\mathbf{X}_t;\boldsymbol{\theta}) + z_{1-\beta/2}\sigma_{s|t}(\mathbf{X}_t;\boldsymbol{\theta})\right)$

Volatility forecasting. Our next example is the volatility forecasting problem explored in Gibbs & Candès (2021). We work with the same daily stock data as in the return forecasting problem. Instead of the log return, we compute the squared volatility as $Y_t = (P_t/P_{t-1} - 1)^2$. We then use the lagged volatility $\{Y_{t-1}, \ldots, Y_{t-100}\}$ as the predictors. Instead of using modern neural networks, we use the classical time series model `GARCH(1, 1)` Engle (1982b); Bollerslev (1986) that is often used to model stock price volatility. As we will see in Section 4.3, this classic model delivers nearly (marginally) calibrated prediction intervals. The `GARCH(1, 1)` model assumes that the squared volatility Y_s conditioned on the current observation follows a non-central χ-squared distribution and we can use its $\beta/2$- and $1-\beta/2$-th quantile to form the prediction interval $C_{s|t}(1-\beta)$. We provide a detailed description of the `GARCH(1, 1)` model in Appendix A.4.

Google trend popularity. Finally, we apply BCI to non-financial data. We consider the daily Google search popularity for the keyword `deep learning` from 2006 to 2011. For this dataset, we use the same neural-network-based fitting procedure as in the return forecasting task, except that we use a 5-layer LSTM Hochreiter & Schmidhuber (1997) recurrent network instead of the transformer to demonstrate the compatibility of BCI with a broader class of forecasters.

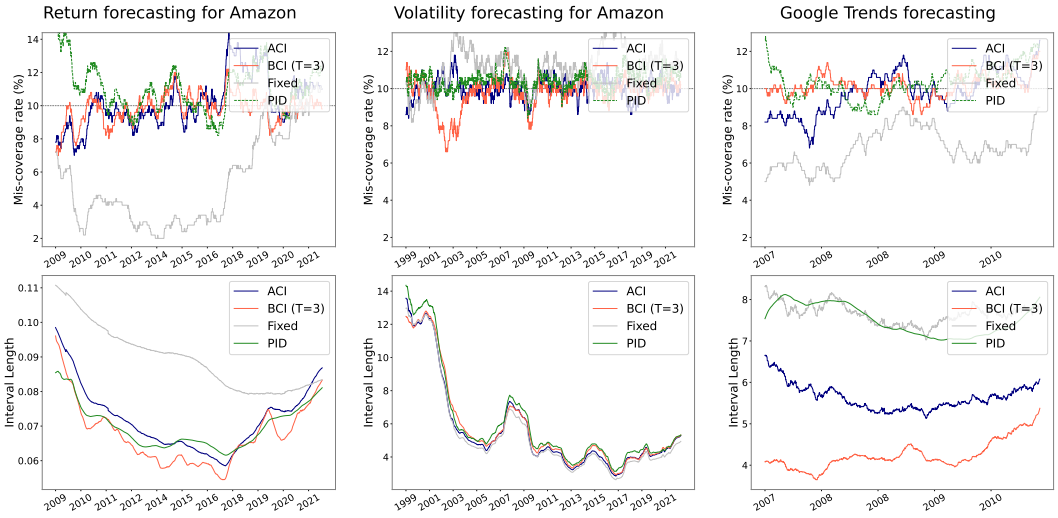

(a) Return forecasting (b) Volatility forecasting (c) Google trend

Figure 2: Same as Figure 1, except that the stepsize γ for ACI is 0.008 for looser control.

4.2 Performance Evaluation

We turn to discussing how we evaluate the performance of BCI. In particular, we point out caveats when comparing the performance of online prediction intervals and propose a fair evaluation.

Evaluation metrics. We evaluate the performance of online prediction intervals using two metrics: average miscoverage rate and average interval length. Following Gibbs & Candès (2021), we

measure the local average of both measures over a moving window of size 500:

$$\begin{bmatrix}\text{LocalMiscov}_t \\ \text{LocalLength}_t\end{bmatrix} = \frac{1}{500}\sum_{s=t-250}^{t+250}\begin{bmatrix}\text{err}_s \\ |C_s(1-\alpha_s)|\end{bmatrix}.$$

We report the standard deviation of the local miscoverage rate (LocalMiscov_t) in Tables 1 and 2.

An approach for fair comparison. For ACI, BCI, and conformal PID, the stepsize γ can trade off the tightness of the coverage and the rate of change of the interval length. Typically, a smaller stepsize means that the local miscoverage rate will exhibit larger excursions away from the target, hence looser control, and that the length of the prediction intervals will be smoother due to the smaller increment in α_t's. We illustrate this tradeoff in Figure 2. Ideally, we want to compare the average interval lengths when both methods achieve similar levels of coverage control. To make the comparison fair, we first choose a set of ACI parameters ($\gamma = 0.1$ for tight control and $\gamma = 0.008$ for loose control). For each choice of γ for ACI, we tune BCI and PID; see Appendix A.3. While the step size parameters are chosen post-hoc, length comparisons can be misleading without matching the tightness of miscoverate control due to the miscoverate-length tradeoff.

4.3 Empirical results

For each forecasting problem, we apply ACI with $\gamma = 0.1$ (tight coverage control) and $\gamma = 0.008$ (loose coverage control). We run BCI and conformal PID with stepsizes chosen through the grid search method discussed in the last subsection. We also include the naive benchmark with $\alpha_t = \bar{\alpha}$.

The time series of LocalMiscov_t and LocalLength_t are plotted in Figure 1 and Figure 2 for a subset of experiments. All three methods achieve comparable levels of coverage control across every combination of dataset and γ (Tables 1, 2). BCI achieves the shortest average interval length in most cases; ACI is shortest on volatility forecasting specifically, where the underlying GARCH model already produces nearly well-calibrated nominal intervals. Conformal PID does not achieve the shortest interval length in any of the tabulated combinations we test (Tables 1, 2); results on the additional datasets in Appendix A are qualitatively consistent.

Uninformative infinite-length intervals. One issue pointed out by Angelopoulos et al. (2023) is that ACI may generate infinite-length intervals. From Table 1, we observe that ACI generates a moderate fraction of infinite-length prediction intervals under tight coverage control. In contrast, BCI avoids infinite intervals entirely at our main horizon $T = 3$ and generally tends to produce shorter intervals (see the horizon-sensitivity discussion below for how this changes at larger T). Under loose control, neither ACI nor BCI generate any uninformative intervals.

Quality of nominal prediction intervals. From Figure 1 and Figure 2, we can see that BCI generally outperforms ACI. For return forecasting and Google trend forecasting, the gain is more prominent than for volatility forecasting. Intuitively, if the nominal prediction intervals are well-calibrated, BCI would not gain much from using the multi-step ahead prediction intervals. We confirm this intuition through the expected calibration curves (ECC) for each of the datasets; see Appendix A.1.

Sensitivity to Prediction Horizon. Finally, we investigate the sensitivity of BCI to the choice of the prediction horizon T. While our main results utilize $T = 3$, we also evaluated the framework across $T \in \{1, 3, 5, 7, 9\}$ and found that larger T yields only marginal further reduction in interval length but a substantially higher fraction of unbounded intervals (up to 5–6% for Google Trends and Volatility-AMD under tight control by $T = 7$–9, versus under 1% at $T = 3$), making $T = 3$ a reasonable balance between horizon length and interval boundedness.

5 Conclusion

We propose Bellman Conformal Inference (BCI) as an extension of ACI that calibrates nominal prediction intervals produced by any forecasting algorithms with coverage guarantees. BCI leverages multi-step ahead prediction intervals and applies Model Predictive Control (MPC) techniques to explicitly optimize the interval lengths. When the nominal prediction intervals are poorly calibrated, BCI improves substantially upon ACI in terms of the average interval lengths; otherwise, BCI has comparable performance to ACI.

AI USE STATEMENT

In this work, we used generative AI tools to aid programming, data visualization, checking the logical steps of our mathematical proofs, and polishing the writing. We did not use generative AI tools for research ideation or to generate synthetic datasets. We have reviewed all AI-assisted work; LLM-generated code was tested and verified for correctness by the authors, and AI-assisted proof-checking was independently verified against the formal proofs in the appendix. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

ETHICS STATEMENT

This research does not involve human subjects or crowdsourcing. All data used are publicly available: stock price data from the Wall Street Journal Market Data and search-popularity data from Google Trends; neither contains personally identifiable information. BCI is a general statistical methodology for uncertainty quantification and does not itself control any safety-critical system. Our coverage guarantees hold on average over time and do not imply uniform reliability across all subgroups or market regimes; practitioners should not rely on BCI as the sole basis for high-stakes decisions without appropriate domain expertise.

REPRODUCIBILITY STATEMENT

All training and evaluation details necessary to reproduce our experimental results are provided in Section 4.1 and Section A.3, including model architectures, dataset descriptions, data splits (the first 100 days of each series are discarded as burn-in), the hyperparameter search grids for BCI and PID, and the criteria used to select hyperparameters for a fair comparison across methods. Standard deviations of the local miscoverage rate are reported alongside all point estimates in Tables 1 and 2 as a measure of variability, computed as described in Section 4.2. All experiments were run on a single CPU; further compute details are given in Section A.3. Code and configuration files sufficient to reproduce all experiments will be made available upon publication.

REFERENCES

Anastasios N Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*, 2021.

Anastasios N Angelopoulos, Stephen Bates, Emmanuel J Candès, Michael I Jordan, and Lihua Lei. Learn then test: Calibrating predictive algorithms to achieve risk control. *arXiv preprint arXiv:2110.01052*, 2021.

Anastasios N Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. Conformal risk control. *arXiv preprint arXiv:2208.02814*, 2022.

Anastasios N Angelopoulos, Emmanuel J Candes, and Ryan J Tibshirani. Conformal PID control for time series prediction. *arXiv preprint arXiv:2307.16895*, 2023.

Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. *The Annals of Statistics*, 51(2):816–845, 2023.

Stephen Bates, Anastasios Angelopoulos, Lihua Lei, Jitendra Malik, and Michael Jordan. Distribution-free, risk-controlling prediction sets. *Journal of the ACM (JACM)*, 68(6):1–34, 2021.

Dimitri P. Bertsekas. *Dynamic Programming and Stochastic Control*. Academic Press, Inc., USA, 1976. ISBN 0120932504.

Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327, 1986.

Francesco Borrelli, Alberto Bemporad, and Manfred Morari. *Predictive Control for Linear and Hybrid Systems*. Cambridge University Press, 2017. doi: 10.1017/9781139061759.

George.E.P. Box and Gwilym M. Jenkins. *Time Series Analysis: Forecasting and Control*. Holden-Day, 1976.

Robert Goodell Brown. Smoothing, forecasting and prediction of discrete time series. In ., 1964. URL https://api.semanticscholar.org/CorpusID:120260777.

Emmanuel Candès, Lihua Lei, and Zhimei Ren. Conformalized survival analysis. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(1):24–45, 2023.

Victor Chernozhukov, Kaspar Wüthrich, and Zhu Yinchu. Exact and robust conformal inference methods for predictive machine learning with dependent data. In *Conference On learning theory*, pp. 732–749. PMLR, 2018.

Thomas G Dietterich and Jesse Hostetler. Conformal prediction intervals for markov decision process trajectories. *arXiv preprint arXiv:2206.04860*, 2022.

Robert Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982a. URL https://EconPapers.repec.org/RePEc:ecm:emetrp:v:50:y:1982:i:4:p:987-1007.

Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pp. 987–1007, 1982b.

Jianqing Fan and Qiwei Yao. *Nonlinear time series: nonparametric and parametric methods*, volume 20. Springer, 2003.

Shai Feldman, Liran Ringel, Stephen Bates, and Yaniv Romano. Achieving risk control in online learning settings, 2023.

Isaac Gibbs and Emmanuel Candès. Adaptive conformal inference under distribution shift. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/forum?id=6vaActvpcp3.

Isaac Gibbs and Emmanuel Candès. Conformal inference for online prediction with arbitrary distribution shifts, 2022.

Chirag Gupta, Arun K Kuchibhotla, and Aaditya Ramdas. Nested conformal prediction and quantile out-of-bag ensemble methods. *Pattern Recognition*, 127:108496, 2022.

Julien Herzen, Francesco LAossig, Samuele Giuliano Piazzetta, Thomas Neuer, LAoo Tafti, Guillaume Raille, Tomas Van Pottelbergh, Marek Pasieka, Andrzej Skrodzki, Nicolas Huguenin, Maxime Dumonal, Jan KoAcisz, Dennis Bader, FrACdAcrick Gusset, Mounir Benheddi, Camila Williamson, Michal Kosinski, Matej Petrik, and GaAl Grosch. Darts: User-friendly modern machine learning for time series. *Journal of Machine Learning Research*, 23(124):1–6, 2022. URL http://jmlr.org/papers/v23/21-1177.html.

Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8): 1735–1780, nov 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.

Oscar Jorda. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, March 2005. doi: 10.1257/0002828053828518. URL https://www.aeaweb.org/articles?id=10.1257/0002828053828518.

Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014.

Jing Lei, Alessandro Rinaldo, and Larry Wasserman. A conformal prediction approach to explore functional data. *Annals of Mathematics and Artificial Intelligence*, 74:29–43, 2015. doi: 10.1007/s10472-013-9366-6.

Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113 (523):1094–1111, 2018.

Lihua Lei and Emmanuel J Candès. Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(5):911–938, 2021.

Jordan Lekeufack, Anastasios N. Angelopoulos, Andrea Bajcsy, Michael I. Jordan, and Jitendra Malik. Conformal decision theory: Safe autonomous decisions from imperfect predictions, 2023.

Spyros Makridakis, Evangelos Spiliotis, and Vassilios Assimakopoulos. Statistical and machine learning forecasting methods: Concerns and ways forward. *PloS one*, 13(3):e0194889, 2018.

Roberto I Oliveira, Paulo Orenstein, Thiago Ramos, and João Vitor Romano. Split conformal prediction for dependent data. *arXiv preprint arXiv:2203.15885*, 2022.

Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *Machine Learning: ECML 2002: 13th European Conference on Machine Learning Helsinki, Finland, August 19–23, 2002 Proceedings 13*, pp. 345–356. Springer, 2002.

Dimitris N Politis and Kejin Wu. Multi-step-ahead prediction intervals for nonparametric autoregressions via bootstrap: Consistency, debiasing, and pertinence. *Stats*, 6(3):839–867, 2023.

David Salinas, Valentin Flunkert, Jan Gasthaus, and Tim Januschowski. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *International Journal of Forecasting*, 36(3):1181–1191, 2020. ISSN 0169-2070. doi: https://doi.org/10.1016/j.ijforecast.2019.07.001. URL https://www.sciencedirect.com/science/article/pii/S0169207019301888.

Craig Saunders, Alexander Gammerman, and Volodya Vovk. Transduction with confidence and credibility. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, IJCAI '99, pp. 722–726, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc. ISBN 1558606130.

Kamile Stankeviciute, Ahmed M Alaa, and Mihaela van der Schaar. Conformal time-series forecasting. *Advances in neural information processing systems*, 34:6216–6228, 2021.

