PARAMETER-FREE AND GROUP CONDITIONAL ONLINE CONFORMAL PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Quantifying the uncertainty in predictions made by machine learning models is challenging in online settings where the characteristics of the data and labels change over time. Online conformal prediction (OCP) methods tackle this problem by producing intervals for each observation and guaranteeing that a proportion of them contain their labels, a property known as marginal coverage. OCP methods, however, may systematically miscover observations belonging to groups defined by attributes like age or race, raising fairness concerns. Group-conditional OCP (G-OCP) methods address this by producing intervals that provide coverage for multiple groups. Existing approaches rely on learning rates, which control how they react to changes in the data stream, but can be hard to tune in online settings. Leveraging advancements in learning-rate-free methods, we develop a learning-rate-free G-OCP algorithm with the strongest group-coverage guarantees among existing methods. Empirical results show that our method provides group coverage, produces interval lengths competitive with those of standard methods, and reacts more effectively to changes in the data than fixed-learning-rate methods.

1 INTRODUCTION

Machine learning models are increasingly deployed in high-stakes domains, such as supporting patient care in hospitals (Hong et al., 2022), informing loan decisions in banking (Bhatore et al., 2020), and enabling real-time perception in autonomous vehicles (Grigorescu et al., 2020). Since these models impact consequential decisions, it is critical to quantify the uncertainty in their predictions.

Conformal prediction (CP) is a versatile framework for uncertainty quantification. Given an input X, its unknown label Y, a predictor for Y, and miscoverage level $\alpha \in (0, 1)$, CP methods construct prediction sets that contain the unknown label with probability $1 - \alpha$ (Angelopoulos & Bates, 2023; Lei & Wasserman, 2014; Papadopoulos et al., 2002; Saunders et al., 1999; Shafer & Vovk, 2008; Vovk et al., 2005). This guarantee, however, only holds when the data is exchangeable, an assumption rarely satisfied in practice. In medical triage, for example, the stream of incoming patients may change due to seasonal events, like the flu (Anderer, 2025), and large transportation accidents which can cause surges in patient admission (Rozenberg et al., 2017).

Online conformal prediction (OCP) methods (Angelopoulos et al., 2023; Areces et al., 2025; Bastani et al., 2022; Bhatnagar et al., 2023; Gibbs & Candès, 2021; 2024; Liu et al., 2026; Podkopaev et al., 2024; Zaffran et al., 2022) address this limitation by generating a sequence of prediction intervals $(\mathcal{I}_t)_{t\geq 1}$ for the labels of an arbitrary data stream $((X_t, Y_t))_{t\geq 1} = (X_1, Y_1), (X_2, Y_2), \ldots$. Most OCP methods provide *marginal* coverage: as the time horizon T grows, approximately $(1 - \alpha)$-fraction of the intervals will contain their labels, i.e. $\sum_{t=1}^{T} \mathbf{1}\{Y_t \in \mathcal{I}_t\} \approx T(1-\alpha)$. This guarantee, however, may harm certain sub-populations in the data. For example, if the samples X_t of a minority group appear with frequency α, an OCP method can systematically fail to cover that group while maintaining marginal coverage. This may lead to harmful downstream decisions for this group.

Accordingly, several works have developed OCP algorithms with group-conditional coverage guarantees (Bastani et al., 2022; Ramalingam et al., 2025). That is, for a set of k groups in the data, these methods generate a single sequence of intervals $(\mathcal{I}_t)_{t\geq 1}$ that provides $(1 - \alpha)$-coverage within each group, i.e., for all $j \in \{1, \ldots, k\}$, $\sum_{t=1}^{T} \mathbf{1}\{Y_t \in \mathcal{I}_t \text{ and } X_t \in \text{group } j\} \approx T_j(1 - \alpha)$, where T_j is the number of samples in group j. However, they are not *parameter-free*: their performance depends on a learning rate which can be problematic for OCP where little to no assumptions are made about the data stream. Since the data may be adversarial, relying on a validation set to tune the learning

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.

rate can be futile: if one identifies a "good" learning rate, a shift in the incoming data could make the choice inadequate (Orabona, 2019). While there has been a rise in parameter-free OCP methods, existing ones only provide marginal coverage (Liu et al., 2026; Podkopaev et al., 2024; Zhang et al., 2024). Thus, there is no parameter-free OCP method that provides group-conditional coverage.

Contributions. In this work, we address this gap. Specifically, we make the following contributions:

1. We introduce the first parameter-free algorithm for group-conditional online conformal prediction, that we dub **P**ortfolios for **O**nline **G**roup C**o**nformal (POGO).
2. We establish rigorous finite-time group-conditional coverage guarantees for POGO with the best-known rates with respect to time T (up to logarithmic factors), that hold under weaker assumptions than those required by prior work.
3. We perform extensive experiments on synthetic and real-world datasets, showing that POGO is competitive with existing approaches: POGO can effectively react to shifts in the data stream that other methods cannot and produce small prediction intervals that provide group coverage.

1.1 Related Work

Closest to our work are two lines of research that draw on classical and parameter-free online learning: marginal OCP and group-conditional coverage in stochastic settings. We discuss each of them.

Online conformal prediction. Gibbs & Candès (2021) introduced Adaptive Conformal Inference (ACI) for marginal OCP. ACI parameterizes the prediction sets' width via a single parameter and updates it with Online Subgradient Descent (OSD) on the quantile loss. To address ACI's sensitivity to its step size, Bastani et al. (2022) introduced MultiValid Predictor, which provides long-term threshold-calibrated coverage but cannot rapidly react to changes in the data. Subsequent approaches–Aggregated ACI (Zaffran et al., 2022), Strongly Adaptive OCP (Bhatnagar et al., 2023), and Dynamically-Tuned ACI (Gibbs & Candès, 2021)–try to avoid manual step-size tuning by using online learning tools to automatically adjust ACI's updates. A few works explore related problems. Angelopoulos et al. (2024) analyzed OSD with decaying step sizes and developed a method for both marginal OCP and CP with stochastic data. Angelopoulos et al. (2025) then introduced a general theory, showing when OSD works for OCP and similar problems. Areces et al. (2025) developed OCP algorithms for a class of stochastic processes with temporal dependence. Others used control theory, recasting OCP as a feedback-control problem (Angelopoulos et al., 2023; Yang et al., 2024). These methods, however, require learning rates and sometimes other parameters. To address this, Zhang et al. (2024) and Podkopaev et al. (2024) used parameter-free online learning (Jun et al., 2017a;b; Orabona & Pál, 2016; 2021) for OCP. Recently, Liu et al. (2026) developed Universal Portfolio OCP (UP-OCP), a parameter-free algorithm that provides marginal coverage. The concepts introduced in Liu et al. (2026) provide the foundation for our method. We elaborate in Sec. 3.

Group-conditional coverage. The goal of providing group-conditional coverage was introduced in the i.i.d. data setting. Vovk et al. (2003) and Romano et al. (2020) proposed algorithms with coverage guarantees for disjoint groups. Shortly after came a conservative method with guarantees for intersecting groups (Foygel Barber et al., 2021). Numerous works followed. Jung et al. (2021; 2023) developed methods that produce smaller prediction sets which provide coverage conditional on group membership and the threshold used to produce the set. Gibbs et al. (2025) provided further improvements, giving tighter finite-sample coverage, allowing for discrete outcomes, and richer conditional guarantees. These works laid the foundation for group-conditional OCP. The method of Gupta et al. (2022) generates prediction sets in an online setting with a group calibration-like guarantees. This was improved by Bastani et al. (2022) and then Ramalingam et al. (2025), who developed group-conditional ACI (GCACI). GCACI parameterizes sets using group-specific parameters and updates them via online quantile regression with follow-the-regularized-leader; this procedure requires a learning rate. Concurrent work (Angelopoulos et al., 2025) derived similar results, extending the theory to decaying learning rates. Recently, Lyu & Cheung (2026) studied how normalizing GCACI's update rule can yield improvements. Many of these approaches represent prediction sets as linear functions of group membership, an idea that we also adopt.

2 Problem Setting

We observe a stream of data $((X_t, Y_t))_{t\geq 1}$ and have a predictor f_t that depends on the data up to time $t-1$. At each time t, $X_t \in \mathbb{R}^d$ is a d-dimensional input, $Y_t \in \mathbb{R}$ is a univariate outcome, and

denote $\hat{Y}_t = f_t(X_t)$ to be the prediction of f_t on X_t. We make no assumptions on the data stream, or f_t which may change arbitrarily over time. For a sequence of radii $(\tau_t)_{t\geq 1}$, $\tau_t \in \mathbb{R}$, we define prediction intervals as $\mathcal{I}_t(\tau_t) = [\hat{Y}_t - \tau_t, \hat{Y}_t + \tau_t]$ with the convention that $\mathcal{I}_t$ is *empty* if $\tau_t < 0$.

Marginal OCP. The goal of OCP is to generate intervals $(\mathcal{I}_t)_{t\geq 1}$ that control the empirical coverage rate at level $1-\alpha \in (0,1)$, i.e. as the time horizon T grows, we require $\frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\{Y_t \in \mathcal{I}_t\} \approx 1-\alpha$. At each time t, the radius τ_t of the interval $\mathcal{I}_t(\tau_t)$ must be selected using only the data up to $t-1$. Denote $S_t = |Y_t - \hat{Y}_t|$ to be *non-conformity scores*. Noting that $\mathbf{1}\{Y_t \in \mathcal{I}_t\} = \mathbf{1}\{S_t \leq \tau_t\}$, the goal of marginal OCP is equivalent to learning a sequence of radii $(\tau_t)_{t\geq 1}$ that satisfies

$$\lim_{T\to\infty} \left| \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\{S_t \leq \tau_t\} - (1-\alpha) \right| = 0. \tag{OCP$_\text{M}$}$$

Unfortunately, radii $(\tau_t)_{t\geq 1}$ can provide (OCP$_\text{M}$) by systematically miscovering the Y_t of certain groups. This raises fairness and safety concerns, motivating group-conditional OCP (G-OCP).

Group-conditional OCP. Consider a collection of k potentially intersecting groups represented by a set of k functions $\mathcal{C} = \{c_j : \mathcal{X} \to \{0,1\} \mid j \in [k]\}$, $[k] \coloneqq \{1,\dots,k\}$. Here $c_j(X) = 1$ indicates that X belongs to group j and $T_j = \sum_{t=1}^{T} c_j(X_t)$ is the count of samples belonging to group j up to time T. As an example, groups may be defined by biological sex, age, or race, and an individual X may belong to, say, both the MALE and UNDER 50 groups. G-OCP extends the objective of marginal OCP, requiring that the learned sequence of radii $(\tau_t)_{t\geq 1}$ satisfies[1]

$$\lim_{T\to\infty} \left| \frac{1}{T_j}\sum_{t=1}^{T} \mathbf{1}\{S_t \leq \tau_t\} c_j(X_t) - (1-\alpha) \right| = 0 \quad \text{for all } j \in [k]. \tag{OCP$_\text{G}$}$$

In this work, we develop a parameter-free algorithm for G-OCP. [2] We begin with an overview of a parameter-free marginal OCP method (Liu et al., 2026) whose key ideas we adapt.

3 Marginal Coverage via Portfolio Optimization

Similar to CP with i.i.d data (Angelopoulos & Bates, 2023; Shafer & Vovk, 2008), online marginal coverage (OCP$_\text{M}$) can be obtained by learning the $(1-\alpha)$-quantile of the scores $(S_t)_{t\geq 1}$. Since OCP considers a stream of data, this is typically done with online quantile regression. At each time t, one considers the quantile loss (Gneiting & Raftery, 2007; Koenker & Hallock, 2001)

$$l_t(\tau_t, S_t) = \max\{(1-\alpha)(S_t - \tau_t), \alpha(\tau_t - S_t)\}, \tag{1}$$

and uses classic online learning methods (Cesa-Bianchi & Lugosi, 2006; Orabona, 2019) to update τ_t. Interestingly, the subgradients of l_t (with respect to τ_t) quantify the miscoverage error (Angelopoulos et al., 2025; Gibbs & Candès, 2021). Observe that

$$g_t(\tau_t) = \mathbf{1}\{S_t \leq \tau_t\} - (1-\alpha) \tag{2}$$

is a subgradient of l_t at τ_t and that the miscoverage error at time T can be expressed as

$$\mathsf{MisCov}_T = \left| \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}\{S_t \leq \tau_t\} - (1-\alpha) \right| = \frac{1}{T}\left| \sum_{t=1}^{T} g_t(\tau_t) \right|. \tag{3}$$

Thus, radii τ_t for which the cumulative subgradient $|\sum_{t=1}^{T} g_t(\tau_t)|$ grows slowly will necessarily control the empirical miscoverage error.

3.1 A Parameter-Free Method for OCP via Universal Portfolio

Liu et al. (2026) proposed UP-OCP, a portfolio-optimization-based method for OCP that uses a sequential betting game to control the cumulative subgradient $|\sum_{t=1}^{T} g_t(\tau_t)|$. The method begins with wealth $W_0 = 1$. At each time t, using the data up to $t-1$, UP-OCP chooses a *signed betting*

[1] Implicit is the requirement that $T_j > 0$ and $T_j \to \infty$.

[2] Our algorithm and all presented analysis is for the more general setting where $c_j(X_t) \in [0,1]$.

fraction $\beta_t \in [-1/1-\alpha, 1/\alpha]$ of the current wealth to bet on the negative subgradient $-g_t(\tau_t) \in \{-\alpha, 1-\alpha\}$. This sequential game induces a wealth process $(W_t)_{t\geq 1}$ given by

$$W_t = W_{t-1}(1 - \beta_t g_t). \tag{4}$$

UP-OCP employs Universal Portfolio (UP) (Cover, 1991; Cover & Ordentlich, 1996) to learn $(\beta_t)_{t\geq 1}$ that promote rapid growth of $\ln W_t$, and sets the radii as $\tau_t = \beta_t W_{t-1}$. Liu et al. (2026) show that for this choice of τ_t, $|\sum_{t=1}^{T} g_t(\tau_t)|$ is bounded by a slowly growing function in T, yielding $\mathsf{MisCov}_T = \mathcal{O}\left((\ln T)/T + \sqrt{(\alpha \ln T)/T}\right)$. Notably, UP-OCP provides this guarantee under a mild polynomial growth condition on the non-conformity scores: a condition weaker than what is assumed in prior works.[3] Moreover, it does not use a learning rate, making it *parameter-free*.

While UP-OCP only provides marginal (not group-conditional) guarantees, it demonstrates that *solving a portfolio optimization problem can produce a parameter-free OCP algorithm*. We show how a suitable, carefully-constructed, generalization of this insight leads to a G-OCP algorithm.

4 Online Group Coverage via Betting & Portfolio Optimization

We now present POGO, a parameter-free algorithm for G-OCP. Like UP-OCP, POGO utilizes betting and portfolio optimization, extending it to G-OCP.

