

Adaptive conformal classification with noisy labels

Matteo Sesia*, Y. X. Rachel Wang†, Xin Tong*

February 23, 2024

Abstract

This paper develops novel conformal prediction methods for classification tasks that can automatically adapt to random label contamination in the calibration sample, leading to more informative prediction sets with stronger coverage guarantees compared to state-of-the-art approaches. This is made possible by a precise characterization of the effective coverage inflation (or deflation) suffered by standard conformal inferences in the presence of label contamination, which is then made actionable through new calibration algorithms. Our solution is flexible and can leverage different modeling assumptions about the label contamination process, while requiring no knowledge of the underlying data distribution or of the inner workings of the machine-learning classifier. The advantages of the proposed methods are demonstrated through extensive simulations and an application to object classification with the CIFAR-10H image data set.

1 Introduction

1.1 Background and motivation

Conformal inference [1] is a versatile and increasingly popular framework for estimating the uncertainty of predictions output by any supervised learning model, including for example modern deep neural networks for multi-class classification. Its two key strengths are that: (1) it requires no parametric assumptions about the data distribution, making it relevant to a variety of real-world applications; and (2) it can accommodate arbitrarily complex *black-box* predictive models, mitigating the risk of early obsolescence in the rapidly evolving field of machine learning. In a nutshell, conformal inference is able to transform the output of any model into a relatively informative and well-calibrated *prediction set* for the unknown label of a future test point, while providing precise coverage guarantees in finite samples. Intuitively, this is achieved by carefully leveraging the empirical distribution of suitable

*Department of Data Sciences and Operations, University of Southern California, Los Angeles, California.

†School of Mathematics and Statistics, University of Sydney, Sydney, Australia.

residuals (or *conformity scores*) evaluated on held-out data that were not utilized for training. Notably, these guarantees can be established assuming only that the calibration data are *exchangeable* (or, for simplicity, independent and identically distributed) random samples from the population of interest. While this framework can provide useful uncertainty estimates while relying on weaker assumptions compared to classical parametric modeling, existing conformal prediction methods are not always fully satisfactory. One limitation that we aim to address in this paper is the reliance on the assumption that the calibration data are all labeled correctly, which is often unrealistic.

In fact, it may be expensive or even outright unfeasible to acquire accurately labeled (or *clean*) data in many applications, even if there is an abundance of lower-quality observations with imperfect labels [2–5], which we also call *noisy* or *contaminated*. For example, the Amazon Mechanical Turk is a crowdsourcing platform that allows researchers and organizations to assign labels to large-scale unsupervised data by leveraging a global workforce of remote human annotators [6]. This platform is utilized across diverse fields, including image recognition and natural language processing. Crowdsourcing tends to be fast and cost-effective, but it raises concerns about poor annotation quality and its impacts on the reliability of downstream analyses [7–10]. Other applications motivating our work are those in which data with labels carrying sensitive personal information are randomly anonymized to safeguard user privacy [11–15].

Although significant efforts have been dedicated to *training* predictive models using data with noisy labels [2–4, 16, 17], the challenge of *calibrating* those models using conformal inference has only recently begun to receive some attention [18–20] and is still understudied. This paper aims to help fill this gap.

1.2 Preview of our contributions and paper outline

This paper addresses the questions of whether and how conformal inference should account for the possible presence of incorrectly labeled calibration samples, focusing on the task of constructing prediction sets for multi-class classification. After recalling the relevant background, we begin in Section 2 by carefully studying the impact of label contamination on the effective coverage achieved by standard conformal prediction sets. Our analysis shows and quantifies precisely how label contamination may lead to prediction sets that are either too liberal or too conservative. The practical significance of these results becomes clear in Section 3, where we develop and study a novel calibration method that can automatically adapt to label contamination, without assuming any knowledge of the data distribution or of the classifier. As we shall see, our method can produce more informative prediction sets with more robust coverage guarantees compared to standard approaches. Initially, we assume the label contamination process is known and belongs to a broad family encompassing most of the typical models from the related literature on learning from noisy labels [2, 21, 22]. Then, we extend our solution to accommodate models that may depend on unknown parameters, and we explain how to estimate those parameters.

In Section 4, we demonstrate the practical performance of our methods through exten-

sive numerical experiments and an analysis of the CIFAR-10H image classification data set [23]. A preview of some results is given here by Figure 1, which demonstrates that our methods are practical and produce smaller (more informative) prediction sets compared to standard conformal inference techniques, while maintaining valid coverage. Finally, Section 5 concludes with a discussion and some ideas for future work.

Some content is deferred to the Appendix, for lack of space. Section A1 reviews the standard conformal prediction approach. Section A2 provides additional algorithmic tables outlining the proposed methods. Section A3 describes some simplified implementations of our methods under more restricted contamination models. Section A4 discusses the relation between our work and existing theoretical worst-case results about the behavior of standard conformal predictions applied to non-exchangeable data [20]. Section A5 extends our theoretical results from Section 2 to study marginal instead of label-conditional coverage. Section A6 extends our methods from Section 3 to construct adaptive prediction sets satisfying other types of theoretical guarantees, such as marginal coverage and calibration-conditional coverage [24]. Further extensions encompassing even stronger guarantees, such as equalized coverage over protected categories [25], would also be possible but are omitted for length-related reasons. Section A7 contains all mathematical proofs. Section A8 describes additional numerical results, and Section A9 provides further technical details about our experiments.

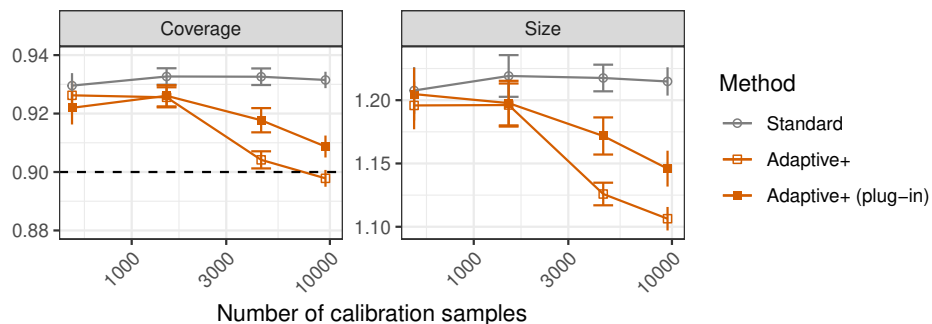


Figure 1: Performance of the proposed adaptive conformal prediction methods on CIFAR-10H image data with imperfect labels. The results are shown as a function of the calibration sample size and compared to the performance of standard conformal predictions, which are too conservative due to the presence of inaccurate labels. The dashed line indicates the nominal 90% marginal coverage level.

1.3 Related work

Conformal inference was pioneered by Vovk and collaborators [1] and has become a very active area of research [26–28], with applications including outlier detection [29–33], regression [34–36], and classification [37–41].

Numerous works have studied the robustness of conformal inferences to the breakdown

of standard data exchangeability assumptions. However, these investigations largely concentrated on different challenges, such as covariate shift [42], repeated sampling [43], time-series dependencies [44], and label shift [45, 46]. We refer to Section 2.4 for further details on the crucial distinction between label *shift* and the label *contamination* problem considered in this paper.

Barber et al. [20] conducted a general theoretical study of conformal inference with non-exchangeable data. Although their goals and approaches differ from ours, some of their insights also apply to scenarios involving label contamination. Our findings are consistent with theirs, albeit somewhat weaker, when viewed through the lens of a *worst-case* analysis of traditional methods. However, our goal is not to refine their *theoretical* analysis. On the contrary, we make a completely separate methodological contribution—we develop novel methods for constructing more informative prediction sets in the presence of *random* label contamination.

Our paper is more closely related to Cauchois et al. [18] and Einbinder et al. [19], which also considered conformal prediction sets calibrated using imperfectly labeled data. However, our research is distinct and involves many novelties. The problem of Cauchois et al. [18] is different because they allow the calibration data to contain imperfect labels but seek to predict analogous imperfect labels for the test data. By contrast, we aim to predict the true (clean) labels. Our perspective is more aligned with that of Einbinder et al. [19], but their efforts are focused on establishing the conservativeness of standard conformal inferences to certain forms of random label contamination, and they do not attempt to mitigate that (often excessive) conservativeness or to protect against other forms of contamination. Our theoretical analyses take a different approach, leading to more general and quantitative results that hold for a broader class of contamination models. Then, our main contribution is to develop novel methods that can automatically adapt to label contamination. To the best of our knowledge, this is the first work to propose conformal prediction methods that are adaptive to label contamination.

2 Preliminaries

2.1 Problem statement

Consider $n + 1$ data points (X_i, Y_i, \tilde{Y}_i) , for $i \in [n + 1] = \{1, \dots, n + 1\}$, where $X_i \in \mathbb{R}^d$ is a feature vector, $Y_i \in [K]$ is a latent categorical label, and $\tilde{Y}_i \in [K]$ is an observable label that we interpret as a contaminated version of Y_i . Assume the data are i.i.d. random samples from some unknown distribution. A weaker *exchangeability* assumption often turns out to be sufficient in the related literature, but we focus on i.i.d. data in this paper. The problem we consider is that of constructing informative conformal prediction sets for the true label Y_{n+1} of a test point with features X_{n+1} , leveraging the available observations (X_i, \tilde{Y}_i) indexed by $[n]$, to which we collectively refer as \mathcal{D} . Of course, in order to establish precise coverage guarantees, it will be necessary to introduce some assumptions about the

relation between the true and contaminated labels, as we shall see. Our setup thus extends the standard conformal prediction framework for uncontaminated data, which corresponds for the special case in which $\tilde{Y} = Y$ almost surely.

2.2 Relevant technical background

This section recalls some relevant background on standard conformal classification methods, which ignore the possibility of label contamination. Our notation is inspired by that of Romano, Sesia, and Candès [40] but involves some helpful adaptations. Note that we will refer to standard conformal prediction approaches as working with observations (X, \tilde{Y}) to keep our notation consistent, although the prior literature typically did not distinguish between Y and \tilde{Y} .

Let \mathcal{A} denote any learning algorithm for classification; e.g., a logistic regression model or a neural network. The role of \mathcal{A} , which is treated as a black-box throughout the paper, is to train a model $\hat{\pi}$ that estimates the distribution of $\tilde{Y} \mid X$ using the data in \mathcal{D} . For any $k \in [K]$ and $x \in \mathbb{R}^d$, let $\hat{\pi}(x, k) \in [0, 1]$ be the estimated probability of $\tilde{Y} = k \mid X = x$. In the following, it will sometimes be convenient to assume that the distribution of $\hat{\pi}(X, k)$ is continuous. This condition could be relaxed at the cost of a more involved notation, but it can be made realistic by simply adding a small amount of independent noise to the output of the model. Further, we assume that $\hat{\pi}$ is normalized; i.e., $\sum_{k=1}^K \hat{\pi}(x, k) = 1$. Aside from these requirements, we allow $\hat{\pi}$ to be anything and do not necessarily expect it to model $\tilde{Y} \mid X$ accurately. Thus, typical off-the-shelf classifiers can provide suitable statistics $\hat{\pi}$ by default. For example, one may choose $\hat{\pi}$ to be the output of the final soft-max layer of a deep neural network.

This paper studies how to translate such a black-box model into a reliable prediction set for Y_{n+1} given X_{n+1} . But first we need to review the standard conformal approach to the special case where $Y = \tilde{Y}$. The first key notion that we recall is that of a *prediction function*.

Definition 1 (Prediction function). *Let \mathcal{C} be a set-valued function, whose form may depend on the model $\hat{\pi}$, that takes as input $x \in \mathbb{R}^d$ and $\tau = (\tau_1, \dots, \tau_K) \in [0, 1]^K$, and returns as output a subset of $[K]$. We say that \mathcal{C} is a prediction function if it is monotone increasing with respect to each element of τ , is such that the event $k \in \mathcal{C}(x, \tau)$ depends on τ only through its k -th element τ_k , and satisfies $k \in \mathcal{C}(X, \tau)$ whenever $\tau_k = 1$, for any $k \in [K]$.*

Note that the dependence of a prediction function \mathcal{C} on $\hat{\pi}$ will typically be kept implicit unless otherwise necessary to avoid ambiguity; i.e., $\mathcal{C}(x, \tau) = \mathcal{C}(x, \tau; \hat{\pi}) \subseteq [K]$.

For any prediction function \mathcal{C} , we define the *conformity score* function $\hat{s} : \mathbb{R}^d \times [K] \mapsto [0, 1]$, also implicitly depending on $\hat{\pi}$, as that function which outputs the smallest value of τ_k allowing the label k to be contained in the set $\mathcal{C}(x, \tau_k)$. That is,

$$\hat{s}(x, k) = \inf \{ \tau_k \in [0, 1] : k \in \mathcal{C}(x, \tau_k) \}. \quad (1)$$

Note that the short-hand $\mathcal{C}(X, \tau_k)$ above is a slight abuse of notation, since \mathcal{C} was defined for a vector-valued input τ , but it does not introduce any ambiguity.

An example of \mathcal{C} is the function that outputs the set of all labels $k \in [K]$ for which the estimated conditional probability of $\tilde{Y} = k \mid X = x$ is sufficiently large; i.e.,

$$\mathcal{C}(x, \tau; \hat{\pi}) := \{k \in [K] : \hat{\pi}(x, k) \geq 1 - \tau_k\}. \quad (2)$$

The associated scores are $\hat{s}(x, k) = 1 - \hat{\pi}(x, k)$. These are sometimes called *homogeneous conformity scores* because the prediction function in (2) is not designed to account for heteroscedasticity in the distribution of $Y \mid X$ [39]. While this choice of \mathcal{C} is convenient to keep the notation simple, all of our results also extend to other prediction functions, including those associated with the more sophisticated *generalized inverse quantile* conformity scores of Romano, Sesia, and Candès [40], which we review in Section A1.

A standard implementation of conformal inference begins by randomly splitting the data into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} , such that $\mathcal{D} = \mathcal{D}^{\text{train}} \cup \mathcal{D}^{\text{cal}}$. The model $\hat{\pi}$ is trained using the data in $\mathcal{D}^{\text{train}}$. The held-out data in \mathcal{D}^{cal} are utilized to compute conformity scores $\hat{s}(X_i, k)$ via (1), for all $k \in [K]$ and $i \in \mathcal{D}^{\text{cal}}$, according to the desired prediction function \mathcal{C} . These scores are then utilized to calibrate a prediction set for a new test point with features X_{n+1} as follows. For each $k \in [K]$, define $\mathcal{D}_k^{\text{cal}} = \{i \in \mathcal{D}^{\text{cal}} : \tilde{Y}_i = k\}$, $n_k = |\mathcal{D}_k^{\text{cal}}|$, and $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_K)$, where

$$\hat{\tau}_k = \lceil (1 + n_k) \cdot (1 - \alpha) \rceil\text{-th smallest value in } \{\hat{s}(X_i, k)\}_{i \in \mathcal{D}_k^{\text{cal}}}, \quad (3)$$

and $\alpha \in (0, 1)$ is the desired significance level. The prediction set for Y_{n+1} is given by $\hat{C}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}; \hat{\pi})$. See Algorithm A1 in Section A1.2 for a summary of this method. This procedure makes it possible to prove that $\hat{C}(X_{n+1})$ has *label-conditional coverage* [47],

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] \geq 1 - \alpha, \quad \text{for all } k \in [K], \quad (4)$$

as long as $Y = \tilde{Y}$ almost surely. See Proposition A1 for a formal statement of this result, which also provides an almost-matching coverage upper bound. Note that the probability in (4) is taken with respect to X_{n+1} and the data in \mathcal{D} , both of which are random. In the following sections, we will study the behavior of this method while allowing $\tilde{Y} \neq Y$.

In the meantime, we recall that an alternative standard approach is to construct prediction sets $\hat{C}(X_{n+1})$ satisfying the following weaker notion of *marginal coverage*:

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \right] \geq 1 - \alpha. \quad (5)$$

As long as $Y = \tilde{Y}$ almost surely, prediction sets with marginal coverage can be obtained by simply replacing the subset $\mathcal{D}_k^{\text{cal}}$ with \mathcal{D}^{cal} in Algorithm A1. See Algorithm A2 and Proposition A2 in Section A1.3 for further details. While marginal coverage (5) is not as strong as label-conditional coverage (4), it is a useful notion because it is easier to achieve using smaller, and hence more informative, prediction sets when limited data are available. Therefore, we will study how to efficiently control both (4) and (5) using contaminated data, leaving it to practitioners to determine which type of guarantee is most appropriate for a given application. However, due to lack of space, most details about marginal coverage are presented in the Appendix.

2.3 General coverage bounds under label contamination

This section analyzes theoretically the behavior of standard conformal classification approaches applied with contaminated calibration data. In particular, we demonstrate how label noise can cause the effective coverage of these methods to be either inflated or deflated, as made precise by an explicit factor that depends on the distribution of the conformity scores. For simplicity, we focus here on label-conditional coverage (4), studying the behavior of Algorithm A1. These results will be extended to study marginal coverage (5) in Section A5.

It is worth emphasizing that the results presented in this section require no assumptions on the contamination process and encompass a wide range of scenarios in which label noise leads to over-coverage or under-coverage; thus, our analysis is more general than that of Einbinder et al. [19], which focused on establishing conservativeness under a narrower random corruption model. That being said, it will become clear in Section 3 that some assumptions about the contamination model are useful to achieve the more ambitious goal of developing practical conformal prediction methods that can automatically adapt to label noise.

The following notation will be helpful. For any $k, l \in [K]$ and $t \in \mathbb{R}$, define

$$\begin{aligned} F_l^k(t) &:= \mathbb{P}[\hat{s}(X, k) \leq t \mid Y = l, \mathcal{D}^{\text{train}}], \\ \tilde{F}_l^k(t) &:= \mathbb{P}[\hat{s}(X, k) \leq t \mid \tilde{Y} = l, \mathcal{D}^{\text{train}}]. \end{aligned} \quad (6)$$

In words, F_l^k is the cumulative distribution function of $\hat{s}(X, k)$, based on a fixed function \hat{s} and a random sample X from the distribution of $X \mid Y = l$, namely P_l . Analogously, \tilde{F}_l^k is the cumulative distribution function corresponding to a random sample X from the distribution of $X \mid \tilde{Y} = l$, namely \tilde{P}_l . For any $k \in [K]$ and $t \in \mathbb{R}$, define also $\Delta_k(t)$ as

$$\Delta_k(t) := F_k^k(t) - \tilde{F}_k^k(t). \quad (7)$$

We refer to $\Delta_k(t)$ as the *coverage inflation factor* because its expected value controls the discrepancy between the real and nominal coverage of the prediction sets output by Algorithm A1, as established by the next result. Note that $\Delta_k(t)$ may be either positive or negative.

Theorem 1. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n+1]$. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}(X_{n+1})$ indicate the prediction set output by Algorithm A1 applied using the corrupted labels \tilde{Y}_i instead of the clean labels Y_i , for all $i \in [n]$. Then,*

$$\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k] \geq 1 - \alpha + \mathbb{E}[\Delta_k(\hat{\tau}_k)]. \quad (8)$$

Further, if the conformity scores $\hat{s}(X_i, \tilde{Y}_i)$ used by Algorithm A1 are almost-surely distinct,

$$\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k] \leq 1 - \alpha + \frac{1}{n_k + 1} + \mathbb{E}[\Delta_k(\hat{\tau}_k)]. \quad (9)$$

Above, $\mathbb{E}[\Delta_k(\hat{\tau}_k)]$ denotes an expected value taken with respect to the randomness in all data in \mathcal{D} , including the calibration samples \mathcal{D}^{cal} upon which $\hat{\tau}_k$ depends. Note that it would also be possible to obtain slightly stronger versions of (8) and (9) in which the data \mathcal{D} are treated as fixed and $\mathbb{E}[\Delta_k(\hat{\tau}_k)]$ is replaced by $\Delta_k(\hat{\tau}_k)$; see the proof of Theorem 1. In any case, this result tells us that standard conformal inferences may be either overly conservative or too liberal when dealing with contaminated data, depending on whether $\mathbb{E}[\Delta_k(\hat{\tau}_k)]$ is positive or negative.

This result, which we extend in Section A5 to study marginal coverage, serves as the starting point of our methodology. As a preliminary step towards the goal of achieving *tight* coverage at level $1 - \alpha$, the next subsection introduces additional modeling assumptions about the label contamination process and sheds more light onto the situations in which standard conformal inferences can be guaranteed to be conservative.

2.4 Coverage lower bounds under a general linear contamination model

To obtain more interpretable and actionable expressions for the coverage bounds presented above, it is necessary to introduce some assumptions about the relation between the latent labels Y and the observable labels \tilde{Y} . Fortunately, significant progress can be made by simply assuming that \tilde{Y} is conditionally independent of X given Y . This corresponds to a widely used class of label contamination models [2, 22].

Assumption 1. $\tilde{Y} \perp\!\!\!\perp X \mid Y$.

Assumption 1 gives rise to a convenient relation between the distribution of $X \mid \tilde{Y} = k$, namely \tilde{P}_k , and the distributions of $X \mid Y = l$, namely P_l , for all $k, l \in [K]$. In plain words, for any $k \in [K]$, the distribution \tilde{P}_k is a linear mixture of the distributions P_l for all $l \in [K]$.

Proposition 1. *Let M be a $K \times K$ matrix such that $M_{kl} = \mathbb{P}[Y = l \mid \tilde{Y} = k]$ for any $k, l \in [K]$. Then, Assumption 1 implies that*

$$\tilde{P}_k = \sum_{l=1}^K M_{kl} P_l. \quad (10)$$

The mixture model in (10) extends the classical binary class-dependent noise model in Scott, Blanchard, and Handy [21] to the multi-class setting. Further, combined with Bayes' theorem, Assumption 1 leads to: $\mathbb{P}[\tilde{Y} = k \mid X, Y = l] = M_{kl} \tilde{\rho}_k / \rho_l$, where $\rho_k := \mathbb{P}[Y = k]$ and $\tilde{\rho}_k = \mathbb{P}[\tilde{Y} = k]$ denote the marginal distributions of the clean and corrupted labels, respectively. If $M_{kl} = \eta / (K - 1)$ for all $k \neq l$, where $0 \leq \eta \leq 1$, and $\rho_k = \tilde{\rho}_k = 1/K$ for all $k \in [K]$, this model reduces to the simple homogeneous noise setting [22].

We remark that it should now be clear how our problem is distinct from that of label shift [45, 46]. The latter refers to situations in which the marginal distribution of \tilde{Y} may differ from that of Y but $\tilde{P}_k = P_k$ for all $k \in [K]$, while label contamination generally leads

to $\tilde{P}_k \neq P_k$. In this sense, label contamination is more challenging because the distribution of $X \mid \tilde{Y}$ used in training and calibration differs from that of $X \mid Y$ used in testing. At the same time, our problem is also distinct from *covariate shift* [42] because in our case the *marginal* distribution of X does not vary.