James Stock and Mark Watson. *Dynamic Factor Models*, chapter ., pp. . Oxford University Press, Oxford, 2010. URL http://www.economics.harvard.edu/faculty/stock/files/dfm_oup_4.pdf.

Sophia Sun and Rose Yu. Copula conformal prediction for multi-step time series forecasting, 2023.

Sean J Taylor and Benjamin Letham. Forecasting at scale. *The American Statistician*, 72(1):37–45, 2018.

Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candes, and Aaditya Ramdas. Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Ł ukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.

Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer-Verlag, Berlin, Heidelberg, 2005. ISBN 0387001522.

Mike West and Jeff Harrison. *Bayesian forecasting and dynamic models*. Springer Science & Business Media, 2006.

Chen Xu and Yao Xie. Conformal prediction interval for dynamic time-series. In *International Conference on Machine Learning*, pp. 11559–11569. PMLR, 2021.

Chen Xu and Yao Xie. Sequential predictive conformal inference for time series. In *International Conference on Machine Learning*, pp. 38707–38727. PMLR, 2023.

Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *International Conference on Machine Learning*, pp. 25834–25866. PMLR, 2022.

A ADDITIONAL EXPERIMENTAL DETAILS

A.1 WHEN DOES BCI IMPROVE MOST UPON ACI?

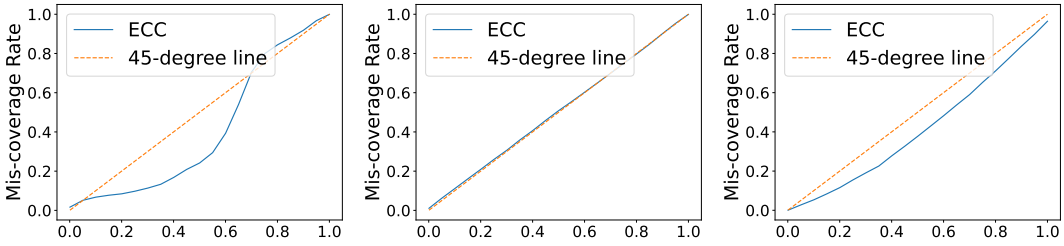

(a) Google trend forecasting (b) Volatility forecasting (c) Return forecasting

Figure 3: Expected calibration curve of nominal prediction intervals for Google trend, volatility, and return forecasting.

To examine the intuition that the amount of improvement of BCI over ACI depends on calibration of nominal prediction intervals, we define the expected calibration curve for nominal prediction intervals $C_t(\cdot)$ as follows:

$$\text{ecc}(\alpha) = \frac{1}{T}\sum_{t=1}^{T} 1\{Y_t \notin C_t(1-\alpha)\}.$$

If $\text{ecc}(\alpha) \approx \alpha$ for all $\alpha \in (0, 1)$, it means that the nominal prediction intervals are already well-calibrated.

From Figure 3, we can see that the Google trend forecasting and return forecasting with LSTM neural network is too conservative for small α and too aggressive for large α. In contrast, the `GARCH(1, 1)` generates nearly calibrated prediction invervals, leaving little possibility for BCI to alter the prediction intervals. This supports our intuition that BCI improves upon ACI more substantially when the nominal prediction intervals are poorly calibrated.

A.2 ADDITIONAL EXPERIMENTS

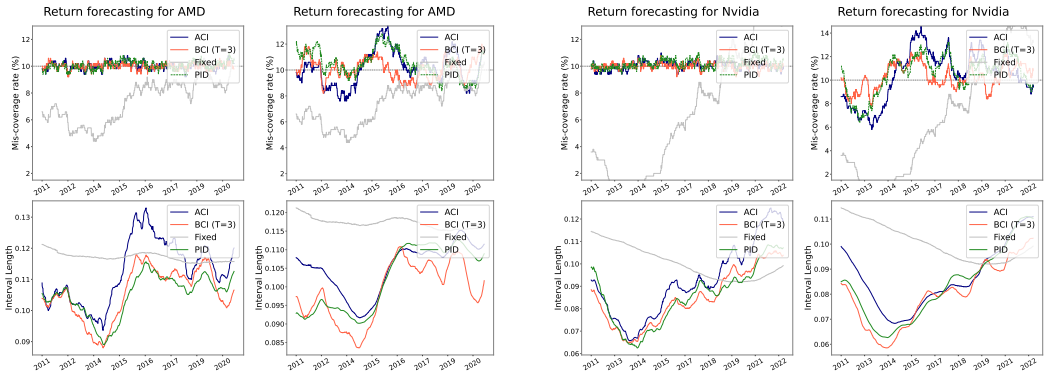

(a) Amazon tight v.s. loose control (b) Nvidia tight v.s. loose control

Figure 4: Additional return forecasting problems.

A.3 BCI AND PID TUNING FOR FAIR COMPARISON

Tuning procedure. A natural approach would be to tune each method to minimize average interval length, but this would unfairly favor methods that achieve coverage with narrower intervals by being less stable. Conversely, tuning purely on coverage stability would ignore interval efficiency entirely.

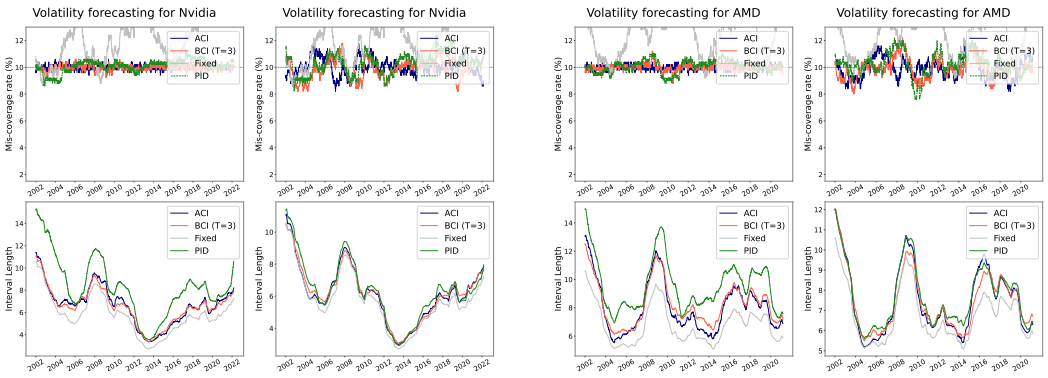

(a) Nvidia tight v.s. loose control

(b) AMD tight v.s. loose control

Figure 5: Additional volatility forecasting problems.

Table 1: Summary statistics for ACI with $\gamma = 0.1$ (tight control), BCI with $T = 3$, and PID. Miscoverage columns report mean $\pm$ std (%); length excludes ∞-sized intervals.

	Miscoverage (%) $\pm$ Std.			Avg. length			Frac. ∞ (%)		
Task/dataset	BCI	ACI	PID	BCI	ACI	PID	BCI	ACI	PID
Return-AMD	9.95 ± 0.256	9.95 ± 0.284	10.06 ± 0.441	0.1051	0.1128	0.1059	0.00	3.37	0.00
Return-Amazon	9.97 ± 0.270	9.99 ± 0.298	10.08 ± 0.246	0.0730	0.0787	0.0891	0.00	2.89	0.00
Return-Nvidia	10.04 ± 0.255	10.04 ± 0.287	10.21 ± 0.330	0.0865	0.0928	0.0886	0.00	4.29	0.00
Volatility-AMD	9.97 ± 0.235	9.99 ± 0.189	10.09 ± 0.498	8.3956	8.2062	9.7662	0.00	10.75	0.00
Volatility-Amazon	9.97 ± 0.168	9.99 ± 0.198	10.31 ± 0.288	6.7463	6.4999	7.0544	0.00	8.85	0.00
Volatility-Nvidia	10.01 ± 0.249	10.01 ± 0.190	10.14 ± 0.485	6.7998	6.9466	8.7391	0.00	10.71	0.00
Google Trend	10.36 ± 0.545	10.15 ± 0.377	10.92 ± 0.695	4.7895	6.3510	7.9425	0.00	2.81	0.00

Instead, we use the standard deviation of the local miscoverage rate as our tuning criterion. We define $\hat{\sigma} = \mathrm{std}(\mathrm{LocalMiscov}_t)$ as our stability metric. For a fixed choice of ACI stepsize γ (either 0.1 for tight control or 0.008 for loose control), we first run ACI and record its $\hat{\sigma}_{\mathrm{ACI}}$. We then tune BCI and PID to match $\hat{\sigma}_{\mathrm{ACI}}$ as closely as possible, subject to the constraint that miscoverage remains within $\pm 0.5\%$ of the target $\bar{\alpha} = 0.1$. This ensures that all methods operate at a comparable level of coverage stability, making interval length the primary axis of comparison.

BCI Tuning. BCI requires two hyperparameters: $\lambda_{\max}$ and c, where $\gamma = c \cdot \lambda_{\max}$.

To find $\lambda_{\max}$, we exploit the theoretical property that the optimal nominal miscoverage rate $\beta^*(\lambda)$, defined as

$$\beta^*(\lambda) = \arg\min_{\beta} \left[\bar{L}(\beta) + \lambda \cdot \max(\bar{m}(\beta) - \bar{\alpha},\, 0)\right],$$

where $\bar{L}(\beta)$ and $\bar{m}(\beta)$ are the average interval length and miscoverage rate over the first 100 days, converges to a floor value as $\lambda \to \infty$ – beyond which increasing λ no longer changes the solution. We conduct a geometric grid search over $\lambda \in [10^{-4}, 10^{6}]$ and identify the smallest λ^* at which $\beta^*(\lambda)$ has plateaued – specifically, the first λ^* such that $\beta^*(\lambda^*) \le 1.01 \cdot \beta^*(\lambda^{\mathrm{grid}}_{\max})$, where $\beta^*(\lambda^{\mathrm{grid}}_{\max})$ is the value at the right boundary of the grid. We set $\lambda_{\max} = \lambda^*$ and $\lambda_{\mathrm{init}} = 0.01 \cdot \lambda_{\max}$. This stage uses only the first 100 days and is discarded from all reported metrics.

To tune c, we conduct a geometric grid search over $c \in [10^{-4}, 1)$ with 40 points on the full dataset. Our tuning criterion is designed to be fair to BCI: we first identify all candidate values

$$\mathcal{C} = \{c : |\hat{m}(c) - \bar{\alpha}| \le 0.5\%\},$$

where $\hat{m}(c)$ is the empirical miscoverage rate on days 100+ and $\bar{\alpha} = 0.1$ is the target. Among $\mathcal{C}$, we select the c whose 50-day rolling miscoverage standard deviation is closest to that of ACI:

$$c^* = \arg\min_{c \in \mathcal{C}} |\hat{\sigma}(c) - \hat{\sigma}_{\mathrm{ACI}}|,$$

Table 2: Same as Table 1, except ACI uses $\gamma = 0.008$ (loose control).

	Miscoverage (%) ± Std.			Avg. length			Frac. ∞ (%)		
Task/dataset	BCI	ACI	PID	BCI	ACI	PID	BCI	ACI	PID
Return-AMD	10.20 ± 0.953	10.13 ± 1.284	10.83 ± 1.463	0.0985	0.1057	0.1017	0.00	0.00	0.00
Return-Amazon	9.53 ± 1.067	10.02 ± 1.541	11.47 ± 2.171	0.0712	0.0756	0.0720	0.00	0.00	0.00
Return-Nvidia	10.28 ± 1.390	10.07 ± 2.109	10.62 ± 2.212	0.0822	0.0896	0.0854	0.00	0.00	0.00
Volatility-AMD	9.92 ± 0.762	10.03 ± 0.803	10.31 ± 0.868	7.6328	7.6393	7.8388	0.00	0.00	0.00
Volatility-Amazon	9.88 ± 1.020	9.93 ± 0.559	10.42 ± 0.505	6.0761	6.0013	6.3920	0.00	0.00	0.00
Volatility-Nvidia	10.01 ± 0.829	9.96 ± 0.669	10.37 ± 0.885	6.4198	6.4208	6.6602	0.00	0.00	0.00
Google Trend	10.41 ± 0.730	10.10 ± 1.593	11.12 ± 1.881	4.4483	5.9183	7.6147	0.00	0.00	0.00

where $\hat{\sigma}(\cdot)$ denotes the standard deviation of the 50-day rolling miscoverage rate on days 100+. This criterion allows BCI to be as stable as it naturally wants to be while remaining calibrated, rather than forcing it into an artificially volatile regime to match ACI's standard deviation. If no candidate passes the coverage filter, the tolerance is relaxed to $\pm 1\%$.

BCI with $T = 3$ runs in approximately $O(T^2)$ operations per time step, making it computationally efficient. The full experimental pipeline, including hyperparameter tuning, completes within a few hours on a standard laptop CPU.

PID Tuning. PID requires two hyperparameters: the learning rate η and the integral gain K_I. The saturation parameter C_{sat} is set following Angelopoulos et al. (2023):

$$C_{\text{sat}} = \frac{2}{\pi}\left(\lceil \log(T) \cdot 0.05 \rceil - \frac{1}{\log T}\right),$$

where T is the dataset length.

A.4 STOCHASTIC MODEL FOR THE VOLATILITY SERIES

In the definition below, we introduce the GARCH(1, 1) process which we will later use as a stochastic model for the return series $r_1, \ldots, r_K$.

Definition A.1 (GARCH(1, 1))**.** Let $p_k(r_1, \ldots, r_k; \mu, \omega, a, b)$ be a joint probability density function of $r_1, \ldots, r_k$ with parameters (μ, ω, a, b). p_k is defined through the following sampling process: For each $k \geq 1$

1. Define ${\sigma_k}^2 = \omega + a{\epsilon_{k-1}}^2 + b{\sigma_{k-1}}^2$ with the convention that $\epsilon_0 = 0$ and $\sigma_0 = 0$.
2. Sample $e_k \sim \mathcal{N}(0, 1)$.
3. Set $\epsilon_k = \sigma_k e_k$ and $r_k = \epsilon_k + \mu$.

Forecasting with GARCH. Now we introduce how to use the GARCH model to generate prediction intervals. The stochastic model specifies the conditional distribution

$$r_{K+1}|r_1, \ldots, r_K \sim \mathcal{N}(\mu, \sigma_{K+1}^2).$$

Therefore, once we know the values of σ_{K+1} and μ, we know the conditional distribution $r_{K+1}|r_1, \ldots, r_K$. With a bit of algebra, we can show that

$$\sigma_{K+1}^2 = \omega \frac{1 - b^{K+1}}{1 - b} + a \sum_{k=1}^{K} b^{K-k}(r_k - \mu)^2$$

The unknown variables in the equation above are (μ, ω, a, b), which can be estimated by applying MLE on historical data $r_1, r_2, \ldots, r_K$:

$$(\hat{\mu}, \hat{\omega}, \hat{a}, \hat{b}) = \arg \max_{(\mu,\omega,a,b)} \log p_K(r_1, \ldots, r_K; \mu, \omega, a, b).$$

We use python package `arch`[2] to perform the fitting. This gives an estimate of the conditional distribution $r_{K+1}|r_1, \ldots, r_K \sim \mathcal{N}(\hat{\mu}, \hat{\sigma}_{K+1}^2)$, where

$$\hat{\sigma}_{K+1}^2 = \hat{\omega} \frac{1 - \hat{b}^{K+1}}{1 - \hat{b}} + \hat{a} \sum_{k=1}^{K} \hat{b}^{K-k}(r_k - \hat{\mu})^2.$$

Let $Q_{\mu,\sigma} : [0, 1] \to [0, \infty]$ be the quantile function of the squared normal distribution $\mathcal{N}(\mu, \sigma)^2$, meaning that

$$\mathbb{P}(\mathcal{N}(\mu, \sigma^2)^2 \leq Q_{\mu,\sigma}(1 - \beta)) = 1 - \beta.$$

Using $Q_{\mu,\sigma}$, a natural prediction interval for r_{K+1} is

$$\hat{C}_{K+1}(\beta) = \big(Q_{\hat{\mu},\hat{\sigma}_{K+1}^2}(\beta/2), Q_{\hat{\mu},\hat{\sigma}_{K+1}^2}(1 - \beta/2)\big).$$

Intuitively, we expect $\hat{C}_{K+1}(\beta)$ to have mis-coverage rate β.