In standard OCP, there is a *single* coverage objective (OCP_M) and the radii τ_t are thus scalars. In G-OCP, the radii must satisfy k coverage guarantees, one for each group (OCP_G). This naturally suggests parameterizing τ_t with k parameters and learning them sequentially. Accordingly, for grouping functions $c_1, \ldots, c_k$, define $\mathbf{c} = (c_1, \ldots, c_k)^T$ to be the group membership vector. POGO parametrizes the radii $\tau_t(X_t)$ as a linear function of the group membership vector[4]

$$\tau_t(X_t) = \langle \boldsymbol{\theta}_t, \mathbf{c}(X_t) \rangle, \quad \boldsymbol{\theta}_t \in \mathbb{R}^k, \tag{5}$$

and learns $(\boldsymbol{\theta}_t)_{t\geq 1}$ that provide group-conditional coverage. This parameterization is simple to implement and empirically very effective. Consequently, it has been used in several works (Areces et al., 2025; Jung et al., 2023; Ramalingam et al., 2025).

More importantly though, it is amenable to an analysis of miscoverage. That is, the linear model provides a link between subgradients and miscoverage. Denote $\mathbf{c}_t = \mathbf{c}(X_t)$ and $l_t(\boldsymbol{\theta}_t, S_t) := \max\{(1-\alpha)(S_t - \langle \boldsymbol{\theta}_t, \mathbf{c}_t \rangle), \alpha(\langle \boldsymbol{\theta}_t, \mathbf{c}_t \rangle - S_t)\}$ to be the quantile loss at $\boldsymbol{\theta}_t$ (i.e., the loss incurred by the radius $\tau_t(X_t)$). A subgradient vector $\mathbf{g}_t \in \mathbb{R}^k$ of l_t at $\boldsymbol{\theta}_t$ is

$$\mathbf{g}_t = [\mathbf{1}\{S_t \leq \langle \boldsymbol{\theta}_t, \mathbf{c}_t \rangle\} - (1-\alpha)]\, \mathbf{c}_t. \tag{6}$$

Analogous to marginal OCP, these subgradients quantify the group miscoverage errors. Namely, the j^{th} entry of the vector $\left|\sum_{t=1}^{T} \mathbf{g}_t\right|$ quantifies the miscoverage error of group j:

$$\mathsf{MisCov}_T(j) := \left| \frac{1}{T_j} \sum_{t=1}^{T} \mathbf{1}\{S_t \leq \langle \boldsymbol{\theta}_t, \mathbf{c}_t \rangle\} c_{t,j} - (1-\alpha) \right| = \frac{1}{T_j} \left| \sum_{t=1}^{T} g_{t,j} \right|. \tag{7}$$

Thus, learning $\boldsymbol{\theta}_t$ such that each coordinate of the cumulative subgradient $|\sum_{t=1}^{T} \mathbf{g}_t|$ grows slowly will yield group coverage. The remaining question is whether this can be achieved with a parameter-free algorithm. Recall from Sec. 3.1, UP-OCP learns scalar radii through a wealth process and betting. POGO generalizes UP-OCP by maintaining k wealth processes and maximizing each via UP. In doing so, it learns k-dimensional parameters $(\boldsymbol{\theta}_t)_{t\geq 1}$ that provide group coverage guarantees.

4.1 Group Conditional Online Conformal via POGO

POGO learns parameters $(\boldsymbol{\theta}_t)_{t\geq 1}$ that control each coordinate of $|\sum_{t=1}^{T} \mathbf{g}_t|$ by playing k sequential betting games in parallel. For each group $j \in [k]$, POGO maintains a wealth process $(W_{t,j})_{t\geq 1}$, initialized at $W_{0,j} = 1/k$. At time t, in the j-th game, POGO chooses a signed betting fraction $\beta_{t,j} \in [-1/1-\alpha, 1/\alpha]$ of the current wealth $W_{t-1,j}$ in that game, and places this bet on $-g_{t,j}$. The fraction

[3] Liu et al. (2026) assume $0 \leq S_t \leq Dt^q$ for $D > 0,\ q \geq 0$. Previous works assume that $S_t \in [0, D]$.

[4] If $c_j(X_t) = 0$ for all j, then $\tau_t = 0$ which will typically miscover. This is prevented by augmenting the group vector to $\mathbf{c}'_t = (1, \mathbf{c}_t)^\top$, which ensures that we provide at least marginal coverage.

Algorithm 1 Portfolios for Online Group Conformal (POGO)

Input: Target miscoverage level α.

Initialize: Wealths $W_{0,j} \leftarrow 1/k$ for all $j \in [k]$ and Jeffrey prior $\mu(\bar{\lambda}) = 1/(\bar{\lambda}(1-\bar{\lambda}))^{1/2}$

for $t = 1, \ldots$ **do**
  for $j = 1, \ldots, k$ **do**
    Define $w_j : \mathbb{R} \to \mathbb{R}$ as $w_j(\bar{\lambda}) = \prod_{i=1}^{t-1} \left(\bar{\lambda}\left(1 - \frac{g_{i,j}}{\alpha}\right) + (1-\bar{\lambda})\left(1 + \frac{g_{i,j}}{1-\alpha}\right)\right)$
    Update portfolio $\lambda_{t,j} \leftarrow \frac{\int_0^1 \bar{\lambda} w_j(\bar{\lambda}) d\mu(\bar{\lambda})}{\int_0^1 w_j(\bar{\lambda}) d\mu(\bar{\lambda})}$
    Update j-th betting fraction $\beta_{t,j} = \frac{\lambda_{t,j} - \alpha}{\alpha(1-\alpha)}$ and j-th entry of $\boldsymbol{\theta}_t$, $\theta_{t,j} \leftarrow \beta_{t,j} W_{t-1,j}$
  end for
  Observe $(X_t, \mathbf{c}(X_t))$, compute radius $\tau_t \leftarrow \langle \boldsymbol{\theta}_t, \mathbf{c}(X_t) \rangle$
  Observe Y_t, compute score $S_t \leftarrow |Y_t - \hat{Y}_t|$ and subgradient $\mathbf{g}_t \leftarrow [\mathbf{1}\{S_t \le \tau_t\} - (1-\alpha)]\mathbf{c}_t$
  for $j = 1, \ldots, k$ **do**
    Update individual wealth processes: $W_{t,j} \leftarrow W_{t-1,j}\left(1 - \left(\frac{\lambda_{t,j} - \alpha}{\alpha(1-\alpha)}\right) g_{t,j}\right)$
  end for
end for

$\beta_{t,j}$ depends only on the history of subgradients until that point, $g_{1,j}, \ldots, g_{t-1,j}$. Accordingly, the j-th wealth process evolves as

$$W_{t,j} = W_{t-1,j}(1 - \beta_{t,j} g_{t,j}). \tag{8}$$

POGO is presented in Algorithm 1. Specifically, POGO uses UP to learn the betting fractions $(\beta_{t,j})_{t \ge 1}$ for each game, ensuring the logarithm of the corresponding wealth process, $\ln W_{t,j}$, grows quickly. The parameter vector $\boldsymbol{\theta}_t$ is defined coordinate-wise by $\theta_{t,j} = \beta_{t,j} W_{t-1,j}$. This construction ensures that each coordinate of the cumulative subgradient $\left|\sum_{t=1}^{T} \mathbf{g}_t\right|$ grows slowly, thereby controlling $\mathsf{MisCov}_T(j)$ for all $j \in [k]$ and yielding the following finite-time coverage guarantee.

Theorem 4.1 (POGO coverage guarantee)**.** *Let $\alpha \in (0,1)$, and let $(S_t, \mathbf{c}_t)_{t=1}^T$ be the sequence of non-conformity scores and group membership vectors observed by Algorithm 1, respectively. Let $D > 0$ and $q \ge 0$, and assume $S_t \le D t^q$, $\forall t \ge 1$ and $T_j > 0$. For every integer $T \ge 1$, define*

$$U_T(k) = \ln\left(1 + (1-\alpha)\frac{D(T+1)^{q+1}}{q+1}\right) + \frac{1}{2}\ln(\pi(T+1)) + \ln(k). \tag{9}$$

Then, POGO guarantees $\mathsf{MisCov}_T(j) \le \frac{1}{T_j}\left[U_T(k) + \sqrt{2T_j\alpha(1-\alpha)U_T(k)}\right]$.

We now compare POGO's guarantee with those of three standard G-OCP algorithms, GCACI (Ramalingam et al., 2025), GCACI with decaying learning rates (D-GCACI) (Angelopoulos et al., 2025), and Normalized-GCACI (N-GCACI) (Lyu & Cheung, 2026). At their core, all three methods use online subgradient descent, updating $\boldsymbol{\theta_t}$ with the rule $\boldsymbol{\theta_t} = \boldsymbol{\theta_{t-1}} - \eta \mathbf{g}_{t-1}$ ($\boldsymbol{\theta_1} = \mathbf{0}$). GCACI and N-GCACI use a fixed η while D-GCACI uses $\eta_t = \eta_0 t^{-a}$ for small $a \in (0, 1/2)$ and $\eta_0 \in \mathbb{R}$.

Table 1 presents the coverage guarantees for various methods. POGO's guarantee holds under a weaker growth condition $S_t \le D t^q$, for $D > 0$ and $q \ge 0$, whereas others rely on uniformly bounded scores, $S_t \in [0, D]$. *Thus, POGO's guarantee applies to a broader class of score sequences*. Nonetheless, in the interest of comparing methods, assume uniformly bounded scores, $S_t \in [0, D]$. Moreover, suppose α is small and each group occurs with non-vanishing frequency, i.e., $T_j \asymp \rho_j T$ for some constant $\rho_j > 0$. Then, all bounds vanish as $T \to \infty$, yielding asymptotic coverage. They differ, however, in their dependence on k and α, as well as in their large-T behavior.

Varying T, fixed α and k. POGO guarantees $\mathsf{MisCov}_T(j) = \mathcal{O}(\sqrt{\ln(T)/T})$, whereas, for fixed η and η_0, GCACI and D-GCACI guarantee rates of $\mathcal{O}(1/\sqrt{\eta T})$ and $\mathcal{O}(1/\sqrt{\eta_0 T^{1-a}})$, respectively. Thus, POGO's miscoverage bound decreases faster than D-GCACI's but slower than GCACI's. However, GCACI's guarantee depends on a fixed learning rate, which can be problematic. Relative to GCACI, POGO incurs an additional factor of $\sqrt{\ln T}$ in exchange for being parameter-free. This factor is *very* mild. D-GCACI's bound has the slowest rate, as expected: using decaying learning

Table 1: Coverage guarantees of different G-OCP methods for $\alpha < 1/2$ (omiting constants).

Method	Parameter-Free	Score Assumption	Coverage Rate
GCACI (Ramalingam et al., 2025)	✗	$S_t \in [0, D]$	$T_j^{-1}\sqrt{T(k(1-\alpha)^2 + 2(1-\alpha)/\eta)}$
D-GCACI (Angelopoulos et al., 2025)	✗	$S_t \in [0, D]$	$T_j^{-1}\sqrt{T^{2a}[(1-\alpha)^2 k \sum_{t=1}^{T} t^{-2a} + (2(1-\alpha)\sqrt{k}/\eta_0)\sum_{t=1}^{T} t^{-a}]}$
N-GCACI (Lyu & Cheung, 2026)	✗	$S_t \in [0, D]$	$T_j^{-1}\min\left\{\sqrt{Tk((1-\alpha)^2 + 2(1-\alpha)/\eta)}, 2^{k+1}(1+\eta)/\eta\right\}$
POGO	✓	$S_t \le Dt^q$	$T_j^{-1}(\ln kT + \sqrt{T_j \alpha \ln(kT)})$

rates trades convergence speed for stability and smaller prediction-set radii (Angelopoulos et al., 2024). N-GCACI's guarantee is better than GCACI's [5] for $T = \Omega(4^k/k)$, yielding $\mathsf{MisCov}_T(j) = \mathcal{O}((1+\eta)/(\eta T))$. However, this improvement occurs only at impractical time horizons (when $k = 25$, the crossover occurs only after approximately $T \approx 4.5 \times 10^{13}$ samples). Thus, it is also important to consider the smaller finite T regime. In the analysis that follows, we focus on the finite-horizon setting $T \lesssim 4^k/k$. Since, in this regime, the N-GCACI and D-GCACI miscoverage bounds decrease at a slower rate than GCACI's, we only compare POGO to GCACI.

Fixed α and varying k. POGO yields $\mathsf{MisCov}_T(j) = \mathcal{O}(\sqrt{(1/T)\ln(kT)})$, while GCACI with fixed η gives $\mathcal{O}(\sqrt{(1/T)(k + (1/\eta))})$. POGO's bound depends on k via $\ln k$ whereas GCACI's scales with $\sqrt{k}$ (and also depends on η). When k is very small (i.e. $k = 1$) GCACI may have a better dependence in T. In contrast, when k is moderate or large, POGO provides better guarantees.

Fixed k and $\alpha \to 0$. When k is fixed and $\alpha \to 0$, POGO's guarantees $\mathsf{MisCov}_T(c_j) = \mathcal{O}((\ln T)/T)$. By contrast, for fixed η, the GCACI bound is $\mathcal{O}(1/\sqrt{\eta T})$. Thus, as $\alpha \to 0$, POGO achieves exact asymptotic coverage at a faster rate than GCACI. Moreover, POGO's rate achieves the fastest known rate for time-averaged constraint violation (up to log factor) (Yu & Neely, 2020).

Lastly, let us consider the rare scenario where a group j does appear with a vanishing frequency in time T. Then, the miscoverage bound of GCACI will go to 0 as $T \to \infty$ as long as $T_j = \omega(\sqrt{T})$. On the other hand, the bound of POGO will approach 0 as $T \to \infty$ as long as $T_j = \omega(\ln T)$. Thus, in this scenario, POGO can provide coverage for groups that are rarer than GCACI can.

In conclusion, POGO incurs a mild $\sqrt{\ln T}$ factor in its finite-horizon coverage guarantee for two advantages. First, its guarantee remains valid under polynomially growing non-conformity scores, rather than requiring uniformly bounded scores. Second, POGO is parameter-free. Although GCACI's miscoverage bound may decrease faster for small T for suitably chosen η, selecting such η is difficult as the data stream may change arbitrarily. We illustrate this issue empirically in Section 5.

4.2 Implementation Details

At each time $t \in \{1, \dots, T\}$, POGO evaluates k integrals, each requiring $\mathcal{O}(T^2)$ operations by time T (Cover, 1991), leading to a total time complexity of $\mathcal{O}(kT^2)$. Such numerical integration is an approximation, whereas the coverage guarantee applies to the exact update. Fortunately, for binary groups $c_j(X_t) \in \{0, 1\}$, the update in Line 4 of Alg. 1 has the closed form (Liu et al., 2026, App. G)

$$\lambda_{t,j} = \frac{1}{\sum_{i=1}^{t-1} c_j(X_i) + 1}\left(\sum_{i=1}^{t-1} \mathbf{1}\{g_{i,j} = \alpha - 1\} + \frac{1}{2}\right). \tag{10}$$

Therefore, our experiments consider only the binary-group setting. For soft grouping functions $c_j(X_t) \in [0, 1]$, we provide details on approximating the integrals in Appendix E.