The first useful implication of Proposition 1 is that standard conformal inferences under label contamination are often conservative.

Corollary 1. *Consider the same setting of Theorem 1 and assume Assumption 1 holds. Suppose also that the cumulative distribution functions of the scores (6) satisfy*

$$\max_{l \neq k} F_l^k(t) \leq F_k^k(t), \quad (11)$$

for all $t \in \mathbb{R}$ and $k \in [K]$. Then, $\Delta_k(\hat{\tau}_k) \geq 0$ almost-surely for any $k \in [K]$, and hence the prediction sets $\hat{C}(X_{n+1})$ output by Algorithm A1 satisfy (4).

The “stochastic dominance” condition in (11) states that the trained model tends to assign smaller scores $\hat{s}(X, k)$ when $Y = k$. If $\hat{s}(X, k) = 1 - \hat{\pi}(X, k)$, this is intuitively equivalent to requiring the model to estimate the distribution of $Y \mid X$ sufficiently accurately as to at least preserve the relative ranking of the most likely labels, on average.

In summary, Corollary 1, together with its extension to marginal coverage presented in Section A5, provides a lower bound that highlights a certain robustness of standard conformal inferences to label contamination, consistently with Einbinder et al. [19]. However, these results are not yet fully satisfactory for at least two reasons. Firstly, it is unclear how to check whether the condition in (11) holds in practice. Secondly, even if (11) is satisfied, one may be concerned that standard conformal inferences can be too conservative under label contamination, leading to unnecessarily large prediction sets. This is why we develop in the next section novel methods that can automatically adapt to label contamination, producing more informative prediction sets that rigorously guarantee coverage at the desired level.

We conclude this section by noting that Theorem 1 could also be applied to derive theoretical *worst-case* coverage bounds for standard conformal prediction sets calibrated with contaminated data, even without any (empirical) information about the conformity score distribution. Section A4 elaborates on this subject and discusses its connection to the elegant theoretical insights of Barber et al. [20]. However, it is worth emphasizing that such worst-case bounds have limited practical relevance in the context of this paper. This is because worst-case bounds, while theoretically interesting, do not offer actionable guidance on how to enhance the informativeness of standard conformal prediction sets in the face of random label contamination.

3 General methodology

3.1 Adaptive coverage under a known label contamination model

We present a method for constructing prediction sets that automatically adapt to label contamination. For simplicity, we begin by focusing on label-conditional coverage assuming that the contamination model is known. Subsequently, we will extend similar ideas to accommodate unknown contamination models, and to provide other types of coverage guarantees. The assumption of a known contamination model is convenient and provides a useful stepping stone for our next developments. Further, this assumption is well justified in several interesting applications, such as those involving controlled label randomization designed to ensure label-differential privacy [12–14].

3.1.1 A plug-in estimate for the coverage inflation factor

Our method leverages Assumption 1 through Proposition 1, which makes it possible to write the inflation factor $\Delta_k(t)$ in (7) in terms of quantities that are either known or estimable. In fact, if M in (10) admits a matrix inverse $V = M^{-1}$, the factor $\Delta_k(t)$ can be expressed as:

$$\Delta_k(t) = (V_{kk} - 1)\tilde{F}_k^k(t) + \sum_{l \neq k} V_{kl}\tilde{F}_l^k(t). \quad (12)$$

This expression only depends on V , which is assumed to be known, and on the distributions of the scores computed from \tilde{Y} , which are observable. This suggests it may be possible to estimate $\Delta_k(t)$ from the data and then leverage Theorem 1 to obtain an adaptive prediction method with tighter coverage guarantees compared to the standard approach studied in Section 2. In the following, we develop such a method and establish both upper and lower bounds for its coverage, assuming that V is known. The problem of estimating V will be addressed later.

Our method begins by randomly splitting the labeled data into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} , similarly to standard conformal inference. The observations in $\mathcal{D}^{\text{train}}$ are used to train the model $\hat{\pi}$, while those in \mathcal{D}^{cal} are used to compute conformity scores $\hat{s}(X_i, k)$ via (1), for all $k \in [K]$ and $i \in \mathcal{D}^{\text{cal}}$, according to the desired prediction function \mathcal{C} . For any $k, l \in [K]$, let \hat{F}_l^k denote the empirical cumulative distribution function of $\hat{s}(X_i, k)$ for $i \in \mathcal{D}_l^{\text{cal}} = \{i \in \mathcal{D}^{\text{cal}} : \tilde{Y}_i = l\}$; i.e.,

$$\hat{F}_l^k(t) := \frac{1}{n_l} \sum_{i \in \mathcal{D}_l^{\text{cal}}} \mathbb{I}[\hat{s}(X_i, k) \leq t], \quad (13)$$

where $n_l = |\mathcal{D}_l^{\text{cal}}|$. In other words, $\hat{F}_l^k(t)$ intuitively estimates \tilde{F}_l^k . If the matrix V is known, one can leverage the \hat{F}_l^k functions to compute a plug-in estimate of $\Delta_k(t)$:

$$\hat{\Delta}_k(t) := (V_{kk} - 1)\hat{F}_k^k(t) + \sum_{l \neq k} V_{kl}\hat{F}_l^k(t). \quad (14)$$

If $\hat{\Delta}_k$ could estimate Δ_k accurately, one would guess from Theorem 1 that Algorithm A1—the standard conformal inference method that ignores label contamination—leads to an effective coverage close to $1 - \alpha + \hat{\Delta}_k(\hat{\tau}_k)$, where $\hat{\tau}_k$ is the data-driven calibration parameter computed via (3). This suggests adjusting the nominal significance level to something close to $\alpha - \hat{\Delta}_k(\hat{\tau}_k)$ to achieve $1 - \alpha$ coverage. We will now translate this intuition into a rigorous method.

3.1.2 The adaptive calibration algorithm

For any $k \in [K]$, define the set $\hat{\mathcal{I}}_k \subseteq [n_k]$ as

$$\hat{\mathcal{I}}_k := \left\{ i \in [n_k] : \frac{i}{n_k} \geq 1 - \alpha - \hat{\Delta}_k(S_{(i)}^k) + \delta(n_k, n_*) \right\}, \quad (15)$$

where $S_{(i)}^k$, for $i \in [n_k]$, are the ascending order statistics of $\{\hat{s}(X_j, k)\}_{j \in \mathcal{D}_j^{\text{cal}}}$, while $n_* := \min_{k \in [K]} n_k$, and $\delta(n_k, n_*)$ is a correction factor specified later. Then, our threshold $\hat{\tau}_k$ is:

$$\hat{\tau}_k = \begin{cases} S_{(\hat{i}_k)}^k \text{ where } \hat{i}_k = \min\{i \in \hat{\mathcal{I}}_k\}, & \text{if } \hat{\mathcal{I}}_k \neq \emptyset, \\ 1, & \text{if } \hat{\mathcal{I}}_k = \emptyset. \end{cases} \quad (16)$$

Finally, the adaptive prediction set output by our proposed method is $\hat{C}(X_{n+1}) = \mathcal{C}(X, \hat{\tau})$, where $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_K)$. This procedure is outlined by Algorithm 1.

If we ignored the noise and finite-sample correction terms (i.e., imagining that $\hat{\Delta}_k(S_{(i)}^k) = 0$ and $\delta(n_k, n_*) = 0$), the threshold $\hat{\tau}_k$ in (16) would intuitively reduce to the $\lceil (1 - \alpha)n_k \rceil$ -th smallest value among the conformity scores for the calibration points with label k , similarly to the standard method reviewed in Section 2.2. In general, though, the more complicated form of $\hat{\tau}_k$ in (16) is designed to approximately cancel the unknown coverage inflation factor $\Delta_k(\hat{\tau}_k)$ arising when the standard conformal inference method is applied to contaminated data, as described by Theorem 1. The main purpose of $\delta(n_k, n_*)$ in (15) is to account for possible random errors in the estimation of the unknown function $\Delta_k(\cdot)$ through $\hat{\Delta}_k(\cdot)$, allowing us to obtain finite-sample guarantees. The exact form of this correction term is discussed next.

For any $k \in [K]$, let U_1, \dots, U_{n_k} be i.i.d. uniform random variables on $[0, 1]$, and denote their order statistics as $U_{(1)}, \dots, U_{(n_k)}$. Then, define

$$c(n_k) := \mathbb{E} \left[\sup_{i \in [n_k]} \left\{ \frac{i}{n_k} - U_{(i)} \right\} \right], \quad (17)$$

and

$$\delta(n_k, n_*) := c(n_k) + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\}. \quad (18)$$

Algorithm 1: Adaptive classification under a known label contamination model

- Input:** Data set $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with corrupted labels $\tilde{Y}_i \in [K]$.
 The inverse V of the matrix M in (10).
 Unlabeled test point with features X_{n+1} .
 Machine learning algorithm \mathcal{A} for training a K -class classifier.
 Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).
 Desired coverage level $1 - \alpha \in (0, 1)$.
- 1 Randomly split $[n]$ into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} .
 - 2 Train the classifier \mathcal{A} on the data in $\mathcal{D}^{\text{train}}$.
 - 3 Compute conformity scores $\hat{s}(X_i, k)$ using (1) for all $i \in \mathcal{D}^{\text{cal}}$ and all $k \in [K]$.
 - 4 Define the empirical CDF \hat{F}_l^k of $\{\hat{s}(X_i, k) : i \in \mathcal{D}_l^{\text{cal}}\}$ for all $l \in [K]$, as in (13).
 - 5 **for** $k = 1, \dots, K$ **do**
 - 6 Define $\mathcal{D}_k^{\text{cal}} = \{i \in \mathcal{D}^{\text{cal}} : \tilde{Y}_i = k\}$ and $n_k = |\mathcal{D}_k^{\text{cal}}|$.
 - 7 Sort $\{\hat{s}(X_i, k) : i \in \mathcal{D}_k^{\text{cal}}\}$ into $(S_{(1)}^k, S_{(2)}^k, \dots, S_{(n_k)}^k)$, in ascending order.
 - 8 Compute $\hat{F}_l^k(S_{(i)}^k)$ for all $i \in [n_k]$ and $l \in [K]$.
 - 9 Compute $\hat{\Delta}_k(S_{(i)}^k)$ for all $i \in [n_k]$, as in (14).
 - 10 Compute $\delta(n_k, n_*)$ using (18), based on a Monte Carlo estimate of $c(n_k)$ in (17).
 - 11 Construct the set $\hat{\mathcal{I}}_k \subseteq [K]$ as in (15).
 - 12 Evaluate $\hat{\tau}_k$ based on $\hat{\mathcal{I}}_k$ as in (16).
 - 13 **end**
 - 14 Evaluate $\hat{C}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}; \hat{\pi})$, where $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_K)$.
- Output:** Conformal prediction set $\hat{C}(X_{n+1})$ for Y_{n+1} .
-

We know from classical results in empirical process theory that $c(n_k)$ in (17) scales as $1/\sqrt{n_k}$ if n_k is large, and thus the overall correction term $\delta(n_k, n_*)$ tends to vanish as $1/\sqrt{n_*}$ in the large-sample limit. Note that the constant $c(n_k)$ in (17) will be assumed henceforth to be known because it can be easily estimated up to arbitrary precision via a Monte Carlo simulation of n_k independent standard uniform random variables. Combined with the adaptive nature of our threshold $\hat{\tau}_k$ in (16), this finite-sample correction allows Algorithm 1 to enjoy a stronger coverage guarantee under label contamination compared to standard conformal prediction.

Theorem 2. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n + 1]$, and that Assumption 1 holds. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}(X_{n+1})$ indicate the prediction set output by Algorithm 1 based on the inverse V of the model matrix M in the label contamination model (10). Then, $\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k] \geq 1 - \alpha$ for all $k \in [K]$.*

Intuitively, this says that Algorithm 1 provides valid prediction sets at level $1 - \alpha$ despite the label noise. Crucially, this does not require any assumptions about the classifier’s accuracy, in contrast with the potentially more delicate behavior of standard conformal inferences (Algorithm A1); i.e., see Theorem 1 and Corollary 1. Further, under some additional regularity conditions, it can be proved that Algorithm 1 is not overly conservative, as discussed next.

3.1.3 A coverage upper bound

The assumptions needed for our coverage upper bound are stated here and explained below.

Assumption 2. For all $k, l \in [K]$, the cumulative distribution functions \tilde{F}_l^k are differentiable on the interval $(0, 1)$, and the corresponding densities \tilde{f}_l^k are uniformly bounded with $\|\tilde{f}_l^k\|_\infty \leq f_{\max}$, for some $f_{\max} > 0$. Further, $f_{\min} := \min_{k \in [K]} \inf_{t \in (0, 1)} \tilde{f}_k^k(t) > 0$.

Assumption 3. For any $k \in [K]$, the cumulative distribution function \tilde{F}_k^k satisfies

$$\max_{l \neq k} \tilde{F}_l^k(t) \leq \tilde{F}_k^k(t), \quad \forall t \in \mathbb{R}.$$

Assumption 4. The coverage inflation factor $\Delta_k(t)$ is bounded from below by:

$$\inf_{t \in (0, 1)} \Delta_k(t) \geq -\alpha + c(n_k) + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2 \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right).$$

Assumption 2 merely requires that the distribution of the conformity scores should be continuous with bounded density; this can be ensured in practice by adding a small amount of random noise to the scores computed by any classifier. Assumption 3 simply states that the classifier tends to assign smaller scores $\hat{s}(X, k)$ to data points with $\tilde{Y} = k$. This may be reminiscent of the stochastic dominance condition in Corollary 1, although it is different and arguably weaker. In fact, the classifier is trained on data with corrupted labels. Therefore, as long as it can achieve non-trivial prediction accuracy, it should assign smaller scores $\hat{s}(X, k)$ when $\tilde{Y} = k$.

Assumption 4 looks slightly more involved, but it is also quite realistic. For example, it is always satisfied in the large-sample limit, $n_* \rightarrow \infty$, if the stochastic dominance condition defined in (11) holds, because in that case $\inf_{t \in (0, 1)} \Delta_k(t) \geq 0$. Further, as discussed in more detail in Section A3, Assumption 4 can also be satisfied if the stochastic dominance condition in (11) does not hold, as long as some additional assumptions are imposed on the label contamination model. Under this setup, a finite-sample upper bound for the coverage of the conformal prediction sets output by Algorithm 1 is established below.

Theorem 3. Under the setup of Theorem 2, let $\hat{C}(X_{n+1})$ be the prediction set output by Algorithm 1 based on the inverse V of the matrix M in (10). Suppose also that Assumptions 2–4 hold. Then, $\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k] \leq 1 - \alpha + \varphi_k(n_k, n_*)$ for all $k \in [K]$, where

$$\varphi_k(n_k, n_*) = 2\delta(n_k, n_*) + \frac{1}{n_*} + \frac{1 + 2 \sum_{l \neq k} |V_{kl}| \frac{f_{\max}}{f_{\min}} \sum_{j=1}^{n_k+1} \frac{1}{j}}{n_k} + \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k + 1}.$$

Thus, the sets output by Algorithm 1 are *asymptotically tight* because $\varphi_k(n_k, n_*) \rightarrow 0$ as $n_* \rightarrow \infty$. While this is already encouraging about the efficiency of Algorithm 1, our method can be further refined to produce even more informative prediction sets that remain valid in those (rather common) scenarios in which standard conformal inferences are too conservative.

3.1.4 Boosting power with more optimistic calibration

We know from Corollary 1 that even the standard conformal inference approach of Algorithm A1 is conservative under the (relatively mild) stochastic dominance condition in (11). This motivates us to devise a hybrid method that can outperform both Algorithm 1 and Algorithm A1, while retaining guaranteed coverage under a slightly stronger version of (11). Intuitively, the idea is to adaptively choose between Algorithm 1 and Algorithm A1 depending on which approach leads to a lower (less conservative) calibrated threshold. In other words, we propose to apply Algorithm 1 with the set $\hat{\mathcal{I}}_k$ in (15) replaced by

$$\hat{\mathcal{I}}_k := \left\{ i \in [n_k] : \frac{i}{n_k} \geq 1 - \alpha - \max \left\{ \hat{\Delta}_k(S_{(i)}^k) - \delta(n_k, n_*), -\frac{1 - \alpha}{n_k} \right\} \right\}. \quad (19)$$

Perhaps surprisingly, this somewhat greedy approach typically produces valid predictions.

Proposition 2. *Under the setup of Theorem 2, assume also that $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta(n_k, n_*) - (1 - \alpha)/n_k$. If $\hat{C}(X_{n+1})$ is the prediction set output by Algorithm 1 applied with the set $\hat{\mathcal{I}}_k$ defined in (19) instead of (15), then $\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k] \geq 1 - \alpha$ for all $k \in [K]$.*

The additional assumption of Proposition 2, $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta(n_k, n_*) - (1 - \alpha)/n_k$, is stronger than the stochastic dominance condition in (11), but it is not unrealistic. When the calibration set size n_k is sufficiently large to make $\delta(n_k, n_*)$ small, this assumption is closely related to (11), which implies $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq 0$; see the proof of Corollary 1. In fact, Section 4 will show that the hybrid method described in this section tends to work very well in practice.

3.2 Adaptive coverage under a bounded label contamination model

We now extend Algorithm 1 by relaxing the assumption that the matrix M in (10) is fully known. In particular, we assume only that M is invertible and a joint confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ is available for the off-diagonal entries of $V = M^{-1}$, with $\hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \in \mathbb{R}^{K \times K}$, such that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ with probability at least $1 - \alpha_V$ simultaneously for all $l \neq k$, at some significance level $\alpha_V \in (0, 1)$. Here, it is understood that the matrices \hat{V}^{low} and \hat{V}^{upp} are independent of the data utilized to calibrate our conformal inferences. It should be anticipated that there will be some trade-offs involved in the choice of α_V , which should generally not exceed the desired level α of the output conformal prediction sets, but this matter will become clearer later.

To simplify the notation in the following, it is helpful to define $\hat{\delta}_{kl}^{(V)} := \hat{V}_{kl}^{\text{upp}} - \hat{V}_{kl}^{\text{low}}$ for all $l \neq k$. Further, it is useful to imagine that a (possibly very conservative) deterministic upper bound $\bar{V}^{\text{upp}} \in \mathbb{R}^{K \times K}$ for the off-diagonal entries of V is also known a priori, such that $\max\{|\hat{V}_{kl}^{\text{low}}|, |\hat{V}_{kl}^{\text{upp}}|, |V_{kl}|\} \leq |\bar{V}_{kl}^{\text{upp}}|$ almost-surely for all $l \neq k$. We refer to Section A3 for concrete examples of $\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}$ and \bar{V}^{upp} corresponding to two special cases in which more specific knowledge about the structure of the matrix V is also available. In the meantime, here we continue describing our method in generality, without additional constraints on V .

For ease of notation, let us define also

$$\hat{\delta}_{k^*}^{(V)} := \max_{l \neq k} \hat{\delta}_{kl}^{(V)}, \quad \hat{\zeta}_k := \max_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) - \min_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right). \quad (20)$$

Intuitively, $\hat{\delta}_{k^*}^{(V)}$ represents the width of the simultaneous confidence band for the off-diagonal entries of V at its widest point, while $\hat{\zeta}_k$ quantifies the uniformity of that confidence band. In particular, $\hat{\zeta}_k = 0$ if $\hat{V}_{kl}^{\text{low}}$, $\hat{V}_{kl}^{\text{upp}}$, and V_{kl} are constant for all $l \neq k$, as one should expect to be true under the special label contamination model discussed in Section A3.1, for example. Further, let $\hat{\zeta}_k^{\text{upp}}$ denote a $1 - \alpha_V$ upper confidence bound for $\hat{\zeta}_k$. This can generally be extracted directly from $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$, since $\hat{\zeta}_k$ is a known functional of \hat{V}_{kl} and V_{kl} for $l \neq k$, but it could also be informed by additional prior knowledge about the structure of the matrix V . We refer to Section A3.2 for some concrete examples on how to compute $\hat{\zeta}_k$ in practice.

Then, our solution consists of applying Algorithm 1 after replacing $\hat{\Delta}_k(t)$ in (14) with

$$\begin{aligned} \hat{\Delta}_k^{\text{ci}}(t) := & \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) - \hat{\delta}_{k^*}^{(V)} (K-1) \left| \hat{F}_k^k(t) - \frac{\sum_{l \neq k} \hat{F}_l^k(t)}{K-1} \right| \\ & - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right|, \end{aligned} \quad (21)$$

the correction factor $\delta(n_k, n_*)$ in (18) with

$$\begin{aligned} \delta^{\text{ci}}(n_k, n_*) := & c(n_k) + \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\ & + 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}|, \end{aligned} \quad (22)$$

and the set $\hat{\mathcal{I}}_k \subseteq [K]$ in (15) with

$$\hat{\mathcal{I}}_k^{\text{ci}} := \left\{ i \in [n_k] : \frac{i}{n_k} \geq 1 - \alpha - \hat{\Delta}_k^{\text{ci}}(S_{(i)}) + \delta^{\text{ci}}(n_k, n_*) \right\}. \quad (23)$$

This method, outlined by Algorithm 2, provably achieves label-conditional coverage.