B THE DYNAMIC PROGRAMMING ALGORITHM FOR $\alpha_{t|t}^*$

B.1 BACKGROUND ON STOCHASTIC CONTROL AND DYNAMIC PROGRAMMING

This is a self-contained section outlining some background on stochastic control Bertsekas (1976) needed for this paper. To be consistent with the control literature, we use the standard notation in this section. Consider a finite horizon stochastic system defined by the update rule

$$x_{t+1} = f_t(x_t, u_t, w_t) \text{ for } t = 1, 2, \ldots, T, \tag{13}$$

where x_t is the state variable, u_t is the control variable chosen by the analyst, and w_t is the random disturbance to the system at time t that affects the next state x_{t+1}. The control u_t can depend on any information available at time t but not on future knowledge. The disturbance w_t are sampled from

[2] `https://arch.readthedocs.io/en/latest/univariate/introduction.html`

$w_t \sim P(\cdot|x_t, u_t)$ and can depend on the state x_t and the control u_t. At each time t, the analyst bears the cost

$$g_t(x_t, u_t, w_t). \tag{14}$$

At the terminal time $T+1$, the analyst bears the cost $g_{T+1}(x_{T+1})$. The planning strategy is described by policy $\pi = \{\mu_t\}_{t=1,\dots,T}$, where each μ_t determines the control at time t through

$$u_t = \mu_t(x_t). \tag{15}$$

Given an initial state x_1, the expected cost under policy π is

$$J_\pi(x_1) = \mathbb{E}_{w_{1:T}}\left[g_{T+1}(x_{T+1}) + \sum_{t=1}^{T} g_t(x_t, \mu_t(x_t), w_t)\right]. \tag{16}$$

The problem of optimal control is to find the optimal policy

$$\pi^*_{x_1} = \arg\min_\pi J_\pi(x_1). \tag{17}$$

The optimal control problem admits exact solution through dynamic programming. To introduce the dynamic programming algorithm, we first define auxiliary functions

$$J_t(x_t) := \text{The optimal expected cost given that we start at state } x_t \text{ at time } t. \tag{18}$$

These functions are called "cost-to-go" function in dynamic programming literature. Under this definition, we have that

$$J_{T+1}(x_{T+1}) = g_{T+1}(x_{T+1}).$$

Now suppose that we have solved function J_{t+1} exactly. We can solve J_t according the update rule

$$J_t(x) = \min_u \mathbb{E}_{w\sim P(\cdot|x,u)}\left[g_t(x, u, w) + J_{t+1}(f_t(x, u, w))\right]. \tag{19}$$

Mathematically, the minimizer u^* of the program in (19) will depend on x, which defines the relation $u^* = \mu_t(x)$. A standard result that establishes the optimality of dynamic programming algorithm dictates the policy $\{\mu_t\}$ find through (19) is the same as the optimal policy in (17).

B.2 Algorithm for $\alpha^*_{t|t}$

Following the above DP terminology, the cost-to-go function at time $t + T$ is

$$J_{t+T|t}(\rho) = \lambda_t \max\left[\rho/T - \bar{\alpha}, 0\right], \rho \in \{0, \dots, T\}.$$

This measures the loss incurred in the last step when $\rho_{t+T|t} = \rho$. The cost-to-go functions for time $s = t + T - 1, t + T - 2, \dots, t$ are recursively defined through the Bellman equation: $J_{s|t}(\rho) =$

$$\min_{\alpha_{s|t}} \mathbb{E}_{\beta_{s|t}\sim F_{s|t}}\left[L_{s|t}(\alpha_{s|t}) + J_{s+1|t}(\rho + 1(\alpha_{s|t} \geq \beta_{s|t}))\right], \tag{20}$$

for $\rho \in \{0, \dots, s - t\}$. It measures the optimal cumulative loss incurred from time $s + 1$ onwards when $\rho_{s|t} = \rho$. Since $\beta_{s|t} \sim F_{s|t}$, we can rewrite $J_{s|t}(\rho)$ as

$$J_{s|t}(\rho) = \min_{\alpha_{s|t}} \Big\{L_{s|t}(\alpha_{s|t}) + J_{s+1|t}(\rho + 1)F_{s|t}(\alpha_{s|t}) + J_{s+1|t}(\rho)\left(1 - F_{s|t}(\alpha_{s|t})\right)\Big\}. \tag{21}$$

To simplify the expression, we define $D_{s|t}(\rho) = J_{s+1|t}(\rho + 1) - J_{s+1|t}(\rho)$. Then

$$J_{s|t}(\rho) = J_{s+1|t}(\rho) + \min_\alpha \left\{L_{s|t}(\alpha) + D_{s|t}(\rho)F_{s|t}(\alpha)\right\},$$

and the optimal policy at time s can be characterized as

$$\tilde{\alpha}_{s|t}(\rho) = \arg\min_\alpha \left\{L_{s|t}(\alpha) + D_{s|t}(\rho)F_{s|t}(\alpha)\right\}. \tag{22}$$

The following Proposition provides a useful property of the objective function in (22).

Proposition B.1. *At any $s \geq t$ and nonnegative $\lambda_t \geq 0$, the cost-to-go function $J_{s|t}(\rho)$ is nonnegative and non-decreasing in ρ.*

Proposition B.1 implies that the objective function $L_{s|t}(\alpha) + D_{s|t}(\rho)F_t(\alpha)$ is a sum of increasing function $D_{s|t}(\rho)F_t(\alpha)$ and a decreasing function $L_{s|t}(\alpha)$, which depicts the efficiency-coverage tradeoff discussed earlier.

Since the optimization problem in (22) is one-dimensional, it can be efficiently solved with a grid search. In addition, when we choose $F_{s|t}$ to be the empirical CDF of $\{\beta_{t-1}, \ldots, \beta_{t-B}\}$ as discussed earlier, it is easy to see that $\alpha^*_{s|t}(\rho) \in \{\beta_{t-1}, \ldots, \beta_{t-B}\}$ if $L_{s|t}(\cdot)$ is continuous. As a result, (22) can be solved exactly.

Finally, since $\rho_{t|t} = 0$, the optimal solution $\alpha^*_{t|t}$ in (7) is given by $\tilde{\alpha}_{t|t}(0)$. Since $\rho_{s|t}$ is integer-valued, the computation of $\tilde{\alpha}_{t|t}(0)$ involves solving $(T+1)+T+\ldots+1 = O(T^2)$ one-dimensional optimization problems in the form of (22). When T is moderate, it is computationally efficient.

B.3 Proof of Proposition B.1

Proof. We shall use the induction on s. When $s = t + T$,

$$J_{s|t}(\rho) = \lambda_t \max\left[\rho/T - \bar{\alpha}, 0\right] \tag{23}$$

is increasing. Now we assume that $J_{s+1}(\rho)$ is non-decreasing in ρ. Then for any ρ

$$\begin{aligned} &J_{s|t}(\rho+1) - J_{s|t}(\rho) \\ &\quad = \left[J_{s+1|t}(\rho+1) - J_{s+1|t}(\rho)\right] + \min_\alpha \left\{L_{s|t}(\alpha) + D_{s|t}(\rho+1)F_{s|t}(\alpha)\right\} - \min_\alpha \left\{L_{s|t}(\alpha) + D_{s|t}(\rho)F_{s|t}(\alpha)\right\}. \end{aligned}$$

Write $\tilde{\alpha}_{s|t}(\rho+1)$ as α' for notational convenience. Then

$$\min_\alpha \left\{L_{s|t}(\alpha) + D_{s|t}(\rho+1)F_{s|t}(\alpha)\right\} = L_{s|t}(\alpha') + D_{s|t}(\rho+1)F_{s|t}(\alpha'),$$

and

$$-\min_\alpha \left\{L_{s|t}(\alpha) + D_{s|t}(\rho)F_{s|t}(\alpha)\right\} \geq -\left\{L_{s|t}(\alpha') + D_{s|t}(\rho)F_{s|t}(\alpha')\right\}.$$

Therefore

$$\begin{aligned} J_{s|t}(\rho+1) - J_{s|t}(\rho) &\geq D_{s|t}(\rho) + \left\{L_{s|t}(\alpha') + D_{s|t}(\rho+1)F_{s|t}(\alpha')\right\} - \left\{L_{s|t}(\alpha') + D_{s|t}(\rho)F_{s|t}(\alpha')\right\}, \\ &\geq D_{s|t}(\rho)\left[1 - F_{s|t}(\alpha')\right] + D_{s|t}(\rho+1)F_{s|t}(\alpha'), \\ &\geq 0. \end{aligned}$$

Since ρ is arbitrary, $J_{s|t}(\rho)$ is increasing in ρ. This completes the proof. □

C Proof of Theorem 3.1

Proof. By (6),

$$\begin{aligned} \lambda_{m+K+1} - \lambda_{m+1} &= \gamma \sum_{t=m+1}^{m+K} (\mathrm{err}_t - \bar{\alpha}), \\ \left|\frac{1}{K}\sum_{t=m+1}^{m+K} \mathrm{err}_t - \bar{\alpha}\right| &= \frac{|\lambda_{m+K+1} - \lambda_{m+1}|}{K\gamma}. \end{aligned} \tag{24}$$

Now we prove that $\lambda_k \in [-\gamma\bar{\alpha}, \lambda_{\max} + \gamma(1-\bar{\alpha})]$ for all k by induction. For $k = 1$, this follows from $\lambda_1 \in [0, \lambda_{\max}]$. Suppose the claim holds for some $k \geq 1$.

1. If $\lambda_k < 0$, by (5), $\alpha_k = 1$ and hence $\mathrm{err}_k = 1$. As a result,
$$-\gamma\bar{\alpha} \leq \lambda_k < \lambda_{k+1} = \lambda_k + \gamma(1-\bar{\alpha}) < \gamma \leq \lambda_{\max}.$$

2. If $\lambda_k > \lambda_{\max}$, then $\alpha_k = 0$ and hence $\mathrm{err}_k = 0$. As a result,
$$\lambda_{\max} + \gamma(1-\bar{\alpha}) \geq \lambda_k > \lambda_{k+1} = \lambda_k - \gamma\bar{\alpha} > \lambda_{\max} - \gamma\bar{\alpha} \geq \lambda_{\max} - \gamma > 0.$$

3. If $\lambda_k \in [0, \lambda_{\max}]$, then we either subtract $\gamma\bar{\alpha}$ from λ_k or add $\gamma(1-\bar{\alpha})$ onto λ_k. This guarantees that $\lambda_{k+1} \in [-\gamma\bar{\alpha}, \lambda_{\max} + \gamma(1-\bar{\alpha})]$.

Thus the induction hypothesis holds for every $k \geq 1$, and

$$|\lambda_{m+K+1} - \lambda_{m+1}| \leq \lambda_{\max} + \gamma(1-\bar{\alpha}) + \gamma\bar{\alpha} = \lambda_{\max} + \gamma. \tag{25}$$

Combining (24) and (25) with $\gamma = c\lambda_{\max}$ gives (8).

□

D Formal statement and proof of Theorem 3.2

D.1 Setting

Fix integers $p \geq 1$, $B \geq 1$, and $\bar{\alpha} = 1/q$ for some integer $q \geq 2$. Fix $\lambda_{\max} > 0$ and $c \in (0,1)$, set $\gamma := c\lambda_{\max}$, and set

$$\alpha := \frac{1}{q}, \quad \Lambda := \left[-\frac{\gamma}{q},\, \lambda_{\max} + \gamma\Big(1 - \frac{1}{q}\Big)\right] = \lambda_{\max} \cdot \left[-\frac{c}{q},\, 1 + c\Big(1 - \frac{1}{q}\Big)\right]$$

Write

$$Z_t := (Y_t, \ldots, Y_{t-p+1}) \in \mathcal{Z} = \mathbb{R}^p.$$

We assume that (Y_t) is Markov of order p with one-step transition density $p_\star(\cdot \mid z)$ relative to Lebesgue measure, meaning that for every Borel set $A \subset \mathbb{R}$,

$$\mathbb{P}\big(Y_{t+1} \in A \mid Y_t, Y_{t-1}, \ldots\big) = \int_A p_\star(x \mid Z_t)\, dx.$$

At time t, the rolling window prediction algorithm uses Z_t and produces a one-step predictive CDF:

$$y \mapsto F(z, y), \quad z \in \mathcal{Z}.$$

Then the nominal prediction interval is given by

$$C(\beta) = \big(F^{-1}(z, \beta/2), F^{-1}(z, 1-\beta/2)\big).$$

Then

$$\beta_t := s(Z_t, Y_{t+1}),$$

where

$$s(z, y) := 2\min\{F(z, y), 1 - F(z, y)\} \in [0, 1], \quad (z, y) \in \mathcal{Z} \times \mathbb{R},$$

Define the sequence of β's as

$$\Gamma_t := (\beta_{t-B}, \ldots, \beta_{t-1}) \in [0,1]^B.$$

Let

$$a : \mathcal{Z} \times [0,1]^B \times \Lambda \to [0,1]$$

be the map from the input $(Z_t, \Gamma_t, \lambda_t)$ to the resulting solution α_{t+1}. Using the fact that both the length and miscoverage terms are monotone in $(\alpha_{s|t})_{s=t}^{t+T-1}$, we have (7):

$$a(z, \gamma, \lambda) = 1 \quad \text{if } \lambda \leq 0, \quad a(z, \gamma, \lambda) = 0 \quad \text{if } \lambda \geq \lambda_{\max}. \tag{26}$$

Define

$$\mathrm{err}_t := 1\{a(Z_t, \Gamma_t, \lambda_t) \geq \beta_t\}, \quad \lambda_{t+1} = \lambda_t + \gamma\left(\mathrm{err}_t - \frac{1}{q}\right), \quad t \geq 1. \tag{27}$$

Consider the concatenated state

$$X_t := (Z_t, \Gamma_t, \lambda_t) \in \mathsf{X} := \mathcal{Z} \times [0,1]^B \times \Lambda.$$

Since Z_t is Markov and (Γ_t, λ_t) is a function of Z_t, X_t is also Markov. Let P denote the transition kernel of $(X_t)_{t \geq 1}$.