4.3 Alternative Naive Approaches

Before presenting our experiments, we address a natural question: *why* use POGO rather than running UP-OCP separately for each group? Concretely, for each group j, one can parameterize a radius as $\tau_{t,j} = \theta_{t,j}\, c_j(X_t)$, for $\theta_{t,j} \in \mathbb{R}$. Since c_j is fixed, learning $(\tau_{t,j})_{t\ge1}$ reduces to learning $(\theta_{t,j})_{t\ge1}$, which can be done with UP-OCP. Specifically, define the quantile loss $l_t(\theta_{t,j}, S_t)$. Then a subgradient of l_t at $\theta_{t,j}$ is $g_{t,j} = [\mathbf{1}\{S_t \le \theta_{t,j} c_j(X_t)\} - (1-\alpha)] \cdot c_j(X_t)$ and $\frac{1}{T_j}|\sum_{t=1}^{T} g_{t,j}|$ is precisely the miscoverage error of group j. With UP-OCP, one can learn a sequence $(\theta_{t,j})_{t\ge1}$ which yields a fast decreasing bound (in T) on $\frac{1}{T_j}|\sum_{t=1}^{T} g_{t,j}|$ and hence a coverage guarantee for group j.

[5] We suppress the dependence on α and η for simplicity.

However, running UP-OCP separately for each group will produce k sequences of radii $\{(\tau_{t,j})_{t=1}^{T}\}_{j=1}^{k}$. This is impractical: given a new sample X_t belonging to multiple groups, which radii should we use? Returning the maximum radius will likely over cover (Foygel Barber et al., 2021), and returning all complicates downstream decision-making. We therefore seek algorithms that return a *single* sequence of radii $(\tau_t)_{t\geq 1}$ that provide coverage for all groups $j \in [k]$. Moreover, if the algorithm is parameter-free, we overcome common issues with tuning learning rates in online settings. POGO satisfies these requirements, making it well-suited for G-OCP.

5 Experiments

We support our theoretical results with experiments on synthetic and real-world datasets.

Baselines. We compare POGO with GCACI, D-GCACI, and N-GCACI. We exclude MVP (Bastani et al., 2022) since GCACI outperforms MVP theoretically and empirically (Ramalingam et al., 2025). For the real-data experiments, we also compare with UP-OCP. We do not compare with UP-OCP in the synthetic experiments because they are designed so that marginal methods perform poorly (e.g., UP-OCP achieves group coverage rates of $\leq 70\%$).

Metrics. Evaluating methods only by whether they provide coverage is insufficient because arbitrarily large radii can achieve coverage. For a target level $1-\alpha$, a G-OCP method should (i) attain $1-\alpha$ coverage for all groups, (ii) produce small radii, and (iii) quickly react to changes in the data. To evaluate coverage, we report the worst observed group coverage: the coverage of the group with the largest absolute deviation from the target. To summarize the tradeoffs between coverage, radii length, and reactivity, we provide Pareto-frontier plots of observed coverage versus radius length and reactivity (Liu et al., 2026). We quantify reactivity by computing, for each group, the longest consecutive run of miscoverage events $\{S_t > \tau_t\}$, and reporting the maximum (lower is better).

5.1 Synthetic experiments

Groups are modeled to have different frequencies. Denote $\mathcal{U} \subset [k]$ to be the set of under-represented groups. At each time t, we generate a group vector $\mathbf{c}(X_t) = (c_1(X_t), \ldots, c_k(X_t)) \in \{0,1\}^k$, where $c_j(X_t) \sim \text{Bernoulli}(p_j)$ independently across $j \in [k]$ with $p_j = 0.05$ if $j \in \mathcal{U}$ and 0.25 if $j \in [k] \setminus \mathcal{U}$. Thus, groups in $\mathcal{U}$ appear infrequently, making it challenging to ensure their coverage.

We generate non-conformity scores following (Angelopoulos et al., 2023; Liu et al., 2026), as:

$$S_t = \max\{0, S_t^{\text{base}} + \langle \mathbf{b}(t), \mathbf{c}(X_t)\rangle\}. \tag{11}$$

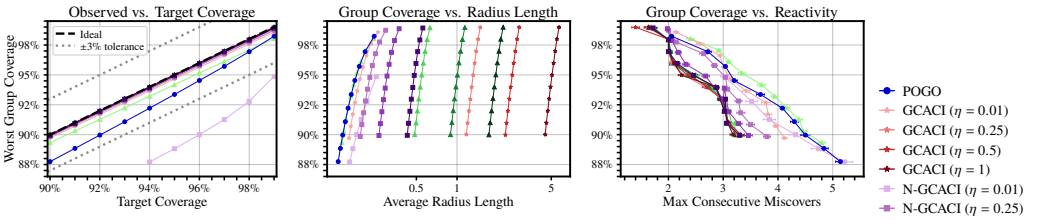
Here, $S_t^{\text{base}} \sim \text{Beta}(1, 20)$ and $\mathbf{b}(t) = (b_1(t), \ldots, b_k(t))$ introduces time-varying, group-specific shifts. Specifically, $\mathbf{b}(t)$ is a first-order auto-regressive (AR(1)) process with bounded uniform noise. The specific parameters of $\mathbf{b}(t)$ are provided in Appendix D.1. The key detail is that $\mathbf{b}(t)$ only alters the scores of the under-represented groups: for $j \in \mathcal{U}$, $b_j(t)$ is an AR(1) process whereas for $j \notin \mathcal{U}$, $b_j(t) \sim \text{Unif}(-1/10, 1/10)$ for all t. We set $|\mathcal{U}| = \lfloor \frac{1}{20} k \rfloor$ and run all methods for target coverage levels $1-\alpha \in \{0.90, 0.91, \ldots, 0.99\}$ and $T = 50{,}000$. The following results are for $k = 50$ groups and presented as averages over 50 independent runs (other k are in Appendix D.2).

Bounded-$[0,1]$ scores. We model time-varying group shifts with no drift. To ensure a fair comparison to the GCACI-based methods, we clip the scores to $[0,1]$. Fig. 1a (left) shows that, with one exception, all methods, including POGO, provide group coverage within 3% of the target. Moreover, POGO produces the smallest average radii at all observed coverage levels (Fig. 1a, middle) and remains competitive in terms of reactivity: its longest sequence of miscovers across all groups is within 2 of that of the best GCACI-based method (Fig. 1a, right).

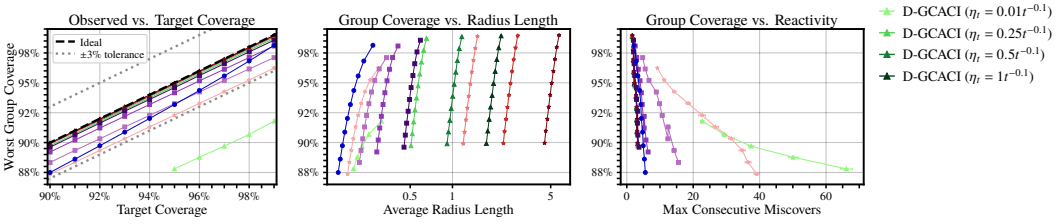
Bounded-$[0,1]$ scores with a shift. In the previous setting, in addition to POGO, GCACI ($\eta = 0.01$), N-GCACI ($\eta = 0.25$), and D-GCACI ($\eta_t = 0.01t^{-0.1}$) achieved favorable tradeoffs between coverage, radius length, and the longest miscoverage sequence. We now model an abrupt shift after a long stable period to illustrate the risks of committing to learning rates. The scores follow the same model until $t = \lfloor T/3 \rfloor$, after which the scores of one group $j^* \in \mathcal{U}$ increase. Specifically,

$$S_t = \max\{0, S_t^{\text{base}} + \langle \mathbf{b}(t), \mathbf{c}(X_t)\rangle + 0.6 \cdot \mathbf{1}\{t \geq \lfloor T/3 \rfloor\} c_{j^*}(X_t)\} \tag{12}$$

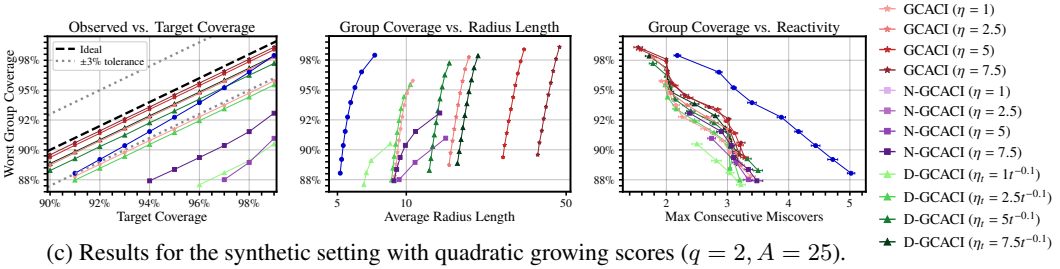
Now, D-GCACI ($\eta_t = 0.01t^{-0.1}$), does not provide adequate group coverage, whereas the other methods do (Fig. 1b (left)). POGO again provides the best coverage–radius tradeoff, outperforming GCACI ($\eta = 0.01$) and N-GCACI ($\eta = 0.25$). Although these two methods still provide coverage



(a) Results for the synthetic setting with bounded, gradually varying scores.



(b) Results for the synthetic setting with bounded scores and a changepoint.



(c) Results for the synthetic setting with quadratic growing scores ($q = 2, A = 25$).

Figure 1: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

and attain small radii (Fig. 1b (middle)), they respond poorly to the shift, especially GCACI ($\eta = 0.01$). They exhibit longer runs of miscovers, particularly at lower coverage levels (Fig. 1b (right)). POGO maintains short runs and competes with the best-performing methods. These results illustrate a key benefit of POGO: a learning rate η selected in a no-shift regime can yield poor reactivity after a shift, whereas POGO remains robust across coverage, radius, and miscoverage-sequence metrics.

Polynomial growing scores. We now select a group $j^* \in \mathcal{U}$ and generate scores according to

$$S_t = \max\{0, S_t^{\text{base}} + \langle \mathbf{b}(t), \mathbf{c}(X_t) \rangle + A(1 + \nu_t)\,(t/T)^q\, c_{j^*}(X_t)\}, \tag{13}$$

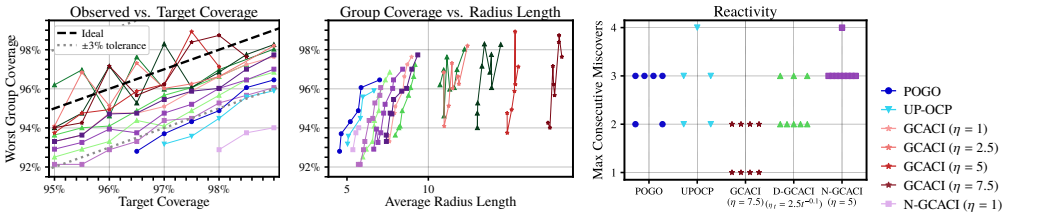
where $\nu_t \sim \text{Unif}(-1/2, 1/2)$, $q = 2$, and $A = 25$. The scores of group j^* grow quadratically, up to the multiplicative noise $1 + \nu_t$, with A controlling the magnitude of the shift. To ensure a fair comparison, we consider $\eta, \eta_0 \in \{1, 2.5, 5, 7.5\}$ for GCACI, N-GCACI, and D-GCACI, so they can accommodate the growing scores. Results for additional A and $q = 3$ are deferred to Appendix D.2.

Several of the GCACI-based methods with smaller η, η_0 fail to provide group coverage within 3% of the target (Fig. 1c (left)). Among the methods that do, D-GCACI ($\eta_t = 5t^{-0.1}$) and GCACI ($\eta = 2.5$) achieve the most favorable coverage-radius length tradeoff. POGO, however, outperforms them: it provides group coverage while producing radii 50% shorter than the best GCACI-based methods (Fig. 1c (middle)). Although the results suggest that a better η for GCACI or η_0 for D-GCACI may lie in $[1, 2.5]$, they also indicate such variants would produce larger radii than POGO.

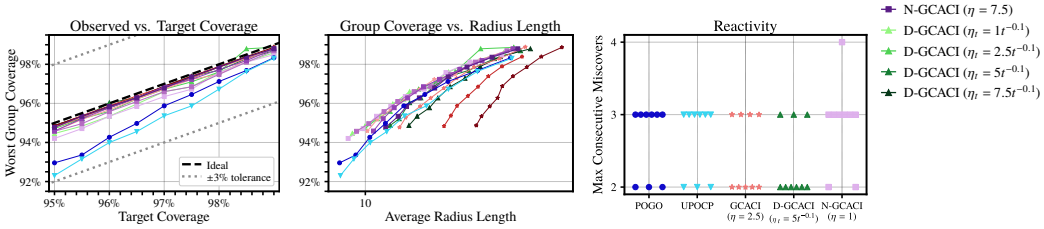
5.2 Real World Data

For our real data experiments, we first learn a base model $\hat{Y} = f(X)$ (training) and then compare POGO, UP-OCP [6], and GCACI on a test dataset. We run all methods for target coverage levels

[6] UP-OCP only provides marginal coverage – to compute group quantities, we run UP-OCP at a target level α to generate radii $(\tau_t)_{t \geq 1}$. Then, for each group $j \in [k]$, we compute $|\frac{1}{T_j} \sum_{t=1}^{T} \mathbf{1}\{S_t \leq \tau_t\} c_j(X_t)|$.



(a) Results for Nvidia (NVDA) stock.



(b) Results for MIMIC-IV.

Figure 2: Performance across target coverage levels $1-\alpha \in \{0.95, 0.955, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle:** worst-group coverage vs. radius length (curves closer to the top-left indicate better performance). **Right:** Max consecutive miscovers across target coverage levels. For clarity, among the GCACI, N-GCACI, and D-GCACI families, we show only the variant with the lowest median max consecutive miscovers. Data point shown if coverage ≥ 0.92.

$1-\alpha \in \{0.95, 0.955, ..., 0.99\}$. Since, a priori, we do not know the bounds of the non-conformity scores, we consider $\eta \in \{1, 2.5, 5, 7.5\}$ for GCACI.

Length-of-Stay in the Intensive Care Unit. We consider an MLP model that predicts the length of stay in the intensive care unit. We use the MIMIC-IV dataset (Johnson et al., 2023), a de-identified dataset of patients at BIDMC, collected within 24 hours of admission. We define groups using self-reported race, insurance status, and biological sex. Additional details are provided in Appendix D.5.

Stock Market. We consider a model that predicts the daily opening price of a stock. We use stock data of Nvidia (NVDA), Meta (META), and others from the last 15 years (Nugent, 2018). We use Prophet (Taylor & Letham, 2018), a popular model to produce daily opening-price forecasts, trained on the log opening prices of the first 100 trading days and then refit daily (NYSE, 2026). Groups are defined using market indicators such as volatility relative to a rolling median, calendar-year markers such as day-of-week, and indices of trading days from the first Monday (Bastani et al., 2022; Ramalingam et al., 2025). Additional details and results for other stocks are in Appendix D.3.

Fig. 2 shows that POGO is competitive with the other methods across both experiments. POGO provides slightly smaller radii but undercovers by 1–2%. Among the GCACI-based methods, those with smaller η and η_0 offer favorable trade-offs among coverage, radius length, and reactivity. Importantly, *these learning rates cannot be selected reliably in advance*. Being learning-rate-free, POGO avoids this problem while remaining competitive with these favorable GCACI variants.