Theorem 4. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n+1]$. Assume the general linear mixture contamination model described in Section 2.4 holds, with $V = M^{-1}$. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}^{\text{ci}}(X_{n+1})$ indicate the prediction set output by Algorithm 2 based on an independent simultaneous confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ such that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$ with probability at least $1 - \alpha_V$. Assume also that $\max\{|\hat{V}_{kl}^{\text{upp}}|, |V_{kl}|\} \leq |\bar{V}_{kl}^{\text{upp}}|$ almost-surely for all $l \neq k$, for some known constants $|\bar{V}_{kl}^{\text{upp}}| > 0$. Then, $\mathbb{P}[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k] \geq 1 - \alpha$ for all $k \in [K]$.*

Algorithm 2: Adaptive classification under a bounded label contamination model

- Input:** Data set $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with corrupted labels $\tilde{Y}_i \in [K]$.
 A $1 - \alpha_V$ joint confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ for the off-diagonal elements of V .
 Constants $\bar{V}_{kl}^{\text{upp}}$: $\max\{|\hat{V}_{kl}^{\text{upp}}|, |V_{kl}|\} \leq |\bar{V}_{kl}^{\text{upp}}|$ almost-surely for all $l \neq k$.
 Unlabeled test point with features X_{n+1} .
 Machine learning algorithm \mathcal{A} for training a K -class classifier.
 Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).
 Desired significance level $\alpha \in (0, 1)$.
- 1 Randomly split $[n]$ into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} .
 - 2 Train the classifier \mathcal{A} on the data in $\mathcal{D}^{\text{train}}$.
 - 3 Compute conformity scores $\hat{s}(X_i, k)$ using (1) for all $i \in \mathcal{D}^{\text{cal}}$ and all $k \in [K]$.
 - 4 Define the empirical CDF \hat{F}_l^k of $\{\hat{s}(X_i, k) : i \in \mathcal{D}_l^{\text{cal}}\}$ for all $l \in [K]$, as in (13).
 - 5 **for** $k = 1, \dots, K$ **do**
 - 6 Define $\mathcal{D}_k^{\text{cal}} = \{i \in \mathcal{D}^{\text{cal}} : \tilde{Y}_i = k\}$ and $n_k = |\mathcal{D}_k^{\text{cal}}|$.
 - 7 Sort $\{\hat{s}(X_i, k) : i \in \mathcal{D}_k^{\text{cal}}\}$ into $(S_{(1)}^k, S_{(2)}^k, \dots, S_{(n_k)}^k)$, in ascending order.
 - 8 Compute $\hat{F}_l^k(S_{(i)}^k)$ for all $i \in [n_k]$ and $l \in [K]$.
 - 9 Compute $\hat{\Delta}_k^{\text{ci}}(S_{(i)}^k)$ for all $i \in [n_k]$, as in (21).
 - 10 Compute $\delta^{\text{ci}}(n_k, n_*)$ using (22), based on a Monte Carlo estimate of $c(n_k)$ in (17).
 - 11 Construct the set $\hat{\mathcal{I}}_k^{\text{ci}} \subseteq [K]$ as in (23), based on $\hat{\Delta}_k^{\text{ci}}$ and δ^{ci} .
 - 12 Evaluate $\hat{\tau}_k$ by applying (16) with $\hat{\mathcal{I}}_k^{\text{ci}}$ instead of $\hat{\mathcal{I}}_k$.
 - 13 **end**
 - 14 Evaluate $\hat{C}^{\text{ci}}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}; \hat{\pi})$, where $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_K)$.
- Output:** Conformal prediction set $\hat{C}^{\text{ci}}(X_{n+1})$ for Y_{n+1} .
-

It is now clear that the level α_V of the confidence region for V affects the magnitude of the finite-sample correction term in (22) through the product $\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}|$. Therefore, there is an important trade-off in the choice of α_V , because smaller values of the latter tend to lead to larger upper bounds $|\bar{V}_{kl}^{\text{upp}}|$ and $|\hat{V}_{kl}^{\text{upp}}|$. In practice, we have observed that choosing a relatively small value such as $\alpha_V = 0.01$ often works well in practice, although it may not be optimal.

The following theorem formalizes the intuition that the prediction sets computed by Algorithm 2 tend to be more informative if the confidence bands for V_{kl} are tighter. This result naturally extends Theorem 3 to provide a coverage upper bound for the conformal prediction sets output by Algorithm 2. Similarly to Theorem 3, three technical conditions are needed: Assumptions 2, 3 and 5, with the latter being a suitable variation of Assumption 4.

Assumption 5. *The coverage inflation factor $\Delta_k(t)$ is almost-surely bounded by:*

$$\begin{aligned} \inf_{t \in (0,1)} \Delta_k(t) \geq & -\alpha + c(n_k) + \frac{2 \sum_{l \neq k} (|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)})}{\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right) \\ & + 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right). \end{aligned} \quad (24)$$

The difference between Assumption 5 and Assumption 4 is that the upper bound for $\Delta_k(t)$ imposed by the latter is a random variable that depends on the confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$. In the limit of $n_* \rightarrow \infty$, Assumption 5 becomes approximately equivalent to

$$\sum_{l \neq k} |V_{kl}| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) \leq \alpha - 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}|. \quad (25)$$

This means that Assumption 5 is often realistic, similarly to Assumption 4, as long as $\alpha_V \leq \alpha$, the contaminated labels are not too different from the true labels, and the confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ is not too wide. We refer to Section A3.1 for further details on the interpretation of (25), which can be simplified under more specific label contamination models. Moreover, we remark that Assumption 5 is always satisfied in the large-sample limit if the stochastic dominance condition in (11) holds (i.e., $\inf_{t \in (0,1)} \Delta_k(t) \geq 0$) and sufficiently tight confidence bounds $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ for a small enough α_V are available. Then, a finite-sample upper bound for the coverage of the conformal prediction sets output by Algorithm 2 is established below.

Theorem 5. *Under the setup of Theorem 4, assume the general contamination model from Section 2.4 holds. Let $\hat{C}^{\text{ci}}(X_{n+1})$ indicate the prediction set output by Algorithm 2 based on an independent simultaneous confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ such that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$ with probability at least $1 - \alpha_V$. Assume also that $\max\{|\hat{V}_{kl}^{\text{upp}}|, |V_{kl}|\} \leq |\bar{V}_{kl}^{\text{upp}}|$ almost-surely for all $l \neq k$, for some known constants $|\bar{V}_{kl}^{\text{upp}}| > 0$. Suppose also that Assumptions 2, 3, and 5 hold. Then, $\mathbb{P}[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k] \leq 1 - \alpha + \varphi_k^{\text{ci}}(n_k, n_*)$, for all $k \in [K]$, where*

$$\begin{aligned} \varphi_k^{\text{ci}}(n_k, n_*) = & \frac{1}{n_*} + \left(1 + 4 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \right) \alpha_V + 2c(n_k) + (K-1) \left(2\mathbb{E} \left[\hat{\delta}_{k*}^{(V)} \right] + \mathbb{E} \left[|\hat{\zeta}_k^{\text{upp}}| \right] \right) \\ & + \frac{4 \sum_{l \neq k} \mathbb{E} \left[|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right]}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\ & + \frac{2}{n_k} \left[1 + \sum_{l \neq k} |V_{kl}| + \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \sum_{j=1}^{n_k+1} \frac{1}{j} \right]. \end{aligned}$$

The interpretation of Theorem 5 is similar to that of Theorem 3, although now the unknown nature of the contamination model necessarily introduces some slack in the coverage

upper bound. In particular, note that $\varphi_k^{\text{ci}}(n_k, n_*)$ converges to a finite quantity as $n_* \rightarrow \infty$, but our prediction sets can still be (approximately) tight if $\bar{V}_{kl}^{\text{upp}}$ is finite for all $l \neq k$, α_V is small, and the expected lengths of all confidence intervals $[\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ are also small for all $l \neq k$.

We conclude this section by noting that the power of Algorithm 2 can be further boosted without losing the coverage guarantee, as long as a relatively mild “optimistic” condition on the coverage inflation factor in (7) holds. Concretely, we propose to apply Algorithm 2 based on

$$\hat{\mathcal{I}}_k^{\text{ci}} := \left\{ i \in [n_k] : \frac{i}{n_k} \geq 1 - \alpha - \max \left\{ \hat{\Delta}_k^{\text{ci}}(S_{(i)}) - \delta^{\text{ci}}(n_k, n_*), -\frac{1 - \alpha}{n_k} \right\} \right\}, \quad (26)$$

instead of the more conservative option described above in (23). This leads to an optimistic variation of Algorithm 2, analogous to the extension of Algorithm 1 presented earlier in Section 3.1.4, which still enjoys similar coverage properties.

Proposition 3. *Under the setup of Theorem 4, assume also that $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta^{\text{ci}}(n_k, n_*) - (1 - \alpha)/n_k$. If $\hat{C}(X_{n+1})$ is the set output by Algorithm 2 applied with the definition of $\hat{\mathcal{I}}_k^{\text{ci}}$ in (26) instead of (23), then $\mathbb{P}[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k] \geq 1 - \alpha$ for any $k \in [K]$.*

3.3 A general method for fitting the label contamination model

We now shift our focus to the estimation of the contamination model. On the one hand, practitioners sometimes have prior information about the matrix M in (10). For example, they might have knowledge of the labeling processes [6] or previous experiences with related data. Further, the contamination model can be known in applications involving differential privacy [11–14]. On the other hand, it is interesting to also consider a data-driven approach that can estimate M leveraging some observations of both contaminated and clean data.

Of course, if clean data are available, one may think of circumventing the problem studied in this paper by calibrating the conformal inferences without using the contaminated samples. However, data with high-quality labels are often scarce, and they may not always be available in sufficient numbers to reliably calibrate the K thresholds $\hat{\tau}_k$ needed to guarantee label-conditional coverage [24], especially if K is large. Further, even more abundant clean calibration data may be needed if one aims to achieve stronger guarantees such as equalized coverage over protected categories [25]. By contrast, our label contamination model can easily incorporate well-justified structural constraints that reduce the number of parameters to be estimated, making our estimation task seem manageable even with a small clean data set.

With this premise, we begin to tackle the estimation of $V = M^{-1}$ from a general perspective, avoiding for the time being any additional assumptions about the structure of M . Subsequently, in Section A3, our attention will turn to more specific contamination models that can be fitted quite accurately even with a very limited amount of clean data.

Let \mathcal{D}^1 denote a random contaminated data set, independent of \mathcal{D} but identically distributed, with $n^1 = |\mathcal{D}^1| < n$. Similarly, let \mathcal{D}^0 denote a (smaller) clean data set containing i.i.d. pairs of observations (X_i^0, Y_i^0) , for $i \in [n^0]$, independently drawn from the same distribution P_{XY} corresponding to the data in \mathcal{D} . First, randomly partition \mathcal{D}^1 into two disjoint subsets, \mathcal{D}_a^1 and \mathcal{D}_b^1 . The data in \mathcal{D}_a^1 are utilized to train a K -class classifier, possibly with the same machine learning algorithm utilized to compute the conformity scores. Let $\hat{f}(X) \in [K]$ denote the most likely label predicted by this classifier for a new sample with features X . Then, define the matrices $\tilde{Q} \in [0, 1]^{K \times K}$ and $Q \in [0, 1]^{K \times K}$ such that, for any $l, k \in [K]$,

$$\tilde{Q}_{lk} := \mathbb{P} \left[\hat{f}(X) = k \mid \tilde{Y} = l, \hat{f} \right], \quad Q_{lk} := \mathbb{P} \left[\hat{f}(X) = k \mid Y = l, \hat{f} \right]. \quad (27)$$

As proved in Section A7, the matrices Q and \tilde{Q} satisfy the following *estimating equation*:

$$\tilde{Q} = MQ. \quad (28)$$

Thus, as long as \tilde{Q} is invertible, $V := M^{-1}$ is given by

$$V = Q\tilde{Q}^{-1}. \quad (29)$$

Equation (29) suggests the following strategy for estimating V . First, note that the data in \mathcal{D}_b^1 can be utilized to compute an intuitive point estimate of \tilde{Q} ; that is, for each $l, k \in [K]$,

$$\tilde{Q}_{lk} \approx \frac{\sum_{i \in \mathcal{D}_b^1} \mathbb{I} \left[\hat{f}(X_i) = k, \tilde{Y} = l \right]}{\sum_{i \in \mathcal{D}_b^1} \mathbb{I} \left[\tilde{Y} = l \right]}. \quad (30)$$

Similarly, the data in \mathcal{D}^0 can be utilized to obtain an intuitive point estimate of Q ; i.e.,

$$Q_{lk} \approx \frac{\sum_{i \in \mathcal{D}^0} \mathbb{I} \left[\hat{f}(X_i) = k, Y = l \right]}{\sum_{i \in \mathcal{D}^0} \mathbb{I} \left[Y = l \right]}. \quad (31)$$

If the contaminated observations are relatively abundant (i.e., $|\mathcal{D}^1| \gg |\mathcal{D}^0|$), Equation (30) provides an estimate of \tilde{Q} with low variance compared to that of Q in (31). Therefore, the leading source of uncertainty in V comes from Q and is due to the unknown joint distribution of $(\hat{f}(X), Y)$ conditional on \hat{f} . This is a multinomial distribution with K^2 categories and event probabilities equal to $\lambda_{lk} = \mathbb{P} \left[\hat{f}(X) = k, Y = l \mid \hat{f} \right]$, for all $l, k \in [K]$. Then, since $Q_{lk} = \lambda_{lk} / \sum_{s=1}^K \lambda_{ls}$ for any $l, k \in [K]$, each element of the matrix V in (29) can be written as a function of the multinomial parameter vector $\lambda = (\lambda_{lk})_{l,k \in [K]}$ and of other quantities (i.e., \tilde{Q}) that are already known with relatively high accuracy:

$$V_{lk} = \frac{\sum_{s=1}^K \lambda_{ls} (\tilde{Q}^{-1})_{sk}}{\sum_{s=1}^K \lambda_{ls}}. \quad (32)$$

Thus, a simultaneous confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ for the off-diagonal entries of V can be obtained by applying standard parametric bootstrap techniques for multinomial parameters; e.g., see Sison and Glaz [48]. Finally, an adaptive prediction set $\hat{C}^{\text{ci}}(X_{n+1})$ for Y_{n+1} can be constructed by applying Algorithm 2 based on $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$. This two-step procedure is summarized by Algorithm A3 in Section A2.

Alternatively, one could seek only a point estimate \hat{V} of V , by replacing the multinomial parameters in (32) with their standard maximum-likelihood estimates. Then, it seems intuitive to construct prediction sets by applying Algorithm 1 with the *plug-in* estimate \hat{V} instead of V . Section 4 will demonstrate that this partly heuristic approach, outlined by Algorithm A4 in Section A2, tends to work well in practice and often leads to more informative prediction sets compared to Algorithm A3. To conclude this section, we note that the estimation process for the label contamination model, as covered here, can be streamlined and made more practically feasible in cases with limited clean data. This simplification is achieved by introducing additional assumptions about the structure of the matrix M , as discussed in Section A3.

4 Empirical demonstrations

Sections 4.1–4.3 demonstrate the usage of our methods on simulated data. Section 4.1 applies methods from Section 3.1 and their extensions from Section A6, assuming a known contamination model. Section 4.2 applies methods from Section 3.2, which are useful when the contamination model is unknown. Section 4.3 focuses on the estimation of the contamination model, applying methods from Section 3.3. Finally, Section 4.4 presents an application to image data.

4.1 Simulations under a known label contamination model

We begin by demonstrating the performance of our methods on synthetic data. For this purpose, we simulate classification data with $K = 4$ labels and $d = 50$ features from a Gaussian mixture distribution using the standard *make_classification* function from the Scikit-Learn Python package [49]. This function creates $2K$ clusters of points normally distributed, with unit variance, about the vertices of a 25-dimensional hypercube with sides of length 2, and then randomly assigns an equal number of clusters to each of the K classes. Note that this leads to uniform label frequencies; i.e., $\rho_k := \mathbb{P}[Y = k] = 1/K$ for all $k \in [K]$. We refer to Pedregosa et al. [49] and Guyon [50] for further details about the data-generating process. The results of additional experiments based on different data distributions are in Section A8. Conditional on the simulated data, the contaminated labels \tilde{Y} are generated following a *randomized response* model [11], an intuitive special case of the linear mixture model from Section 2.4. Specifically, $\mathbb{P}[\tilde{Y} = k \mid X, Y = l] = (1 - \epsilon)\mathbb{I}[k = l] + \epsilon/K$, which corresponds to $M_{kl} = \mathbb{P}[\tilde{Y} = k \mid X, Y = l] \cdot \rho_l / \tilde{\rho}_k$, for all $l, k \in [K]$, considering a range of values for $\epsilon \in [0, 0.2]$. Additional experiments based on different contamination processes

will be presented later.

A random forest classifier implemented by Scikit-Learn is trained on 10,000 independent observations with contaminated labels generated as described above. The classifier is then applied to an independent and identically distributed calibration data set, whose labels are also similarly contaminated, in order to construct generalized inverse quantile conformity scores with the recipe of Romano, Sesia, and Candès [40], reviewed in Section A1.1. These scores are transformed into prediction sets for 2000 independent unlabeled test points following three alternative approaches. The first one is the standard conformal inference approach, which seeks 90% label-conditional coverage while ignoring the presence of label contamination. The second approach, called *Adaptive*, is the method outlined by Algorithm 1, which we also apply with $\alpha = 0.1$. The third approach, called *Adaptive+*, is the optimistic variation of the *Adaptive* approach, as described in Section 3.1.4. Both the *Adaptive* and *Adaptive+* methods are applied assuming perfect knowledge of M .




Figure 2: Performances of different conformal prediction methods on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. The dashed horizontal line indicates the 90% nominal label-conditional coverage level.

Figure 2 compares the performances of the prediction sets obtained with the three alternative methods, measured in terms empirical coverage—the average proportion of test points for which the true label is contained in the prediction set—and average size. The results are shown as a function of the number of calibration data points and of the contamination parameter ϵ , averaging over 25 independent repetitions of each experiment. Unsurprisingly, the standard conformal prediction sets are overly conservative if $\epsilon > 0$ and their size does not change significantly as the number of calibration samples increases. By contrast, the *Adaptive* and *Adaptive+* methods tend to produce more informative prediction sets as the calibration sample grows. Further, the *Adaptive+* sets can be smaller than both the standard and *Adaptive* sets, as long as the number of calibration samples is large enough. Overall, these experiments demonstrate that our *Adaptive* and *Adaptive+* methods are effective at constructing more informative prediction sets with valid coverage in the presence of label contamination, and that the optimistic *Adaptive+* version is preferable in practice,

even though its theoretical guarantee relies on slightly stronger technical assumptions on the coverage inflation factor (ref. Section 3.1.4).

Additional numerical results. Figures A5–A17 in Section A8.1 present the results of further experiments with similar conclusions. Figure A5 reports additional performance metrics from the experiments of Figure 2, stratifying the results based on the true label of the test point. This confirms that all methods under comparison achieve 90% label-conditional coverage. Figure A6 gives an alternative view of these experiments, by plotting the results as a function of ϵ separately for different calibration sample sizes. Figure A7 presents results from experiments analogous to those in Figure A6, but fixing $\epsilon = 0.1$ and varying instead the number of labels K . Figure A8 presents additional results from experiments analogous to those in Figure 2, fixing $\epsilon = 0.1$ but utilizing different types of classifiers, namely a support vector machine and a neural network.

The effect of the data distribution. The robustness of our results to different data distributions is demonstrated by Figures A9–A10, which report on experiments similar to those of Figure 2. Figure A9 considers data with $K = 4$ classes simulated from a logistic model with random parameters, which is inspired by Romano, Sesia, and Candès [40] and described as follows. The features follow a standard multivariate Gaussian distribution of dimensions $d = 50$, and $Y | X$ is multinomial with weights proportional to $\exp[(X^T W)_k]$, for all $k \in [K]$, where $W \in \mathbb{R}^{d \times K}$ is an independent standard normal random vector. Figure A10 also involves synthetic data with $K = 4$ classes, but those data are generated from the heteroscedastic decision-tree model defined in Section A9.1, which is borrowed from Romano, Sesia, and Candès [40].

The effect of the label contamination process. The robustness of our results to different data distributions is demonstrated by Figures A11–A12, which report on experiments based on data simulated from a logistic model with random parameters, as in Figure A9. However, now the labels are contaminated differently. In Figure A11, the contamination process is described by a transition matrix $T \in [0, 1]^{K \times K}$, with $T_{kl} = \mathbb{P}[\tilde{Y} = k | X, Y = l]$ for all $l, k \in [K]$, given by $T = (1 - \epsilon)I_K + \epsilon/K \cdot B_{K,2}$, where $B_{K,2}$ is a block-diagonal matrix with $K/2$ constant blocks equal to J_2 —the 2×2 matrix of ones. In Figure A12, $T = (1 - \epsilon)I_K + \epsilon/K \cdot U_K$, where U_K is a matrix of i.i.d. uniform random numbers on $[0, 1]$, standardized to have its columns sum to one.

Prediction sets targeting other notions of coverage. Figures A13–A14 present results from experiments analogous to those in Figures 2 and A6, respectively, with the difference that all methods under comparison are applied to seek 90% marginal coverage instead of 90% label-conditional coverage. Within our adaptive framework, this is obtained by replacing Algorithm 1 with Algorithm A5, as explained in Section A6.1. Then, Figures A15–A16 demonstrate the robustness of Algorithm A5 to the empirical estimation of the contaminated label frequencies \tilde{p} from the available data. Finally, Figure A17 reports on experiments in which the goal is to achieve valid coverage conditional on the calibration data. Within our adaptive framework, this is obtained by replacing Algorithm 1 with Algorithm A6, as explained in Section A6.2.

4.2 Simulations under a bounded label contamination model

We now apply Algorithm 2, which does not require perfect knowledge of the model matrix M . As a starting point, we focus on a contamination process described by the randomized response model [11] with an unknown parameter. As explained in Section A3.1, this model streamlines the implementation of Algorithm 2, which in this case requires as input only a confidence interval for a scalar parameter $\epsilon \in [0, 1]$; see Section A3.1 for further details.

Figure 3 compares the performance of Algorithm 2 (*Adaptive*) and its optimistic variation (*Adaptive+*), from Section 3.2, to that of the standard conformal method that ignores label contamination, using synthetic data similar to those in Figure 2. The differences are that now the number of possible labels is varied, $K \in \{2, 4, 8\}$, the number of calibration samples is 10,000, and the true noise parameter is $\epsilon = 0.2$. The *Adaptive* (and *Adaptive+*) prediction sets are constructed by applying Algorithm 2 (and its optimistic variation) based on a 99% confidence interval for ϵ whose lower bound is varied, while the upper bound is fixed to 0.2. The results show that our prediction sets always achieve valid label-conditional coverage and become increasingly informative as the lower bound for ϵ increases, as anticipated by our theory.

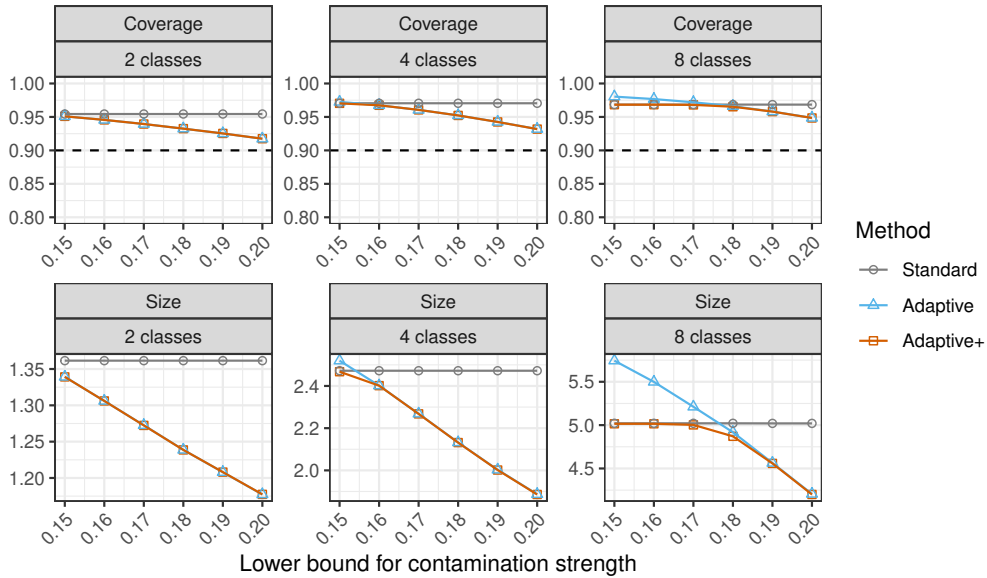


Figure 3: Performances of different conformal prediction methods on simulated data with contaminated labels from a randomized response model. The results are shown as a function of the known lower bound for the label noise parameter $\epsilon = 0.2$ and of the number of possible labels. The number of calibration samples is 10,000. Other details are as in Figure 2.