For a weight $\lambda \in \Lambda$, define

$$\mathcal{E}(\lambda) := \begin{cases} \{1\}, & \lambda \le 0, \\ \{0,1\}, & 0 < \lambda < \lambda_{\max}, \\ \{0\}, & \lambda \ge \lambda_{\max}. \end{cases}$$

We say that a sequence $w_1, \dots, w_N \in \{0,1\}$ is admissible from λ if, after setting $\lambda_0 = \lambda$ and

$$\lambda_j = \lambda_{j-1} + \gamma \left(w_j - \frac{1}{q} \right), \qquad j = 1, \dots, N,$$

one has

$$w_j \in \mathcal{E}(\lambda_{j-1}) \text{ for every } j.$$

Equivalently, the sequence never asks for a bit that is forbidden by the clipped boundary rule.

Fix an initial value $\lambda_1 \in \Lambda$. For $r = 1, \dots, q$, define

$$\Lambda_r := \Lambda \cap \big(\lambda_1 - (r-1)\gamma/q + \gamma\mathbb{Z}\big), \tag{28}$$

$$\mathsf{X}_r := \mathcal{Z} \times [0,1]^B \times \Lambda_r, \tag{29}$$

where $\mathbb{Z}$ is the set of integers. We will show in Lemma D.2 that $\lambda_t \in \Lambda_r$ if $t - r$ is divisible by q.

For each $r \in \{1, \dots, q\}$, define the kernel Q_r on X_r by

$$Q_r(x, A) := P^q(x, A), \quad x \in \mathsf{X}_r, \quad A \subseteq \mathsf{X}_r \text{ measurable}.$$

Lastly, we define

$$\begin{aligned} T_* &:= \left\lceil q \left(\frac{1}{c} + 1 \right) \right\rceil, \quad L = \left\lceil \frac{1}{q}(p + B + T_*) \right\rceil, \\ n &= Lq - (p + B), \quad N := Lq = n + p + B. \end{aligned} \tag{30}$$

D.2 Assumptions

Assumption D.1. For every nonempty compact set $A \subset \mathbb{R}$,

$$\inf_{x \in A,\ z \in \mathcal{Z}} p_\star(x \mid z) > 0.$$

Assumption D.2. There exist a continuous strictly increasing CDF F_0 on $\mathbb{R}$ with $F_0(0) = 1/2$, functions $\widehat{m} : \mathcal{Z} \to \mathbb{R}$ and $\widehat{\sigma} : \mathcal{Z} \to (0, \infty)$, and constants

$$m_- \le m_+, \quad 0 < \sigma_- \le \sigma_+,$$

such that

$$F(z, y) = F_0 \left(\frac{y - \widehat{m}(z)}{\widehat{\sigma}(z)} \right), \quad m_- \le \widehat{m}(z) \le m_+, \quad \sigma_- \le \widehat{\sigma}(z) \le \sigma_+, \tag{31}$$

for every $z \in \mathcal{Z}$ and $y \in \mathbb{R}$.

Assumption D.3. There exists $0 < \ell_a \le h_a < 1$ such that

$$a(z, \gamma, \lambda) \in [\ell_a, h_a] \quad \text{for every } z \in \mathcal{Z},\ \gamma \in [0,1]^B,\ 0 < \lambda < \lambda_{\max}.$$

Remark D.1. Assumption D.3 can be satisfied by forcing α_t to be ℓ_a when $\lambda_t < \lambda_{\max}$ but the resulting DP solution is less than ℓ_a and h_a when $\lambda_t > 0$ but the resulting DP solution is less than ℓ_a. In real implementations, we can simply round λ to achieve it.

Assumption D.4. With the constants from Assumption D.2 and the value h_a from Assumption D.3,

$$m_+ + \sigma_- F_0^{-1}(h_a/2) < m_- + \sigma_- F_0^{-1}(1 - h_a/2). \tag{32}$$

D.3 Main result

With the assumptions, we state the formal version of Theorem 3.2.

Theorem D.1. *Assume Assumptions D.1, D.2, D.3, D.4, and a constant B. Let*

$$\eta := \mathrm{Law}(Z_1, \Gamma_1) \otimes \delta_{\lambda_1}, \quad \eta_r := \eta P^{r-1}, \quad r = 1, \dots, q.$$

Then the following hold.

1. *For each $r = 1, \dots, q$, there exists a measure μ_r on X_r such that*

   $$\mathrm{Law}(X_{mq+r}) \longrightarrow \mu_r \text{ in total variation} \quad \text{as } m \to \infty$$

   Moreover, these measures form a cyclic group:

   $$\mu_{r+1} = \mu_r P, \quad r = 1, \dots, q-1, \quad \mu_1 = \mu_q P.$$

   Consequently,

   $$\overline{\mu} := \frac{1}{q} \sum_{r=1}^{q} \mu_r$$

   is invariant for P.

2. *Let $\mathbb{E}_\lambda$ denote expectation with respect to X_t with $\lambda_1 = \lambda$. Then, for any $\lambda_1 \in (0, \lambda_{\max})$,*

   $$\lim_{t \to \infty} \frac{1}{q} \sum_{j=0}^{q-1} \mathbb{E}_{\lambda_1}[\mathrm{err}_{t+j}] = \frac{1}{q}.$$

   If $X_1 = (Z_1, \Gamma_1, \lambda_1) \sim \overline{\mu}$, then $(X_t)_{t \geq 1}$ is stationary. Under this stationary distribution,

   $$\mathbb{E}[\mathrm{err}_t] = \frac{1}{q} \quad \text{for every } t \geq 1.$$

D.4 Proof of Theorem D.1

D.4.1 Lemmas

We first show that $X_t \in \mathsf{X}_r$ whenever $t - r$ is divisible by q.

Lemma D.2. *Assume $\lambda_1 \in \Lambda$ is fixed. Then*

$$[\lambda_{t+1}]_\gamma = [\lambda_t - \gamma/q]_\gamma, \quad t \geq 1,$$

where $[\cdot]_\gamma$ denotes residual modulo γ. In particular,

$$[\lambda_{t+q}]_\gamma = [\lambda_t]_\gamma, \quad t \geq 1.$$

For the sets Λ_r and X_r in (28) *and* (29)*, each Λ_r is finite and*

$$P(x, \mathsf{X}_{r+1}) = 1 \quad \text{for } x \in \mathsf{X}_r, \quad r = 1, \dots, q-1, \quad P(x, \mathsf{X}_1) = 1 \quad \text{for } x \in \mathsf{X}_q. \tag{33}$$

Proof. Reducing (27) modulo γ gives

$$[\lambda_{t+1}]_\gamma = [\lambda_t + \gamma \mathrm{err}_t - \gamma/q]_\gamma = [\lambda_t - \gamma/q]_\gamma,$$

because $\mathrm{err}_t \in \{0, 1\}$. Iterating q times gives

$$[\lambda_{t+q}]_\gamma = [\lambda_t - \gamma]_\gamma = [\lambda_t]_\gamma.$$

Since Λ is bounded, each intersection Λ_r is finite. (33) follows from (27) and the definitions in (28) and (29). □

The next result introduces a useful formulation for the dynamic of (Z_t, Γ_t).

Lemma D.3. *There is a measurable map*

$$\Theta : \mathbb{R}^{p+B} \to \mathcal{Z} \times [0,1]^B$$

such that, for every integer $k \geq p + B$ and every starting state $x \in \mathsf{X}$,

$$(Z_{k+1}, \Gamma_{k+1}) = \Theta(Y_{k-p-B+2}, \ldots, Y_{k+1})$$

for every value of the displayed observations.

Proof. For $y = (y_1, \ldots, y_{p+B}) \in \mathbb{R}^{p+B}$ and $j = 1, \ldots, B+1$, define

$$z_j(y) := (y_{p+j-1}, y_{p+j-2}, \ldots, y_j) \in \mathcal{Z}.$$

For $j = 1, \ldots, B$, define

$$\theta_j(y) := s\big(z_j(y), y_{p+j}\big).$$

Set

$$\Theta(y) := \big(z_{B+1}(y), (\theta_1(y), \ldots, \theta_B(y))\big).$$

This map is measurable.

Now fix $k \geq p + B$. The terminal forecast state is

$$Z_{k+1} = (Y_{k+1}, \ldots, Y_{k-p+2}),$$

which is the last p entries of $(Y_{k-p-B+2}, \ldots, Y_{k+1})$. Thus,

$$\Gamma_{k+1} = (\beta_{k-B+1}, \ldots, \beta_k).$$

For $j = 1, \ldots, B$, the PIT value β_{k-B+j} is obtained by applying s to the rolling p-window immediately before $Y_{k-B+j+1}$ and to the next observation $Y_{k-B+j+1}$. These quantities are exactly encoded by $z_j(y)$ and y_{p+j}. Therefore the stated identity holds. □

The next result identifies two intervals of Y_{t+1} on which $s(z, y)$ is guaranteed to be small and large, respectively. This serves as a building block for the Doeblin minorization condition.

Lemma D.4. *Assume Assumptions D.2 and D.4. Then there exist compact nondegenerate intervals $G_1, G_0 \subset \mathbb{R}$ such that*

$$s(z, y) < \ell_a \quad \textit{for every } (z, y) \in \mathcal{Z} \times G_1,$$

and

$$s(z, y) > h_a \quad \textit{for every } (z, y) \in \mathcal{Z} \times G_0.$$

Proof. Define the open interval

$$\mathcal{C}_a := \left(m_+ + \sigma_- F_0^{-1}(h_a/2), m_- + \sigma_- F_0^{-1}(1 - h_a/2)\right).$$

Assumption D.4 says exactly that $\mathcal{C}_a$ is nonempty. Choose a compact interval of positive length

$$G_0 \subset \mathcal{C}_a.$$

Also choose a compact interval of positive length

$$G_1 \subset \left(m_+ + \sigma_+ F_0^{-1}(1 - \ell_a/2), \infty\right).$$

Let $y \in G_0$. By the definition of $\mathcal{C}_a$, for every $z \in \mathcal{Z}$,

$$y - \widehat{m}(z) > \sigma_- F_0^{-1}(h_a/2) \geq \widehat{\sigma}(z) F_0^{-1}(h_a/2),$$

because $F_0^{-1}(h_a/2) < 0$ and $\widehat{\sigma}(z) \geq \sigma_-$. Similarly,

$$y - \widehat{m}(z) < \sigma_- F_0^{-1}(1 - h_a/2) \leq \widehat{\sigma}(z) F_0^{-1}(1 - h_a/2),$$

because $F_0^{-1}(1 - h_a/2) > 0$ and $\widehat{\sigma}(z) \geq \sigma_-$. Hence

$$h_a/2 < F(z, y) < 1 - h_a/2, \quad z \in \mathcal{Z},$$

and therefore

$$s(z, y) = 2\min\{F(z, y), 1 - F(z, y)\} > h_a, \quad z \in \mathcal{Z},\ y \in G_0.$$

Now let $y \in G_1$. For every $z \in \mathcal{Z}$,

$$y - \widehat{m}(z) > \sigma_+ F_0^{-1}(1 - \ell_a/2) \geq \widehat{\sigma}(z) F_0^{-1}(1 - \ell_a/2),$$

because $F_0^{-1}(1 - \ell_a/2) > 0$ and $\widehat{\sigma}(z) \leq \sigma_+$. Thus $F(z, y) > 1 - \ell_a/2$, and hence

$$s(z, y) < \ell_a, \quad z \in \mathcal{Z},\ y \in G_1.$$

□

For the remainder of the proof, fix compact intervals $G_1, G_0 \subset \mathbb{R}$ supplied by Lemma D.4.1. The following lemma shows that G_1 and G_0 are able to pin down err_t.

Lemma D.5. *Assume Assumptions D.2, D.3, and D.4. Let $e \in \{0, 1\}$. Then*

$$Y_{t+1} \in G_e \Longrightarrow \mathrm{err}_t = e.$$

Proof. If $e = 1$, then $Y_{t+1} \in G_1$. If $\lambda_t \leq 0$, then $a(Z_t, \Gamma_t, \lambda_t) = 1$ by clipping and

$$\beta_t = s(Z_t, Y_{t+1}) \leq h_a < 1,$$

so $\mathrm{err}_t = 1$. If $e = 1$ and $0 < \lambda_t < \lambda_{\max}$, then Lemma D.4.1 implies

$$\beta_t < \ell_a \leq a(Z_t, \Gamma_t, \lambda_t),$$

so again $\mathrm{err}_t = 1$.

If $e = 0$, then $Y_{t+1} \in G_0$. If $\lambda_t \geq \lambda_{\max}$, clipping gives $a(Z_t, \Gamma_t, \lambda_t) = 0 < \beta_t$, hence $\mathrm{err}_t = 0$. If $0 < \lambda_t < \lambda_{\max}$, then

$$\beta_t = s(Z_t, Y_{t+1}) > h_a \geq a(Z_t, \Gamma_t, \lambda_t),$$

so $\mathrm{err}_t = 0$. □

Lemma D.4.1 also implies a weaker version of the Doeblin minorization condition for (Y_t).

Lemma D.6. *Under Assumption D.1, there exists a constant $c_G > 0$ such that, for every $z \in \mathcal{Z}$ and every Borel set $A \subset G_0 \cup G_1$,*

$$\mathbb{P}\big(Y_{t+1} \in A \mid Z_t = z\big) \geq c_G \mathrm{Leb}(A).$$

Proof. The set $G_0 \cup G_1$ is compact and nonempty. By Assumption D.1,

$$c_G := \inf_{x \in G_0 \cup G_1,\ z \in \mathcal{Z}} p_\star(x \mid z) > 0.$$

For every Borel set $A \subset G_0 \cup G_1$ and every $z \in \mathcal{Z}$,

$$\mathbb{P}\big(Y_{t+1} \in A \mid Z_t = z\big) = \int_A p_\star(x \mid z)\, dx \geq c_G \mathrm{Leb}(A),$$

which proves the claim. □

The one-step lower bound extends to blocks by repeated conditioning. This result gives a product Lebesgue lower bound for any finite block restricted to the $(G_0 \cup G_1)^k$.

Lemma D.7. *Under Assumption D.1, for every integer $k \geq 1$, every starting state $x \in \mathsf{X}$, and every Borel set $A \subset (G_0 \cup G_1)^k$,*

$$\mathbb{P}_x\big((Y_2, \ldots, Y_{k+1}) \in A\big) \geq c_G^k\, \mathrm{Leb}_k(A).$$

Proof. For rectangles $A = A_1 \times \cdots \times A_k$ with each $A_j \subset G_0 \cup G_1$, condition successively on the past before generating $Y_2, Y_3, \ldots, Y_{k+1}$. Lemma D.6 gives

$$\mathbb{P}(Y_{j+1} \in A_j \mid Y_j, \ldots, Y_2) \geq c_G \mathrm{Leb}(A_j)$$

at every step. Multiplying these conditional lower bounds yields

$$\mathbb{P}_x\big((Y_2, \ldots, Y_{k+1}) \in A\big) \geq c_G^k \prod_{j=1}^{k} \mathrm{Leb}(A_j) = c_G^k \mathrm{Leb}_k(A).$$

A monotone-class argument extends the bound to all Borel subsets of $(G_0 \cup G_1)^k$. □

The following two lemmas provide another building block for the Doeblin minorization condition.