6 Conclusion

We introduced POGO, the first parameter-free algorithm for group-conditional online conformal prediction. By combining ideas from online conformal prediction and parameter-free online learning, POGO provides rigorous coverage guarantees for all groups without tuning learning rates. Our analysis establishes the strongest known finite-time group-coverage bounds, and experiments on synthetic and real-world data show that POGO consistently attains the target coverage level for every group while maintaining competitive prediction interval lengths and reactivity relative to non–parameter-free baselines. Looking ahead, it remains unknown if stronger group-coverage guarantees are possible using richer models for the radius beyond the linear parameterization. Extending our theory and tools to other forms of online uncertainty quantification, such as calibration, is also an interesting direction. Finally, studying how POGO can support online decision-making in safety-critical applications like clinical decision support systems is a promising avenue for impact.

AI USE STATEMENT

In this work, we used generative AI tools for generating synthetic data sets, assisting in the writing of proofs, and implementing methods. We have not used generative AI tools for helping develop theoretical models or conceptual frameworks, formulating mathematical claims, providing critical ingredients for proving mathematical claims, proposing or refining hypotheses, and designing or providing feedback on research methodology or experiments. The rest of the required disclosure takes, including assistance with translation, cleaning and reformatting datasets, supporting qualitative and thematic data analysis, and interpreting results are not applicable to this work. Additionally, we used generative AI tools for editing particular sentences for readability, and creating code to run experiments and generate scientific figures. We have reviewed all AI-assisted work. We checked an LLM-generated list of works related to ours to ensure we are up to date and made all necessary comparisons. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Samantha Anderer. 2024-2025 flu season marked highest hospitalization rate in nearly 15 years, 2025.

Anastasios Angelopoulos, Emmanuel Candes, and Ryan J Tibshirani. Conformal pid control for time series prediction. *Advances in neural information processing systems*, 36:23047–23074, 2023.

Anastasios N Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023.

Anastasios N. Angelopoulos, Michael I. Jordan, and Ryan J. Tibshirani. Gradient equilibrium in online learning: Theory and applications. *Journal of Machine Learning Research*, 26(305):1–68, 2025. URL http://jmlr.org/papers/v26/25-0356.html.

Anastasios Nikolas Angelopoulos, Rina Barber, and Stephen Bates. Online conformal prediction with decaying step sizes. In *International Conference on Machine Learning*, pp. 1616–1630. PMLR, 2024.

Felipe Areces, Christopher Mohri, Tatsunori Hashimoto, and John Duchi. Online conformal prediction via online optimization. In *International Conference on Machine Learning*, pp. 1604–1649. PMLR, 2025.

Osbert Bastani, Varun Gupta, Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Practical adversarial multivalid conformai prediction. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, pp. 29362–29373, 2022.

Aadyot Bhatnagar, Huan Wang, Caiming Xiong, and Yu Bai. Improved online conformal prediction via strongly adaptive online learning. In *International Conference on Machine Learning*, pp. 2337–2363. PMLR, 2023.

Siddharth Bhatore, Lalit Mohan, and Y Raghu Reddy. Machine learning techniques for credit risk evaluation: a systematic literature review. *Journal of Banking and Financial Technology*, 4(1): 111–138, 2020.

Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

Thomas M Cover. Universal portfolios. *Mathematical finance*, 1(1):1–29, 1991.

Thomas M Cover and Erik Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2):348–363, 1996.

Rina Foygel Barber, Emmanuel J Candes, Aaditya Ramdas, and Ryan J Tibshirani. The limits of distribution-free conditional predictive inference. *Information and Inference: A Journal of the IMA*, 10(2):455–482, 2021.

Isaac Gibbs and Emmanuel J Candès. Adaptive conformal inference under distribution shift. In *Proceedings of the 35th International Conference on Neural Information Processing Systems*, pp. 1660–1672, 2021.

Isaac Gibbs and Emmanuel J Candès. Conformal inference for online prediction with arbitrary distribution shifts. *Journal of Machine Learning Research*, 25(162):1–36, 2024.

Isaac Gibbs, John J Cherian, and Emmanuel J Candès. Conformal prediction with conditional guarantees. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 87(4):1100–1126, 2025.

Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102(477):359–378, 2007.

Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. A survey of deep learning techniques for autonomous driving. *Journal of field robotics*, 37(3):362–386, 2020.

Varun Gupta, Christopher Jung, Georgy Noarov, Mallesh M Pai, and Aaron Roth. Online multivalid learning: Means, moments, and prediction intervals. *Innovations in Theoretical Computer Science (ITCS)*, 2022.

Na Hong, Chun Liu, Jianwei Gao, Lin Han, Fengxiang Chang, Mengchun Gong, and Longxiang Su. State of the art of machine learning–enabled clinical decision support in intensive care units: literature review. *JMIR medical informatics*, 10(3):e28781, 2022.

Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics: Methodology and distribution*, pp. 492–518. Springer, 1992.

Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. Mimic-iv, a freely accessible electronic health record dataset. *Scientific data*, 10(1):1, 2023.

Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Improved strongly adaptive online learning using coin betting. In *Artificial Intelligence and Statistics*, pp. 943–951. PMLR, 2017a.

Kwang-Sung Jun, Francesco Orabona, Stephen Wright, and Rebecca Willett. Online learning for changing environments using coin betting. *Electronic Journal of Statistics*, 11:5282–5310, 2017b.

Christopher Jung, Changhwa Lee, Mallesh Pai, Aaron Roth, and Rakesh Vohra. Moment multicalibration for uncertainty estimation. In *Conference on Learning Theory*, pp. 2634–2678. PMLR, 2021.

Christopher Jung, Georgy Noarov, Ramya Ramalingam, and Aaron Roth. Batch multivalid conformal prediction. In *International Conference on Learning Representations (ICLR)*, 2023.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Roger Koenker and Kevin F Hallock. Quantile regression. *Journal of economic perspectives*, 15(4): 143–156, 2001.

Jing Lei and Larry Wasserman. Distribution-free prediction bands for non-parametric regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 76(1):71–96, 2014.

Tuo Liu, Edgar Dobriban, and Francesco Orabona. Online conformal prediction via universal portfolio algorithms. *arXiv preprint arXiv:2602.03168*, 2026.

Lixing Lyu and Wang Chi Cheung. Normalization is enough: An improved algorithm for online group-wise conformal prediction. In *ICML 2026 Statistical Frameworks for Uncertainty in Agentic Systems*, 2026.

Cam Nugent. S&p 500 stock data, Feb 2018. URL https://www.kaggle.com/datasets/camnugent/sandp500.

NYSE. Holidays & trading hours, 2026. URL https://www.nyse.com/trade/hours-calendars.

Francesco Orabona. A modern introduction to online learning. *arXiv preprint arXiv:1912.13213*, 2019.

Francesco Orabona and Dávid Pál. Coin betting and parameter-free online learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 577–585, 2016.

Francesco Orabona and Dávid Pál. Parameter-free stochastic optimization of variationally coherent functions. *arXiv preprint arXiv:2102.00236*, 2021.

Harris Papadopoulos, Kostas Proedrou, Volodya Vovk, and Alex Gammerman. Inductive confidence machines for regression. In *European conference on machine learning*, pp. 345–356. Springer, 2002.

Aleksandr Podkopaev, Darren Xu, and Kuang-chih Lee. Adaptive conformal inference by betting. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 40886–40907, 2024.

Ramya Ramalingam, Shayan Kiyani, and Aaron Roth. The relationship between no-regret learning and online conformal prediction. *arXiv preprint arXiv:2502.10947*, 2025.

Yaniv Romano, Rina Foygel Barber, Chiara Sabatti, and Emmanuel Candès. With Malice Toward None: Assessing Uncertainty via Equalized Coverage. *Harvard Data Science Review*, 2(2), apr 30 2020. https://hdsr.mitpress.mit.edu/pub/qedrwcz3.

Aleksandr Rozenberg, Timothy Danish, Viktor Y Dombrovskiy, and Todd R Vogel. Outcomes after motor vehicle trauma: transfers to level i trauma centers compared with direct admissions. *The Journal of Emergency Medicine*, 53(3):295–301, 2017.

C Saunders, A Gammerman, and V Vovk. Transduction with confidence and credibility. In *Proceedings of the 16th international joint conference on Artificial intelligence-Volume 2*, pp. 722–726, 1999.

Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of machine learning research*, 9(3), 2008.

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

Graham Teasdale and Bryan Jennett. Assessment of coma and impaired consciousness: a practical scale. *The lancet*, 304(7872):81–84, 1974.

J-L Vincent, Rui Moreno, Jukka Takala, Sheila Willatts, Arnaldo De Mendonça, Hajo Bruining, C Kathryn Reinhart, PeterM Suter, and Lambertius G Thijs. The sofa (sepsis-related organ failure assessment) score to describe organ dysfunction/failure: On behalf of the working group on sepsis-related problems of the european society of intensive care medicine (see contributors to the project in the appendix). *Intensive care medicine*, 22(7):707–710, 1996.

Vladimir Vovk, David Lindsay, Ilia Nouretdinov, Alex Gammerman, et al. Mondrian confidence machine. *Technical Report*, 2003.

Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*. Springer, 2005.

Zitong Yang, Emmanuel Candès, and Lihua Lei. Bellman conformal inference: Calibrating prediction intervals for time series. *arXiv preprint arXiv:2402.05203*, 2024.

Hao Yu and Michael J Neely. A low complexity algorithm with o ($\sqrt{}$ t) regret and o (1) constraint violations for online convex optimization with long term constraints. *Journal of Machine Learning Research*, 21(1):1–24, 2020.

Margaux Zaffran, Olivier Féron, Yannig Goude, Julie Josse, and Aymeric Dieuleveut. Adaptive conformal predictions for time series. In *ICML 2022-39th International Conference on Machine Learning*, volume 162, 2022.

Zhiyu Zhang, David Bombara, and Heng Yang. Discounted adaptive online learning: towards better regularization. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 58631–58661, 2024.

A Appendix

B Proofs

B.1 Proof of Theorem 4.1

Theorem 4.1 (POGO coverage guarantee)**.** *Let $\alpha \in (0,1)$, and let $(S_t, \mathbf{c}_t)_{t=1}^T$ be the sequence of non-conformity scores and group membership vectors observed by Algorithm 1, respectively. Let $D > 0$ and $q \geq 0$, and assume $S_t \leq Dt^q$, $\forall t \geq 1$ and $T_j > 0$. For every integer $T \geq 1$, define*

$$U_T(k) = \ln\left(1 + (1-\alpha)\frac{D(T+1)^{q+1}}{q+1}\right) + \frac{1}{2}\ln(\pi(T+1)) + \ln(k). \tag{9}$$

Then, POGO guarantees $\mathsf{MisCov}_T(j) \leq \frac{1}{T_j}\left[U_T(k) + \sqrt{2T_j\alpha(1-\alpha)U_T(k)}\right]$.

Proof. Recall $\mathbf{c}_t = (c_{t,1}, \ldots, c_{t,k}) \in [0,1]^k$ and the subgradient vector $\mathbf{g}_t$ at time t is defined as

$$\mathbf{g}_t = Z_t\mathbf{c}_t, \quad \text{where} \quad Z_t = \mathbf{1}\{S_t \leq \langle\theta_t, \mathbf{c}_t\rangle\} - (1-\alpha). \tag{14}$$

Let $g_{t,j}$ denote the j^{th} entry of the subgradient vector $\mathbf{g}_t$, which, by definition above, is $g_{t,j} = c_{t,j}Z_t$. Define $w_{t,j,1} = 1 - \frac{g_{t,j}}{\alpha}$ and $w_{t,j,2} = 1 + \frac{g_{t,j}}{1-\alpha}$ to be the returns of two stocks and the j^{th} wealth process

$$W_{t,j} = W_{t-1,j}\left(\lambda_{t,j}w_{t,j,1} + (1-\lambda_{t,j})w_{t,j,2}\right) \quad \text{with} \quad W_{0,j} = \frac{1}{k}. \tag{15}$$

where $(\lambda_{t,j})_{t=1}^T$ is a sequence of portfolios in $[0,1]$. By this recursive definition of $W_{t,j}$, it is clear that $W_{T,j} = \frac{1}{k}\prod_{t=1}^T\left(\lambda_{t,j}w_{t,j,1} + (1-\lambda_{t,j})w_{t,j,2}\right)$. As a result,

$$\ln(W_{T,j}) = \ln\left(\frac{1}{k}\prod_{t=1}^T\left(\lambda_{t,j}w_{t,j,1} + (1-\lambda_{t,j})w_{t,j,2}\right)\right) \tag{16}$$

$$= \ln\left(\prod_{t=1}^T\left(\lambda_{t,j}w_{t,j,1} + (1-\lambda_{t,j})w_{t,j,2}\right)\right) + \ln\left(\frac{1}{k}\right). \tag{17}$$

Since the $\lambda_{t,j}$ are determined by Universal Portfolio (Cover, 1991; Cover & Ordentlich, 1996) with Jeffrey's prior, we have

$$\ln\left(\prod_{t=1}^T\left(\lambda_{t,j}w_{t,j,1} + (1-\lambda_{t,j})w_{t,j,2}\right)\right) \geq \max_{\lambda\in[0,1]}\ln\left(\prod_{s=1}^T\left(\lambda w_{t,j,1} + (1-\lambda)w_{t,j,2}\right)\right) \tag{18}$$

$$- \frac{1}{2}\ln(\pi(T+1)). \tag{19}$$

By the definitions of $w_{t,j,1}$ and $w_{t,j,2}$, along with noting that $g_{t,j} = c_{t,j}Z_t$, where $c_{t,j} \in [0,1]$ and $Z_t \in \{\alpha - 1, \alpha\}$, by Lemma C.1

$$\max_{\lambda\in[0,1]}\ln\left(\prod_{t=1}^T\left(\lambda w_{t,j,1} + (1-\lambda)w_{t,j,2}\right)\right) \geq \frac{\left(\sum_{t=1}^T g_{t,j}\right)^2}{2T_j\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^T g_{t,j}\right|}, \tag{20}$$

where $T_j = \sum_{t=1}^{T} c_{t,j}$. Therefore,

$$\ln(W_{T,j}) \geq \frac{\left(\sum_{t=1}^{T} g_{t,j}\right)^2}{2T_j\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} g_{t,j}\right|} + \ln\left(\frac{1}{k}\right) - \frac{1}{2}\ln(\pi(T+1)). \tag{21}$$

Applying $\exp(\cdot)$ to both sides and summing the k wealth processes $(W_{t,1})_{t=1}^{T}, \ldots (W_{t,k})_{t=1}^{T}$, gives

$$\sum_{j=1}^{k} W_{T,j} \geq \sum_{j=1}^{k} \frac{1}{k\sqrt{\pi}(T+1)^{1/2}} \exp\left(\frac{\left(\sum_{t=1}^{T} g_{t,j}\right)^2}{2T_j\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} g_{t,j}\right|}\right). \tag{22}$$

Thus we have a lower bound on $\sum_{j=1}^{k} W_{T,j}$.