Additional results with qualitatively consistent conclusions are presented by Figures A18–A20 in Section A8.2.1. Figure A18 reports on experiments that differ from those in Figure 3 in that the upper bound for ϵ is varied while the lower bound is fixed equal to the true parameter value. Figures A19–A20 report on the results of experiments similar to those of

Figure 3, respectively fixing $\epsilon = 0.2$ and $\epsilon = 0.1$, while varying the corresponding lower confidence bound and the calibration set size. Overall, these results confirm that tighter bounds for ϵ generally allow Algorithm 2 to construct more informative prediction sets, and that the advantage of our adaptive method is more noticeable when the contaminated data are abundant.

Section A8.2.2 presents the results of similar experiments in which the label contamination process is more complex, following a *two-level* extension of the randomized response model. This model is discussed in Section A3.2 and can accommodate potential label hierarchies using two distinct parameters, $\epsilon \in [0, 1)$ and $\nu \in [0, 1]$. As explained in Section A3.2, Algorithm 2 also simplifies under a two-level randomized response model, and in this case it requires as input only a pair of simultaneously valid confidence intervals for ϵ and ν . Figures A21–A24 report on the performance of our method as a function of the width of the confidence interval for ϵ , using different values of ν , confidence intervals for ν , and calibration set sizes. Figures A25–A28 report on similar experiments in which the width of the confidence interval for ν is varied, using different values of ν , confidence intervals for ϵ , and calibration set sizes. Overall, these results support the previous conclusions: Algorithm 2 leads to more informative prediction sets compared to standard conformal methods even if the label contamination model is unknown. Further, its advantage tends to grow as the contaminated calibration set becomes larger.

4.3 Robustness to model estimation and mis-specification

We now examine the performance of the methods described in Section 3.3 for estimating the contamination model using an independent “model-fitting” data set containing both clean and contaminated labels. Again, we begin by focusing on the randomized response model. More challenging estimation settings will be considered subsequently.

To begin, we focus on a randomized response model with an unknown scalar parameter ϵ . We compare the performance of three alternative implementations of our *Adaptive+* method to the standard conformal approach, on synthetic data with $K = 2$ labels, similar to those utilized for Figure 3. The first implementation of our method (*Adaptive+*) consists of applying the optimistic version of Algorithm 1 using perfect *oracle* knowledge of the matrix M , as written explicitly in Section A3.1, based on the correct value $\epsilon = 0.2$. This corresponds to the general method presented in Section 4.1. The second implementation of our method, which we call *Adaptive+ (plug-in)*, consists of applying the optimistic version of Algorithm 1 using an approximate version of M obtained by replacing the unknown noise parameter ϵ with an intuitive point estimate $\hat{\epsilon}$ calculated from the model-fitting data as explained in Section A3.1—the latter simplifies the more general procedure described in Section 3.3 for the special case of the randomized response model. The third implementation of our method, which we call *Adaptive+ (CI)*, consists of applying the optimistic version of Algorithm 2 using a 99% bootstrap confidence interval for ϵ . This confidence interval is produced by a specialized version of the general method from Section 3.3, as explained in Section A3.1. For simplicity, Algorithm 2 is applied assuming a known fixed upper bound

for $\epsilon = 0.2$ equal to $\bar{\epsilon} = 0.2$, so that the bootstrap is effectively only needed to estimate the lower confidence bound.

Figure 4 reports on the performance of all methods as a function of the size and composition of the model-fitting data set. The results show that the heuristic *Adaptive+* (*plug-in*) method performs very similarly to the ideal *Adaptive+* method based on oracle knowledge of the true contamination parameter ϵ . By contrast, the *Adaptive+* (*CI*) tends to be more conservative and can lead to prediction sets that are significantly more informative compared to the standard conformal inference benchmark only if the number of clean samples in the model-fitting data set is large. Figure A29 in Section A8.3.1 plots explicitly the average upper and lower bounds of the bootstrap confidence intervals estimated by the *Adaptive+* (*CI*) in the experiments of Figure 4.




Figure 4: Performances of different conformal prediction methods, as a function of the numbers of clean and corrupted samples used to fit the unknown parameter of a randomized response model for the contamination process. The number of possible classes is $K = 2$. Other details are as in Figure 3.

Further results from related experiments, with qualitatively similar conclusions, are presented by Figures A30–A34 in Section A8.3.1. Figures A30 and A31 report on results analogous to those in Figures 4 and A29, respectively, with the only difference that now the fixed upper bound for $\epsilon = 0.2$ utilized by the *Adaptive+* (*CI*) method is set equal to $\bar{\epsilon} = 0.25$. Figures A32–A34 further demonstrate the power and robustness of the *Adaptive+* (*plug-in*) method in experiments with different values of the true noise parameter ϵ and different numbers of possible labels.

Section A8.3.2 details further experiments conducted within a more general contamination process, described by a two-level randomized response model. Specifically, Figure A35 focuses on synthetic data with $K = 4$ classes and summarizes the performances of prediction sets obtained with different methods, while Figures A36–A37 plot the corresponding estimated confidence intervals for the unknown parameters of the contamination model.

Figures A38–A40 report on similar results for data with $K = 8$ classes. Overall, our *Adaptive+* (*plug-in*) method performs well across all scenarios considered, even though it ignores the true contamination model parameters. Notably, this partly heuristic implementation of Algorithm 1 attains valid coverage empirically and leads to more informative predictions compared to the more rigorous *Adaptive+* (*CI*) approach of Algorithm 2. Moreover, our estimators of the contamination model parameters (Section 3.3) are confirmed to be consistent, in the sense that their estimated confidence intervals become narrower as the number of samples available for estimation increases.

Finally, Figures A41–A42 in Section A8.4 describe experiments demonstrating the robustness of our methods to possible mis-specification of the label contamination process, going beyond the estimation of the parameters in a (generalized) randomized response model.

4.4 Demonstrations with CIFAR-10 image data

This section demonstrates the use of our methods in a object classification application based on real-world 32x32 color images. As anticipated in Section 1, we focus on the CIFAR-10H data set [23], a variation of the larger CIFAR-10 data set [51] that includes imperfect labels assigned by approximately 50 independent human annotators via the Amazon Mechanical Turk, for a subset of 10,000 images. Each image depicts an object belonging to one of 10 possible classes: airplane, car, bird, cat, deer, dog, frog, horse, ship, or truck. Since the individual annotators do not always agree on the content of each image, we can think of their labels as being a randomly contaminated version of the corresponding “true” labels contained in the original CIFAR-10 data set. Our goal is to construct informative prediction sets for the true labels, using a conformal predictor calibrated on the contaminated data. For simplicity, we work with a slightly modified version of the CIFAR-10H data in which each image has a single corrupted label \tilde{Y} , randomly sampled from a multinomial distribution whose weights are equal to the relative label frequencies assigned to that image by different human annotators. Note that these corrupted labels coincide with the true CIFAR-10 labels approximately 95.4% of the time; see Figure A43 (a) for a visualization of some images for which the labels do not match.

A ResNet-18 convolutional neural network serves as base classifier; this was implemented by the PyTorch Python package [52] and pre-trained using the 50,000 CIFAR-10 images excluded from the CIFAR-10H data set. The output of the final soft-max layer of the deep neural network provides estimates of the class probabilities for any new given image, and from that we calculate conformity scores with the recipe reviewed in Section A1.1. The conformal predictor is then calibrated using three alternative methods, based on a random subset of the 10,000 CIFAR-10H images whose size is varied as a control parameter.

Due to the larger number of labels, $K = 10$, and to help highlight the flexibility of our methods, we focus first on marginal (5) instead of label-conditional (4) coverage. The first method considered is Algorithm A2, the standard conformal approach that seeks marginal coverage while ignoring label contamination. The second method, *Adaptive+*, is the opti-

mistic variation of our Algorithm A5 from Section A6.1, which we apply imagining that the label contamination process follows the randomized response model described in Section A3.1. This model assumes that the observed labels are conditionally independent of the image features, given the true labels, and that an incorrect label is as likely as any other incorrect label. Despite the potential misalignment with the true data generating process, this model is useful for our purposes. Specifically, it simplifies the matrix M , allowing us to characterize it using a single noise parameter, $\epsilon \in [0, 1)$, which is fixed here as $\epsilon = 0.051$. This ϵ value is selected to match the mean fraction of CIFAR-10H samples where $Y \neq \tilde{Y}$, roughly $0.046 = \epsilon(1 - 1/K)$. The third method, *Adaptive+* (*plug-in*), differs from *Adaptive+* in that it utilizes a plug-in estimate of ϵ obtained via maximum-likelihood. This estimate is evaluated as explained in Section A3.1, by applying the same pre-trained ResNet-18 convolutional neural network to a smaller independent data set containing both clean and corrupted data in equal proportions. In particular, the number of clean images used to estimate ϵ is 10% of the total number of corrupted calibration images.

Figure 1, previewed in Section 1.2, reports on the prediction sets constructed by the three methods for a random test set 500 CIFAR-10H images, varying the size of the calibration sample between 500 and 9500. All experiments are independently repeated 50 times, using different random splits of the CIFAR-10H data into calibration, model-fitting, and test subsets. The results show that the standard method is overly conservative, while our adaptive approaches are able to achieve valid coverage with increasingly more informative prediction sets as the size of the calibration sample grows. See Figure A43 (b) for a visualization of some concrete examples in which our *Adaptive+* method leads to more informative prediction sets compared to the standard conformal inference benchmark. Finally, Figure A44 in Section A8.5 presents analogous results from similar experiments in which we target label-conditional coverage, using Algorithms 1 and A1 instead of Algorithms A5 and A2, respectively. Note that Algorithm 1 requires a larger data set in this case, compared to the marginal coverage setting, in order to produce prediction sets that are significantly more informative compared to those of Algorithm A1. The reason for this is the stronger nature of the label-conditional guarantee, which effectively diminishes the usable sample size for both Algorithm 1 and Algorithm A1 by a factor of $K = 10$.

5 Discussion

This paper studied in-depth the problem of conformal classification in the presence of calibration data with contaminated labels. Our research contributes to the growing literature on conformal inference beyond exchangeability [20], addressing an open practical problem. A key innovation of our approach is its capacity to automatically adjust to random label contamination, resulting in prediction sets that offer more robust coverage guarantees and are often more informative compared to those given by state-of-the-art approaches. Additionally, our framework is highly flexible, enabling several variations of our methodology that target different coverage metrics and can accommodate varying degrees of understanding of the label contamination process. These adaptations hint at potential for further

expansion in future research. Both theoretical and practical findings presented in this paper underscore the enhanced benefits of our adaptive approach when employed with larger sets of contaminated calibration data. This aspect renders our research particularly applicable to real-world situations where there are abundant data of modest quality, but precise labels are scarce.

This research opens several opportunities for future work. For example, it may be interesting to study possible extensions of our methods that can be applied with regression data, or even with other types of more complex data for which conformal inference has already been utilized, including causal inference [53], survival analysis [54], and matrix completion [55]. Alternatively, it may be possible to account for label contamination in the context of more sophisticated conformal prediction frameworks such as full-conformal inference [1] and cross-validation+ [56], which are more computationally expensive but can make more efficient use of limited observations. Finally, future research might be able to uncover further methodological opportunities by combining the ideas presented in this paper with the theoretical analyses of [20].

A software implementation of the methods presented in this paper is available online at <https://github.com/mnesia/conformal-label-noise>.

Acknowledgements

M. S. was supported by NSF grant DMS 2210637 and by an Amazon Research Award. We are grateful to two anonymous referees for their constructive feedback about an earlier version of this manuscript. We also thank Adel Javanmard for helpful suggestions.

References

- [1] V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic learning in a random world*. Springer, 2005.
- [2] N. Natarajan, I. S. Dhillon, P. K. Ravikumar, and A. Tewari. “Learning with noisy labels”. In: *Adv. Neural Inf. Process. Syst.* 26 (2013).
- [3] S. Sukhbaatar, J. Bruna, M. Paluri, L. Bourdev, and R. Fergus. “Training convolutional networks with noisy labels”. In: *preprint at arXiv:1406.2080* (2014).
- [4] H. Song, M. Kim, D. Park, Y. Shin, and J.-G. Lee. “Learning from noisy labels with deep neural networks: A survey”. In: *IEEE Transactions on Neural Networks and Learning Systems* (2022).
- [5] S. Yao, B. Rava, X. Tong, and G. James. “Asymmetric Error Control Under Imperfect Supervision: A Label-Noise-Adjusted Neyman-Pearson Umbrella Algorithm”. In: *J. Am. Stat. Assoc.* (2022), pp. 1–13.
- [6] A. Sorokin and D. Forsyth. “Utility data annotation with Amazon mechanical turk”. In: *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*. IEEE, 2008, pp. 1–8.
- [7] P. G. Ipeirotis, F. Provost, and J. Wang. “Quality management on Amazon mechanical turk”. In: *Proceedings of the ACM SIGKDD Workshop on Human Computation*. 2010, pp. 64–67.
- [8] R. Snow, B. O’connor, D. Jurafsky, and A. Y. Ng. “Cheap and fast—but is it good? evaluating non-expert annotations for natural language tasks”. In: *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*. 2008, pp. 254–263.
- [9] R. Kennedy, S. Clifford, T. Burleigh, P. D. Waggoner, R. Jewell, and N. J. Winter. “The shape of and solutions to the MTurk quality crisis”. In: *Political Science Research and Methods* 8.4 (2020), pp. 614–629.
- [10] H. Aguinis, I. Villamor, and R. S. Ramani. “MTurk research: Review and recommendations”. In: *Journal of Management* 47.4 (2021), pp. 823–837.
- [11] S. L. Warner. “Randomized response: A survey technique for eliminating evasive answer bias”. In: *J. Am. Stat. Assoc.* 60.309 (1965), pp. 63–69.
- [12] A. Evfimievski, J. Gehrke, and R. Srikant. “Limiting privacy breaches in privacy preserving data mining”. In: *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of Database Systems*. 2003, pp. 211–222.
- [13] S. P. Kasiviswanathan, H. K. Lee, K. Nissim, S. Raskhodnikova, and A. Smith. “What can we learn privately?” In: *SIAM Journal on Computing* 40.3 (2011), pp. 793–826.
- [14] B. Ghazi, N. Golowich, R. Kumar, P. Manurangsi, and C. Zhang. “Deep learning with label differential privacy”. In: *Adv. Neural Inf. Process. Syst.* 34 (2021), pp. 27131–27145.
- [15] A. N. Angelopoulos, S. Bates, T. Zrnica, and M. I. Jordan. “Private prediction sets”. In: *Harvard Data Science Review* (2022).

- [16] C. G. Northcutt, A. Athalye, and J. Mueller. “Pervasive label errors in test sets destabilize machine learning benchmarks”. In: *preprint at arXiv:2103.14749* (2021).
- [17] G. Algan and I. Ulusoy. “Image classification with deep learning in the presence of noisy labels: A survey”. In: *Knowledge-Based Systems* 215 (2021), p. 106771.
- [18] M. Cauchois, S. Gupta, A. Ali, and J. Duchi. “Predictive inference with weak supervision”. In: *preprint at arXiv:2201.08315* (2022).
- [19] B.-S. Einbinder, S. Bates, A. N. Angelopoulos, A. Gendler, and Y. Romano. “Conformal Prediction is Robust to Label Noise”. In: *preprint at arXiv:2209.14295* (2022).
- [20] R. F. Barber, E. J. Candès, A. Ramdas, and R. J. Tibshirani. “Conformal prediction beyond exchangeability”. In: *Ann. Stat.* 51.2 (2023), pp. 816–845.
- [21] C. Scott, G. Blanchard, and G. Handy. “Classification with asymmetric label noise: Consistency and maximal denoising”. In: *Conference on Learning Theory*. PMLR, 2013, pp. 489–511.
- [22] A. Ghosh, H. Kumar, and P. S. Sastry. “Robust loss functions under label noise for deep neural networks”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 31. 1. 2017.
- [23] J. C. Peterson, R. M. Battleday, T. L. Griffiths, and O. Russakovsky. “Human uncertainty makes classification more robust”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2019, pp. 9617–9626.
- [24] V. Vovk. “Conditional Validity of Inductive Conformal Predictors”. In: *Proceedings of the Asian Conference on Machine Learning*. Vol. 25. 2012, pp. 475–490.
- [25] Y. Romano, R. F. Barber, C. Sabatti, and E. Candès. “With Malice Toward None: Assessing Uncertainty via Equalized Coverage”. In: *Harvard Data Science Review* (2020).
- [26] J. Lei, J. Robins, and L. Wasserman. “Distribution-Free Prediction Sets”. In: *J. Am. Stat. Assoc.* 108.501 (2013), pp. 278–287.
- [27] J. Lei and L. Wasserman. “Distribution-free prediction bands for non-parametric regression”. In: *J. R. Stat. Soc. (B)* 76.1 (2014), pp. 71–96.
- [28] J. Lei, M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman. “Distribution-free predictive inference for regression”. In: *J. Am. Stat. Assoc.* 113.523 (2018), pp. 1094–1111.
- [29] J. Smith, I. Nouretdinov, R. Craddock, C. Offer, and A. Gammerman. “Conformal anomaly detection of trajectories with a multi-class hierarchy”. In: *International symposium on statistical learning and data sciences*. Springer, 2015, pp. 281–290.
- [30] L. Guan and R. Tibshirani. “Prediction and outlier detection in classification problems”. In: *J. R. Stat. Soc. (B)* 84.2 (2022), pp. 524–546.
- [31] Z. Liang, M. Sesia, and W. Sun. “Integrative conformal p-values for out-of-distribution testing with labelled outliers”. In: *J. R. Stat. Soc. (B)* (2024), qkad138.
- [32] S. Bates, E. Candès, L. Lei, Y. Romano, and M. Sesia. “Testing for outliers with conformal p-values”. In: *Ann. Stat.* 51.1 (2023), pp. 149–178.

- [33] M. Bashari, A. Epstein, Y. Romano, and M. Sesia. “Derandomized novelty detection with FDR control via conformal e-values”. In: *Adv. Neural Inf. Process. Syst.* Vol. 36. 2023.
- [34] Y. Romano, E. Patterson, and E. Candès. “Conformalized quantile regression”. In: *Adv. Neural Inf. Process. Syst.* Vol. 32. 2019, pp. 3538–3548.
- [35] M. Sesia and E. Candès. “A comparison of some conformal quantile regression methods”. In: *Stat 9.1* (2020).
- [36] M. Sesia and Y. Romano. “Conformal Prediction using Conditional Histograms”. In: *Adv. Neural Inf. Process. Syst.* 34 (2021).
- [37] Y. Hechtlinger, B. Póczos, and L. Wasserman. “Cautious Deep Learning”. In: *preprint at arXiv:1805.09460* (2018).
- [38] M. Sadinle, J. Lei, and L. Wasserman. “Least ambiguous set-valued classifiers with bounded error levels”. In: *J. Am. Stat. Assoc.* 114.525 (2019), pp. 223–234.
- [39] M. Cauchois, S. Gupta, and J. C. Duchi. “Knowing what You Know: valid and validated confidence sets in multiclass and multilabel prediction.” In: *J. Mach. Learn. Res.* 22 (2021), pp. 81–1.
- [40] Y. Romano, M. Sesia, and E. Candès. “Classification with Valid and Adaptive Coverage”. In: *Adv. Neural Inf. Process. Syst.* 33 (2020).
- [41] B.-S. Einbinder, Y. Romano, M. Sesia, and Y. Zhou. “Training Uncertainty-Aware Classifiers with Conformalized Deep Learning”. In: *Adv. Neural Inf. Process. Syst.* Vol. 35. 2022.
- [42] R. J. Tibshirani, R. Foygel Barber, E. Candès, and A. Ramdas. “Conformal prediction under covariate shift”. In: *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [43] M. Sesia, S. Favaro, and E. Dobriban. “Conformal frequency estimation using discrete sketched data with coverage for distinct queries”. In: *J. Mach. Learn. Res.* 24.348 (2023), pp. 1–80.
- [44] I. Gibbs and E. Candès. “Adaptive conformal inference under distribution shift”. In: *Adv. Neural Inf. Process. Syst.* 34 (2021), pp. 1660–1672.
- [45] A. Podkopaev and A. Ramdas. “Distribution-free uncertainty quantification for classification under label shift”. In: *Uncertainty in Artificial Intelligence*. PMLR. 2021, pp. 844–853.
- [46] W. Si, S. Park, I. Lee, E. Dobriban, and O. Bastani. “PAC Prediction Sets Under Label Shift”. In: *arXiv preprint arXiv:2310.12964* (2023).
- [47] V. Vovk, D. Lindsay, I. Nourtdinov, and A. Gammerman. *Mondrian Confidence Machine*. Technical Report. On-line Compression Modelling project. Royal Holloway, University of London, 2003.
- [48] C. P. Sison and J. Glaz. “Simultaneous confidence intervals and sample size determination for multinomial proportions”. In: *J. Am. Stat. Assoc.* 90.429 (1995), pp. 366–369.

- [49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [50] I. Guyon. “Design of experiments for the NIPS 2003 variable selection benchmark 2003 [Available from: <http://clopinet.com/isabelle/Projects>”]. In: *NIPS2003* (2003).
- [51] A. Krizhevsky, V. Nair, and G. Hinton. “The CIFAR-10 dataset (2014)”. In: *Online: <http://www.cs.toronto.edu/kriz/cifar.html>* 55 (2020).
- [52] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al. “Pytorch: An imperative style, high-performance deep learning library”. In: *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [53] L. Lei and E. Candès. “Conformal inference of counterfactuals and individual treatment effects”. In: *J. R. Stat. Soc. (B)* 83.5 (2021), pp. 911–938.
- [54] E. Candès, L. Lei, and Z. Ren. “Conformalized survival analysis”. In: *J. R. Stat. Soc. (B)* 85.1 (2023), pp. 24–45.
- [55] Y. Gui, R. F. Barber, and C. Ma. “Conformalized matrix completion”. In: *preprint at arXiv:2305.10637* (2023).
- [56] R. F. Barber, E. Candès, A. Ramdas, and R. J. Tibshirani. “Predictive inference with the jackknife+”. In: *Ann. Stat.* 49.1 (2021), pp. 486–507.
- [57] J. Duchi, M. J. Wainwright, and M. I. Jordan. “Local privacy and minimax bounds: Sharp rates for probability estimation”. In: *Adv. Neural Inf. Process. Syst.* 26 (2013).
- [58] P. Kairouz, K. Bonawitz, and D. Ramage. “Discrete distribution estimation under local privacy”. In: *International Conference on Machine Learning*. PMLR, 2016, pp. 2436–2444.