Lemma D.8. *For $\lambda \in \Lambda$, define*

$$d(\lambda) := \begin{cases} 1, & \lambda \le 0, \\ 0, & \lambda > 0, \end{cases} \quad D(\lambda) := \lambda + \gamma \left(d(\lambda) - \frac{1}{q} \right).$$

For every Λ_r and every $\lambda, \lambda' \in \Lambda_r$,

$$D^t(\lambda) = D^t(\lambda') \quad \textit{for every } t \ge T_*.$$

Proof. First, D maps Λ into itself. If $\lambda > 0$, then $D(\lambda) = \lambda - \gamma/q \ge -\gamma/q$. If $\lambda \le 0$, then $\lambda \in [-\gamma/q, 0]$ and

$$D(\lambda) \in \left[\gamma \left(1 - \frac{2}{q} \right), \gamma \left(1 - \frac{1}{q} \right) \right] \subset \Lambda.$$

Normalize λ by setting

$$w := \frac{q\lambda}{\gamma}.$$

The interval Λ becomes $[-1, U]$, where

$$U := \frac{q}{\gamma} \left\{ \lambda_{\max} + \gamma \left(1 - \frac{1}{q} \right) \right\} = q \left(\frac{1}{c} + 1 \right) - 1.$$

Define the following map

$$w \mapsto \begin{cases} w + q - 1, & w \le 0, \\ w - 1, & w > 0. \end{cases}$$

The half-open interval $(-1, q-1]$ is invariant under this map. Also, every trajectory starting in $[-1, U]$ enters $(-1, q-1]$ by time $\lceil U \rceil + 1 = \left\lceil q \left(\frac{1}{c} + 1 \right) \right\rceil = T_*$: while $w > 0$ the map subtracts one, and after the first nonpositive value one more step sends the point into $(-1, q-1]$.

If $\lambda, \lambda' \in \Lambda_r$, then their normalized values differ by an integer multiple of q. The above map preserves the property that this difference is an integer multiple of q. Once both normalized trajectories lie in $(-1, q-1]$, two different points in that interval cannot differ by a nonzero integer multiple of q, because the interval has length q and excludes its left endpoint. Therefore the two normalized values must be equal. This proves the claim. □

Lemma D.9. *For each $r \in \{1, \dots, q\}$, there exist a common entrance level $\widehat{\lambda}_r \in \Lambda$, a reference level $\overline{\lambda}_r \in \Lambda_r$, and a common final sequence*

$$c_{r,1}, \dots, c_{r,p+B} \in \{0, 1\}$$

with the following property. For every starting level $\lambda \in \Lambda_r$, there is a prefix sequence

$$d_{r,\lambda,1}, \dots, d_{r,\lambda,n} \in \{0, 1\}$$

such that the full sequence

$$w_{r,\lambda} := (d_{r,\lambda,1}, \dots, d_{r,\lambda,n}, c_{r,1}, \dots, c_{r,p+B})$$

is admissible from λ. If

$$\lambda_0^{r,\lambda} := \lambda, \quad \lambda_j^{r,\lambda} := \lambda_{j-1}^{r,\lambda} + \gamma \left(w_{r,\lambda,j} - \frac{1}{q} \right), \quad j = 1, \dots, N,$$

then

$$w_{r,\lambda,j} \in \mathcal{E}(\lambda_{j-1}^{r,\lambda}) \quad \textit{for every } j = 1, \dots, N,$$

and

$$\lambda_n^{r,\lambda} = \widehat{\lambda}_r, \quad \lambda_N^{r,\lambda} = \overline{\lambda}_r.$$

Proof. Fix $r \in \{1, \dots, q\}$. For the prefix, set

$$\lambda_0^{r,\lambda} := \lambda, \quad d_{r,\lambda,j} := d(\lambda_{j-1}^{r,\lambda}), \quad \lambda_j^{r,\lambda} := D(\lambda_{j-1}^{r,\lambda}), \quad j = 1, \dots, n.$$

Set $w_{r,\lambda,j} := d_{r,\lambda,j}$ for $j = 1, \dots, n$. Each prefix bit is admissible by definition of d. Since $n \geq T_*$, Lemma D.8 implies that $\lambda_n^{r,\lambda}$ is the same for every starting $\lambda \in \Lambda_r$. Denote this common value by $\widehat{\lambda}_r$.

Starting from $\widehat{\lambda}_r$, continue the same rule for the final $p + B$ steps:

$$\widehat{\lambda}_{r,0} := \widehat{\lambda}_r, \quad c_{r,j} := d(\widehat{\lambda}_{r,j-1}), \quad \widehat{\lambda}_{r,j} := D(\widehat{\lambda}_{r,j-1}), \quad j = 1, \dots, p + B.$$

Set $\overline{\lambda}_r := \widehat{\lambda}_{r,p+B}$. For $j = n + 1, \dots, N$, define

$$w_{r,\lambda,j} := c_{r,j-n}, \quad \lambda_j^{r,\lambda} := \widehat{\lambda}_{r,j-n}.$$

The final sequence $c_{r,1}, \dots, c_{r,p+B}$ is common across all starting levels in Λ_r. The full sequence is admissible at every step, reaches $\widehat{\lambda}_r$ after the prefix, and reaches $\overline{\lambda}_r$ after $N = Lq$ steps. Since Lq is a multiple of q, Lemma D.2 implies $\overline{\lambda}_r \in \Lambda_r$. □

The preceding ingredients combine to give the uniform small-set condition on each X_r. This proposition constructs the minorizing measure for $Q_r^L = P^{Lq}$.

Proposition D.1. *Assume Assumptions D.1, D.2, D.3, and D.4. Fix $r \in \{1, \dots, q\}$. Use the levels $\lambda_j^{r,\lambda}$ and bits $w_{r,\lambda,j}$ supplied by Lemma D.9. Define*

$$S_{r,\lambda} := \prod_{j=1}^{n} G_{w_{r,\lambda,j}} \subset (G_0 \cup G_1)^n,$$

with the convention that $\mathrm{Leb}_0(S_{r,\lambda}) = 1$ when $n = 0$. Define the common tail levels by

$$\widehat{\lambda}_{r,0} := \widehat{\lambda}_r, \quad \widehat{\lambda}_{r,j} := \widehat{\lambda}_{r,j-1} + \gamma \left(c_{r,j} - \frac{1}{q} \right), \quad j = 1, \dots, p + B,$$

and the common tail box

$$T_r := \prod_{j=1}^{p+B} G_{c_{r,j}} \subset (G_0 \cup G_1)^{p+B}.$$

Let ν_r be the image of normalized Lebesgue measure on T_r under

$$y \mapsto (\Theta(y), \overline{\lambda}_r).$$

Equivalently,

$$\nu_r(A) := \frac{1}{\mathrm{Leb}_{p+B}(T_r)} \mathrm{Leb}_{p+B}\big(\{y \in T_r : (\Theta(y), \overline{\lambda}_r) \in A\}\big).$$

Set

$$\delta_r := c_G^N \, \mathrm{Leb}_{p+B}(T_r) \min_{\lambda \in \Lambda_r} \mathrm{Leb}_n(S_{r,\lambda}).$$

Then $\delta_r > 0$ and

$$Q_r^L(x, A) = P^{Lq}(x, A) \geq \delta_r \, \nu_r(A) \quad \textit{for every } x \in \mathsf{X}_r \textit{ and measurable } A \subseteq \mathsf{X}_r. \tag{34}$$

Proof. Fix $r \in \{1, \dots, q\}$ and $x = (z, \gamma, \lambda) \in \mathsf{X}_r$. Let

$$D_{r,\lambda} := S_{r,\lambda} \times T_r \subset (G_0 \cup G_1)^N.$$

On the event

$$(Y_2, \dots, Y_{N+1}) \in D_{r,\lambda},$$

the observation at step j belongs to $G_{w_{r,\lambda,j}}$ for $j = 1, \dots, N$. By Lemma D.5 and the admissibility of the sequence,

$$\mathrm{err}_j = w_{r,\lambda,j} \quad \text{and} \quad \lambda_{j+1} = \lambda_j^{r,\lambda}, \quad j = 1, \dots, N.$$

Thus $\lambda_{N+1} = \lambda_N^{r,\lambda} = \overline{\lambda}_r$.

Since $N \geq p + B$, Lemma D.3 gives

$$(Z_{N+1}, \Gamma_{N+1}) = \Theta(Y_{N-p-B+2}, \ldots, Y_{N+1}).$$

The last $p + B$ observations are exactly the coordinates in the common tail box T_r. Therefore, for any measurable $A \subseteq \mathsf{X}_r$,

$$P^N(x, A) \geq \mathbb{P}_x\Big((Y_2, \ldots, Y_{N+1}) \in E_{r,\lambda}(A)\Big),$$

where

$$E_{r,\lambda}(A) := \Big\{(u, y) \in S_{r,\lambda} \times T_r : (\Theta(y), \overline{\lambda}_r) \in A\Big\}.$$

By Lemma D.7,

$$P^N(x, A) \geq c_G^N \mathrm{Leb}_N(E_{r,\lambda}(A)).$$

Fubini's theorem gives

$$\mathrm{Leb}_N(E_{r,\lambda}(A)) = \mathrm{Leb}_n(S_{r,\lambda})\mathrm{Leb}_{p+B}\big(\{y \in T_r : (\Theta(y), \overline{\lambda}_r) \in A\}\big).$$

Using the definition of ν_r and taking the minimum over the finite set Λ_r yields

$$P^N(x, A) \geq c_G^N \mathrm{Leb}_{p+B}(T_r) \min_{\lambda' \in \Lambda_r} \mathrm{Leb}_n(S_{r,\lambda'})\, \nu_r(A) = \delta_r \nu_r(A).$$

Because G_0 and G_1 are nondegenerate intervals, every factor in each box has positive Lebesgue measure. Since Λ_r is finite, $\delta_r > 0$. □

D.4.2 Proof of Theorem D.1

For Part 1, fix $r \in \{1, \ldots, q\}$. Let δ_r be the minorization constant from Proposition D.1. Proposition D.1 gives

$$Q_r^L(x, \cdot) \geq \delta_r \nu_r(\cdot) \quad \text{for every } x \in \mathsf{X}_r.$$

Lemma D.10, applied to $Q = Q_r$, gives a unique invariant probability measure μ_r for Q_r and the stated total-variation bound. Since

$$\mathrm{Law}(X_{mq+r}) = \eta_r Q_r^m,$$

we can define μ_r as its limit in total variation.

When $r < q$, Lemma D.2 implies that $\mu_r P$ is supported on X_{r+1}. Moreover,

$$(\mu_r P)Q_{r+1} = \mu_r P^{q+1} = \mu_r Q_r P = \mu_r P,$$

so $\mu_r P$ is invariant for Q_{r+1}. By uniqueness of the invariant law on slice $r + 1$,

$$\mu_r P = \mu_{r+1}.$$

The same argument with $r = q$ gives $\mu_q P = \mu_1$. Averaging the cyclic identities gives

$$\overline{\mu}P = \frac{1}{q}\sum_{r=1}^{q} \mu_r P = \frac{1}{q}\sum_{r=1}^{q} \mu_r = \overline{\mu}.$$

Thus $\overline{\mu}$ is invariant for P.

For Part 2, start the chain with $X_1 \sim \overline{\mu}$. Part 1 implies that $(X_t)_{t \geq 1}$ is stationary. Since λ_t is bounded almost surely, taking expectations in the bounded update

$$\lambda_{t+1} - \lambda_t = \gamma\left(\mathrm{err}_t - \frac{1}{q}\right)$$

under this stationary distribution gives

$$0 = \mathbb{E}[\lambda_{t+1} - \lambda_t] = \gamma\left(\mathbb{E}[\mathrm{err}_t] - \frac{1}{q}\right),$$

so $\mathbb{E}[\mathrm{err}_t] = 1/q$ for every t.

D.4.3 DOEBLIN MINORIZATION CONDITION

In this subsection, we revisit Doeblin minorization condition and its implication on the convergence to the stationary distribution with any initialization. We provide a self-contained proof for completeness.

Lemma D.10. *Let Q be a Markov kernel on a measurable space E. Suppose that for some integer $L \geq 1$, some $\delta \in (0, 1]$, and some probability measure ν on E,*

$$Q^L(x, A) \geq \delta\, \nu(A) \quad \textit{for every } x \in E \textit{ and measurable } A \subseteq E. \tag{35}$$

Then Q has a unique invariant probability measure μ, and for every probability measure η on E,

$$\|\eta Q^n - \mu\|_{\mathrm{TV}} \leq (1 - \delta)^{\lfloor n/L \rfloor}, \quad n \geq 0.$$

Proof. First apply the usual Doeblin decomposition to the kernel $\widetilde{Q} := Q^L$. If $\delta = 1$, then $\widetilde{Q}(x, \cdot) = \nu(\cdot)$ for every x and the contraction below is immediate. If $0 < \delta < 1$, define

$$R(x, A) := \frac{\widetilde{Q}(x, A) - \delta\nu(A)}{1 - \delta}.$$

Then R is a Markov kernel and

$$\widetilde{Q}(x, \cdot) = \delta\nu(\cdot) + (1 - \delta)R(x, \cdot).$$

Therefore, for any two probability measures η_1, η_2,

$$\eta_1\widetilde{Q} - \eta_2\widetilde{Q} = (1 - \delta)(\eta_1 R - \eta_2 R),$$

so

$$\|\eta_1\widetilde{Q} - \eta_2\widetilde{Q}\|_{\mathrm{TV}} \leq (1 - \delta)\|\eta_1 - \eta_2\|_{\mathrm{TV}}. \tag{36}$$

Thus $\eta \mapsto \eta\widetilde{Q}$ is a strict contraction on probability measures with the total-variation metric. Hence $\widetilde{Q}$ has a unique invariant probability measure μ and

$$\|\eta\widetilde{Q}^k - \mu\|_{\mathrm{TV}} \leq (1 - \delta)^k, \quad k \geq 0.$$

It remains to check that this same μ is invariant for Q. Since $\mu\widetilde{Q} = \mu$,

$$(\mu Q)\widetilde{Q} = \mu Q^{L+1} = \mu Q^L Q = \mu Q.$$

Thus μQ is also invariant for $\widetilde{Q}$. By uniqueness for $\widetilde{Q}$, $\mu Q = \mu$. Therefore μ is invariant for Q. If $\pi Q = \pi$, then $\pi\widetilde{Q} = \pi$, so uniqueness for $\widetilde{Q}$ also gives $\pi = \mu$. Thus μ is the unique invariant law for Q.

Finally write $n = kL + j$ with $0 \leq j < L$. Since μ is invariant for Q and hence for $\widetilde{Q}$,

$$\eta Q^n = (\eta Q^j)\widetilde{Q}^k,$$

and the contraction bound for $\widetilde{Q}$ gives

$$\|\eta Q^n - \mu\|_{\mathrm{TV}} = \|(\eta Q^j)\widetilde{Q}^k - \mu\|_{\mathrm{TV}} \leq (1 - \delta)^k.$$

Since $k = \lfloor n/L \rfloor$, the claim follows. □

E PROOF OF THEOREM 3.3

E.1 NOTATION

Fix a time t, a planning horizon $T \geq 1$, and a historical window length $B \geq 4$ with $t \geq B + 1$. For a planned action vector

$$a = (a_0, \ldots, a_{T-1}) \in [0, 1]^T,$$

define the deterministic length term

$$\ell_t(a) := \sum_{h=0}^{T-1} L_{t+h|t}(a_h),$$

where $L_{t+h|t}(a_h)$ denotes the nominal prediction interval length used by BCI at horizon h.