Now, we have that for any $j \in [k]$, by the definitions of $w_{t,j,1}$ and $w_{t,j,2}$,

$$\lambda_{t,j} w_{t,j,1} + (1-\lambda_{t,j}) w_{t,j,2} = \lambda_{t,j} - \frac{\lambda_{t,j}}{\alpha} g_{t,j} + 1 - \lambda_{t,j} + \frac{1-\lambda_{t,j}}{1-\alpha} g_{t,j} \tag{23}$$

$$= 1 - \left(\frac{\lambda_{t,j}}{\alpha} - \frac{1-\lambda_{t,j}}{1-\alpha}\right) g_{t,j} \tag{24}$$

$$= 1 - \left(\frac{\lambda_{t,j} - \alpha}{\alpha(1-\alpha)}\right) g_{t,j}. \tag{25}$$

By defining $\theta_{t,j} = W_{t-1,j}\left(\frac{\lambda_{t,j}-\alpha}{\alpha(1-\alpha)}\right)$, we have

$$W_{T,j} = W_{T-1,j}(\lambda_{t,j} w_{t,j,1} + (1-\lambda_{t,j}) w_{t,j,2}) \tag{26}$$

$$= W_{T-1,j}\left(1 - \left(\frac{\lambda_{t,j} - \alpha}{\alpha(1-\alpha)}\right) g_{t,j}\right) \tag{27}$$

$$= W_{T-1,j} - \theta_{T,j} g_{T,j} \tag{28}$$

$$= W_{T-2,j} - \theta_{T-1,j} g_{T-1,j} - \theta_{T,j} g_{T,j} \tag{29}$$

$$= W_{0,j} - \sum_{t=1}^{T} \theta_{t,j} g_{t,j} = \frac{1}{k} - \sum_{t=1}^{T} \theta_{t,j} g_{t,j}. \tag{30}$$

Summing the individual $W_{T,j}$, we obtain

$$\sum_{j=1}^{k} W_{T,j} = \sum_{j=1}^{k}\left(\frac{1}{k} - \sum_{t=1}^{T} \theta_{t,j} g_{t,j}\right) \tag{31}$$

$$= 1 - \sum_{j=1}^{k}\left(\sum_{t=1}^{T} \theta_{t,j} g_{t,j}\right) \tag{32}$$

$$= 1 - \sum_{t=1}^{T}\sum_{j=1}^{k} \theta_{t,j} g_{t,j} = 1 - \sum_{t=1}^{T} \langle \boldsymbol{\theta}_t, \mathbf{g}_t \rangle, \tag{33}$$

where we define $(\boldsymbol{\theta}_t = (\theta_{t,1}, \ldots \theta_{t,k}))$. By Lemma C.2, and the assumption $S_t \leq Dt^q$ for $D > 0$ and $q \geq 0$, we obtain

$$1 - \sum_{t=1}^{T} \langle \theta_t, \mathbf{g}_t \rangle \leq 1 + (1-\alpha)\sum_{t=1}^{T} S_t \leq 1 + (1-\alpha)\frac{D(T+1)^{q+1}}{q+1}, \tag{34}$$

where we have used the integral inequality $\sum_{t=1}^{T} t^q \cdot 1 \leq \int_0^{T+1} t^q dt = \frac{(T+1)^{q+1}}{q+1}$. Therefore, we have an upper bound for $\sum_{j=1}^{k} W_{T,j}$.

Comparing the upper and the lower bounds,

$$\sum_{j=1}^{T} \frac{1}{k\sqrt{\pi}(T+1)^{1/2}} \exp\left(\frac{\left(\sum_{t=1}^{T} g_{t,j}\right)^2}{2T_j\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} g_{t,j}\right|}\right) \leq 1 + (1-\alpha)\frac{D(T+1)^{q+1}}{q+1} \tag{35}$$

Every term in the left hand side equation is non-negative. Therefore, for any $j \in [k]$,

$$\frac{1}{k\sqrt{\pi}(T+1)^{1/2}} \exp\left(\frac{\left(\sum_{t=1}^{T} g_{t,j}\right)^2}{2T_j\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} g_{t,j}\right|}\right) \leq 1 + (1-\alpha)\frac{D(T+1)^{q+1}}{q+1}. \tag{36}$$

Multiplying both sides of (36) by $k\sqrt{\pi}(T+1)^{1/2}$ and then applying $\ln(\cdot)$ to both sides, we obtain

$$\frac{\left(\sum_{t=1}^{T} g_{t,j}\right)^2}{2T_j\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} g_{t,j}\right|} \leq U_T(k), \tag{37}$$

where we define

$$U_T(k) \coloneqq \ln\left(1 + (1-\alpha)\frac{D(T+1)^{q+1}}{q+1}\right) + \frac{1}{2}\ln(\pi(T+1)) + \ln(k). \tag{38}$$

Equivalently,

$$\left(\sum_{t=1}^{T} g_{t,j}\right)^2 - \frac{2}{3}U_T(k)\left|\sum_{t=1}^{T} g_{t,j}\right| - 2U_T(k)T_j\alpha(1-\alpha) \leq 0. \tag{39}$$

The inequality above is of the form $z^2 + Q|z| + R \leq 0$, and the LHS might be non-convex depending on Q, R. Fortunately, we only require an upper bound on $|z|$, and we can look at the part of the curve in $z \geq 0$, which is a quadratic, and remains non-positive in $0 \leq z \leq z^\star$, where $z^\star = \frac{-Q+\sqrt{Q^2-4R}}{2}$.

$$\left|\sum_{t=1}^{T} g_{t,j}\right| \leq \frac{1}{3}U_T(k) + \sqrt{2U_T(k)T_j\alpha(1-\alpha) + \frac{1}{9}U_T(k)^2} \tag{40}$$

$$\leq U_T(k) + \sqrt{2T_j\alpha(1-\alpha)U_T(k)}, \tag{41}$$

where we have used the inequality $\sqrt{Q+R} \leq \sqrt{Q} + \sqrt{R}$, with $Q = 2T_j\alpha(1-\alpha)U_T(k)$, $R = \frac{1}{9}U_T(k)^2$. Finally, dividing both sides by T_j obtains the desired result. □

C Useful Lemmas and Inequalities

In this section, we present results that are used throughout the proofs in Section B. For cited results, their proofs can be found in the referenced works; otherwise, proofs are provided here.

C.1 Lower bound on the optimal wealth process

Lemma C.1. *Let $\alpha \in (0,1)$. Consider any integer $T \geq 1$ and a sequence $(c_t Z_t)_{t=1}^T$ where $c_t \in [0,1]$ and $Z_t \in \{\alpha - 1, \alpha\}$. Define $T_c = \sum_{t=1}^T c_t$ and assume $T_c > 0$. Then*

$$\max_{\lambda\in[0,1]} \ln\left(\prod_{t=1}^{T} \lambda\left(1 - \frac{c_t Z_t}{\alpha}\right) + (1-\lambda)\left(1 + \frac{c_t Z_t}{1-\alpha}\right)\right) \geq \frac{\left(\sum_{t=1}^{T} c_t Z_t\right)^2}{2T_c\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} c_t Z_t\right|}. \tag{42}$$

Proof. First, observe

$$\lambda\left(1 - \frac{c_t Z_t}{\alpha}\right) + (1-\lambda)\left(1 + \frac{c_t Z_t}{1-\alpha}\right) = \lambda - \frac{\lambda}{\alpha}c_t Z_t + 1 - \lambda + \frac{(1-\lambda)}{1-\alpha}c_t Z_t \tag{43}$$

$$= 1 - \left(\frac{\lambda}{\alpha} - \frac{1-\lambda}{1-\alpha}\right)c_t Z_t \tag{44}$$

$$= 1 - \left(\frac{\lambda - \alpha}{\alpha(1-\alpha)}\right)c_t Z_t. \tag{45}$$

Then, rearrange to obtain

$$1 - \frac{\lambda - \alpha}{\alpha(1-\alpha)} c_t Z_t = 1 - c_t + c_t - \frac{\lambda - \alpha}{\alpha(1-\alpha)} c_t Z_t \tag{46}$$

$$= 1 - c_t + \left(1 + \frac{\lambda - \alpha}{\alpha(\alpha - 1)} Z_t\right) \cdot c_t \tag{47}$$

$$= 1 - c_t + u_t \cdot c_t, \tag{48}$$

where $u_t = \frac{\lambda}{\alpha}$ when $Z_t = \alpha - 1$, and $u_t = \frac{1-\lambda}{1-\alpha}$ when $Z_t = \alpha$.

Since $c_t \in [0, 1]$, by the concavity of $\ln(x)$,

$$\ln((1 - c_t)(1) + c_t u_t) \geq (1 - c_t) \ln(1) + c_t \ln(u_t) = c_t \ln(u_t). \tag{49}$$

Therefore, for any $\lambda \in [0, 1]$

$$\max_{\lambda \in [0,1]} \ln \left(\prod_{t=1}^{T} \lambda \left(1 - \frac{c_t Z_t}{\alpha}\right) + (1 - \lambda) \left(1 + \frac{c_t Z_t}{1 - \alpha}\right) \right) \tag{50}$$

$$\geq \ln \left(\prod_{t=1}^{T} \lambda \left(1 - \frac{c_t Z_t}{\alpha}\right) + (1 - \lambda) \left(1 + \frac{c_t Z_t}{1 - \alpha}\right) \right) \tag{51}$$

$$= \sum_{t=1}^{T} \ln \left(1 - c_t + c_t u_t\right) \tag{52}$$

$$\geq \sum_{t=1}^{T} c_t \ln(u_t) \tag{53}$$

$$= \sum_{t: Z_t = \alpha} c_t \ln(u_t) + \sum_{t: Z_t = \alpha - 1} c_t \ln(u_t) \tag{54}$$

$$= \ln \left(\frac{1 - \lambda}{1 - \alpha}\right) T_1 + \ln \left(\frac{\lambda}{\alpha}\right) T_2, \tag{55}$$

where $T_1 = \sum_{\{t:\ Z_t = \alpha\}} c_t$ and $T_2 = \sum_{\{t:\ Z_t = \alpha - 1\}} c_t$. Recall that $T_c = T_1 + T_2$.

The lower bound in (55) holds for any $\lambda \in [0, 1]$: we instantiate it for λ^*, which is the maximizer of (55). To maximize (55), set the partial derivative with respect to λ to 0, to obtain that

$$\lambda^\star = \frac{T_2}{T_c} \in [0, 1]. \tag{56}$$

Now $T_1 = (1 - \lambda^\star) T_c$ and $T_2 = \lambda^\star T_c$. Substituting this, (55) can be rewritten as

$$T_c \left((1 - \lambda^*) \ln \left(\frac{1 - \lambda^*}{1 - \alpha}\right) + \lambda^* \ln \left(\frac{\lambda^*}{\alpha}\right) \right) = T_c \cdot \mathrm{KL}\left(\mathrm{Ber}(\lambda^*) || \mathrm{Ber}(\alpha)\right). \tag{57}$$

Further, we can express $\lambda^\star$ as a function of α and $\sum c_t Z_t$ using (56):

$$\sum_{t=1}^{T} c_t Z_t = \alpha T_1 + (\alpha - 1) T_2 = \alpha T_c - T_2 = \alpha T_c - \lambda^\star T_c. \tag{58}$$

Thus

$$T_c \cdot \mathrm{KL}\left(\mathrm{Ber}(\lambda^*) \,||\mathrm{Ber}(\alpha)\right) = T_c \cdot \mathrm{KL}\left(\mathrm{Ber}\left(\alpha - \frac{\sum_{t=1}^{T} c_t Z_t}{T_c}\right) \middle|\middle| \mathrm{Ber}(\alpha)\right) \tag{59}$$

By Lemma C.3

$$T_c \cdot \mathrm{KL}\left(\mathrm{Ber}\left(\alpha - \frac{\sum_{t=1}^{T} c_t Z_t}{T_c}\right) ||\mathrm{Ber}(\alpha)\right) \geq T_c \frac{\left(\frac{\sum_{t=1}^{T} c_t Z_t}{T_c}\right)^2}{2\alpha(1-\alpha) + \frac{2}{3}\left|\frac{\sum_{t=1}^{T} c_t Z_t}{T_c}\right|} \tag{60}$$

$$= \frac{\left(\sum_{t=1}^{T} c_t Z_t\right)^2}{2T_c\alpha(1-\alpha) + \frac{2}{3}\left|\sum_{t=1}^{T} c_t Z_t\right|} \tag{61}$$

Thus we have shown the result

$$\max_{\lambda\in[0,1]} \ln\left(\prod_{t=1}^{T} \lambda\left(1-\frac{c_t Z_t}{\alpha}\right)+(1-\lambda)\left(1+\frac{c_t Z_t}{1-\alpha}\right)\right) \geq \frac{\left(\sum_{t=1}^{T} c_t Z_t\right)^2}{2T_c\alpha(1-\alpha)+\frac{2}{3}\left|\sum_{t=1}^{T} c_t Z_t\right|} \tag{62}$$

□

C.2 Upper bound on the wealth process

Lemma C.2. *Let $\alpha \in (0,1)$. Consider any integer $T \geq 1$ and let $(S_t)_{t=1}^T$ and $(\mathbf{c}_t)_{t=1}^T$ be a sequence of non-conformity scores and group membership vectors. At every time t consider a vector $\boldsymbol{\theta}_t \in \mathbb{R}^k$ and define*

$$\mathbf{g}_t = [\mathbf{1}\{S_t \leq \langle \theta_t, \mathbf{c}_t\rangle\} - (1-\alpha)]\mathbf{c}_t. \tag{63}$$

Then

$$1 - \sum_{t=1}^{T} \langle \boldsymbol{\theta}_t, \mathbf{g}_t\rangle \leq 1 + (1-\alpha)\sum_{t=1}^{T} S_t. \tag{64}$$

Proof. It suffices to show

$$-\langle \boldsymbol{\theta}, \mathbf{g}_t\rangle \leq (1-\alpha)S_t, \quad \forall t \geq 1 \tag{65}$$

because, if true, then $-\sum_{t=1}^{T}\langle \boldsymbol{\theta}, \mathbf{g}_t\rangle \leq \sum_{t=1}^{T}(1-\alpha)S_t$, and the result is immediate.

By definition of $\mathbf{g}_t$, we have

$$\langle \theta_t, \mathbf{g}_t\rangle = \langle \theta_t, \mathbf{c}_t\rangle \cdot [\mathbf{1}\{S_t \leq \langle \theta_t, \mathbf{c}_t\rangle\} - (1-\alpha)]. \tag{66}$$

The range of $\langle \theta_t, \mathbf{g}_t\rangle$ is $[-\infty, \infty]$. We will analyze the three exhaustive cases of where this quantity is less than 0, in between 0 and S_t, and greater than zero. We will show that no matter the scenario, the inequality $\langle \boldsymbol{\theta}, \mathbf{g}_t\rangle \leq (1-\alpha)S_t$ holds $\forall t \geq 1$.

Case 1: Suppose $\langle \theta_t, \mathbf{c}_t\rangle < 0$. Then, since $S_t \geq 0$ because it is a non-conformity score, we have

$$[\mathbf{1}\{S_t \leq \langle \theta_t, \mathbf{c}_t\rangle\} - (1-\alpha)] = -(1-\alpha). \tag{67}$$

Thus, $-\langle \theta_t, \mathbf{g}_t\rangle = \langle \theta_t, \mathbf{c}_t\rangle(1-\alpha) \leq 0 \leq (1-\alpha)S_t$.

Case 2: Suppose $0 \leq \langle \theta_t, \mathbf{c}_t\rangle < S_t$. Then,

$$[\mathbf{1}\{S_t \leq \langle \theta_t, \mathbf{c}_t\rangle\} - (1-\alpha)] = -(1-\alpha) \tag{68}$$

and, as a result, $-\langle \theta_t, \mathbf{g}_t\rangle = \langle \theta_t, \mathbf{c}_t\rangle(1-\alpha) \leq (1-\alpha)S_t$.