A1 Review of standard conformal classification methods

A1.1 Conformity scores based on generalized inverse quantiles

We briefly review here the construction of the generalized inverse quantile conformity scores proposed by [40], upon which we rely in the empirical demonstrations of Section 4. These conformity scores are more involved compared to the classical homogeneous scores defined in (2), but they have the advantage of leading to more flexible prediction sets that can account for possible heteroscedasticity in the distribution of $Y \mid X$. We refer to [40] or [39] for further information about the limitations of the scores in (2).

For any $x \in \mathbb{R}^d$ and $t \in [0, 1]$, define

$$\hat{Q}(x, \hat{\pi}, t) = \min\{k \in \{1, \dots, K\} : \hat{\pi}_{(1)}(x) + \hat{\pi}_{(2)}(x) + \dots + \hat{\pi}_{(k)}(x) \geq t\}, \quad (\text{A33})$$

where $\hat{\pi}_{(1)}(x) \geq \dots \geq \hat{\pi}_{(K)}(x)$ are the descending order statistics of $\hat{\pi}(x, 1), \dots, \hat{\pi}(x, K)$. Intuitively, $\hat{Q}(x, \hat{\pi}, \cdot)$ may be seen as a generalized quantile function. Similarly, let $\hat{r}(x, \hat{\pi}, k)$ denote the rank of $\hat{\pi}(x, k)$ among $\hat{\pi}(x, 1), \dots, \hat{\pi}(x, K)$. With this notation, one can also define a corresponding generalized cumulative distribution function:

$$\hat{\Pi}(x, \hat{\pi}, k) = \hat{\pi}_{(1)}(x) + \hat{\pi}_{(2)}(x) + \dots + \hat{\pi}_{(\hat{r}(x, \hat{\pi}, k))}(x).$$

Then, the function \mathcal{C} proposed by [40] can be written as:

$$\mathcal{C}(x; \hat{\pi}, \tau) = \{k \in [K] : \hat{r}(x, \hat{\pi}, k) \leq \hat{Q}(x, \hat{\pi}, \tau_k)\}, \quad (\text{A34})$$

and the corresponding conformity scores defined in (1) can be evaluated efficiently by noting that $\hat{s}(x, k) = \hat{\Pi}(x, k)$; see [40] for further details.

The prediction function defined in (A34) may be understood by noting that, if $\tau = (\tau_0, \dots, \tau_0)$ for some $\tau_0 \in [0, 1]$, the output of $\mathcal{C}(x; \hat{\pi}, \tau)$ is the list of most likely classes according to $\hat{\pi}(x)$ up until the first label l for which $\hat{\Pi}(x, \hat{\pi}, l) \geq \tau_0$. Therefore, in the ideal case where $\hat{\pi}(x, k) = \mathbb{P}[Y = k \mid X = x]$, one can verify that $\mathcal{C}(x; \hat{\pi}, \tau)$ is the smallest possible (deterministic) prediction set for Y with perfect object-conditional coverage at level τ_0 , i.e., satisfying $\mathbb{P}[Y \in \mathcal{C}(x; \hat{\pi}, \tau) \mid X = x] \geq \tau_0$. Note that [40] also developed a more powerful randomized version of (A34) that enjoys similar theoretical properties while being able to produce even more informative prediction sets. The results of this paper can also seamlessly accommodate such additional randomness in \mathcal{C} , and indeed that is the practical approach followed in the empirical demonstrations of Section 4, but we choose not to review such extension explicitly here to avoid making the notation too cumbersome.

A1.2 Prediction sets with label-conditional coverage

Algorithm A1: Standard conformal classification with label-conditional coverage.

- Input:** Data set $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$, with observable labels $\tilde{Y}_i \in [K]$.
 Unlabeled test point with features X_{n+1} .
 Machine learning algorithm \mathcal{A} for training a K -class classifier.
 Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).
 Desired coverage parameter α .
- 1 Randomly split $[n]$ into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} .
 - 2 Train the classifier \mathcal{A} on the data in $\mathcal{D}^{\text{train}}$.
 - 3 **for** $k = 1, \dots, K$ **do**
 - 4 Define $\mathcal{D}_k^{\text{cal}} = \{i \in \mathcal{D}^{\text{cal}} : Y_i = k\}$ and $n_k = |\mathcal{D}_k^{\text{cal}}|$.
 - 5 Compute $\hat{s}(X_i, k)$ using (1), for all $i \in \mathcal{D}_k^{\text{cal}}$.
 - 6 Define $\hat{\tau}_k$ as the $\lceil (1 + n_k) \cdot (1 - \alpha) \rceil$ smallest value in $\{\hat{s}(X_i, k)\}_{i \in \mathcal{D}_k^{\text{cal}}}$.
 - 7 **end**
 - 8 Evaluate $\hat{C}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}; \hat{\pi})$, where $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_K)$.
- Output:** Conformal prediction set $\hat{C}(X_{n+1})$ for Y_{n+1} , satisfying (4).
-

Proposition A1 (e.g., from [47] or [40]). *If the data pairs (X_i, \tilde{Y}_i) , for all $i \in [n + 1]$, are exchangeable random samples from some joint distribution, the prediction set $\hat{C}(X_{n+1})$ output by Algorithm A1 has label-conditional coverage (4) for the observable labels \tilde{Y} ; i.e.,*

$$\mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] \geq 1 - \alpha, \quad \text{for all } k \in [K].$$

Further, if all scores $\hat{s}(X_i, \tilde{Y}_i)$ computed by Algorithm A1 are almost-surely distinct,

$$\mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] \leq 1 - \alpha + \frac{1}{n_k + 1}, \quad \text{for all } k \in [K],$$

where $n_k = |\mathcal{D}_k^{\text{cal}}|$ is the number of data points with label k in the calibration set.

In words, this result tells us that Algorithm A1 is always guaranteed to achieve valid label-conditional coverage for the observable labels \tilde{Y} , regardless of which prediction function \mathcal{C} is employed. Note that in the standard setting without label contamination, the observable labels \tilde{Y} are simply assumed to be always equal to Y .

A1.3 Prediction sets with marginal coverage

Algorithm A2: Standard conformal classification with marginal coverage.

Input: Data set $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with observable labels $\tilde{Y}_i \in [K]$.

Unlabeled test point with features X_{n+1} .

Machine learning algorithm \mathcal{A} for training a K -class classifier.

Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).

Desired coverage parameter α .

- 1 Randomly split $[n]$ into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} , defining $n_{\text{cal}} = |\mathcal{D}^{\text{cal}}|$.
- 2 Train the classifier \mathcal{A} on the data in $\mathcal{D}^{\text{train}}$.
- 3 Compute $\hat{s}(X_i, \tilde{Y}_i)$ using (1), for all $i \in \mathcal{D}^{\text{cal}}$.
- 4 Define $\hat{\tau}$ as the $\lceil (1 + n_{\text{cal}}) \cdot (1 - \alpha) \rceil$ smallest value in $\{\hat{s}(X_i, \tilde{Y}_i)\}_{i \in \mathcal{D}^{\text{cal}}}$.
- 5 Evaluate $\hat{C}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}; \hat{\pi})$, where $\hat{\tau} = (\hat{\tau}_0, \dots, \hat{\tau}_0)$.

Output: Conformal prediction set $\hat{C}(X_{n+1})$ for \tilde{Y}_{n+1} , satisfying (5).

Proposition A2 (e.g., from [26] or [40]). *If the data pairs (X_i, \tilde{Y}_i) , for all $i \in [n + 1]$, are exchangeable random samples from some joint distribution, the prediction set $\hat{C}(X_{n+1})$ output by Algorithm A1 has marginal coverage (5) for the observable labels \tilde{Y} ; i.e.,*

$$\mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \right] \geq 1 - \alpha.$$

Further, if all scores $\hat{s}(X_i, \tilde{Y}_i)$ computed by Algorithm A1 are almost-surely distinct,

$$\mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \right] \leq 1 - \alpha + \frac{1}{n_{\text{cal}} + 1},$$

where $n_{\text{cal}} = |\mathcal{D}^{\text{cal}}|$ is the number of data points in the calibration set.

A2 Additional methodological details

Algorithm A3: Adaptive classification with contamination model fitting (CI)

- Input:** Data set $\mathcal{D} = \{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with corrupted labels $\tilde{Y}_i \in [K]$.
 Independent data set $\mathcal{D}^1 = \{(X_i^1, \tilde{Y}_i^1)\}_{i=1}^{n^1}$ with corrupted labels $\tilde{Y}_i^1 \in [K]$.
 Clean data set $\mathcal{D}^0 = \{(X_i^0, Y_i^0)\}_{i=1}^{n^0}$ with labels $Y_i^0 \in [K]$.
 Unlabeled test point with features X_{n+1} .
 Machine learning algorithm \mathcal{A} for training a K -class classifier.
 Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).
 Desired significance level $\alpha \in (0, 1)$.
 Desired significance level $\alpha_V \in (0, \alpha)$ for the estimation of V .
- 1 Randomly split the data in \mathcal{D}^1 into two disjoint subsets, \mathcal{D}_a^1 and \mathcal{D}_b^1 .
 - 2 Train the classifier \mathcal{A} on the data in \mathcal{D}_a^1 .
 - 3 Assess the accuracy of \mathcal{A} to predict \tilde{Y} in \mathcal{D}_b^1 and Y in \mathcal{D}^0 , using (30) and (31).
 - 4 Construct a $1 - \alpha_V$ joint confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ for the off-diagonal entries of $V := M^{-1}$ using (32) and the parametric bootstrap, as explained in Section 3.3.
 - 5 Apply Algorithm 2 to construct a conformal prediction set $\hat{C}^{\text{ci}}(X_{n+1})$ for Y_{n+1} .
- Output:** Conformal prediction set $\hat{C}^{\text{ci}}(X_{n+1})$ for Y_{n+1} .
-

Algorithm A4: Adaptive classification with contamination model fitting (plug-in)

- Input:** Data set $\mathcal{D} = \{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with corrupted labels $\tilde{Y}_i \in [K]$.
 Independent data set $\mathcal{D}^1 = \{(X_i^1, \tilde{Y}_i^1)\}_{i=1}^{n^1}$ with corrupted labels $\tilde{Y}_i^1 \in [K]$.
 Clean data set $\mathcal{D}^0 = \{(X_i^0, Y_i^0)\}_{i=1}^{n^0}$ with labels $Y_i^0 \in [K]$.
 Unlabeled test point with features X_{n+1} .
 Machine learning algorithm \mathcal{A} for training a K -class classifier.
 Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).
 Desired significance level $\alpha \in (0, 1)$.
- 1 Randomly split the data in \mathcal{D}^1 into two disjoint subsets, \mathcal{D}_a^1 and \mathcal{D}_b^1 .
 - 2 Train the classifier \mathcal{A} on the data in \mathcal{D}_a^1 .
 - 3 Assess the accuracy of \mathcal{A} to predict \tilde{Y} in \mathcal{D}_b^1 and Y in \mathcal{D}^0 , using (30) and (31).
 - 4 Calculate a point-estimate for the off-diagonal entries of $V := M^{-1}$ using (32), as explained in Section 3.3.
 - 5 Apply Algorithm 1 to construct a conformal prediction set $\hat{C}(X_{n+1})$ for Y_{n+1} .
- Output:** Conformal prediction set $\hat{C}(X_{n+1})$ for Y_{n+1} .
-

A3 Simplified methods for special contamination models

This section illustrates two specific instances of the general label contamination model introduced in Section 2.4. These are particular cases in which the implementation of the adaptive conformal prediction methods developed in this paper can be significantly streamlined, and the estimation of their parameters, as discussed in Section 3.3, is also simplified. We begin in Section A3.1 with an exploration of label contamination processes described by the one-parameter randomized response model [11]. Following this, Section A3.2 expands to consider more complex processes described by a two-parameter extension of the randomized response model, which is specifically designed to accommodate potential label hierarchies.

A3.1 The randomized response model

A3.1.1 Model description

Suppose that the relation between X, Y and \tilde{Y} satisfies

$$\mathbb{P}[\tilde{Y} = k \mid X, Y = l] = (1 - \epsilon)\mathbb{I}[k = l] + \frac{\epsilon}{K}, \quad (\text{A35})$$

for all $l, k \in [K]$, where $\epsilon \in [0, 1)$ is a scalar parameter controlling the amount of random label *noise*. This setup corresponds to the classical *randomized response* model of Warner [11], and it has recently found many relevant applications in the context of differential privacy [14, 57, 58]. In particular, a well-known technique for achieving $(\epsilon, 0)$ -*label differential privacy* consists of replacing each individual observation of Y with a noisy label \tilde{Y} according to the randomized response model defined in (A35), with the parameter ϵ given by

$$\epsilon = \frac{K}{e^\epsilon + K - 1}.$$

We refer to [14] for a formal definition of label differential privacy and a proof of this result.

In the notation of Section 2.4, the randomized response model described above corresponds to a matrix M (10) with the form

$$M_{kl} = \frac{(1 - \epsilon)\rho_k}{(1 - \epsilon)\rho_k + \epsilon/K}\mathbb{I}[k = l] + \frac{\rho_l \cdot \epsilon/K}{(1 - \epsilon)\rho_k + \epsilon/K}, \quad (\text{A36})$$

for any $k, l \in [K]$, while the contaminated label frequencies become

$$\tilde{\rho}_k = (1 - \epsilon)\rho_k + \frac{\epsilon}{K}. \quad (\text{A37})$$

Note that Equation (A36) follows from (A35) with a straightforward application of Bayes' rule, and then Equation (A37) is easily obtained by recalling that $\tilde{\rho}_k := \mathbb{P}[\tilde{Y} = k]$ for all $k \in [K]$.

As long as $\epsilon < 1$ and $\rho_k > 0$ for all $k \in [K]$, this matrix M can be inverted analytically by applying the Sherman-Morrison formula, which leads to $V = M^{-1}$ with

$$\begin{aligned} V_{kl} &= \left(1 + \frac{\epsilon}{1-\epsilon} \cdot \frac{1/K}{\rho_k}\right) \mathbb{I}[k=l] - \frac{\epsilon}{(1-\epsilon)K} \cdot \frac{(1-\epsilon)\rho_l + \epsilon/K}{\rho_k} \\ &= \frac{\tilde{\rho}_k}{\tilde{\rho}_k - \epsilon/K} \mathbb{I}[k=l] - \frac{\epsilon}{K} \cdot \frac{\tilde{\rho}_l}{\tilde{\rho}_k - \epsilon/K}. \end{aligned} \quad (\text{A38})$$

The relatively simple structure of this matrix V makes the randomized response model particularly interesting to focus on. In particular, we will see below how this model leads to an easier-to-interpret version of the general methodology presented in Section 3. Further, the scalar nature of the unknown parameter $\epsilon \in [0, 1)$ simplifies the task of empirically fitting the model given a limited amount of clean data, thus overcoming the main practical limitation of the general estimation approach described in Section 3.3.

A3.1.2 Adaptive coverage under known label noise

We begin by showing how the adaptive conformal prediction methodology from Section 3.1 simplifies under a randomized response model for the label contamination process. First, note that the plug-in estimate $\hat{\Delta}_k(t)$ of the coverage inflation factor $\Delta_k(t)$ utilized by Algorithm 1, originally defined in (14), simplifies to:

$$\hat{\Delta}_k(t) = \frac{\epsilon(1-\tilde{\rho}_k)}{K\tilde{\rho}_k - \epsilon} \left[\hat{F}_k^k(t) - \frac{\sum_{l \neq k} \tilde{\rho}_l \hat{F}_l^k(t)}{\sum_{l \neq k} \tilde{\rho}_l} \right], \quad (\text{A39})$$

while the finite-sample correction term $\delta(n_k, n_*)$ defined in (18) simplifies to

$$\delta(n_k, n_*) = c(n_k) + \frac{2\epsilon(1-\tilde{\rho}_k)}{(K\tilde{\rho}_k - \epsilon)\sqrt{n_*}} \min \left\{ K\sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\}. \quad (\text{A40})$$

Similarly, under the randomized response model, Assumption 4 becomes

$$\inf_{t \in (0,1)} \Delta_k(t) \geq -\alpha + c(n_k) + \frac{2\epsilon(1-\tilde{\rho}_k)}{(K\tilde{\rho}_k - \epsilon)\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2\sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right),$$

which is always satisfied in the large-sample limit, $n_* \rightarrow \infty$, as long as

$$\epsilon \leq \min \left\{ K\tilde{\rho}_k, \frac{\alpha K\tilde{\rho}_k}{1 + \alpha - \tilde{\rho}_k} \right\}, \quad (\text{A41})$$

regardless of whether the stochastic dominance condition in (11) holds.

Combined with Theorem 3, these expressions tell us that the prediction sets output by Algorithm 1 have asymptotically tight coverage if the noise parameter ϵ is not too large and the regularity conditions of Assumptions 2–3 hold. For example, if $\alpha = 0.1$ and $\tilde{\rho}_k = 1/K$

for all $k \in [K]$, the upper bound in (A41) becomes $\epsilon \leq 0.167$ if $K = 2$, and $\epsilon \leq 0.111$ if $K = 5$.

Further, this model makes it easy to bound from above the term $\varphi_k(n_k, n_*)$ in Theorem 3, with a bound that only increases with K at rate $\sqrt{\log K}$. Intuitively, this means that the asymptotic tightness of the prediction sets output by Algorithm 1 also holds for classification problems with many possible classes.

A3.1.3 Adaptive coverage under a bounded label contamination model

The randomized response model also allows simplifying our general method for constructing adaptive prediction sets under imperfect knowledge of the label contamination process. Since the label frequencies $\tilde{\rho}_k$ are easy to estimate accurately from the available contaminated data for all $k \in [K]$, the expression for the matrix V in (A38) involves only one possibly unknown quantity, the scalar noise parameter $\epsilon \in [0, 1)$. In fact, it is easy to verify that, for any $l, k \in [K]$, the matrix entry V_{kl} can be equivalently rewritten as

$$V_{kl} = \frac{\tilde{\rho}_k}{\tilde{\rho}_k - \epsilon/K} \mathbb{I}[k = l] - \frac{\xi \tilde{\rho}_l}{K \tilde{\rho}_k + \xi(K \tilde{\rho}_k - 1)}, \quad (\text{A42})$$

where $\xi := \epsilon/(1 - \epsilon) > 0$ is a monotone increasing transformation of ϵ .

Therefore, in this special case, implementing the methods from Section 3.2 only requires a valid confidence interval for ξ , in lieu of a joint confidence region for all off-diagonal elements of V . In particular, if a confidence interval $[\hat{\xi}^{\text{low}}, \hat{\xi}^{\text{upp}}]$ at level $1 - \alpha_V$ is available for ξ , then it follows immediately from (A42) and (A37) that a valid joint confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ for the off-diagonal elements of V is given by:

$$\hat{V}_{kl}^{\text{upp}} = -\frac{\hat{\xi}^{\text{low}} \tilde{\rho}_l}{K \tilde{\rho}_k + \hat{\xi}^{\text{upp}}(K \tilde{\rho}_k - 1)}, \quad \hat{V}_{kl}^{\text{low}} = -\frac{\hat{\xi}^{\text{upp}} \tilde{\rho}_l}{K \tilde{\rho}_k + \hat{\xi}^{\text{low}}(K \tilde{\rho}_k - 1)}, \quad \forall l \neq k. \quad (\text{A43})$$

To simplify the following notation as much as possible, but without much loss of generality, let us make the additional assumption that the noisy label frequencies are uniform: $\tilde{\rho}_k = 1/K$ for all $k \in [K]$. Note that (A37) tells us this is always the case if the true labels are uniform. Then, Equation (A43) implies that, for all $l \neq k$,

$$\hat{V}_{kl}^{\text{upp}} = -\frac{\hat{\xi}^{\text{low}}}{K}, \quad \hat{V}_{kl}^{\text{low}} = -\frac{\hat{\xi}^{\text{upp}}}{K}, \quad \hat{\delta}_{kl}^{(V)} := \hat{V}_{kl}^{\text{upp}} - \hat{V}_{kl}^{\text{low}} = \frac{\hat{\xi}^{\text{upp}} - \hat{\xi}^{\text{low}}}{K},$$

and therefore $\hat{\delta}_{k*}^{(V)} := \max_{l \neq k} \hat{\delta}_{kl}^{(V)} = \delta_{\hat{\xi}}/K$, where $\delta_{\hat{\xi}} := \hat{\xi}^{\text{upp}} - \hat{\xi}^{\text{low}}$. Further, $\hat{\zeta}_k$ from (20) is equal to zero because $\hat{V}_{kl}^{\text{upp}} = \hat{V}_{kl'}^{\text{upp}}$ and $V_{kl} = V_{kl'}$ for all $l, l' \neq k$. Thus, in this special case, our method from Section 3.2 can be implemented using the following simplified versions of the plug-in estimator $\hat{\Delta}_k^{\text{ci}}(t)$ defined in (21):

$$\hat{\Delta}_k^{\text{ci}}(t) = \hat{\xi}^{\text{low}} \left(1 - \frac{1}{K}\right) \left[\hat{F}_k^k(t) - \frac{\sum_{l \neq k} \hat{F}_l^k(t)}{K-1} \right] - \delta_{\hat{\xi}} \left(1 - \frac{1}{K}\right) \left| \hat{F}_k^k(t) - \frac{\sum_{l \neq k} \hat{F}_l^k(t)}{K-1} \right|, \quad (\text{A44})$$

and $\delta^{\text{ci}}(n_k, n_*)$ in (22),

$$\delta^{\text{ci}}(n_k, n_*) = c(n_k) + \frac{2\hat{\xi}^{\text{upp}}(1 - 1/K)}{\sqrt{n_*}} \min \left\{ K\sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} + 2\alpha_V \bar{\xi}^{\text{upp}}(1 - 1/K),$$

where $\bar{\xi}^{\text{upp}}$ is a (possibly very conservative) deterministic upper bound on ξ .