Let

$$N_a(z) := \sum_{h=0}^{T-1} 1\{z_h \le a_h\}, \qquad z = (z_0, \dots, z_{T-1}) \in [0,1]^T,$$

and define the horizon penalty

$$\phi_{a,\bar{\alpha}}(z) := \Big(\frac{N_a(z)}{T} - \bar{\alpha}\Big)_+,$$

Then we can write $J_t^P(a)$ as

$$J_t^P(a) := \ell_t(a) + \lambda_t \int \phi_{a,\bar{\alpha}}(z)\, P(dz). \tag{37}$$

For $h = 0, \dots, T-1$, define the future predictive CDF

$$G^\star_{t+h|t}(u) := \mathbb{P}(\beta_{t+h} \le u \mid \mathcal{F}_{t-1}).$$

We also write

$$\widehat{G}_t^{(B)}(u) := \frac{1}{B} \sum_{j=1}^{B} 1\{\beta_{t-j} \le u\}$$

for the empirical CDF of $\hat{\mu}_t^{(B)}$, and

$$\overline{G}_{t,B}(u) := \frac{1}{B} \sum_{j=1}^{B} G^\circ_{t-j}(u)$$

for the averaged one-step conditional CDF over the same window.

E.2 Lemmas

The first lemma is the standard argument in the empirical process theory. It converts uniform closeness of a surrogate objective to the true objective into a regret bound for any exact or approximate surrogate optimizer.

Lemma E.1. *Let $\mathcal{A}$ be a nonempty action set and let $J, \widehat{J} : \mathcal{A} \to \mathbb{R}$ be two objectives. Suppose $a^\star \in \operatorname{argmin}_{a\in\mathcal{A}} J(a)$, and suppose $\widetilde{a} \in \mathcal{A}$ satisfies*

$$\widehat{J}(\widetilde{a}) \le \inf_{a\in\mathcal{A}} \widehat{J}(a) + \eta$$

for some $\eta \ge 0$. Then

$$0 \le J(\widetilde{a}) - J(a^\star) \le 2 \sup_{a\in\mathcal{A}} |J(a) - \widehat{J}(a)| + \eta. \tag{38}$$

Proof. The lower bound follows from the optimality of $a^\star$ for J. For the upper bound, write

$$\begin{aligned} J(\widetilde{a}) - J(a^\star) &= \big[J(\widetilde{a}) - \widehat{J}(\widetilde{a})\big] + \big[\widehat{J}(\widetilde{a}) - \widehat{J}(a^\star)\big] + \big[\widehat{J}(a^\star) - J(a^\star)\big] \\ &\le \sup_{a\in\mathcal{A}} |J(a) - \widehat{J}(a)| + \eta + \sup_{a\in\mathcal{A}} |J(a) - \widehat{J}(a)|. \end{aligned}$$

This proves (38). □

The next lemma shows that, for the BCI horizon penalty, replacing the true joint future law by the product of its predictive marginals which ignores the dependence, costs at most the accumulated future one-step calibration error.

Lemma E.2. *Define the product law of the true predictive marginals by*

$$\overline{P}_t^\star := \bigotimes_{h=0}^{T-1} \mu_{t+h|t}^\star,$$

where $\mu_{t+h|t}^\star$ denotes the one-dimensional marginal law of β_{t+h} under $P_t^\star$. Assume (10)*. Then*

$$\Delta_t^{\mathrm{dep}} := \sup_{a\in[0,1]^T} \left| \int \phi_{a,\bar\alpha}(z)\, P_t^\star(dz) - \int \phi_{a,\bar\alpha}(z)\, \overline{P}_t^\star(dz) \right| \le 2\kappa_{t+1}. \tag{39}$$

Proof. Fix $a = (a_0, \dots, a_{T-1}) \in [0,1]^T$, and define

$$I_h := 1\{\beta_{t+h} \le a_h\}, \qquad h = 0, \dots, T-1.$$

Write

$$q_h := \mathbb{P}(I_h = 1 \mid \mathcal{F}_{t-1}) = G_{t+h|t}^\star(a_h).$$

For $h \ge 1$ and for each realized history $i_{0:h-1} = (i_0, \dots, i_{h-1}) \in \{0,1\}^h$ with positive conditional probability, define

$$p_h(i_{0:h-1}) := \mathbb{P}\big(I_h = 1 \mid \mathcal{F}_{t-1}, I_0 = i_0, \dots, I_{h-1} = i_{h-1}\big),$$

and define it arbitrarily on null histories. Set $p_0 := q_0$.

Let $U_0, \dots, U_{T-1}$ be i.i.d. $\mathrm{Unif}[0,1]$, independent of everything else conditionally on $\mathcal{F}_{t-1}$. Define a coupled independent Bernoulli sequence

$$J_h := 1\{U_h \le q_h\}, \qquad h = 0, \dots, T-1,$$

and recursively define a coupled version of the true indicator process by

$$\widetilde{I}_0 := 1\{U_0 \le p_0\}, \qquad \widetilde{I}_h := 1\big\{U_h \le p_h(\widetilde{I}_0, \dots, \widetilde{I}_{h-1})\big\}, \quad h \ge 1.$$

By construction, $(J_0, \dots, J_{T-1})$ has the product law $\bigotimes_{h=0}^{T-1} \mathrm{Bernoulli}(q_h)$, while $(\widetilde{I}_0, \dots, \widetilde{I}_{T-1})$ has the same law as $(I_0, \dots, I_{T-1})$ under $P_t^\star$.

Let

$$\widetilde{N} := \sum_{h=0}^{T-1} \widetilde{I}_h, \qquad N' := \sum_{h=0}^{T-1} J_h.$$

Then

$$\int \phi_{a,\bar\alpha}(z)\, P_t^\star(dz) = \mathbb{E}\Big(\frac{\widetilde{N}}{T} - \bar\alpha\Big)_+, \qquad \int \phi_{a,\bar\alpha}(z)\, \overline{P}_t^\star(dz) = \mathbb{E}\Big(\frac{N'}{T} - \bar\alpha\Big)_+.$$

Since $x \mapsto (x - \bar\alpha)_+$ is 1-Lipschitz, the map $n \mapsto (n/T - \bar\alpha)_+$ is $1/T$-Lipschitz. Therefore

$$\begin{aligned}
\left| \int \phi_{a,\bar\alpha}\, dP_t^\star - \int \phi_{a,\bar\alpha}\, d\overline{P}_t^\star \right| &= \left| \mathbb{E}\Big(\frac{\widetilde{N}}{T} - \bar\alpha\Big)_+ - \mathbb{E}\Big(\frac{N'}{T} - \bar\alpha\Big)_+ \right| \\
&\le \frac{1}{T}\mathbb{E}|\widetilde{N} - N'| \\
&\le \frac{1}{T} \sum_{h=0}^{T-1} \mathbb{P}(\widetilde{I}_h \ne J_h).
\end{aligned} \tag{40}$$

We have $\widetilde{I}_0 = J_0$ almost surely because $p_0 = q_0$. For $h \ge 1$,

$$\mathbb{P}(\widetilde{I}_h \ne J_h \mid \mathcal{F}_{t-1}) = \mathbb{E}\Big[\big|p_h(\widetilde{I}_0, \dots, \widetilde{I}_{h-1}) - q_h\big| \,\Big|\, \mathcal{F}_{t-1}\Big].$$

It remains to bound $|p_h(i_{0:h-1}) - q_h|$ on histories with positive conditional probability.

Fix $h \geq 1$ and such a history $i_{0:h-1}$. Let $A_{i_{0:h-1}}$ denote the event that the threshold indicators $I_0, \ldots, I_{h-1}$ equal $i_0, \ldots, i_{h-1}$. Then $A_{i_{0:h-1}} \in \sigma(\beta_t, \ldots, \beta_{t+h-1}) \subseteq \mathcal{F}_{t+h-1}$. By the tower property and (10),

$$\left|p_h(i_{0:h-1}) - a_h\right| = \left|\mathbb{E}\Big[G^{\circ}_{t+h}(a_h) - a_h \mid \mathcal{F}_{t-1}, A_{i_{0:h-1}}\Big]\right| \leq \kappa_{t+h},$$

and similarly

$$|q_h - a_h| = \left|\mathbb{E}\Big[G^{\circ}_{t+h}(a_h) - a_h \mid \mathcal{F}_{t-1}\Big]\right| \leq \kappa_{t+h}.$$

Therefore

$$|p_h(i_{0:h-1}) - q_h| \leq 2\kappa_{t+h}.$$

Substituting into (40) yields

$$\left|\int \phi_{a,\bar{\alpha}}\, dP_t^{\star} - \int \phi_{a,\bar{\alpha}}\, d\overline{P}_t^{\star}\right| \leq \frac{2}{T}\sum_{h=1}^{T-1} \kappa_{t+h} \leq 2\kappa_{t+1}.$$

Taking the supremum over $a \in [0,1]^T$ proves (39). □

The following grid-reduction lemma turns a finite-grid concentration bound into a uniform bound over $[0,1]$. The Lipschitz property is needed only for the comparison function, not for the empirical CDF.

Lemma E.3. *Let $H : [0,1] \to \mathbb{R}$ be nondecreasing and let $F : [0,1] \to \mathbb{R}$ be nondecreasing and M-Lipschitz. Then for every integer $m \geq 1$,*

$$\|H - F\|_{\infty} \leq \max_{0 \leq \ell \leq m} \big|H(\ell/m) - F(\ell/m)\big| + \frac{M}{m}. \tag{41}$$

Proof. The claim is immediate for $u = 1$, so fix $u \in [0,1)$ and let $\ell \in \{0, \ldots, m-1\}$ be such that $\ell/m \leq u < (\ell+1)/m$. Since H and F are nondecreasing,

$$H(u) - F(u) \leq H\big((\ell+1)/m\big) - F(\ell/m).$$

Add and subtract $F((\ell+1)/m)$ to obtain

$$\begin{aligned} H(u) - F(u) &\leq \big[H\big((\ell+1)/m\big) - F\big((\ell+1)/m\big)\big] + \big[F\big((\ell+1)/m\big) - F(\ell/m)\big] \\ &\leq \max_{0 \leq r \leq m} |H(r/m) - F(r/m)| + \frac{M}{m}. \end{aligned}$$

Similarly,

$$F(u) - H(u) \leq F\big((\ell+1)/m\big) - H(\ell/m) \leq \max_{0 \leq r \leq m} |H(r/m) - F(r/m)| + \frac{M}{m}.$$

Taking the supremum over $u \in [0,1]$ proves (41). □

The next lemma controls the uniform deviation of the empirical PIT CDF from the average of the recent conditional CDFs.

Lemma E.4. *Assume* (11) *and $B \geq 4$. Then for every $\delta \in (0,1)$,*

$$\mathbb{P}\left(\left\|\widehat{G}_t^{(B)} - \overline{G}_{t,B}\right\|_{\infty} \leq \frac{2\big(\sqrt{\log(B/\delta)} + M\big)}{\sqrt{B}}\right) \geq 1 - \delta. \tag{42}$$

Proof. Set $m := \sqrt{B} - 1$, so $m \geq 1$. Use the grid

$$\mathcal{G}_m := \{\ell/m : 0 \leq \ell \leq \lfloor m \rfloor\} \cup \{1\}.$$

This grid has at most $\sqrt{B} + 1$ points and mesh size at most $1/m$. For each $u \in \mathcal{G}_m$, set

$$D_j(u) := 1\{\beta_j \leq u\} - G_j^{\circ}(u), \qquad j = t - B, \ldots, t-1.$$

Because $G_j^\circ(u) = \mathbb{E}[1\{\beta_j \le u\} \mid \mathcal{F}_{j-1}]$, the sequence $\{D_j(u), \mathcal{F}_j\}_{j=t-B}^{t-1}$ is a martingale-difference sequence, and $|D_j(u)| \le 1$ almost surely. Therefore Azuma–Hoeffding gives

$$\mathbb{P}\left(\left|\frac{1}{B}\sum_{j=t-B}^{t-1} D_j(u)\right| > x\right) \le 2\exp\left(-\frac{Bx^2}{2}\right), \qquad x > 0.$$

Applying the union bound over $u \in \mathcal{G}_m$,

$$\mathbb{P}\left(\max_{u\in\mathcal{G}_m}\left|\frac{1}{B}\sum_{j=t-B}^{t-1} D_j(u)\right| > x\right) \le 2(\sqrt{B}+1)\exp\left(-\frac{Bx^2}{2}\right).$$

Choosing

$$x = \sqrt{\frac{2\log\big(2(\sqrt{B}+1)/\delta\big)}{B}}$$

shows that, with probability at least $1-\delta$,

$$\max_{u\in\mathcal{G}_m}\big|\widehat{G}_t^{(B)}(u) - \overline{G}_{t,B}(u)\big| \le \sqrt{\frac{2\log\big(2(\sqrt{B}+1)/\delta\big)}{B}}. \tag{43}$$

The same proof as Lemma E.3, using the mesh size of $\mathcal{G}_m$, gives

$$\big\|\widehat{G}_t^{(B)} - \overline{G}_{t,B}\big\|_\infty \le \max_{u\in\mathcal{G}_m}\big|\widehat{G}_t^{(B)}(u) - \overline{G}_{t,B}(u)\big| + \frac{M}{m}.$$

Since $B \ge 4$ and $\delta \in (0,1)$,

$$\sqrt{\frac{2\log\big(2(\sqrt{B}+1)/\delta\big)}{B}} \le \frac{2\sqrt{\log(B/\delta)}}{\sqrt{B}}, \qquad \frac{M}{\sqrt{B}-1} \le \frac{2M}{\sqrt{B}}.$$

Combining these displays yields (42). □

The next lemma compares each future predictive marginal with the recent conditional average. Under monotone calibration errors, both CDFs are close to the identity map with error controlled by κ_{t-B}.

Lemma E.5. *Assume* (10)*. Then for every* $h = 0, \dots, T-1$*,*

$$\big\|\overline{G}_{t,B} - G^\star_{t+h|t}\big\|_\infty \le 2\kappa_{t-B}. \tag{44}$$

Proof. For any $u \in [0,1]$, conditions (10) give

$$\big|\overline{G}_{t,B}(u) - u\big| \le \frac{1}{B}\sum_{j=1}^{B}\big|G^\circ_{t-j}(u) - u\big| \le \frac{1}{B}\sum_{j=1}^{B}\kappa_{t-j} \le \kappa_{t-B}.$$

On the other hand, by the tower property,

$$G^\star_{t+h|t}(u) = \mathbb{E}\big[G^\circ_{t+h}(u) \mid \mathcal{F}_{t-1}\big],$$

so again by (10),

$$\big|G^\star_{t+h|t}(u) - u\big| \le \mathbb{E}\big[|G^\circ_{t+h}(u) - u| \mid \mathcal{F}_{t-1}\big] \le \kappa_{t+h} \le \kappa_{t-B}.$$

The triangle inequality yields

$$\big|\overline{G}_{t,B}(u) - G^\star_{t+h|t}(u)\big| \le 2\kappa_{t-B}.$$

Taking the supremum over $u \in [0,1]$ proves (44). □

The final lemma combines the empirical-process bound and the marginal comparison. On a single high-probability event, the empirical PIT CDF is uniformly close to every future predictive marginal needed in the horizon objective.