Case 3: Suppose $\langle \theta_t, \mathbf{c}_t\rangle \geq S_t$. Then,

$$[\mathbf{1}\{S_t \leq \langle \theta_t, \mathbf{c}_t\rangle\} - (1-\alpha)] = \alpha. \tag{69}$$

Thus, $-\langle \theta_t, \mathbf{g}_t\rangle = -\langle \theta_t, \mathbf{c}_t\rangle\alpha \leq 0 \leq (1-\alpha)S_t$. □

C.3 A lower bound on the KL divergence between two Bernoulli random variables

Lemma C.3 ((Liu et al., 2026) Lemma F.1)**.** *Let $p, q \in [0,1]$. Then*

$$\mathrm{KL}(\mathrm{Ber}(p)\,||\mathrm{Ber}(q)) = p\ln\frac{p}{q} + (1-p)\ln\frac{1-p}{1-q} \geq \frac{(p-q)^2}{2q(1-q)+\frac{2}{3}|p-q|}. \tag{70}$$

Proof. See Liu et al. (2026) Appendix F.1. □

D Additional Experimental Details and Results

D.1 Additional synthetic experiment details

In the synthetic experiments presented in Section 5.1, we directly generate non-conformity scores. We model the scores as

$$S_t = \max\{0, S_t^{\text{base}} + \langle \mathbf{b}(t), \mathbf{c}(X_t) \rangle\}. \tag{71}$$

Here, $S_t^{\text{base}} \sim \text{Beta}(1, 20)$ and the vector $\mathbf{b}(t) = (b_1(t), \ldots, b_k(t))$ introduces time-varying, group-dependent shifts. Specifically, $\mathbf{b}(t)$ is modeled as a first-order autoregressive (AR(1)) process with bounded uniform noise:

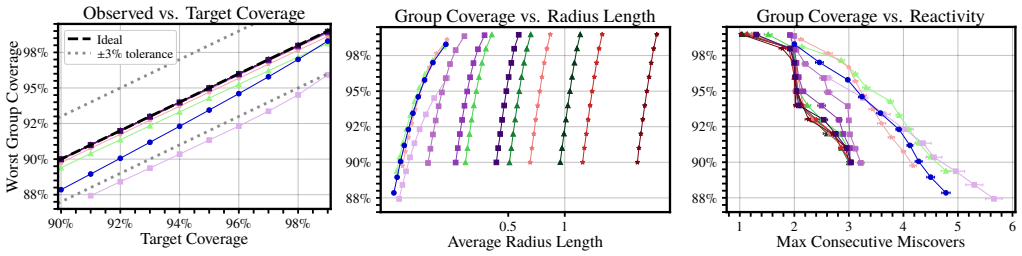
$$b_j(t) = (1 - \gamma_j)\, b_j(t-1) + \gamma_j\, m_j + \mu_j + \sigma_j\, \xi_{t,j}, \qquad b_j(0) = 0, \;\; \xi_{t,j} \overset{iid}{\sim} \text{Unif}(-1, 1), \tag{72}$$

At a high level, μ_j controls the overall drift, σ_j controls the volatility, and γ_j governs persistence. When $\gamma_j = 0$, (72) reduces to a random walk, whereas $\gamma_j > 0$ yields a shift that stabilizes toward m_j. With this model, scores belonging to different groups exhibit different time-varying shifts.

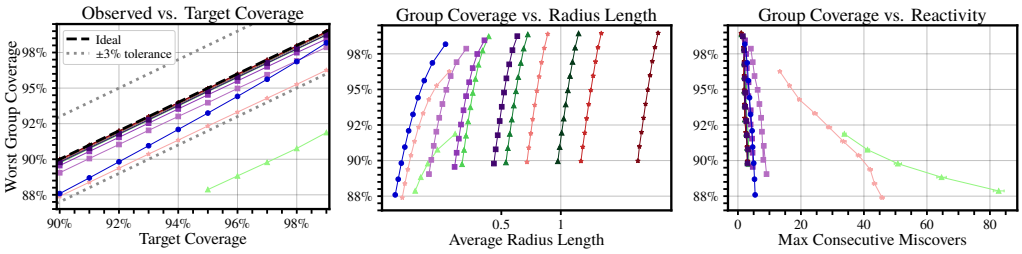
In the experiments, we set $\gamma_j = 0.1$, $m_j = 0.2$, $\mu_j = 10^{-3}$, and $\sigma_j = 10^{-2}$, yielding stochastic shifts that start at 0 and fluctuate around 0.2. Time-varying biases are applied *only* to the under-represented groups: for $j \in \mathcal{U}$, $b_j(t)$ follows (72), whereas for $j \notin \mathcal{U}$, $b_j(t) = \sigma_j \xi_{t,j}\ \forall t$.

D.2 Additional synthetic experiment results

Results for $k = 25$ number of groups.



(a) Results for synthetic setting with bounded, gradually varying scores.



(b) Results for synthetic setting with bounded scores with a changepoint.

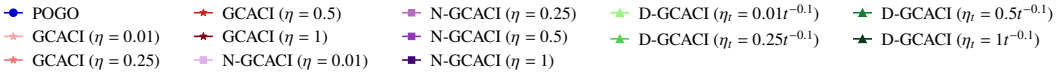
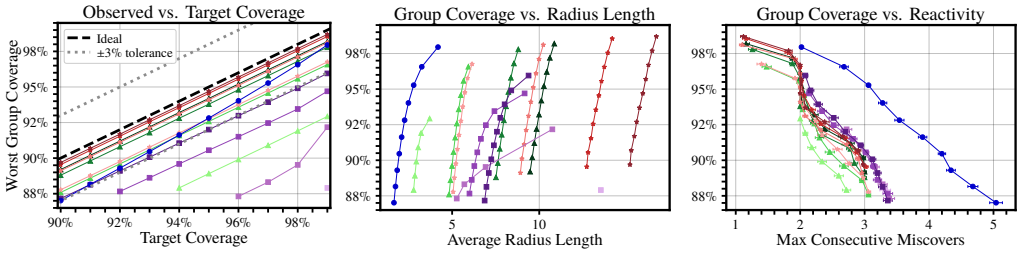
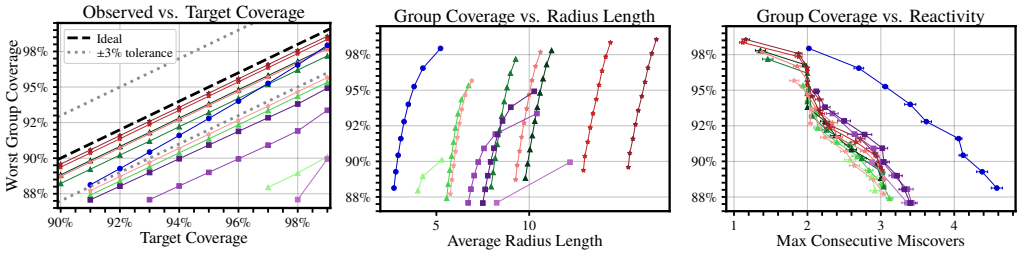


Figure 3: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

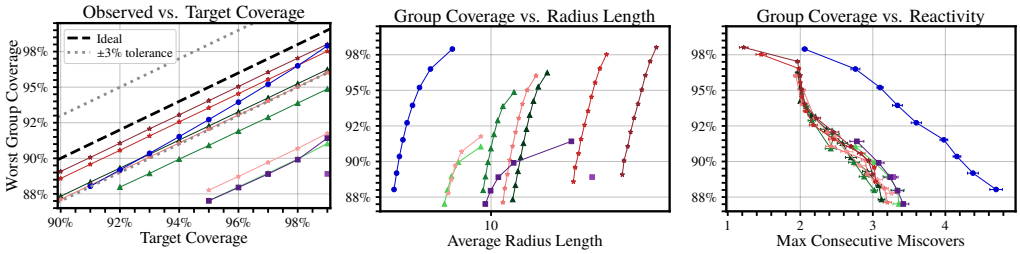
Results for $k = 25$ number of groups, quadratic growing scores.



(a) Results for the synthetic setting with slow growing quadratic scores ($q = 2, A = 5$).



(b) Results for the synthetic setting with moderately growing quadratic scores ($q = 2, A = 25$).



(c) Results for the synthetic setting with extreme growing quadratic scores ($q = 2, A = 100$).

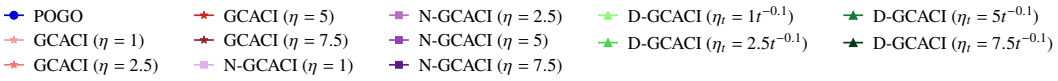
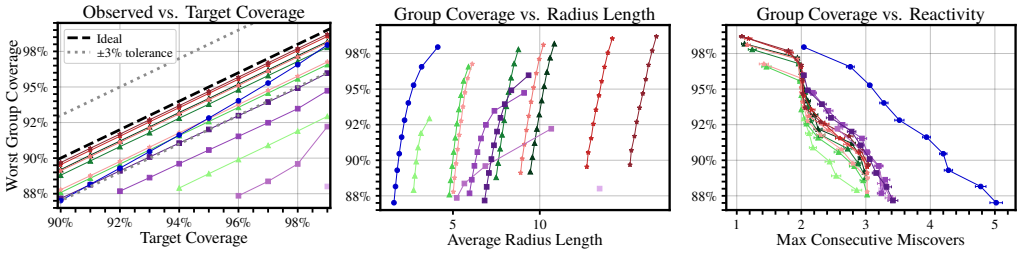
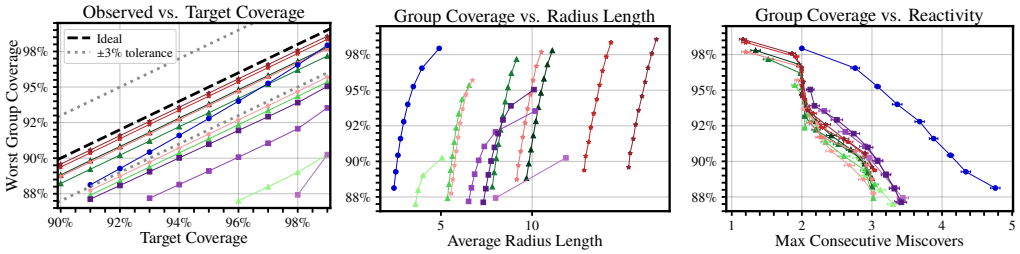


Figure 4: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

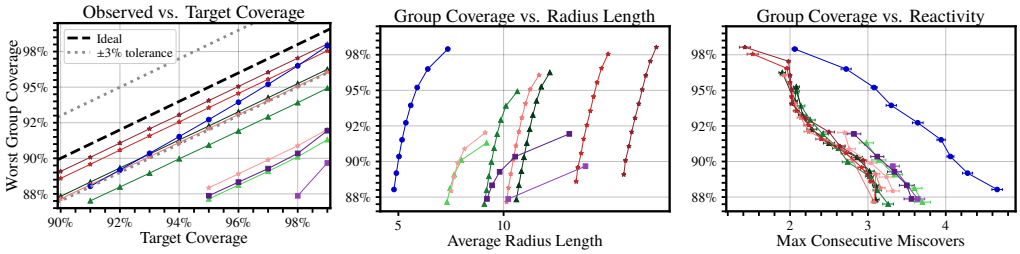
Results for $k = 25$ number of groups, cubic growing scores.



(a) Results for the synthetic setting with slow growing cubic scores ($q = 3, A = 5$).



(b) Results for the synthetic setting with moderately growing cubic scores ($q = 3, A = 25$).



(c) Results for the synthetic setting with extreme growing cubic scores ($q = 3, A = 100$).

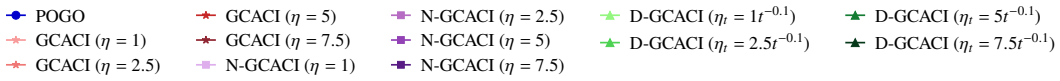
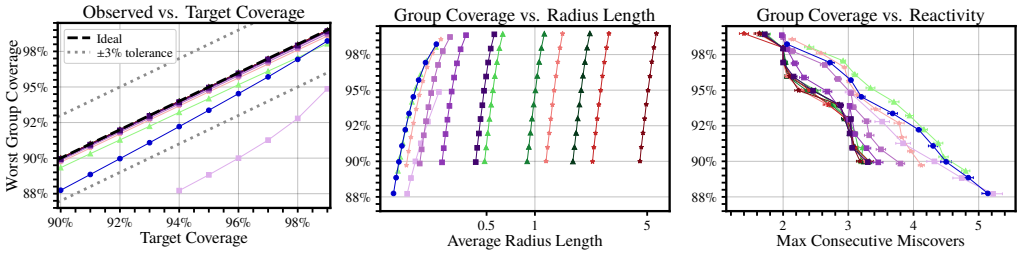
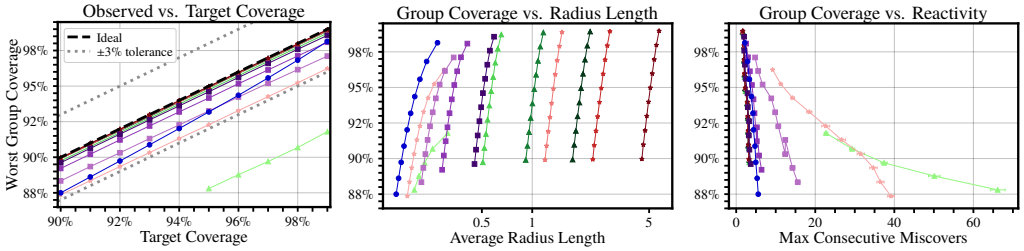


Figure 5: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

Results for $k = 50$ number of groups, bounded scores



(a) Results for synthetic setting with bounded, gradually varying scores.



(b) Results for synthetic setting with bounded scores with a changepoint.