These simplified expressions highlight that, in principle, Algorithm 2 only requires a one-sided confidence interval $[0, \hat{\xi}^{\text{upp}}]$ for ξ in order to achieve valid coverage, because one could always evaluate $\hat{\Delta}_k^{\text{ci}}(t)$ in (A44) using $\hat{\xi}^{\text{low}} = 0$ and $\delta_\xi = \hat{\xi}^{\text{upp}}$. However, we already know from Section 3.2 that the prediction sets computed by Algorithm 2 tend to be more informative if the confidence bounds are tighter. In particular, the $1 - \alpha + \varphi_k^{\text{ci}}(n_k, n_*)$ coverage upper bound given by Theorem 5 can be interpreted even more intuitively in this special case, since

$$\varphi_k^{\text{ci}}(n_k, n_*) \leq \frac{1}{n_*} + 2c(n_k) + \frac{2}{n_k} \cdot \frac{1}{1 - \epsilon} \left[1 + \epsilon \cdot \frac{f_{\max}}{f_{\min}} \cdot \sum_{j=1}^{n_k+1} \frac{1}{j} \right] + (1 + 4\bar{\xi}^{\text{upp}}) \alpha_V + 2\mathbb{E}[\delta_\xi] + \frac{4\mathbb{E}[\hat{\xi}^{\text{upp}}]}{\sqrt{n_*}} \min \left\{ K\sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\}.$$

This highlights that the prediction sets output by Algorithm 2 can be (approximately) asymptotically tight, under Assumption 5, if $\bar{\xi}^{\text{upp}}$ is finite, α_V is small, and the expected $\mathbb{E}[\delta_\xi]$ length of the confidence interval for ξ is also small. Further, under this model, Assumption 5 becomes approximately equivalent, in the limit of $n_* \rightarrow \infty$, to

$$\xi + 2\delta_\xi \leq \alpha - 2\alpha_V \bar{\xi}^{\text{upp}} \quad (\text{A45})$$

This means that Assumption 5 is often realistic, similarly to Assumption 4, as long as $\alpha_V \leq \alpha$, the noise parameter ϵ is not too large, and the confidence interval $[\hat{\xi}^{\text{low}}, \hat{\xi}^{\text{upp}}]$ is not too wide. For example, suppose $K = 2$, $\alpha = 0.1$, $\bar{\xi}^{\text{upp}} = 0.125$, and the confidence interval $[\hat{\xi}^{\text{low}}, \hat{\xi}^{\text{upp}}] = [0.075, 0.125]$ has level $\alpha_V = 0.01$ and width $\delta_\xi = 0.05$. Then, the upper bound for the true noise parameter ξ in (25) is $\xi \leq 0.0975$, which corresponds to $\epsilon \leq 0.089$. Additionally, Assumption 5 is also guaranteed to be satisfied in the large-sample limit if the stochastic dominance condition in (11) holds (i.e., $\inf_{t \in (0,1)} \Delta_k(t) \geq 0$) and a sufficiently tight confidence interval $[\hat{\xi}^{\text{low}}, \hat{\xi}^{\text{upp}}]$ for a small enough α_V is available.

A3.1.4 Estimating the parameter ϵ

The scalar noise parameter $\epsilon \in [0, 1)$ in the randomized response model described in Section A3.1.1 can be empirically estimated, from a small amount of clean data, by applying a simplified version of the general procedure described in Section 3.3.

Recall that the matrices \tilde{Q} and Q , defined in (27), are such that, for any $l, k \in [K]$,

$$\tilde{Q}_{kl} := \mathbb{P} \left[\hat{f}(X) = l \mid \tilde{Y} = k, \hat{f} \right], \quad Q_{kl} := \mathbb{P} \left[\hat{f}(X) = l \mid Y = k, \hat{f} \right].$$

We prove in Section A7.4 that, under the randomized response model, the estimating equation $\tilde{Q} = MQ$ stated in (28) implies

$$\epsilon = \frac{\psi - \tilde{\psi}}{\psi - 1/K}, \quad (\text{A46})$$

where $\tilde{\psi}$ and ψ are the probabilities that the classifier guesses correctly the corrupted and true label, respectively, of a new independent data point,

$$\tilde{\psi} := \mathbb{P} \left[\hat{f}(X) = \tilde{Y} \mid \hat{f} \right], \quad \psi := \mathbb{P} \left[\hat{f}(X) = Y \mid \hat{f} \right]. \quad (\text{A47})$$

This result suggests the following simplified method for estimating ϵ . The clean data in \mathcal{D}^0 can be used to compute an intuitive empirical estimate of ψ :

$$\psi \approx \frac{1}{|\mathcal{D}^0|} \sum_{i \in \mathcal{D}^0} \mathbb{I} \left[Y_i = \hat{f}(X_i) \right]. \quad (\text{A48})$$

Similarly, $\tilde{\psi}$ can be estimated using the held-out contaminated data in \mathcal{D}_b^1 :

$$\tilde{\psi} \approx \frac{1}{|\mathcal{D}_b^1|} \sum_{i \in \mathcal{D}_b^1} \mathbb{I} \left[\tilde{Y}_i = \hat{f}(X_i) \right]. \quad (\text{A49})$$

If the contaminated observations are relatively abundant (i.e., $|\mathcal{D}^1| \gg |\mathcal{D}^0|$), one should expect (A49) to provide an empirical estimate of $\tilde{\psi}$ with low variance compared to that of ψ in (A48). Therefore, (A46) tells us that the leading source of uncertainty in ϵ is due to ψ , which depends on the unknown joint distribution of $(\hat{f}(X), Y)$ conditional on the trained classifier \hat{f} . The latter is a multinomial distribution with K^2 categories and event probabilities equal to

$$\lambda_{lk} = \mathbb{P} \left[\hat{f}(X) = k, Y = l \mid \hat{f} \right], \quad \forall l, k \in [K].$$

Then, since $\psi = \sum_{l=1}^K \lambda_{ll}$, it is easy to see that ϵ in (A46) can be written as a function of the multinomial parameter vector $\lambda = (\lambda_{lk})_{l,k \in [K]}$, as well as of other quantities ($\tilde{\psi}$ and $\tilde{\rho}_k$) that are already known with relatively high accuracy:

$$\epsilon = \frac{\sum_{l=1}^K \lambda_{ll} - \tilde{\psi}}{\sum_{l=1}^K \lambda_{ll} - 1/K}. \quad (\text{A50})$$

Thus, a confidence interval for ϵ can be directly obtained by applying standard parametric bootstrap techniques for multinomial parameters [48], similarly to the more general approach presented in Section 3.3. In turn, this immediately translates into a confidence interval $[\hat{\xi}^{\text{low}}, \hat{\xi}^{\text{upp}}]$ for ξ , which is a monotone increasing function of ϵ , at any desired significance level $\alpha_V \in (0, 1)$. Alternatively, one could consider seeking only a point estimate $\hat{\epsilon}$ of ϵ , by replacing the multinomial parameters in (A50) with their standard maximum-likelihood estimates.

A3.2 The two-level randomized response model

A3.2.1 Model description

This section delves into a different special case of the general label contamination model introduced in Section 2.4, extending the streamlined methods presented in Section A3.1 for the randomized response model [11] to a two-level hierarchical setting. The model considered here describes a natural label contamination process involving two clearly defined groups of labels. To illustrate, consider an animal image recognition task: one label group could represent various dog breeds, while the other could denote different cat breeds. The objective is to accurately identify not only the species but also the specific breed for each new image. In this context, one would often expect that a realistic data annotation process may lead to more frequent mislabeling of dog (or cat) breeds rather than mistakenly identifying a dog as a cat or vice-versa.

The two-level label contamination scenario mentioned above can be formalized using an intuitive model with two scalar parameters: $\epsilon \in [0, 1)$ and $\nu \in [0, 1]$. The parameter ϵ influences the likelihood that the observed label \tilde{Y} deviates from the true label Y . This mirrors the role of ϵ in the randomized response model outlined in Section A3.1. Conversely, the parameter ν governs the interaction between the two distinct label groups. In the special case of $\nu = 0$, this model will reduce to the randomized response model [11], indicating an equal probability of mislabeling across breeds or species. On the other hand, in the special case of $\nu = 1$, this model will describe a scenario in which two separate species-specific randomized response models operate independently of one another, and accurate species labeling is always ensured. The most interesting cases will of course be those in between of these two extremes, for values of $\nu \in (0, 1)$. The specifics of this model are presented next.

For simplicity, let us assume that the total number of possible labels, K , is even. We describe the relation between X, Y and \tilde{Y} with

$$\mathbb{P}[\tilde{Y} = k \mid X, Y = l] = T_{kl}, \quad (\text{A51})$$

for all $k, l \in [K]$, where $T \in [0, 1]^{K \times K}$ is a (2×2) block matrix defined as

$$T = \begin{pmatrix} D^{\text{BRR}} + B^{\text{BRR}} & C^{\text{BRR}} \\ C^{\text{BRR}} & D^{\text{BRR}} + B^{\text{BRR}} \end{pmatrix}.$$

Above, the matrix D^{BRR} is diagonal and such that, for any $k, l \in [K/2]$,

$$D_{kl}^{\text{BRR}} = (1 - \epsilon)\mathbb{I}[k = l],$$

while the matrices B^{BRR} and C^{BRR} are constant and such that, for any $k, l \in [K/2]$,

$$B_{kl}^{\text{BRR}} = \frac{\epsilon}{K}(1 + \nu), \quad C_{kl}^{\text{BRR}} = \frac{\epsilon}{K}(1 - \nu).$$

For example, in the special case of $K = 4$, the matrix T would look like

$$T^{(4)} = \begin{pmatrix} 1 - \epsilon + \frac{(1+\nu)\epsilon}{4} & \frac{(1+\nu)\epsilon}{4} & \frac{(1-\nu)\epsilon}{4} & \frac{(1-\nu)\epsilon}{4} \\ \frac{(1+\nu)\epsilon}{4} & 1 - \epsilon + \frac{(1+\nu)\epsilon}{4} & \frac{(1-\nu)\epsilon}{4} & \frac{(1-\nu)\epsilon}{4} \\ \frac{(1-\nu)\epsilon}{4} & \frac{(1-\nu)\epsilon}{4} & 1 - \epsilon + \frac{(1+\nu)\epsilon}{4} & \frac{(1+\nu)\epsilon}{4} \\ \frac{(1-\nu)\epsilon}{4} & \frac{(1-\nu)\epsilon}{4} & \frac{(1+\nu)\epsilon}{4} & 1 - \epsilon + \frac{(1+\nu)\epsilon}{4} \end{pmatrix}.$$

It is easy to verify that this two-level randomized response model generally leads to the following contaminated label frequencies:

$$\tilde{\rho}_k = (1 - \epsilon)\rho_k + \frac{\epsilon}{K} + \nu \frac{\epsilon}{K} \left(2 \sum_{l=1}^{K/2} \rho_l - 1 \right), \quad \forall k \in [K].$$

However, we will assume henceforth that $\rho_k = 1/K$ for all $k \in [K]$, which implies $\tilde{\rho}_k = 1/K$ for all $k \in [K]$. This simplification is not crucial in principle, but it is useful to make the following computations less tedious. In particular, under the assumption of uniform label frequencies, the matrix M in (1) is simply equal to T :

$$M = T = \begin{pmatrix} D^{\text{BRR}} + B^{\text{BRR}} & C^{\text{BRR}} \\ C^{\text{BRR}} & D^{\text{BRR}} + B^{\text{BRR}} \end{pmatrix}.$$

Having a relatively simple expression for M is useful because it simplifies the implementation of our method, which relies directly on the inverse matrix $V = M^{-1}$. The latter can now be obtained analytically by combining the Sherman-Morrison formula with standard techniques for block-matrix inversion. This leads to the following close-formula expressions. For any $k \in [K]$,

$$V_{kk} = \frac{1}{1 - \epsilon} \left(1 - \frac{\epsilon}{K} \right) - \frac{\epsilon\nu}{K(1 - \epsilon)[1 - \epsilon(1 - \nu)]}.$$

For any $l \in \mathcal{B}_k \setminus \{k\}$, where \mathcal{B}_k indicates the block to which label k belongs—that is, $\mathcal{B}_k = \{1, \dots, K/2\}$ if $k \leq K/2$ and $\mathcal{B}_k = \{K/2 + 1, \dots, K\}$ otherwise—the term V_{kl} is

$$V_{kl} = -\frac{\epsilon}{K(1 - \epsilon)} \cdot \left(1 + \frac{\nu}{1 - \epsilon(1 - \nu)} \right).$$

Finally, for any $l \in \mathcal{B}_k^c$, where $\mathcal{B}_k^c := [K] \setminus \mathcal{B}_k$,

$$V_{kl} = -\frac{\epsilon}{K(1 - \epsilon)} \cdot \left(1 - \frac{\nu}{1 - \epsilon(1 - \nu)} \right).$$

For example, in the special case of $K = 4$, the matrix V would look like

$$V^{(4)} = \begin{pmatrix} \frac{\epsilon^2 - \nu\epsilon^2 + 3\nu\epsilon - 5\epsilon + 4}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} & -\frac{\epsilon(1-\epsilon+\nu+\nu\epsilon)}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} & -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} & -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} \\ -\frac{\epsilon(1-\epsilon+\nu+\nu\epsilon)}{4(1-\epsilon)(\nu\epsilon-\epsilon+1)} & \frac{\epsilon^2 - \nu\epsilon^2 + 3\nu\epsilon - 5\epsilon + 4}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} & -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} & -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} \\ -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} & -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} & \frac{\epsilon^2 - \nu\epsilon^2 + 3\nu\epsilon - 5\epsilon + 4}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} & -\frac{\epsilon(1-\epsilon+\nu+\nu\epsilon)}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} \\ -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} & -\frac{(1-\nu)\epsilon}{4(1-\epsilon+\nu\epsilon)} & -\frac{\epsilon(\nu+\nu\epsilon-\epsilon+1)}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} & \frac{\epsilon^2 - \nu\epsilon^2 + 3\nu\epsilon - 5\epsilon + 4}{4(1-\epsilon)(1-\epsilon+\nu\epsilon)} \end{pmatrix}.$$

As detailed below, the relatively tractable structure of the matrix V under this two-level randomized response model leads to an easier-to-interpret version of the general methodology presented in Section 3. Further, the scalar nature of the unknown parameters $\epsilon \in [0, 1]$ $\nu \in [0, 1]$ simplifies the task of empirically fitting the model given a limited amount of clean data, thus overcoming the main practical limitation of the general estimation approach described in Section 3.3.

A3.2.2 Adaptive coverage under a known label contamination model

We begin by showing how the adaptive conformal prediction methodology from Section 3.1 simplifies under a two-level randomized response model for the label contamination process. First, note that the plug-in estimate $\hat{\Delta}_k(t)$ of the coverage inflation factor $\Delta_k(t)$ utilized by Algorithm 1, originally defined in (14), simplifies to:

$$\begin{aligned} \hat{\Delta}_k(t) = & \frac{\epsilon}{1-\epsilon} \left(1 - \frac{1}{K}\right) \left[\hat{F}_k^k(t) - \frac{1}{K-1} \sum_{l \neq k} \hat{F}_l^k(t) \right] \\ & - \frac{\epsilon}{1-\epsilon} \cdot \frac{\nu}{1-\epsilon(1-\nu)} \left[\frac{1}{K} \sum_{l \in \mathcal{B}_k} \hat{F}_l^k(t) - \frac{1}{K} \sum_{l \in \mathcal{B}_k^c} \hat{F}_l^k(t) \right]. \end{aligned} \quad (\text{A52})$$

In the special case of $\nu = 0$, this recovers to same expression obtained under the standard randomized response model in (A39). In the other extreme case, if $\nu = 1$, the expression for $\hat{\Delta}_k(t)$ in (A52) reduces to

$$\hat{\Delta}_k(t) = \frac{\epsilon}{1-\epsilon} \left(1 - \frac{2}{K}\right) \left[\hat{F}_k^k(t) - \frac{1}{K/2-1} \sum_{l \in \mathcal{B}_l \setminus \{k\}} \hat{F}_l^k(t) \right], \quad (\text{A53})$$

consistently with two block-specific randomized response models operating independently of one another.

In the general case, Equation (A52) implies that

$$\sum_{l \neq k} |V_{kl}| = \frac{\epsilon}{1-\epsilon} \left(1 - \frac{1}{K} - \frac{\nu}{K[1-\epsilon(1-\nu)]}\right). \quad (\text{A54})$$

Therefore, under a two-level randomized response model, the finite-sample correction term $\delta(n_k, n_*)$ defined in (18) simplifies to

$$\begin{aligned} \delta(n_k, n_*) = & c(n_k) + \frac{2\epsilon}{(1-\epsilon)\sqrt{n_*}} \left(1 - \frac{1}{K} - \frac{\nu}{K[1-\epsilon(1-\nu)]}\right) \\ & \cdot \min \left\{ K\sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\}. \end{aligned}$$

Similarly, under this two-level randomized response model, Assumption 4 becomes

$$\inf_{t \in (0,1)} \Delta_k(t) \geq -\alpha + c(n_k) + \frac{2\epsilon \left(1 - \frac{1}{K} - \frac{\nu}{K[1-\epsilon(1-\nu)]}\right)}{(1-\epsilon)\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2\sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right).$$

This is always satisfied in the limit of $n_* \rightarrow \infty$, as long as the following inequality holds:

$$\epsilon \leq \frac{\alpha}{1 + \alpha - \frac{1}{K} + \frac{1}{2} \frac{\nu}{1-\epsilon(1-\nu)}}.$$

Instead of solving this quadratic inequality, it suffices here to note that a stricter condition, for any value of $\nu \in [0, 1]$, is

$$\epsilon \leq \frac{\alpha}{2 + \alpha - \frac{1}{K}}. \quad (\text{A55})$$

Combined with Theorem 3, these expressions tell us that the prediction sets output by Algorithm 1 have asymptotically tight coverage if the noise parameter ϵ is not too large and the regularity conditions of Assumptions 2–3 hold, consistently with the simpler special case of the standard randomized response model (Section A3.1). For example, if $\alpha = 0.1$, the upper bound in (A55) becomes $\epsilon \leq 0.0625$ if $K = 2$, and $\epsilon \leq 0.0526$ if $K = 5$. Further, also consistently with Section A3.1, this model makes it easy to bound from above the term $\varphi_k(n_k, n_*)$ in Theorem 3, with a bound that only increases with K at rate $\sqrt{\log K}$.

A3.2.3 Adaptive coverage under a bounded label contamination model

We now turn our attention to the problem in which the parameters ϵ and ν of the two-level randomized response model are not known exactly. Under this model, applying the adaptive conformal prediction methods presented in Section 3.2 only requires a simultaneously valid pair of confidence intervals for ϵ and ν , in lieu of a joint confidence region for all off-diagonal elements of V .

In particular, if two simultaneously valid confidence intervals $[\hat{\xi}^{\text{low}}, \hat{\xi}^{\text{upp}}]$ and $[\hat{\nu}^{\text{low}}, \hat{\nu}^{\text{upp}}]$ at level $1 - \alpha_V$ are available for $\xi := \epsilon/(1-\epsilon)$ and ν , then it follows immediately that a valid joint confidence region $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ for the off-diagonal elements of V is given as follows.

For any $k \in [K]$ and $l \in \mathcal{B}_k \setminus \{k\}$,

$$V_{kl} = -\frac{\xi}{K} \cdot \frac{1 + \nu(1 + 2\xi)}{1 + \nu\xi}.$$

It is easy to verify by taking partial derivatives that this is a monotone decreasing function of ν for any fixed $\xi \geq 0$, and a monotone decreasing function of ξ for any fixed $\nu \in [0, 1]$. Therefore, for any $k \in [K]$ and $l \in \mathcal{B}_k \setminus \{k\}$,

$$\hat{V}_{kl}^{\text{upp}} = -\frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{1 + \hat{\nu}^{\text{low}}(1 + 2\hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}}, \quad \hat{V}_{kl}^{\text{low}} = -\frac{\hat{\xi}^{\text{upp}}}{K} \cdot \frac{1 + \hat{\nu}^{\text{upp}}(1 + 2\hat{\xi}^{\text{upp}})}{1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{upp}}}.$$

Similarly, for any $k \in [K]$ and $l \in \mathcal{B}_k^c$,

$$V_{kl} = -\frac{\xi}{K} \cdot \frac{1 - \nu}{1 + \nu\xi}.$$

It is easy to verify by taking partial derivatives that this is a monotone increasing function of ν for any fixed $\xi \geq 0$, and a monotone decreasing function of ξ for any fixed $\nu \in [0, 1]$. Therefore, for any $k \in [K]$ and $l \in \mathcal{B}_k^c$,

$$\hat{V}_{kl}^{\text{upp}} = -\frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{1 - \hat{\nu}^{\text{upp}}}{1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}}}, \quad \hat{V}_{kl}^{\text{low}} = -\frac{\hat{\xi}^{\text{upp}}}{K} \cdot \frac{1 - \hat{\nu}^{\text{low}}}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{upp}}}.$$

Next, we require an upper confidence bound for the parameter $\hat{\zeta}_k$ defined in (20). Leveraging the confidence bounds for ξ and ν as well as amenable structure of our V matrix, we prove in Section A7.4 that a valid upper bound for $\hat{\zeta}_k$ is given by

$$\begin{aligned} \hat{\zeta}_k^{\text{upp}} = & \frac{2}{K} \cdot \frac{(\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}}) + \left[\hat{\nu}^{\text{upp}} (\hat{\xi}^{\text{upp}})^2 - \hat{\nu}^{\text{low}} (\hat{\xi}^{\text{low}})^2 \right] + \hat{\nu}^{\text{low}} \hat{\nu}^{\text{upp}} \hat{\xi}^{\text{low}} \hat{\xi}^{\text{upp}} (\hat{\xi}^{\text{upp}} - \hat{\xi}^{\text{low}})}{\left(1 + \hat{\nu}^{\text{upp}} \hat{\xi}^{\text{upp}}\right) \left(1 + \hat{\nu}^{\text{low}} \hat{\xi}^{\text{low}}\right)} \\ & + \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{(\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}}) (1 + \hat{\xi}^{\text{low}})}{\left(1 + \hat{\nu}^{\text{low}} \hat{\xi}^{\text{low}}\right) \left(1 + \hat{\nu}^{\text{upp}} \hat{\xi}^{\text{low}}\right)}. \end{aligned} \tag{A56}$$

Note that, in the special case of $\hat{\nu}^{\text{upp}} = 0$ and $\hat{\nu}^{\text{low}} = 0$, which corresponds to a standard randomized response model with $\nu = 0$, we recover that $\hat{\zeta}_k^{\text{upp}} = 0$, consistently with the results presented earlier in Section A3.1.

Leveraging the explicit expressions for \hat{V}^{low} , \hat{V}^{upp} and $\hat{\zeta}_k^{\text{upp}}$ derived above, one can then directly apply the general conformal prediction methodology detailed in Section 3.2. Note that further simplifications of our method's remaining components within this two-tier randomized response model, though possible in principle, are not explicitly shown here due to the increasingly tedious nature of such analytical calculations beyond this point. Instead, we will later investigate the performance of our method as a function of the confidence intervals for the parameters ξ and ν through numerical experiments.