Lemma E.6. *Assume $B \geq 4$ and* (10)*,* (11)*. Then for every $\delta \in (0,1)$,*

$$\mathbb{P}\left(\forall h \in \{0,\dots,T-1\} : \left\|\widehat{G}_t^{(B)} - G^\star_{t+h|t}\right\|_\infty \leq \frac{2\big(\sqrt{\log(B/\delta)}+M\big)}{\sqrt{B}} + 2\kappa_{t-B}\right) \geq 1-\delta. \tag{45}$$

Proof. By the triangle inequality,

$$\left\|\widehat{G}_t^{(B)} - G^\star_{t+h|t}\right\|_\infty \leq \left\|\widehat{G}_t^{(B)} - \overline{G}_{t,B}\right\|_\infty + \left\|\overline{G}_{t,B} - G^\star_{t+h|t}\right\|_\infty.$$

Apply Lemma E.4 and Lemma E.5. The same event from Lemma E.4 works simultaneously for every h, because the first term does not depend on h. □

E.3 Proof of Theorem 3.3

Apply Lemma E.1 with $\mathcal{A} = [0,1]^T$, $J = J_t^{P_t^\star}$, $\widehat{J} = J_t^{\widehat{P}_t^{(B)}}$, $\widetilde{a} = \widehat{a}_t$, and $\eta = 0$. This gives

$$0 \leq J_t^{P_t^\star}(\widehat{a}_t) - J_t^{P_t^\star}(a_t^\star) \leq 2 \sup_{a\in[0,1]^T} \left|J_t^{P_t^\star}(a) - J_t^{\widehat{P}_t^{(B)}}(a)\right|. \tag{46}$$

It remains to control the uniform objective mismatch.

Fix $a \in [0,1]^T$. Because the deterministic length term $\ell_t(a)$ is identical in both objectives,

$$\begin{aligned}
\left|J_t^{P_t^\star}(a) - J_t^{\widehat{P}_t^{(B)}}(a)\right| &= \lambda_t \left|\int \phi_{a,\bar{\alpha}}(z)\, P_t^\star(dz) - \int \phi_{a,\bar{\alpha}}(z)\, \widehat{P}_t^{(B)}(dz)\right| \\
&\leq \lambda_t \left|\int \phi_{a,\bar{\alpha}}(z)\, P_t^\star(dz) - \int \phi_{a,\bar{\alpha}}(z)\, \overline{P}_t^\star(dz)\right| \\
&\quad + \lambda_t \left|\int \phi_{a,\bar{\alpha}}(z)\, \overline{P}_t^\star(dz) - \int \phi_{a,\bar{\alpha}}(z)\, \widehat{P}_t^{(B)}(dz)\right|.
\end{aligned} \tag{47}$$

By Lemma E.2, the first term on the right-hand side of (47) is bounded uniformly over a by

$$\lambda_t \cdot 2\kappa_{t+1}. \tag{48}$$

For the second term, use a direct product-coupling argument. Set

$$q_h := G^\star_{t+h|t}(a_h), \qquad \widehat{q}_h := \widehat{G}_t^{(B)}(a_h), \qquad h = 0,\dots,T-1.$$

Let $U_0,\dots,U_{T-1}$ be i.i.d. $\mathrm{Unif}[0,1]$, and define

$$I_h := 1\{U_h \leq q_h\}, \qquad J_h := 1\{U_h \leq \widehat{q}_h\}.$$

Then $(I_0,\dots,I_{T-1})$ has the threshold-indicator law induced by $\overline{P}_t^\star$, and $(J_0,\dots,J_{T-1})$ has the threshold-indicator law induced by $\widehat{P}_t^{(B)}$. With $N := \sum_{h=0}^{T-1} I_h$ and $M := \sum_{h=0}^{T-1} J_h$, the map $n \mapsto (n/T - \bar{\alpha})_+$ is $1/T$-Lipschitz, so

$$\begin{aligned}
\left|\int \phi_{a,\bar{\alpha}}\, d\overline{P}_t^\star - \int \phi_{a,\bar{\alpha}}\, d\widehat{P}_t^{(B)}\right| &\leq \frac{1}{T}\mathbb{E}|N-M| \leq \frac{1}{T}\sum_{h=0}^{T-1}\mathbb{P}(I_h \neq J_h) \\
&= \frac{1}{T}\sum_{h=0}^{T-1}|q_h - \widehat{q}_h| \leq \frac{1}{T}\sum_{h=0}^{T-1}\left\|G^\star_{t+h|t} - \widehat{G}_t^{(B)}\right\|_\infty.
\end{aligned} \tag{49}$$

Now invoke Lemma E.6. On an event of probability at least $1-\delta$,

$$\left\|G^\star_{t+h|t} - \widehat{G}_t^{(B)}\right\|_\infty \leq \frac{2\big(\sqrt{\log(B/\delta)}+M\big)}{\sqrt{B}} + 2\kappa_{t-B}, \qquad h = 0,\dots,T-1.$$

Hence, on that same event,

$$\left|\int \phi_{a,\bar{\alpha}}\, d\overline{P}_t^\star - \int \phi_{a,\bar{\alpha}}\, d\widehat{P}_t^{(B)}\right| \leq \frac{2\big(\sqrt{\log(B/\delta)}+M\big)}{\sqrt{B}} + 2\kappa_{t-B}. \tag{50}$$

Combining (47), (48), and (50), we obtain that, with probability at least $1-\delta$,

$$\sup_{a\in[0,1]^T}\left|J_t^{P_t^\star}(a)-J_t^{\widehat{P}_t^{(B)}}(a)\right|\le\lambda_t\left[\frac{2\big(\sqrt{\log(B/\delta)}+M\big)}{\sqrt{B}}+2\kappa_{t-B}+2\kappa_{t+1}\right]. \tag{51}$$

Finally, since $\kappa_{t+1}\le\kappa_{t-B}$, the bracket in (51) is at most

$$\frac{2\big(\sqrt{\log(B/\delta)}+M\big)}{\sqrt{B}}+4\kappa_{t-B}.$$

Substituting this bound into (46) proves (12).

F A THEORETICAL RESULT ON THE LENGTH DOMINANCE OF BCI

In this section, we present a setting in which the expected interval lengths of ACI is provably longer than those of BCI.

F.1 SETTING

We use the time indexing in Section 2, so that C_t predicts Y_t. Let V_t be the covariate vector observed before Y_t, and define

$$\begin{aligned} W_t &:= \big(V_t,(V_{t-1},Y_{t-1}),\dots,(V_{t-p},Y_{t-p})\big),\\ \mathcal{H}_{t-1} &:= \sigma\big((V_s,Y_s):s<t\big), \qquad \mathcal{G}_{t-1}:=\mathcal{H}_{t-1}\vee\sigma(V_t). \end{aligned} \tag{52}$$

Use the interval $[\ell_a,h_a]$ from Assumption D.3.

For ACI, we consider a truncated version as in Assumption D.3 so that the update can be made stationary. With step size $\gamma_A>0$, the update in (2) becomes

$$\begin{aligned} \alpha_t^A &:= \begin{cases} 0, & a_t<\ell_a,\\ a_t, & \ell_a\le a_t\le h_a,\\ 1, & a_t>h_a, \end{cases}\\ \mathrm{err}_t^A &:= 1\{\beta_t\le\alpha_t^A\}, \qquad a_{t+1}:=a_t+\gamma_A(\bar\alpha-\mathrm{err}_t^A), \end{aligned} \tag{53}$$

where β_t is the PIT in (3).

For BCI, consider $\tau>0$ and

$$K(b,a):=b\log\frac{b}{a}+(1-b)\log\frac{1-b}{1-a}, \qquad b\in[0,1],\quad a\in(0,1), \tag{54}$$

with $0\log 0:=0$. Similar to (22), define

$$\begin{aligned} &\alpha_t^B\in\mathrm{argmin}_{a\in[\ell_a,h_a]}\left\{|C_t(1-a)|+D_{t|t}(0)a+\tau K(a,\bar\alpha)\right\}, && 0\le\lambda_t\le\lambda_{\max},\\ &\alpha_t^B=1, && \lambda_t<0,\\ &\alpha_t^B=0, && \lambda_t>\lambda_{\max}. \end{aligned} \tag{55}$$

In particular, (55) uses the minimizer at $\lambda_t=0$ and $\lambda_t=\lambda_{\max}$. Retain $\gamma=c\lambda_{\max}$ with $c\in(0,1)$ and the update in (6):

$$\mathrm{err}_t^B:=1\{\beta_t\le\alpha_t^B\}, \qquad \lambda_{t+1}:=\lambda_t+c\lambda_{\max}(\mathrm{err}_t^B-\bar\alpha). \tag{56}$$

F.2 ASSUMPTIONS

Assumption F.1. Y_t given $\mathcal{G}_{t-1}$ has a continuous CDF $F^\star(W_t,\cdot)$ and the (C_t,α_t^B) are $\mathcal{G}_{t-1}$-measurable, and

$$\mathrm{Law}(\beta_t\mid\mathcal{G}_{t-1})=\mathrm{Unif}[0,1]. \tag{57}$$

Further, $(\alpha_t,\lambda_t,D_{t|t}(0))$ are $\mathcal{H}_{t-1}$-measurable.

Remark F.1. If the conditional CDF is known, then Assumption 2.1 and (57) hold with

$$\begin{aligned} C_t(1-a) &= \left[(F^\star)^{-1}(W_t, a/2), (F^\star)^{-1}(W_t, 1-a/2)\right], \qquad 0 < a < 1, \\ \beta_t &= 2\min\left\{F^\star(W_t, Y_t), 1 - F^\star(W_t, Y_t)\right\}. \end{aligned} \tag{58}$$

If the multi-step ahead forecast intervals and λ_t are $\mathcal{H}_{t-1}$-measurable, (20) also makes $D_{t|t}(0)$ $\mathcal{H}_{t-1}$-measurable.

Assumption F.2. There are constants $\bar{s} < \infty$ and $\mu_S > 0$, a $\mathcal{G}_{t-1}$-measurable variable $S_t \in [0, \bar{s}]$, and a nonincreasing function $w : [0,1] \to [0,\infty)$, differentiable at $\bar{\alpha}$, such that

$$|C_t(1-a)| = S_t w(a), \qquad a \in [0,1], \qquad \mathbb{E}[S_t \mid \mathcal{H}_{t-1}] = \mu_S. \tag{59}$$

For constants $0 < v_- \le v_+ < \infty$,

$$v_-(a-\bar{\alpha})^2 \le w(a) - w(\bar{\alpha}) - w'(\bar{\alpha})(a-\bar{\alpha}) \le v_+(a-\bar{\alpha})^2, \qquad a \in [0,1]. \tag{60}$$

Remark F.2. The requirement $\mathbb{E}[S_t \mid \mathcal{H}_{t-1}] = \mu_S$ in (59) holds, for example, if $S_t = \sigma_0(V_t)$ for a measurable function σ_0 taking values in $[0, \bar{s}]$, the covariates are stationary, V_t is independent of $\mathcal{H}_{t-1}$, and $\mathbb{E}\sigma_0(V_t) > 0$.

Assumption F.3. There are finite constants $D_- \le D_+$ such that

$$D_- \le D_{t|t}(0) \le D_+. \tag{61}$$

Define

$$r_I := \max\{h_a, 1-\ell_a\}, \qquad C_\star := \max\left\{\bar{s}[-w'(\bar{\alpha})] - D_-, D_+\right\}. \tag{62}$$

The step sizes and regularization parameter satisfy

$$\begin{aligned} \gamma_A(1-\bar{\alpha}) \le \ell_a, \qquad \gamma_A\bar{\alpha} \le 1 - h_a, \qquad \ell_a \le \bar{\alpha} \le h_a, \\ v_-\gamma_A > v_+c, \qquad \tau > \frac{v_+ r_I C_\star}{\bar{\alpha}(1-\bar{\alpha})(v_-\gamma_A - v_+c)}. \end{aligned} \tag{63}$$

F.3 Main result

With these assumptions, we can show that BCI has a strictly shorter expected interval length than ACI.

Theorem F.1. *Under Assumptions 2.1, F.1, F.2, and F.3, suppose that the ACI path $(V_t, Y_t, a_t, \alpha_t^A)_{t\ge1}$ and the BCI path $(V_t, Y_t, \lambda_t, \alpha_t^B)_{t\ge1}$ are stationary. Then, for every t,*

$$\mathbb{E}|C_t(1-\alpha_t^B)| < \mathbb{E}|C_t(1-\alpha_t^A)|. \tag{64}$$

Remark F.3. For ACI, Gibbs & Candès (2021, Appendix A.5) give sufficient conditions for the existence of an invariant distribution in a finite-state hidden Markov model. For BCI, Theorem 3.2, formalized in Theorem D.1, establishes an invariant distribution under its stated assumptions. Initializing the respective joint processes from these invariant distributions makes their paths stationary. Therefore, the stationarity of ACI and BCI paths is possible.