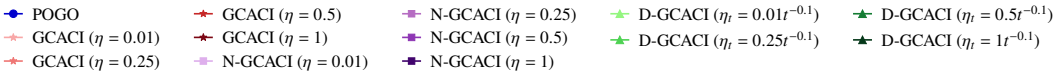
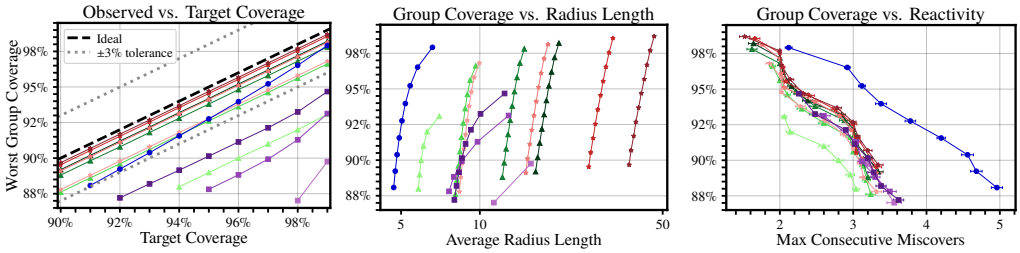
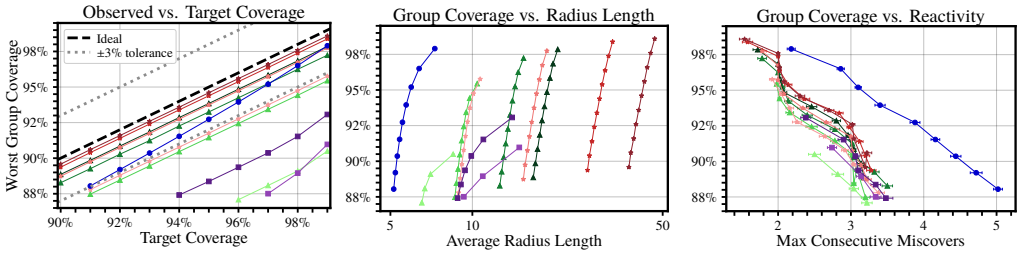


Figure 6: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

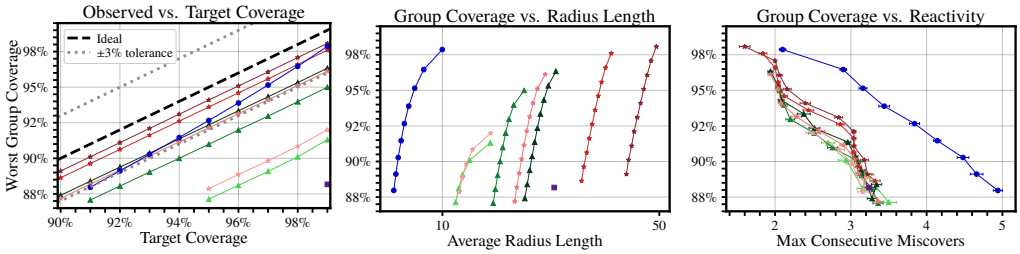
Results for $k = 50$ number of groups, quadratic growing scores.



(a) Results for the synthetic setting with slow growing quadratic scores ($q = 2, A = 5$).



(b) Results for the synthetic setting with moderately growing quadratic scores ($q = 2, A = 25$).



(c) Results for the synthetic setting with extreme growing quadratic scores ($q = 2, A = 100$).

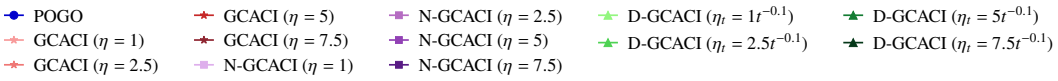
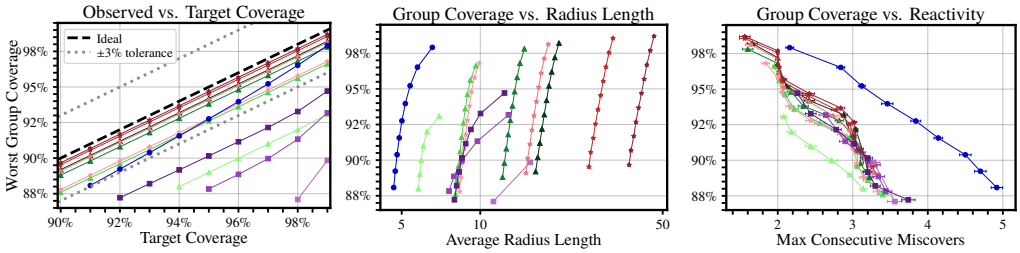
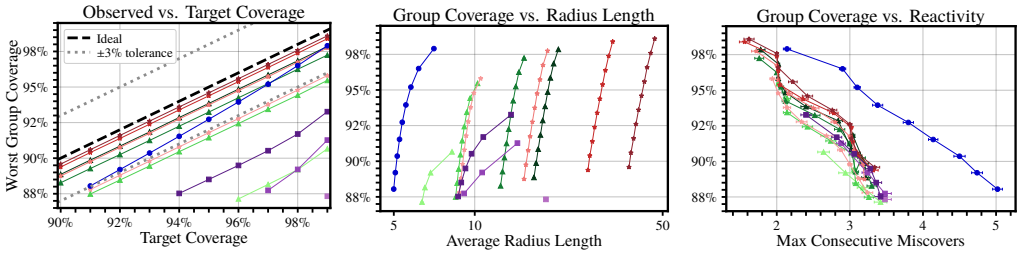


Figure 7: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

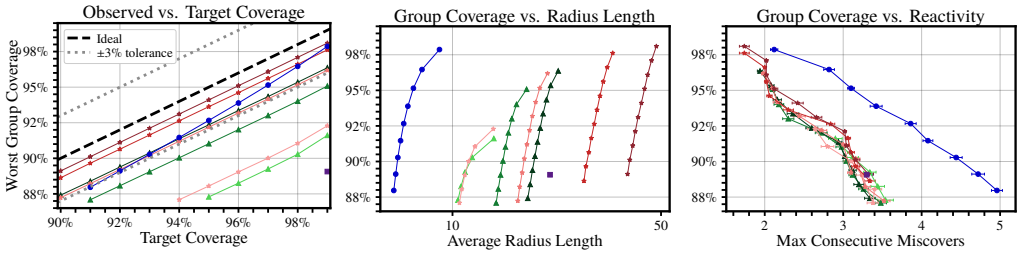
Results for $k = 50$ number of groups, cubic growing scores.



(a) Results for the synthetic setting with slow growing cubic scores ($q = 3, A = 5$).



(b) Results for the synthetic setting with moderately growing cubic scores ($q = 3, A = 25$).



(c) Results for the synthetic setting with extreme growing cubic scores ($q = 3, A = 100$).

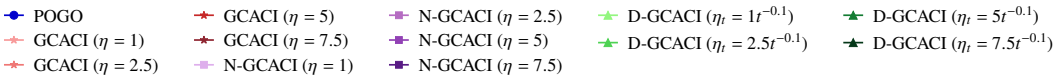
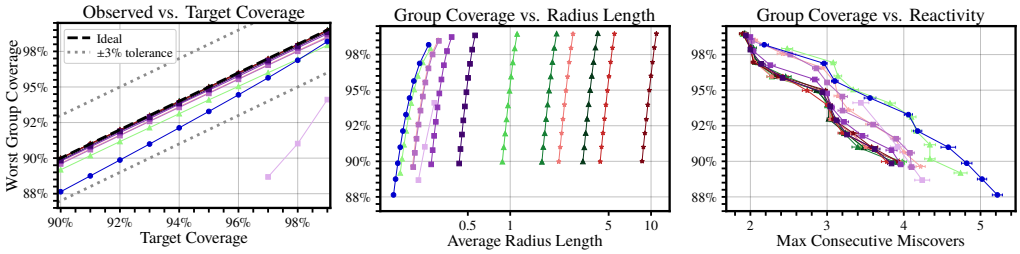
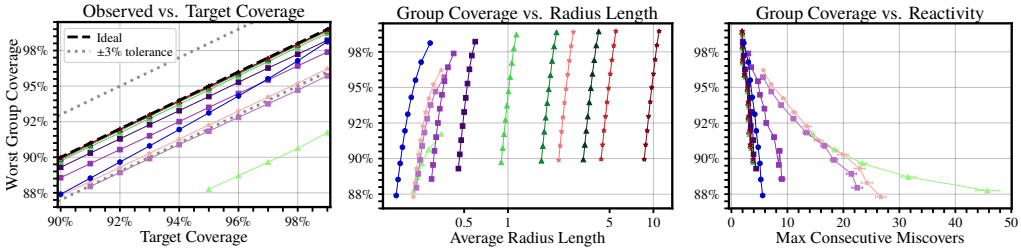


Figure 8: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

Results for $k = 100$ number of groups, bounded scores



(a) Results for synthetic setting with bounded, gradually varying scores.



(b) Results for synthetic setting with bounded scores with a changepoint.

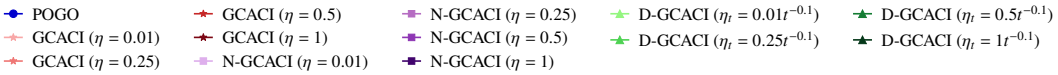
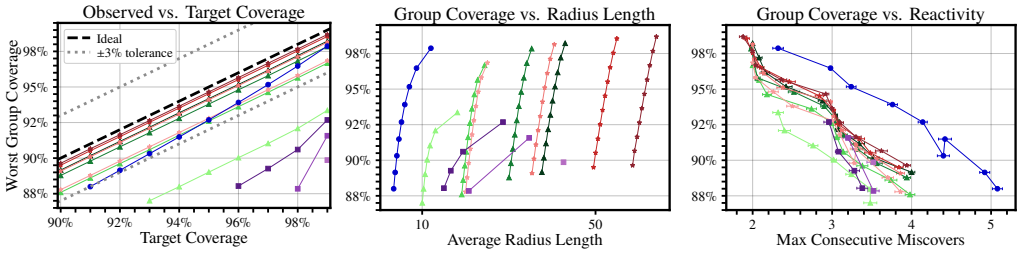
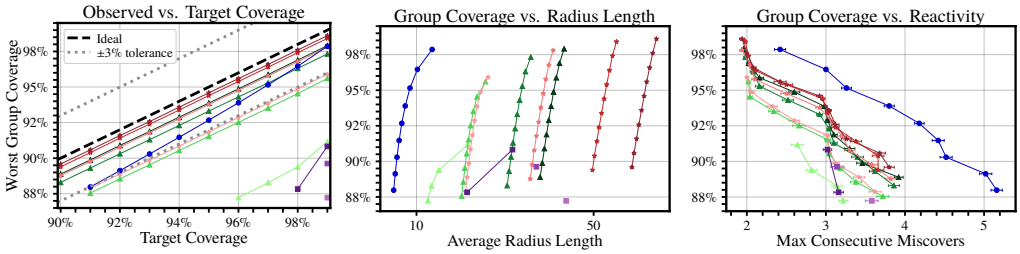


Figure 9: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

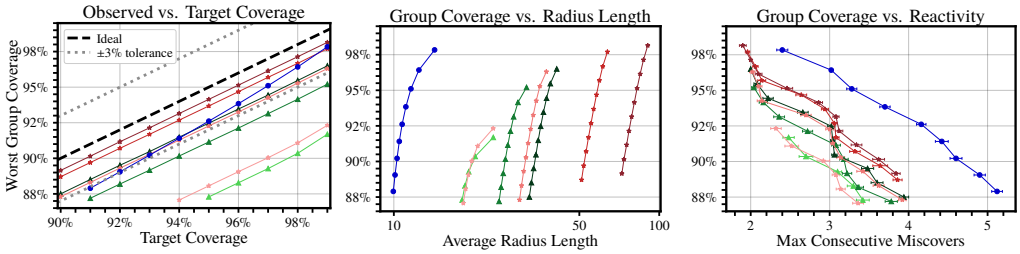
Results for $k = 100$ number of groups, quadratic growing scores.



(a) Results for the synthetic setting with slow growing quadratic scores ($q = 2, A = 5$).



(b) Results for the synthetic setting with moderately growing quadratic scores ($q = 2, A = 25$).



(c) Results for the synthetic setting with extreme growing quadratic scores ($q = 2, A = 100$).

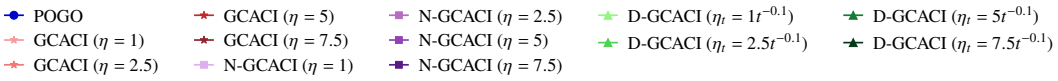
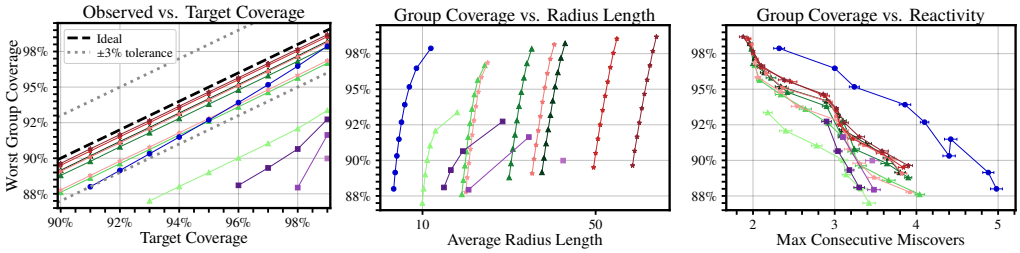
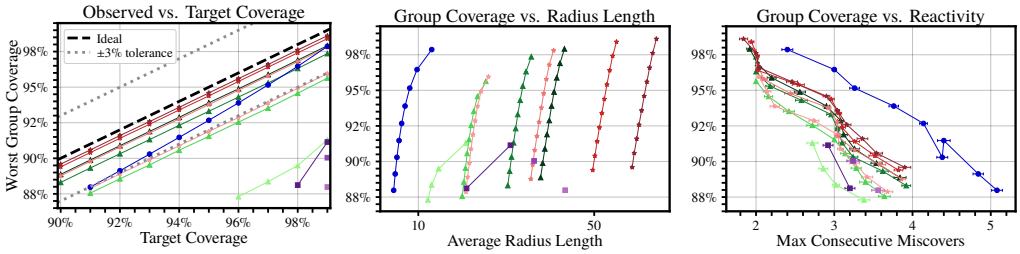


Figure 10: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

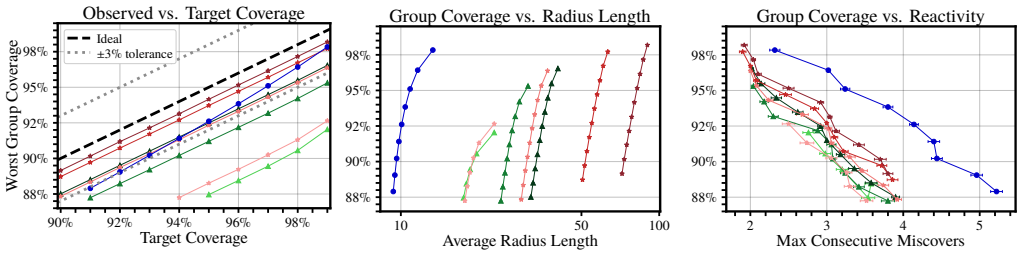
Results for $k = 100$ number of groups, cubic growing scores.



(a) Results for the synthetic setting with slow growing cubic scores ($q = 3, A = 5$).



(b) Results for the synthetic setting with moderately growing cubic scores ($q = 3, A = 25$).



(c) Results for the synthetic setting with extreme growing cubic scores ($q = 3, A = 100$).

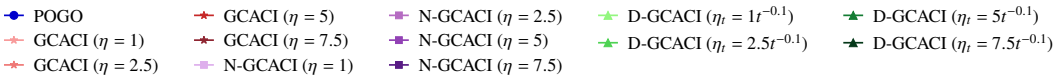


Figure 11: Performance across target coverage levels $1 - \alpha \in \{0.90, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle/right:** worst-group coverage vs. radius length / max consecutive miscovers across groups (curves closer to the top-left indicate better performance; data point shown if observed coverage ≥ 0.87). Error bars show x-axis standard error of the mean (SEM).