A3.2.4 Estimating the parameters ϵ and ν

Both parameters $\epsilon \in [0, 1)$ and $\nu \in [0, 1)$ in the two-level randomized response model described in Section A3.2.1 can be empirically estimated, from a small amount of clean data, by applying a simplified version of the general procedure described in Section 3.3.

Recall that the matrices \tilde{Q} and Q , defined in (27), are such that, for any $l, k \in [K]$,

$$\tilde{Q}_{kl} := \mathbb{P} \left[\hat{f}(X) = l \mid \tilde{Y} = k, \hat{f} \right], \quad Q_{kl} := \mathbb{P} \left[\hat{f}(X) = l \mid Y = k, \hat{f} \right].$$

Recall also that $\tilde{\psi}$ and ψ are the probabilities that the classifier guesses correctly the corrupted and true label, respectively, of a new independent data point, as defined in (A47). Then, let us define two new quantities, $\tilde{\phi}$ and ϕ . These are the probabilities that the classifier guesses correctly the group to which the corrupted and true label, respectively, of a new independent data point belong. That is,

$$\tilde{\phi} := \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = l, \tilde{Y} = k \mid \hat{f} \right], \quad \phi := \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = l, Y = k \mid \hat{f} \right].$$

We prove in Section A7.4 that, under the two-level randomized response model, the estimating equation $\tilde{Q} = MQ$ stated in (28) implies the following system of two equations:

$$\begin{cases} \tilde{\psi} &= (1 - \epsilon)\psi + \frac{\epsilon}{K} + \frac{\epsilon\nu}{K} (2\phi - 1), \\ \tilde{\phi} &= \phi - \epsilon(1 - \nu) \left(\phi - \frac{1}{2} \right). \end{cases} \quad (\text{A57})$$

It is easy to verify that the solution of this system is:

$$\epsilon = \frac{\frac{K}{2} (\psi - \tilde{\psi}) - (\phi - \tilde{\phi})}{\frac{K}{2} \psi - \phi}, \quad \nu = 1 - \frac{(\phi - \tilde{\phi}) \left(\frac{K}{2} \psi - \phi \right)}{(\phi - \frac{1}{2}) \left[\frac{K}{2} (\psi - \tilde{\psi}) - (\phi - \tilde{\phi}) \right]}. \quad (\text{A58})$$

This result suggests the following simplified method for estimating ϵ and ν .

The clean data in \mathcal{D}^0 can be used to compute an intuitive empirical estimate of ψ using (A48), and an estimate of ϕ as:

$$\phi \approx \frac{1}{|\mathcal{D}^0|} \sum_{i \in \mathcal{D}^0} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \mathbb{I} \left[\hat{f}(X_i) = l, Y_i = k \right]. \quad (\text{A59})$$

Similarly, $\tilde{\psi}$ can be estimated using the held-out contaminated data in \mathcal{D}_b^1 using (A49), and $\tilde{\phi}$ can be estimated using

$$\tilde{\phi} \approx \frac{1}{|\mathcal{D}_b^1|} \sum_{i \in \mathcal{D}_b^1} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \mathbb{I} \left[\hat{f}(X_i) = l, \tilde{Y}_i = k \right]. \quad (\text{A60})$$

If the contaminated observations are relatively abundant (i.e., $|\mathcal{D}^1| \gg |\mathcal{D}^0|$), one should expect (A49) and (A60) to provide empirical estimates of $\tilde{\psi}$ and $\tilde{\phi}$, respectively, with low variance compared to those of ψ and ϕ . Therefore, the leading source of uncertainty in ϵ and ν is due to ψ and ϕ , which depend on the unknown joint distribution of $(\hat{f}(X), Y)$ conditional on the trained classifier \hat{f} . We already know that the latter is a multinomial distribution with K^2 categories and event probabilities equal to

$$\lambda_{lk} = \mathbb{P} \left[\hat{f}(X) = k, Y = l \mid \hat{f} \right], \quad \forall l, k \in [K].$$

Further,

$$\psi = \sum_{l=1}^K \lambda_l, \quad \phi = \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \lambda_{kl}. \quad (\text{A61})$$

Therefore, the parameters ϵ and ν in (A58) can be written as a function of the multinomial parameter vector $\lambda = (\lambda_{lk})_{l,k \in [K]}$, as well as of other quantities ($\tilde{\psi}$ and $\tilde{\phi}$) that are already known with relatively high accuracy.

Thus, confidence intervals for ϵ (or, equivalently, ξ) and ν can be directly obtained by applying standard parametric bootstrap techniques for multinomial parameters [48], similarly to the case of the standard randomized response model discussed in Section A3.1. Alternatively, one could consider seeking only a point estimate $\hat{\epsilon}$ and $\hat{\nu}$, by replacing the multinomial parameters in (A58) with their standard maximum-likelihood estimates.

A4 Comparison to worst-case coverage bounds

A4.1 Theoretical bounds under a general linear contamination model

As previously mentioned in Section 2.4, it is notable that Theorem 1 could also be utilized to derive worst-case coverage bounds for standard conformal prediction sets calibrated by Algorithm A1 using contaminated data. This short digression highlights an interesting link with the sophisticated theoretical results of [20], although it does not directly add to the methodological advancements presented in this paper.

Corollary 2. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n+1]$, and assume also that Assumption 1 holds. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}(X_{n+1})$ indicate the prediction set output by Algorithm A1 applied using the corrupted labels \tilde{Y}_i instead of the clean labels Y_i , for all $i \in [n]$. Then,*

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] \geq 1 - \alpha - \sum_{l \neq k} |V_{kl}|. \quad (\text{A62})$$

Further, if the conformity scores $\hat{s}(X_i, \tilde{Y}_i)$ used by Algorithm A1 are almost-surely distinct,

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] \leq 1 - \alpha + \frac{1}{n_k + 1} + \sum_{l \neq k} |V_{kl}|. \quad (\text{A63})$$

The upper and lower coverage bounds established by Corollary 2 are independent of the conformity scores employed by Algorithm A1, consistently with a worst-case type of analysis. However, a practical drawback of this generality is the presence of a constant term in the gap between the upper (A63) and lower (A62) bounds. Therefore, this theoretical worst-case gap cannot vanish asymptotically as the sample size increases:

$$\text{Worst-Case (WC) Gap} = \text{WC-UB} - \text{WC-LB} = \frac{1}{n_k + 1} + 2 \sum_{l \neq k} |V_{kl}|. \quad (\text{A64})$$

This is why worst-case coverage bounds such as those provided by Corollary 2 are not as practically relevant for our objectives as the novel methodologies developed in Section 3. Indeed, our methods are carefully designed to adaptively learn the unknown distributions of the conformity scores from the existing data. This allows us to obtain more informative prediction sets with asymptotically tight coverage guarantees, under the assumption that the label contamination process is random as opposed to worst-case. For example, Theorems 2 and 3 provide coverage lower and upper bounds equal to $1 - \alpha$ and $1 - \alpha + \varphi_k(n_k, n_*)$, respectively, for the prediction sets output by our Algorithm 1, where $\varphi_k(n_k, n_*)$ is a known sequence converging to zero as $n_k \rightarrow \infty$ and $n_* \rightarrow \infty$.

It is important to highlight that the presence of a constant theoretical gap impacts not only the practical utility of our Corollary 2 but also that of the more nuanced theoretical worst-case bounds previously derived by [20]. This observation sets the stage for a detailed comparison of these worst-case bounds, which we present in the following subsection along with more detailed evidence of their practical limitations in our context. For clarity but without much loss of generality, the comparison will be undertaken under the simpler label contamination model previously outlined in Section A3.1.

A4.2 Theoretical bounds under a randomized-response model

Assume for simplicity that the label contamination process follows the randomized response model defined in Section A3.1.1, with noise parameter $\epsilon \in [0, 1)$, and uniform label frequencies $\rho_k = 1/K = \tilde{\rho}_k$ for all $k \in [K]$. Then, it is not difficult to see that the worst-case coverage lower and upper bounds from Corollary 2 become:

$$\begin{aligned} \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] &\geq 1 - \alpha - \frac{\epsilon}{1 - \epsilon} \left(1 - \frac{1}{K} \right), \\ \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] &\leq 1 - \alpha + \frac{1}{n_k + 1} + \frac{\epsilon}{1 - \epsilon} \left(1 - \frac{1}{K} \right). \end{aligned} \tag{A65}$$

Let us now compare these theoretical bounds with those found by [20]. We refer in particular to Appendix C of [20], which studies the behavior of standard conformal predictions under a *Huber contamination* model. Recall that a ϵ' -Huber contamination model is a standard model for describing situations in which a fraction $\epsilon' \in [0, 1)$ of data points are contaminated by outliers from an unknown distribution. Thus, our randomized response model with noise parameter ϵ intuitively corresponds to a special case of a Huber contamination model with parameter $\epsilon' = \epsilon(1 - 1/K)$, because the latter is the expected proportion of contaminated labels \tilde{Y} that differ from the true Y . Then, Theorems 2 and 3 in [20] provide the following worst-case coverage bounds:

$$\begin{aligned} \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] &\geq 1 - \alpha - \epsilon \left(1 - \frac{1}{K} \right), \\ \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] &\leq 1 - \alpha + \frac{1}{n_k + 1} + \epsilon \left(1 - \frac{1}{K} \right). \end{aligned} \tag{A66}$$

These theoretical bounds are similar to ours but tend to be relatively tighter, especially if the label noise parameter is large, because $\epsilon/(1-\epsilon) > \epsilon$ for all $\epsilon > 0$. Further, Theorem 6 in [20] also provides an alternative multiplicative lower bound for the worst-case coverage that can be even tighter than that in (A66) if the significance level α is small:

$$\mathbb{P}\left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k\right] \geq 1 - \frac{\alpha}{1 - \epsilon\left(1 - \frac{1}{K}\right)} \approx 1 - \alpha - \alpha\epsilon\left(1 - \frac{1}{K}\right), \quad (\text{A67})$$

where the approximation above holds for small values of ϵ .

In summary, our comparison of (A65) with (A66) and (A67) indicates that a byproduct of Theorem 1 is a worst-case coverage analysis of standard conformal prediction methods in scenarios with label contamination. This analysis aligns with the more general findings of [20], though certainly with less elegance and precision. It is however worth repeating that the intended function of Theorem 1 is not to conduct a worst-case theoretical analysis of standard conformal prediction methods. On the contrary, our objective is to develop a new method for adaptively correcting the potential over-coverage or under-coverage of standard conformal predictions, thereby generating more insightful prediction sets. Theorem 1, as shown in this paper, paves the path towards effectively accomplishing this goal.




Figure A1: Visualization of theoretical worst-case bounds for the coverage achieved by standard conformal prediction sets, calibrated for label-conditional coverage, on the CIFAR-10H data with noisy human-assigned labels. The thick black lines correspond to the theoretical bounds, whereas the light gray line denotes the average empirical coverage achieved by the standard conformal prediction approach in our experiments. The dashed line indicates the nominal 90% coverage level.

In contrast, theoretical worst-case analyses such as those of [20] are not useful for our aims because they offer no actionable guidance on how to adjust the standard conformal prediction sets. The fixed and considerable gap between their lower and upper bounds prevents us from understanding how to correct the standard conformal prediction sets in the presence of random label contamination. As an example, Figure A1 plots different types of theoretical worst-case upper and lower bounds for the coverage achieved by standard conformal prediction sets, calibrated for label-conditional coverage, on the CIFAR-10H data set studied in Section 4.4. See also Figure A2 for similar results in the context of calibration for marginal instead of label-conditional coverage (see also Section A5). Overall, these results

demonstrate that the theoretical bounds may be technically valid but do not provide actionable insight, particularly because they do not even tell us whether the standard conformal prediction sets are too wide or too narrow in the presence of label contamination!

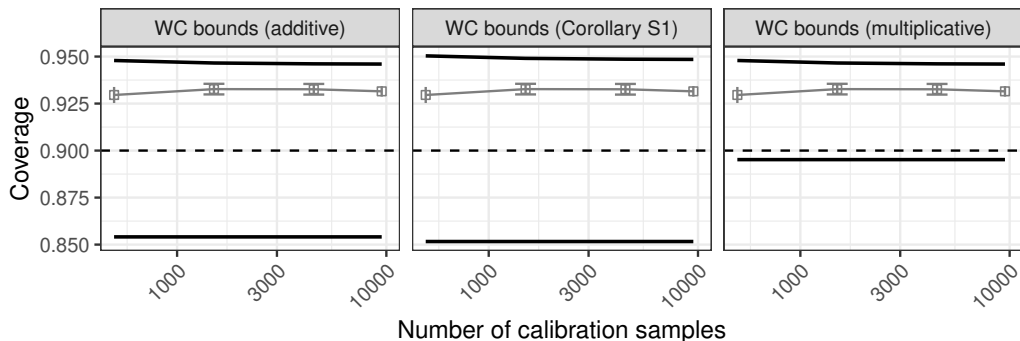


Figure A2: Performance of theoretical worst-case bounds for the coverage achieved by standard conformal prediction sets, calibrated for marginal coverage, on the CIFAR-10H data. Other details are as in Figure A1.

While the current worst-case analysis suggests that the sole principled approach in the context of label contamination may be to enlarge the standard conformal prediction sets, given that the theoretical lower bounds fall below the nominal $1 - \alpha$ level, that is not a satisfactory solution. It lacks adaptability to specific characteristics of the data and classifier, always resulting in prediction sets that are more conservative than the standard ones—which this paper shows already tend to be overly pessimistic. This issue is exemplified in Figures A3 and A4, which compare the performance of our adaptive method to that of a naive alternative approach inspired by the worst-case theoretical analysis. The latter simply consists of applying standard conformal prediction methods at a modified significance level α' defined in such a way that the tightest available (multiplicative) coverage lower bound is exactly equal to the nominal $1 - \alpha$ coverage level.

As expected, the naive theoretical benchmark in a certain sense worsens rather than mitigates the problem considered in this paper. In fact, it leads to even less informative prediction sets compared to the standard approach ignoring the presence of label contamination. By contrast, our adaptive method needs no worst-case analysis; instead, leveraging some knowledge of the random label contamination process, it is able to learn from the available data how to correct the standard conformal prediction sets. Thus, our method produces more informative prediction sets that achieve the desired 90% coverage tightly.




Figure A3: Performance of the proposed adaptive conformal method on CIFAR-10H image data with noisy human-assigned labels, compared to two benchmarks. All methods are calibrated to seek 90% label-conditional coverage. The first benchmark is the standard conformal prediction method ignoring the presence of label contamination. The second benchmark is the standard method applied at a more conservative level, chosen such that the best available worst-case theoretical lower coverage bound matches the desired coverage level $1 - \alpha$. Other details are as in Figure A44.




Figure A4: Performance of the proposed adaptive conformal method on CIFAR-10H image data with noisy human-assigned labels, compared to two benchmarks. All methods are calibrated to seek 90% marginal coverage. Other details are as in Figure 1 and Figure A3.

A5 Extensions of preliminary theoretical results

A5.1 General marginal coverage bounds under label contamination

A result similar to Theorem 1, from Section 2.3, can be reached about the marginal coverage of the prediction sets produced by Algorithm A2. In this case, a useful quantity to define, for any $t \in [0, 1]$, is the *marginal coverage inflation factor*,

$$\Delta(t) := \sum_{k=1}^K \left[\rho_k F_k^k(t) - \tilde{\rho}_k \tilde{F}_k^k(t) \right] = F(t) - \tilde{F}(t), \quad (\text{A68})$$

where $\rho_k := \mathbb{P}[Y = k]$ and $\tilde{\rho}_k = \mathbb{P}[\tilde{Y} = k]$ are the expected proportions of clean and corrupted labels equal to k , respectively. Above, $F(t)$ and $\tilde{F}(t)$ are the marginal cumulative distribution functions of $\hat{s}(X, Y)$ and $\hat{s}(X, \tilde{Y})$, respectively, given \mathcal{D} ; i.e.,

$$F(t) := \mathbb{P}[\hat{s}(X, Y) \leq t \mid \mathcal{D}^{\text{train}}], \quad \tilde{F}(t) := \mathbb{P}[\hat{s}(X, \tilde{Y}) \leq t \mid \mathcal{D}^{\text{train}}].$$

Theorem A1. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n+1]$. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}(X_{n+1})$ indicate the prediction set output by Algorithm A2 applied using the corrupted labels \tilde{Y}_i instead of the clean labels Y_i , for all $i \in [n]$. Then,*

$$\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1})] \geq 1 - \alpha + \mathbb{E}[\Delta(\hat{\tau})]. \quad (\text{A69})$$

Further, if the scores $\hat{s}(X_i, \tilde{Y}_i)$ used by Algorithm A2 are almost-surely distinct,

$$\mathbb{P}[Y_{n+1} \in \hat{C}(X_{n+1})] \leq 1 - \alpha + \frac{1}{n_{\text{cal}} + 1} + \mathbb{E}[\Delta(\hat{\tau})]. \quad (\text{A70})$$

A5.2 Coverage lower bounds under a linear contamination model

A similar stochastic dominance condition as in Corollary 1 also implies that the prediction sets output by Algorithm A2 are conservative in the marginal coverage sense of (5).

Corollary 3. *Consider the same setting of Theorem A1 and assume Assumption 1 holds. Suppose also that the cumulative distribution functions of the scores (6) satisfy*

$$\max_{l \neq k} F_k^l(t) \leq F_k^k(t), \quad (\text{A71})$$

for all $t \in \mathbb{R}$ and $k \in [K]$. Then, $\Delta(\hat{\tau}) \geq 0$ almost-surely, and hence the predictions $\hat{C}(X_{n+1})$ of Algorithm A2 satisfy (5).

Equation (A71) is similar to (11), although the two conditions are not exactly equivalent. Intuitively, (A71) states that the scores $\hat{s}(X, k)$ assigned by the machine learning model tend to be smaller than any other scores $\hat{s}(X, l)$ for $l \neq k$ among data points with true label $Y = k$. In other words, this could be interpreted as saying that the correct label is the most likely point prediction of the machine learning model, for each possible class $k \in [K]$.

A6 Methodology extensions

A6.1 Adaptive prediction sets with marginal coverage

While this paper has so far focused on achieving tight label-conditional coverage (4), the proposed methods can be adapted to alternatively control the weaker notion of marginal coverage (5). Concretely, we present here Algorithm A5, which extends for that purpose Algorithm 1 from Section 3.1. It easy to see that the analogous extensions of the methods described in Sections 3.2–3.3 would also follow similarly.

Algorithm A5 differs from Algorithm 1 in that it calculates a single threshold

$$\hat{\tau}^{\text{marg}} = \begin{cases} S_{(\hat{i}^{\text{marg}})} \text{ where } \hat{i}^{\text{marg}} = \min\{i \in \hat{\mathcal{I}}^{\text{marg}}\}, & \text{if } \hat{\mathcal{I}}^{\text{marg}} \neq \emptyset, \\ 1, & \text{if } \hat{\mathcal{I}}^{\text{marg}} = \emptyset, \end{cases} \quad (\text{A72})$$

where the set $\hat{\mathcal{I}}^{\text{marg}}$ is defined as

$$\hat{\mathcal{I}}^{\text{marg}} := \left\{ i \in [n_{\text{cal}}] : \frac{i}{n_{\text{cal}}} \geq 1 - \alpha - \hat{\Delta}(S_{(i)}) + \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right\}, \quad (\text{A73})$$

for an empirical estimate $\hat{\Delta}(t)$ of the factor $\Delta(t)$ in (A68) given by

$$\hat{\Delta}(t) := \sum_{k=1}^K \left[(\rho_k V_{kk} - \tilde{\rho}_k) \hat{F}_k^k(t) + \rho_k \sum_{l \neq k} V_{kl} \hat{F}_l^k(t) \right], \quad (\text{A74})$$

where, for any $k \in [K]$,

$$\rho_k = \sum_{l=1}^K M_{lk} \tilde{\rho}_l, \quad (\text{A75})$$

and a finite-sample correction term taking the form

$$\delta^{\text{marg}}(n_{\text{cal}}, n_*) := c(n_{\text{cal}}) + \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\}. \quad (\text{A76})$$

It is worth pointing out that, unlike our adaptive methods for label-conditional coverage presented in Section 3, Algorithm A5 requires knowledge of the expected class frequencies for the contaminated labels; i.e., $\tilde{\rho}_k$ for all $k \in [K]$. This additional assumption could be relaxed by replacing $\tilde{\rho}_k$ with an empirical estimate obtained from the available contaminated data, but for simplicity we choose to treat $\tilde{\rho}_k$ as known. Fortunately, this does not prevent our solution from being practical because $\tilde{\rho}_k$ is easy to estimate.

Algorithm A5: Contamination-adaptive classification with marginal coverage

- Input:** Data set $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with corrupted labels $\tilde{Y}_i \in [K]$.
 The inverse V of the matrix M in (10).
 The expected contaminated label frequencies $\tilde{\rho}_k$, for all $k \in [K]$.
 Unlabeled test point with features X_{n+1} .
 Machine learning algorithm \mathcal{A} for training a K -class classifier.
 Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).
 Desired coverage level $1 - \alpha \in (0, 1)$.
- 1 Randomly split $[n]$ into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} .
 - 2 Train the classifier \mathcal{A} on the data in $\mathcal{D}^{\text{train}}$.
 - 3 Compute conformity scores $\hat{s}(X_i, k)$ using (1) for all $i \in \mathcal{D}^{\text{cal}}$ and all $k \in [K]$.
 - 4 Define the empirical CDF \hat{F}_l^k of $\{\hat{s}(X_i, k) : i \in \mathcal{D}^{\text{cal}}\}$ for all $k, l \in [K]$, as in (13).
 - 5 Sort $\{\hat{s}(X_i, Y_i) : i \in \mathcal{D}^{\text{cal}}\}$ into $(S_{(1)}, S_{(2)}, \dots, S_{(n_{\text{cal}})})$, in ascending order.
 - 6 Compute $\hat{\Delta}(S_{(i)})$ for all $i \in [n_{\text{cal}}]$, as in (A74).
 - 7 Compute $\delta^{\text{marg}}(n_{\text{cal}}, n_*)$ using (A76), based on a Monte Carlo estimate of $c(n_{\text{cal}})$.
 - 8 Construct the set $\hat{\mathcal{I}}^{\text{marg}} \subseteq [K]$ as in (A73).
 - 9 Evaluate $\hat{\tau}^{\text{marg}}$ based on $\hat{\mathcal{I}}^{\text{marg}}$ as in (A72).
 - 10 Evaluate $\hat{C}^{\text{marg}}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}^{\text{marg}}; \hat{\pi})$.
- Output:** Conformal prediction set $\hat{C}^{\text{marg}}(X_{n+1})$ for Y_{n+1} .
-

Below, Theorem A2 establishes that the marginal coverage of the prediction sets output by Algorithm A5 is bounded from below by $1 - \alpha$, as long as our method is applied based on the true model matrix M in (10) and the correct contaminated label frequencies $\tilde{\rho}_k$ for all $k \in [K]$.