Proof. We first verify that a_t and λ_t have finite second moments. For ACI, define

$$L_A := \ell_a - \gamma_A(1-\bar{\alpha}), \qquad U_A := h_a + \gamma_A\bar{\alpha}, \qquad \mathcal{I}_A := [L_A, U_A] \subseteq [0,1], \tag{65}$$

where the inclusion follows from (63). For $a_t \in \mathcal{I}_A$, (53) gives, almost surely,

$$\begin{aligned} a_{t+1} &\in \begin{cases} [L_A, \ell_a + \gamma_A\bar{\alpha}], & a_t < \ell_a, \\ [a_t - \gamma_A(1-\bar{\alpha}), a_t + \gamma_A\bar{\alpha}], & \ell_a \le a_t \le h_a, \\ [h_a - \gamma_A(1-\bar{\alpha}), U_A], & a_t > h_a, \end{cases} \\ &\subseteq \mathcal{I}_A. \end{aligned} \tag{66}$$

If $a_1 < L_A$, (53) gives

$$\begin{aligned} n_- &:= \left\lceil \frac{L_A - a_1}{\gamma_A\bar{\alpha}} \right\rceil < \infty, \\ a_{1+k} &= a_1 + k\gamma_A\bar{\alpha} < L_A, \qquad 0 \le k < n_-, \\ a_{1+n_-} &= a_1 + n_-\gamma_A\bar{\alpha} \in [L_A, L_A + \gamma_A\bar{\alpha}) \subseteq \mathcal{I}_A. \end{aligned} \tag{67}$$

If $a_1 > U_A$, (53) gives, almost surely,

$$\begin{aligned} n_+ &:= \left\lceil \frac{a_1 - U_A}{\gamma_A(1-\bar{\alpha})} \right\rceil < \infty, \\ a_{1+k} &= a_1 - k\gamma_A(1-\bar{\alpha}) > U_A, \qquad 0 \le k < n_+, \\ a_{1+n_+} &= a_1 - n_+\gamma_A(1-\bar{\alpha}) \in (U_A - \gamma_A(1-\bar{\alpha}), U_A] \subseteq \mathcal{I}_A. \end{aligned} \tag{68}$$

By (66)–(68), $1\{a_t \notin \mathcal{I}_A\} \to 0$ almost surely. Stationarity and bounded convergence give

$$\mathbb{P}(a_1 \notin \mathcal{I}_A) = \lim_{t\to\infty} \mathbb{P}(a_t \notin \mathcal{I}_A) = 0. \tag{69}$$

For BCI, the interval $\Lambda = [-\gamma\bar{\alpha}, \lambda_{\max} + \gamma(1-\bar{\alpha})]$ is invariant under (56), as shown in the proof of Theorem 3.1: (55) gives $\alpha_t^B = 1$ for $\lambda_t < 0$ and $\alpha_t^B = 0$ for $\lambda_t > \lambda_{\max}$. Write $L_B := -\gamma\bar{\alpha}$ and $U_B := \lambda_{\max} + \gamma(1-\bar{\alpha})$. By (55)–(57), almost surely, if $\lambda_1 < L_B$,

$$\begin{aligned} n_{B,-} &:= \left\lceil \frac{L_B - \lambda_1}{\gamma(1-\bar{\alpha})} \right\rceil < \infty, \\ \lambda_{1+k} &= \lambda_1 + k\gamma(1-\bar{\alpha}) < L_B, \qquad 0 \le k < n_{B,-}, \\ \lambda_{1+n_{B,-}} &= \lambda_1 + n_{B,-}\gamma(1-\bar{\alpha}) \in [L_B, L_B + \gamma(1-\bar{\alpha})) \subseteq \Lambda; \end{aligned} \tag{70}$$

if $\lambda_1 > U_B$,

$$\begin{aligned} n_{B,+} &:= \left\lceil \frac{\lambda_1 - U_B}{\gamma\bar{\alpha}} \right\rceil < \infty, \\ \lambda_{1+k} &= \lambda_1 - k\gamma\bar{\alpha} > U_B, \qquad 0 \le k < n_{B,+}, \\ \lambda_{1+n_{B,+}} &= \lambda_1 - n_{B,+}\gamma\bar{\alpha} \in (U_B - \gamma\bar{\alpha}, U_B] \subseteq \Lambda. \end{aligned} \tag{71}$$

By invariance and (70)–(71), $1\{\lambda_t \notin \Lambda\} \to 0$ almost surely. Stationarity gives $\mathbb{P}(\lambda_t \in \Lambda) = 1$ by the limit argument in (69). Both a_t and λ_t are consequently bounded.

By Assumption F.1, for $j \in \{A, B\}$,

$$\mathbb{E}[\mathrm{err}_t^j \mid \mathcal{G}_{t-1}] = \alpha_t^j, \qquad \mathbb{E}\mathrm{err}_t^j = \mathbb{E}\alpha_t^j. \tag{72}$$

Taking expectations in the two stationary updates gives

$$\begin{aligned} 0 &= \mathbb{E}(a_{t+1} - a_t) = \gamma_A(\bar{\alpha} - \mathbb{E}\alpha_t^A), \\ 0 &= \mathbb{E}(\lambda_{t+1} - \lambda_t) = \gamma(\mathbb{E}\alpha_t^B - \bar{\alpha}). \end{aligned} \tag{73}$$

Hence $\mathbb{E}\alpha_t^A = \mathbb{E}\alpha_t^B = \bar{\alpha}$. Since $(\mathrm{err}_t^j)^2 = \mathrm{err}_t^j$, (72) and (73) imply

$$\mathbb{E}(\mathrm{err}_t^j - \bar{\alpha})^2 = (1-2\bar{\alpha})\mathbb{E}\mathrm{err}_t^j + \bar{\alpha}^2 = \bar{\alpha}(1-\bar{\alpha}), \qquad j \in \{A, B\}. \tag{74}$$

As a result,

$$\begin{aligned} 0 = \mathbb{E}(a_{t+1}^2 - a_t^2) &= 2\gamma_A\mathbb{E}[a_t(\bar{\alpha} - \alpha_t^A)] + \gamma_A^2\bar{\alpha}(1-\bar{\alpha}), \\ \Longrightarrow \quad \mathbb{E}[a_t(\alpha_t^A - \bar{\alpha})] &= \frac{\gamma_A}{2}\bar{\alpha}(1-\bar{\alpha}). \end{aligned} \tag{75}$$

Because $0 \le a_t \le 1$ and $\bar{\alpha} \in [\ell_a, h_a]$, (53) implies

$$(\alpha_t^A - a_t)(\alpha_t^A - \bar{\alpha}) = \begin{cases} a_t\bar{\alpha} \ge 0, & a_t < \ell_a, \\ 0, & \ell_a \le a_t \le h_a, \\ (1-a_t)(1-\bar{\alpha}) \ge 0, & a_t > h_a. \end{cases} \tag{76}$$

Using (73), (75), and (76),

$$\mathbb{E}(\alpha_t^A - \bar{\alpha})^2 = \mathbb{E}[a_t(\alpha_t^A - \bar{\alpha})] + \mathbb{E}[(\alpha_t^A - a_t)(\alpha_t^A - \bar{\alpha})] \ge \frac{\gamma_A}{2}\bar{\alpha}(1-\bar{\alpha}). \tag{77}$$

For BCI, we have

$$\begin{aligned} 0 = \mathbb{E}(\lambda_{t+1}^2 - \lambda_t^2) &= 2\gamma\mathbb{E}[\lambda_t(\alpha_t^B - \bar{\alpha})] + \gamma^2\bar{\alpha}(1-\bar{\alpha}), \\ \Longrightarrow \quad \mathbb{E}[\lambda_t(\alpha_t^B - \bar{\alpha})] &= -\frac{c\lambda_{\max}}{2}\bar{\alpha}(1-\bar{\alpha}). \end{aligned} \tag{78}$$

Define $z_t^B := 1 - \lambda_t/\lambda_{\max}$. By (73) and (78),

$$\mathbb{E}[z_t^B(\alpha_t^B - \bar{\alpha})] = \mathbb{E}(\alpha_t^B - \bar{\alpha}) - \frac{\mathbb{E}[\lambda_t(\alpha_t^B - \bar{\alpha})]}{\lambda_{\max}} = \frac{c}{2}\bar{\alpha}(1-\bar{\alpha}). \tag{79}$$

We next bound the deviation $\alpha_t^B - \bar{\alpha}$ when $0 \le \lambda_t \le \lambda_{\max}$. The feasible action $\bar{\alpha}$ in (55) gives

$$S_t\{w(\alpha_t^B) - w(\bar{\alpha})\} + D_{t|t}(0)(\alpha_t^B - \bar{\alpha}) + \tau K(\alpha_t^B, \bar{\alpha}) \le 0. \tag{80}$$

For the KL divergence in (54), it is well-known that

$$\begin{aligned} K(\bar{\alpha}, \bar{\alpha}) = 0, \qquad & \frac{\partial K(b,\bar{\alpha})}{\partial b}\Big|_{b=\bar{\alpha}} = 0, \\ \frac{\partial^2 K(b,\bar{\alpha})}{\partial b^2} = \frac{1}{b(1-b)} \ge 4, \qquad & b \in (0,1), \\ K(b,\bar{\alpha}) \ge 2(b-\bar{\alpha})^2, \qquad & b \in [0,1], \end{aligned} \tag{81}$$

Since $w'(\bar{\alpha}) \le 0$, (61) implies

$$-D_+ \le S_t[-w'(\bar{\alpha})] - D_{t|t}(0) \le \bar{s}[-w'(\bar{\alpha})] - D_-. \tag{82}$$

Combining (60), (80), (81), and (82) gives

$$\begin{aligned} 2\tau(\alpha_t^B - \bar{\alpha})^2 &\le \tau K(\alpha_t^B, \bar{\alpha}) \\ &\le \{S_t[-w'(\bar{\alpha})] - D_{t|t}(0)\}(\alpha_t^B - \bar{\alpha}) - S_t v_-(\alpha_t^B - \bar{\alpha})^2 \\ &\le C_\star |\alpha_t^B - \bar{\alpha}|. \end{aligned} \tag{83}$$

If $\alpha_t^B - \bar{\alpha} \ne 0$, division by $2\tau|\alpha_t^B - \bar{\alpha}|$ yields

$$|\alpha_t^B - \bar{\alpha}| \le \frac{C_\star}{2\tau}, \qquad 0 \le \lambda_t \le \lambda_{\max}; \tag{84}$$

the inequality also holds when $\alpha_t^B - \bar{\alpha} = 0$.

Decompose the BCI variance using (73):

$$\mathbb{E}(\alpha_t^B - \bar{\alpha})^2 = \mathbb{E}[z_t^B(\alpha_t^B - \bar{\alpha})] + \mathbb{E}[(\alpha_t^B - z_t^B)(\alpha_t^B - \bar{\alpha})]. \tag{85}$$

For $\lambda_t < 0$, $\alpha_t^B = 1$ and $z_t^B > 1$; for $\lambda_t > \lambda_{\max}$, $\alpha_t^B = 0$ and $z_t^B < 0$. Therefore

$$(\alpha_t^B - z_t^B)(\alpha_t^B - \bar{\alpha}) \le 0, \qquad \lambda_t \notin [0, \lambda_{\max}]. \tag{86}$$

For $0 \le \lambda_t \le \lambda_{\max}$, $z_t^B \in [0,1]$ and $\alpha_t^B \in [\ell_a, h_a]$, so

$$|\alpha_t^B - z_t^B| \le \max\{h_a, 1 - \ell_a\} = r_I. \tag{87}$$

Substituting (79), (84), (86), and (87) into (85) yields

$$\mathbb{E}(\alpha_t^B - \bar{\alpha})^2 \le \frac{c}{2}\bar{\alpha}(1-\bar{\alpha}) + \frac{r_I C_\star}{2\tau}. \tag{88}$$

It remains to account for the dependence between S_t and α_t^B. By (60), for $a \in [0,1]$,

$$\begin{aligned} 0 \le w(a) &\le w(\bar{\alpha}) + w'(\bar{\alpha})(a - \bar{\alpha}) + v_+(a-\bar{\alpha})^2 \\ &\le w(\bar{\alpha}) + |w'(\bar{\alpha})| + v_+ =: C_w < \infty. \end{aligned} \tag{89}$$

Since $0 \le S_t \le \bar{s}$, (59) and (89) give, for $j \in \{A, B\}$,

$$\begin{aligned} \mathbb{E}|C_t(1-\alpha_t^j)| &= \mathbb{E}[S_t w(\alpha_t^j)] \le \bar{s} C_w < \infty, \\ \mathbb{E}\big|\mathrm{Cov}(S_t, w(\alpha_t^j) \mid \mathcal{H}_{t-1})\big| &\le \mathbb{E}[S_t w(\alpha_t^j)] + \mathbb{E}\big\{\mathbb{E}[S_t \mid \mathcal{H}_{t-1}]\mathbb{E}[w(\alpha_t^j) \mid \mathcal{H}_{t-1}]\big\} \\ &\le 2\bar{s} C_w < \infty. \end{aligned} \tag{90}$$

For ACI, α_t^A is $\mathcal{H}_{t-1}$-measurable, so (59) gives

$$\mathbb{E}[S_t w(\alpha_t^A)] = \mathbb{E}\{w(\alpha_t^A)\mathbb{E}[S_t \mid \mathcal{H}_{t-1}]\} = \mu_S \mathbb{E} w(\alpha_t^A). \tag{91}$$

For BCI, condition on $\mathcal{H}_{t-1}$, so that λ_t and $D_{t|t}(0)$ are fixed. When $0 \le \lambda_t \le \lambda_{\max}$, define

$$L_t(a, s) := sw(a) + D_{t|t}(0)a + \tau K(a, \bar{\alpha}). \tag{92}$$

For $s_1 < s_2$, let b_i be any minimizer of $L_t(\cdot, s_i)$ on $[\ell_a, h_a]$, for $i = 1, 2$. Optimality of b_i implies

$$\begin{aligned} L_t(b_1, s_1) \le L_t(b_2, s_1), \qquad L_t(b_2, s_2) \le L_t(b_1, s_2), \\ (s_2 - s_1)\{w(b_2) - w(b_1)\} \le 0. \end{aligned} \tag{93}$$

Thus $w(b_2) \le w(b_1)$, including when minimizers are nonunique. If $\lambda_t \notin [0, \lambda_{\max}]$, α_t^B is fixed given $\mathcal{H}_{t-1}$. Let (S_t', b_t') be a conditionally independent copy of (S_t, α_t^B) given $\mathcal{H}_{t-1}$. By (93)

$$\begin{aligned} 2\operatorname{Cov}(S_t, w(\alpha_t^B) \mid \mathcal{H}_{t-1}) &= \mathbb{E}\left[(S_t - S_t')\{w(\alpha_t^B) - w(b_t')\} \mid \mathcal{H}_{t-1}\right] \\ &\le 0. \end{aligned} \tag{94}$$

By (59) and (94),

$$\begin{aligned} &\mathbb{E}[S_t w(\alpha_t^B)] \\ &= \mathbb{E}\operatorname{Cov}(S_t, w(\alpha_t^B) \mid \mathcal{H}_{t-1}) + \mathbb{E}\left\{\mathbb{E}[S_t \mid \mathcal{H}_{t-1}]\mathbb{E}[w(\alpha_t^B) \mid \mathcal{H}_{t-1}]\right\} \\ &\le \mu_S \mathbb{E} w(\alpha_t^B). \end{aligned} \tag{95}$$

Finally, (60) and $\mathbb{E}\alpha_t^A = \mathbb{E}\alpha_t^B = \bar{\alpha}$ give

$$\begin{aligned} \mathbb{E} w(\alpha_t^A) &\ge w(\bar{\alpha}) + v_- \mathbb{E}(\alpha_t^A - \bar{\alpha})^2, \\ \mathbb{E} w(\alpha_t^B) &\le w(\bar{\alpha}) + v_+ \mathbb{E}(\alpha_t^B - \bar{\alpha})^2. \end{aligned} \tag{96}$$

Combining (91), (95), and (96) with the variance bounds (77) and (88),

$$\begin{aligned} \mathbb{E}|C_t(1 - \alpha_t^A)| - \mathbb{E}|C_t(1 - \alpha_t^B)| &\ge \mu_S\{\mathbb{E} w(\alpha_t^A) - \mathbb{E} w(\alpha_t^B)\} \\ &\ge \mu_S\left\{v_- \mathbb{E}(\alpha_t^A - \bar{\alpha})^2 - v_+ \mathbb{E}(\alpha_t^B - \bar{\alpha})^2\right\} \\ &\ge \frac{\mu_S}{2}\left\{\bar{\alpha}(1 - \bar{\alpha})(v_- \gamma_A - v_+ c) - \frac{v_+ r_I C_\star}{\tau}\right\} > 0, \end{aligned} \tag{97}$$

where strict positivity follows from (63). This proves (64). □