D.3 Additional experimental details: Stock Market

Group definitions. We define groups in the following manner:

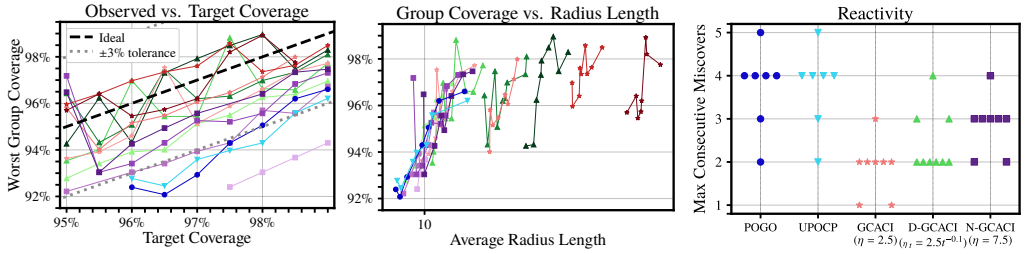
- **Market-regime groups (computed from past returns; no leakage)**
  First, we compute the following quantities.
  - Daily returns $r_t = \frac{Y_t - Y_{t-1}}{Y_{t-1}}$ (we use lagged returns Y_{t-1} to not use data from time t)
  - Over a rolling window of length W (with full-window requirement), we compute
    - Volatility: $\sigma_t = \mathrm{Std}\big(r_{t-1}, r_{t-2}, \ldots, r_{t-W}\big)$
    - Trend: $\mu_t = \mathrm{Mean}\big(r_{t-1}, r_{t-2}, \ldots, r_{t-W}\big)$
  - The rolling median volatility over a window of length W_m: $\tilde{\sigma}_t = \mathrm{Median}\big(\sigma_{t-1}, \ldots, \sigma_{t-W_m}\big)$.

  Then, we define binary groups for each time t
  - $\texttt{is_high_vol}_t = \mathbb{I}\{\sigma_t > \tilde{\sigma}_t\}$, $\texttt{is_low_vol}_t = \mathbb{I}\{\sigma_t \leq \tilde{\sigma}_t\}$

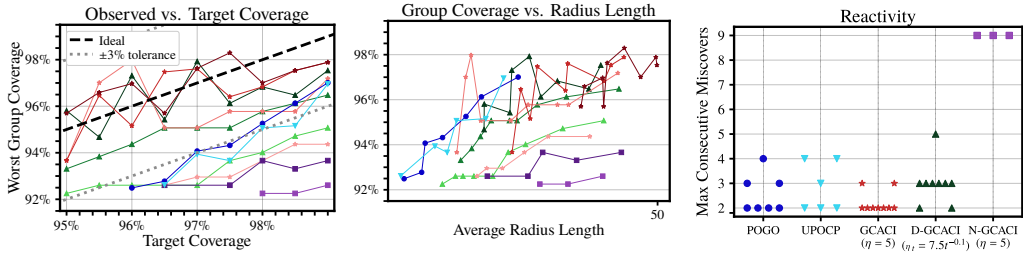
- $\texttt{is_uptrend}_t = \mathbb{I}\{\mu_t > 0\}, \quad \texttt{is_downtrend}_t = \mathbb{I}\{\mu_t \le 0\}$

- **Calendar groups (derived from the forecast date).**
  - Day-of-week one-hot indicators: `is_monday`, ..., `is_friday`.
  - Quarter one-hot indicators: `is_q1`, ..., `is_q4`.
  - Month one-hot indicators: `is_january`, ..., `is_december`.

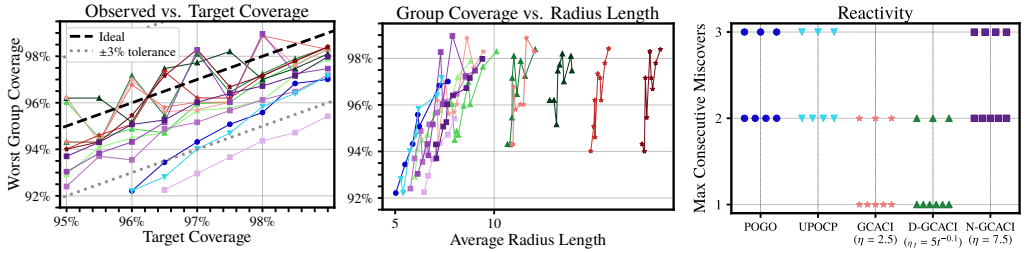
D.4 Additional results for Stock Market Experiment



(a) Results for Apple (`AAPL`) stock.



(b) Results for Boeing Company (`BA`) stock.



(c) Results for Delta Air Lines Inc (`DAL`) stock.

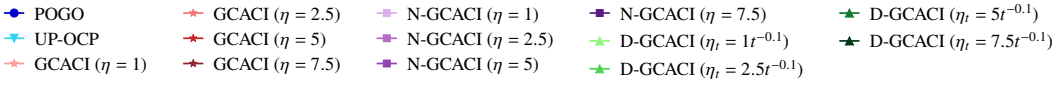
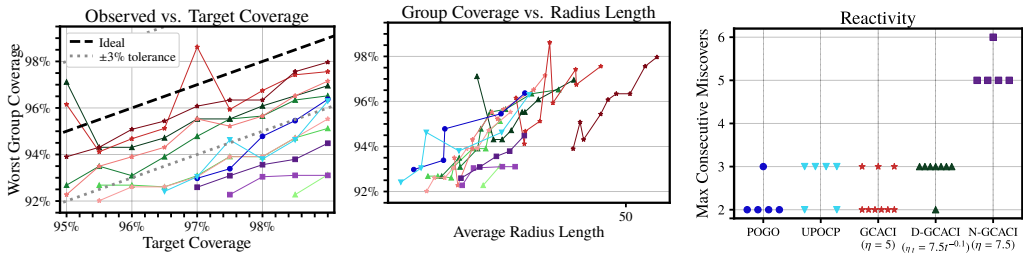
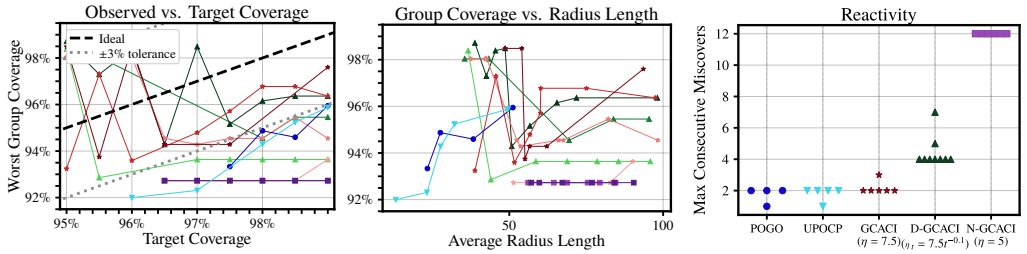


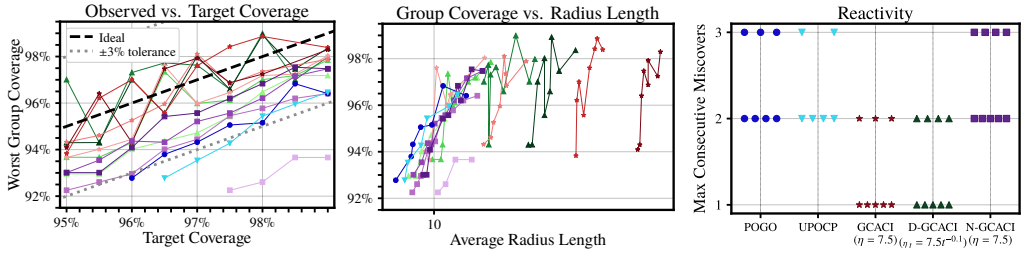
Figure 12: Performance across target coverage levels $1-\alpha \in \{0.95, 0.955, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle:** worst-group coverage vs. radius length (curves closer to the top-left indicate better performance). **Right:** Max consecutive miscovers across target coverage levels. For clarity, among the GCACI, N-GCACI, and D-GCACI families, we show only the variant with the lowest median max consecutive miscovers. Data point shown if coverage ≥ 0.92.



(a) Results for Meta (META) stock.



(b) Results for Moderna Inc (MRNA) stock.



(c) Results for United Airlines Holdings Inc (UAL) stock.

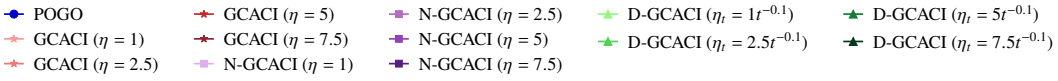


Figure 13: Performance across target coverage levels $1-\alpha \in \{0.95, 0.955, \ldots, 0.99\}$. **Left:** worst-group coverage vs. target coverage. **Middle:** worst-group coverage vs. radius length (curves closer to the top-left indicate better performance). **Right:** Max consecutive miscovers across target coverage levels. For clarity, among the GCACI, N-GCACI, and D-GCACI families, we show only the variant with the lowest median max consecutive miscovers. Data point shown if coverage ≥ 0.92.

D.5 Additional experimental details: Length-of-Stay in the ICU.

Data pre-processing. We selected ICU stays that were at least one day long, and whose patient's date of death was unknown or after discharge ($\approx 66 \times 10^3$ total stays). As features, we used the precomputed MIMIC-IV first day concepts including labs, vital signs, the SOFA score Vincent et al. (1996), and the GCS score Teasdale & Jennett (1974) (169 features total). We imputed missing values with the mean and standardized each feature using statistics from the training set.

Training details. We trained a 4-layer MLP with hidden size of 128 and dropout rate of 0.3 on 50% of the data for 500 epochs. We used the Huber loss (Huber, 1992) on the log-transformed ICU LOS (in days) and Adam optimizer (Kingma & Ba, 2014) (learning rate of 10^{-4}, weight decay of 10^{-3}). We used 10% of the data for model selection, and left the remaining 40% ($\approx 26 \times 10^3$ stays) for online calibration. The trained network achieves an MAE (in days) of 2.23 on the test set, which is compatible with the difficulty of the task using information collected during the first day only. We remark that the MIMIC-IV dataset does not include true admission dates to protect patient privacy. Instead, patient dates are shifted to fictional future dates. Consequently, we cannot guarantee true chronological ordering of patients, so we assign patients to the train, validation, and test splits at random.

Group definitions. We define groups in the following manner:

- **Race (8 groups).** Given the heterogeneity of self-reported race data in MIMIC-IV, we grouped entries as described in Table 1.
- **Insurance (5 groups).** We kept 4 original MIMIC-IV entries as groups: `Medicare`, `Medicaid`, `Private`, `No charge`, and mapped `Other` to an "unknown" group.
- **Biological sex (2 groups).** We kept the original binary `F`, `M` entries in MIMIC-IV.

Table 1: Mapping between race groups and raw self-reported race in the MIMIC-IV dataset.

Race group	MIMIC-IV Entries
Unknown	`UNKNOWN, OTHER, UNABLE TO OBTAIN, PATIENT DECLINED TO ANSWER`
Hispanic	`HISPANIC OR LATINO, PUERTO RICAN, DOMINICAN, GUATEMALAN, SALVADORAN, COLUMBIAN, CUBAN, MEXICAN, HONDURAN, CENTRAL AMERICAN, SOUTH AMERICAN`
White	`WHITE, OTHER EUROPEAN, RUSSIAN, EASTERN EUROPEAN, BRAZILIAN, PORTUGUESE`
Asian	`ASIAN, CHINESE, SOUTH EAST ASIAN, ASIAN INDIAN, KOREAN`
Black	`AFRICAN AMERICAN, CAPE VERDEAN, CARIBBEAN ISLAND, AFRICAN`
Mult	`MULTIPLE RACE/ETHNICITY`
AIAN	`AMERICAN INDIAN/ALASKA NATIVE`
NHPI	`NATIVE HAWAIIAN OR OTHER PACIFIC ISLANDER`

Table 2: Group proportions in the train ($\approx 39,000$) and test ($\approx 26,000$) splits of MIMIC-IV.

	Proportion	
Race	Train	Test
White	67.55%	67.50%
Unknown	14.34%	14.46%
Black	10.73%	10.68%
Hispanic	4.01%	3.87%
Asian	2.94%	3.07%
AIAN	0.21%	0.22%
NHPI	0.13%	0.11%
Mult	0.09%	0.09%

	Proportion	
Insurance	Train	Test
Medicare	54.30%	54.26%
Private	27.08%	26.79%
Medicaid	14.95%	15.09%
Unknown	3.66%	3.85%
No charge	0.01%	0.02%

	Proportion	
Biological sex	Train	Test
Male	56.79%	57.20%
Female	43.21%	42.80%

E POGO Implementation Details for Soft Group Memberships

At iteration t, POGO defines k functions

$$w_j(\bar{\lambda}) = \prod_{i=1}^{t-1} \left[\bar{\lambda} \left(1 - \frac{g_{i,j}}{\alpha}\right) + (1-\bar{\lambda}) \left(1 + \frac{g_{i,j}}{1-\alpha}\right) \right], \qquad j \in \{1, \ldots, k\}, \tag{73}$$

and updates the k portfolios $\lambda_{t,1}, \ldots, \lambda_{t,k}$ via

$$\lambda_{t,j} \leftarrow \frac{\int_0^1 \bar{\lambda} w_j(\bar{\lambda})\, d\mu(\bar{\lambda})}{\int_0^1 w_j(\bar{\lambda})\, d\mu(\bar{\lambda})}, \qquad j \in \{1, \ldots, k\}. \tag{74}$$

When the grouping functions c_j take values in $[0, 1]$, these integrals generally do not admit closed-form solutions. They can be approximated numerically using NumPy's `trapezoid` function on a uniform grid of 20,001 points over $(10^{-12}, 1 - 10^{-12})$, with spacing $h = \frac{1-2\times10^{-12}}{20,000}$ The grid density should be initialized using Jeffreys prior $p_0(\bar{\lambda}) \propto [\bar{\lambda}(1-\bar{\lambda})]^{-1/2}$. At each iteration, the trapezoidal rule can be used to normalize the updated grid density and compute its mean, yielding the next portfolio weight.

F ALGORITHMS

Algorithm 2 UP-OCP

Input: Target miscoverage level α.
Initialize: Wealth $W_0 \leftarrow 1$ and Jeffrey prior $\mu(\bar{\lambda}) = 1/(\bar{\lambda}(1-\bar{\lambda}))^{1/2}$

for $t = 1, \ldots$ **do**
  Observe feature vector X_t and compute prediction $\hat{Y}_t$
  Define $w\colon \mathbb{R} \to \mathbb{R}$ as $w(\bar{\lambda}) = \prod_{i=1}^{t-1} \bar{\lambda}\left(1 - \frac{g_i}{\alpha}\right) + (1-\bar{\lambda})\left(1 + \frac{g_i}{1-\alpha}\right)$
  Update portfolio: $\lambda_t \leftarrow \frac{\int_0^1 \bar{\lambda} w(\bar{\lambda}) d\mu(\bar{\lambda})}{\int_0^1 w(\bar{\lambda}) d\mu(\bar{\lambda})}$
  Update radius: $\tau_t \leftarrow W_{t-1}\left(\frac{\lambda_t - \alpha}{\alpha(1-\alpha)}\right)$
  Observe label Y_t
  Compute non-conformity score, $S_t \leftarrow |Y_t - \hat{Y}_t|$
  $g_t \leftarrow [\mathbf{1}\{S_t \leq \tau_t\} - (1-\alpha)]$
  $W_t \leftarrow W_{t-1}(1 - \frac{\lambda_t - \alpha}{\alpha(1-\alpha)} g_t)$
end for