Theorem A2. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n + 1]$. Assume the label contamination model in Assumption 1 holds. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}^{\text{marg}}(X_{n+1})$ indicate the prediction set output by Algorithm A5 based on the inverse V of the model matrix M in the label contamination model (10) and the true values of the contaminated label frequencies $\tilde{\rho}_k$ for all $k \in [K]$. Then,*

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] \geq 1 - \alpha.$$

Next, we prove that Algorithm A5 is not overly conservative, following an approach similar to that of Theorem 3 for Algorithm 1. This requires a technical lower bound for the marginal coverage inflation factor $\Delta(t)$, whose interpretation is similar to that of Assumption 4.

Assumption 6. *The factor $\Delta(t)$ in (A74) is bounded from below by:*

$$\inf_{t \in (0,1)} \Delta(t) \geq -\alpha + c(n_{\text{cal}}) + \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \frac{1 + 2\sqrt{\log(2K^2) + \log(n_*)}}{\sqrt{n_*}}.$$

Under this setup, an upper bound for the marginal coverage of the conformal prediction sets output by Algorithm A5 is established below.

Theorem A3. *Under the setup of Theorem A2, let $\hat{C}^{\text{marg}}(X_{n+1})$ indicate the prediction set output by Algorithm A5 based on the inverse V of the model matrix M in (10) and the true values of the label frequencies $\rho_k, \tilde{\rho}_k$ for all $k \in [K]$. Suppose also that Assumptions 2, 3, and 6 hold, and that $\tilde{\rho}_k > 0$ for all $k \in [K]$. Then,*

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] \leq 1 - \alpha + \varphi_k^{\text{marg}}(n_k, n_*),$$

where

$$\begin{aligned} \varphi^{\text{marg}}(n_{\text{cal}}, n_*) &= 2\delta^{\text{marg}}(n_{\text{cal}}, n_*) + \frac{1}{n_{\text{cal}}} + \frac{1}{n_*} + \frac{1}{n_{\text{cal}}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \\ &+ \frac{\sum_{j=1}^{n_{\text{cal}}+1} 1/j}{n_{\text{cal}}} \cdot \left[\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} + \frac{f_{\text{max}}}{f_{\text{min}}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right]. \end{aligned} \quad (\text{A77})$$

The interpretation of Theorem A3 is analogous to that of Theorem 3: the marginal coverage of the prediction sets output by Algorithm A5 is guaranteed to be asymptotically tight because $\varphi^{\text{marg}}(n_k, n_*) \rightarrow 0$ as $n_* \rightarrow \infty$.

We conclude this section by noting that the power of Algorithm A5 can be further boosted without losing marginal coverage, as long as a relatively mild “optimistic” assumption on the marginal coverage inflation factor defined in (A68) holds. Concretely, we propose to apply Algorithm A5 with the set $\hat{\mathcal{I}}^{\text{marg}}$ in (A73) replaced by:

$$\hat{\mathcal{I}}^{\text{marg}} := \left\{ i \in [n_{\text{cal}}] : \frac{i}{n_{\text{cal}}} \geq 1 - \alpha - \max \left\{ \hat{\Delta}(S_{(i)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*), -\frac{(1-\alpha)}{n_{\text{cal}}} \right\} \right\}. \quad (\text{A78})$$

This optimistic variation of Algorithm A5 is analogous to the extension of Algorithm 1 presented earlier in Section 3.1.4, and it enjoys a similar coverage guarantee.

Proposition A3. *Under the same setup as Theorem A2, assume also that $\inf_{t \in \mathbf{R}} \Delta(t) \geq \delta^{\text{marg}}(n_{\text{cal}}, n_*) - (1 - \alpha)/n_{\text{cal}}$. If $\hat{C}^{\text{marg}}(X_{n+1})$ is the prediction set output by Algorithm A5 applied with $\hat{\mathcal{I}}$ in (A78) instead of (A73), then $\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] \geq 1 - \alpha$.*

A6.2 Adaptive prediction sets with calibration-conditional coverage

The methods presented in this paper can also be extended to construct prediction sets $\hat{C}^{\text{cc}}(X_{n+1})$ guaranteeing the following notion of calibration-conditional coverage [24]:

$$\mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \geq 1 - \alpha \right] \geq 1 - \gamma, \quad (\text{A79})$$

for any given $\alpha, \gamma \in (0, 1)$. Concretely, we present here Algorithm A6, which extends for the aforementioned purpose Algorithm 1 from Section 1. Analogous extensions of the methods

described in Sections 3.2–3.3 and Section A6.1 would also follow similarly. In a nutshell, Algorithm A6 differs from Algorithm 1 only in that it uses a finite-sample correction factor

$$\delta^{\text{cc}}(n_k, n_*, \gamma) := \sqrt{\frac{\log(1/\gamma_1)}{2n_k}} + 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_2)}{2n_*}}, \quad (\text{A80})$$

where

$$\gamma_1 = \gamma \left(1 - \frac{1}{2} \cdot \frac{\sum_{l \neq k} |V_{kl}|}{\sum_{l=1}^K |V_{kl}|} \right), \quad \gamma_2 = \frac{\gamma}{2} \cdot \frac{\sum_{l \neq k} |V_{kl}|}{\sum_{l=1}^K |V_{kl}|},$$

instead of the factor $\delta(n_k, n_*)$ defined in (18). Note that the second term on the right-hand-side of (A80) simply vanishes in the special case where $\sum_{l=1}^K |V_{kl}| = 0$, which corresponds to the absence of label contamination.

Algorithm A6: Extension of Algorithm 1 with calibration-conditional coverage

Input: Data set $\{(X_i, \tilde{Y}_i)\}_{i=1}^n$ with corrupted labels $\tilde{Y}_i \in [K]$.

The inverse V of the matrix M in (10).

Unlabeled test point with features X_{n+1} .

Machine learning algorithm \mathcal{A} for training a K -class classifier.

Prediction function \mathcal{C} satisfying Definition 1; e.g., (2).

Desired coverage level $1 - \alpha \in (0, 1)$

Desired calibration-conditional significance level $\gamma \in (0, 1)$.

- 1 Randomly split $[n]$ into two disjoint subsets, $\mathcal{D}^{\text{train}}$ and \mathcal{D}^{cal} .
- 2 Train the classifier \mathcal{A} on the data in $\mathcal{D}^{\text{train}}$.
- 3 Compute conformity scores $\hat{s}(X_i, k)$ using (1) for all $i \in \mathcal{D}^{\text{cal}}$ and all $k \in [K]$.
- 4 Define the empirical CDF \hat{F}_l^k of $\{\hat{s}(X_i, k) : i \in \mathcal{D}_l^{\text{cal}}\}$ for all $l \in [K]$, as in (13).
- 5 **for** $k = 1, \dots, K$ **do**
- 6 Define $\mathcal{D}_k^{\text{cal}} = \{i \in \mathcal{D}^{\text{cal}} : \tilde{Y}_i = k\}$ and $n_k = |\mathcal{D}_k^{\text{cal}}|$.
- 7 Sort $\{\hat{s}(X_i, k) : i \in \mathcal{D}_k^{\text{cal}}\}$ into $(S_{(1)}^k, S_{(2)}^k, \dots, S_{(n_k)}^k)$, in ascending order.
- 8 Compute $\hat{F}_l^k(S_{(i)}^k)$ for all $i \in [n_k]$ and $l \in [K]$.
- 9 Compute $\hat{\Delta}_k(S_{(i)}^k)$ for all $i \in [n_k]$, as in (14).
- 10 Compute $\delta^{\text{cc}}(n_k, n_*, \gamma)$ using (A80).
- 11 Construct the set $\hat{\mathcal{I}}_k \subseteq [K]$ as in (15), using $\delta^{\text{cc}}(n_k, n_*, \gamma)$ instead of $\delta(n_k, n_*)$.
- 12 Evaluate $\hat{\tau}_k$ based on $\hat{\mathcal{I}}_k$ as in (16).
- 13 **end**
- 14 Evaluate $\hat{C}^{\text{cc}}(X_{n+1}) = \mathcal{C}(X_{n+1}, \hat{\tau}; \hat{\pi})$, where $\hat{\tau} = (\hat{\tau}_1, \dots, \hat{\tau}_K)$.

Output: Conformal prediction set $\hat{C}^{\text{cc}}(X_{n+1})$ for Y_{n+1} .

Theorem A4. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n + 1]$. Assume the label contamination model in Assumption 1 holds. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}^{\text{cc}}(X_{n+1})$ indicate the prediction set output by Algorithm A6 based on the inverse V of the model matrix M in the label contamination model (10). Then, for any $k \in [K]$,*

$$\mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \geq 1 - \alpha \right] \geq 1 - \gamma.$$

It is interesting to compare the finite-sample correction factor in (A80) to the standard approach for constructing conformal prediction sets with calibration-conditional coverage. In fact, Proposition 2a in Vovk [24] implies that calibration-conditional coverage (A79) can be achieved in the absence of label contamination by simply applying Algorithm A1, in Section A1.2, with the nominal level α replaced by

$$\alpha' = \alpha - \sqrt{\frac{\log(1/\gamma)}{2n_k}}.$$

Intuitively, this correction is similar to the first term on the right-hand-side of (A80), and it becomes equivalent in the case that the matrix V is diagonal, which corresponds to the absence of label contamination.

Following an approach similar to that of Theorem 3, it can be proved that Algorithm A6 is not overly conservative. This requires a technical lower bound for the coverage inflation factor $\Delta_k(t)$, whose interpretation is similar to that of Assumption 4.

Assumption 7. *The coverage inflation factor $\Delta_k(t)$ is bounded from below by:*

$$\inf_{t \in (0,1)} \Delta_k(t) \geq -\alpha + \sqrt{\frac{\log(1/\gamma_1)}{2n_k}} + 4 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_2)}{2n_*}}.$$

Under this setup, a finite-sample upper bound for the marginal coverage of the conformal prediction sets output by Algorithm A6 is established below.

Theorem A5. *Suppose (X_i, Y_i, \tilde{Y}_i) are i.i.d. for all $i \in [n+1]$. Assume the label contamination model in Assumption 1 holds. Fix any prediction function \mathcal{C} satisfying Definition 1, and let $\hat{C}^{\text{cc}}(X_{n+1})$ indicate the prediction set output by Algorithm A6 based on the inverse V of the model matrix M in the label contamination model (10). Suppose also that Assumptions 2, 3, and 7 hold. Then, for any $k \in [K]$,*

$$\mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \leq 1 - \alpha + \varphi_k^{\text{cc}}(n_k, n_*, \gamma) \right] \geq 1 - \gamma.$$

where

$$\begin{aligned} \varphi_k^{\text{cc}}(n_k, n_*, \gamma) &= \delta^{\text{cc}}(n_k, n_*, \gamma) + \frac{1}{n_k} \left(1 + \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{\gamma \bar{v}} \right) + \sqrt{\frac{\log[3/(\gamma \bar{v})]}{2n_k}} \\ &\quad + 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log[3/(\gamma \bar{v})]}{2n_*}} \\ &\quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k} \cdot \left[\log(n_k + 1) + \frac{3}{\gamma \bar{v}} \sum_{j=1}^{n_k+1} \frac{1}{j} \right], \end{aligned}$$

and

$$\bar{v} = \frac{1}{2} \left(1 - \frac{1}{2} \cdot \frac{\sum_{l \neq k} |V_{kl}|}{\sum_{l=1}^K |V_{kl}|} \right).$$

The interpretation of Theorem A5 is similar to that of Theorem 3, because $\varphi_k^{\text{cc}}(n_k, n_*, \gamma) \rightarrow 0$ as $n_* \rightarrow \infty$, for any fixed $\gamma > 0$.

We conclude this section by noting that the power of Algorithm A6 can also be further boosted without losing the calibration-conditional coverage guarantee, as long as a relatively mild “optimistic” assumption on the coverage inflation factor defined in (7) holds. Concretely, we propose to apply Algorithm A6 with the set $\hat{\mathcal{I}}_k$ in (15) replaced by:

$$\hat{\mathcal{I}}_k := \left\{ i \in [n_k] : \frac{i}{n_k} \geq 1 - \alpha - \max \left\{ \hat{\Delta}_k(S_{(i)}) - \delta^{\text{cc}}(n_k, n_*, \gamma), -\sqrt{\frac{\log(1/\gamma)}{2n_k}} \right\} \right\}. \quad (\text{A81})$$

The motivation behind this approach is that, in the absence of label contamination, valid calibration-conditional coverage (A79) can be achieved by simply applying Algorithm A1 with the nominal level α replaced by $\alpha' = \alpha - \sqrt{\log(1/\gamma)/(2n_k)}$ [24]. Therefore, in analogy with the optimistic variation of Algorithm A6 presented earlier in Section 3.1.4, it is now intuitive to propose an optimistic extension of Algorithm A6 that is never more conservative than the standard benchmark for clean data. The following result establishes that this optimistic approach still guarantees the desired calibration-conditional coverage (A79), as long as some relatively mild assumption on the coverage inflation factor holds.

Proposition A4. *Under the setup of Theorem A4, assume also that $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta^{\text{cc}}(n_k, n_*) - \sqrt{\log(1/\gamma)/(2n_k)}$. If $\hat{C}^{\text{cc}}(X_{n+1})$ is the set output by Algorithm A6 applied with $\hat{\mathcal{I}}_k$ in (A81), then, $\mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \geq 1 - \alpha \right] \geq 1 - \gamma$ for any $k \in [K]$.*

A7 Mathematical proofs

A7.1 Preliminaries

Proof of Theorem 1. By definition of the conformity score function in (1), $Y_{n+1} \in \hat{C}(X_{n+1})$ if and only if $\hat{s}(X_{n+1}, Y_{n+1}) \leq \hat{\tau}_{Y_{n+1}}$. Therefore,

$$\begin{aligned}
& \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] \\
&= \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] + \\
&\quad \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] - \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right], \\
&= \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] \\
&\quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k, \mathcal{D} \right] - \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k, \mathcal{D} \right] \right] \\
&= \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] \\
&\quad + \mathbb{E} \left[\mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq \hat{\tau}_k \mid Y_{n+1} = k, \mathcal{D} \right] - \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq \hat{\tau}_k \mid \tilde{Y}_{n+1} = k, \mathcal{D} \right] \right] \\
&= \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] + \mathbb{E} \left[F_k^k(\hat{\tau}_k) - \tilde{F}_k^k(\hat{\tau}_k) \right] \\
&= \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] + \mathbb{E} [\Delta_k(\hat{\tau}_k)].
\end{aligned}$$

Above, the notation $\mathbb{E}[\cdot \mid \mathcal{D}]$ indicates the expected value conditional on the training and calibration data sets. From this, Equations (8) and (9) follow directly because

$$1 - \alpha \leq \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] \leq 1 - \alpha + \frac{1}{n_k + 1},$$

by Proposition A1, which applies here since the data pairs (X_i, \tilde{Y}_i) are i.i.d. random samples. \square

Proof of Proposition 1. By definition of \tilde{P}_k and P_l ,

$$\begin{aligned}
\tilde{P}_k &= \mathbb{P} \left[X \mid \tilde{Y} = k \right] = \sum_{l=1}^K \mathbb{P} \left[X, Y = l \mid \tilde{Y} = k \right] \\
&= \sum_{l=1}^K \frac{\mathbb{P} \left[X, \tilde{Y} = k \mid Y = l \right] \cdot \mathbb{P} [Y = l]}{\mathbb{P} \left[\tilde{Y} = k \right]} \\
&= \sum_{l=1}^K \frac{\mathbb{P} [X \mid Y = l] \cdot \mathbb{P} \left[\tilde{Y} = k \mid Y = l \right] \cdot \mathbb{P} [Y = l]}{\mathbb{P} \left[\tilde{Y} = k \right]} = \sum_{l=1}^K P_l \cdot M_{kl}.
\end{aligned}$$

Above, the fourth equality follows directly from Assumption 1. \square

Proof of Corollary 1. By Theorem 1, it suffices to prove that $\mathbb{E}[\Delta_k(t)] \geq 0$ for all $t \in \mathbb{R}$ and $k \in [K]$. To establish that, note that combining (7) with Proposition 1 gives:

$$\begin{aligned}\Delta_k(t) &= F_k^k(t) - \tilde{F}_k^k(t) = (1 - M_{kk})F_k^k(t) - \sum_{j \neq k} M_{kj}F_j^k(t) \\ &\geq (1 - M_{kk})F_k^k(t) - \left(\sum_{j \neq k} M_{kj} \right) \cdot \max_{j \neq k} F_j^k(t).\end{aligned}$$

Further, by (11),

$$\Delta_k(t) \geq (1 - M_{kk})F_k^k(t) - \left(\sum_{j \neq k} M_{kj} \right) \cdot F_k^k(t) = \left(1 - \sum_{j=1}^K M_{kj} \right) \cdot F_k^k(t) = 0,$$

where the last equality follows from the fact that $\sum_{j=1}^K M_{kj} = 1$ for all $k \in [K]$. This implies that $\mathbb{E}[\Delta_k(\hat{\tau}_k)] \geq 0$, completing the proof. \square

A7.2 Adaptive coverage under a known label contamination model

Proof of Theorem 2. Suppose $Y_{n+1} = k$, for some $k \in [K]$. By definition of the conformity score function in (1), the event $k \notin \hat{C}(X_{n+1})$ occurs if and only if $\hat{s}(X_{n+1}, k) > \hat{\tau}_k$. We will assume without loss of generality that $\hat{\mathcal{I}}_k \neq \emptyset$ and $\hat{i}_k = \min\{i \in \hat{\mathcal{I}}_k\}$; otherwise, $\hat{\tau}_k = 1$ and the result trivially holds. As in the proof of Theorem 1, the probability of miscoverage conditional on $Y_{n+1} = k$ and on the labeled data in \mathcal{D} can be decomposed as:

$$\begin{aligned}& \mathbb{P}\left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D}\right] \\ &= \mathbb{P}\left[\tilde{Y} \notin C(X, S_{(\hat{i}_k)}^k) \mid \tilde{Y} = k, \mathcal{D}\right] \\ &\quad + \mathbb{P}\left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D}\right] - \mathbb{P}\left[\tilde{Y} \notin C(X, S_{(\hat{i}_k)}^k) \mid \tilde{Y} = k, \mathcal{D}\right] \\ &= 1 - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \\ &\quad + \mathbb{P}\left[\hat{s}(X_{n+1}, k) > S_{(\hat{i}_k)}^k \mid Y = k, \mathcal{D}\right] - \mathbb{P}\left[\hat{s}(X_{n+1}, k) > S_{(\hat{i}_k)}^k \mid \tilde{Y} = k, \mathcal{D}\right] \\ &= 1 - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) + \left[1 - F_k^k(S_{(\hat{i}_k)}^k)\right] - \left[1 - \tilde{F}_k^k(S_{(\hat{i}_k)}^k)\right] \\ &= 1 - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k)}^k) \\ &= 1 - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k)}^k) \\ &= \left[1 - \frac{\hat{i}_k}{n_k} - \Delta_k(S_{(\hat{i}_k)}^k)\right] + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \\ &= \left[1 - \frac{\hat{i}_k}{n_k} - \hat{\Delta}_k(S_{(\hat{i}_k)}^k) + \delta(n_k, n_*)\right] - \delta(n_k, n_*) + \\ &\quad + \hat{\Delta}_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k)}^k) + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k)\end{aligned}$$

$$\begin{aligned} &\leq \sup_{i \in \hat{\mathcal{I}}_k} \left\{ 1 - \frac{i}{n_k} - \hat{\Delta}_k(S_{(i)}^k) + \delta(n_k, n_*) \right\} - \delta(n_k, n_*) \\ &\quad + \sup_{t \in \mathbb{R}} [\hat{\Delta}_k(t) - \Delta_k(t)] + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k). \end{aligned}$$

By definition of $\hat{\mathcal{I}}_k$, for all $i \in \hat{\mathcal{I}}_k$,

$$1 - \frac{i}{n_k} - \hat{\Delta}_k(S_{(i)}^k) + \delta(n_k, n_*) \leq \alpha,$$

which implies a.s.

$$\sup_{i \in \hat{\mathcal{I}}_k} \left\{ 1 - \frac{i}{n_k} - \hat{\Delta}_k(S_{(i)}^k) + \delta(n_k, n_*) \right\} \leq \alpha.$$

Therefore,

$$\begin{aligned} &\mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k \right] \\ &\leq \alpha + \mathbb{E} \left[\sup_{t \in \mathbb{R}} [\hat{\Delta}_k(t) - \Delta_k(t)] \right] + \mathbb{E} \left[\hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \right] - \delta(n_k, n_*). \end{aligned}$$

The second term on the right-hand-side above can be bound using the DKW inequality, as made precise by Lemma A1. This leads to:

$$\begin{aligned} &\mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k \right] \\ &\leq \alpha + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\ &\quad + \mathbb{E} \left[\hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \right] - \delta(n_k, n_*). \end{aligned}$$

In order to bound the last expected value, let U_1, \dots, U_{n_k} be i.i.d. uniform random variables on $[0, 1]$, and denote their order statistics as $U_{(1)}, \dots, U_{(n_k)}$. This implies that

$$\hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \stackrel{d}{=} \frac{\hat{i}_k}{n_k} - U_{(\hat{i}_k)}.$$

Therefore, it follows from (17) that

$$\mathbb{E} \left[\hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \right] \leq \mathbb{E} \left[\sup_{i \in [n_k]} \left\{ \frac{i}{n_k} - U_{(i)} \right\} \right] = c(n_k). \quad (\text{A82})$$

We can thus conclude that:

$$\begin{aligned} &\mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k \right] \\ &\leq \alpha + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} + c(n_k) - \delta(n_k, n_*) \\ &= \alpha. \end{aligned}$$

□

Lemma A1. *Under the assumptions of Theorem 2, for any $k \in [K]$,*

$$\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| \right] \leq \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\}.$$

Proof of Lemma A1. Note that, for any $k \in [K]$ and $t \in \mathbb{R}$,

$$\begin{aligned} \hat{\Delta}_k(t) - \Delta_k(t) &= (V_{kk} - 1) \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] + \sum_{l \neq k} V_{kl} \left[\hat{F}_l^k(t) - \tilde{F}_l^k(t) \right] \\ &= \sum_{l \neq k} V_{kl} \left\{ \left[\hat{F}_l^k(t) - \tilde{F}_l^k(t) \right] - \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] \right\} \\ &\leq 2 \left(\sum_{l \neq k} |V_{kl}| \right) \max_{l \in [K]} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)|, \end{aligned}$$

where in the second-to-last line we used the fact that V has row sums equal to 1 because it is the inverse of M , which has row sums equal to 1. Therefore,

$$\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| \right] \leq 2 \sum_{l \neq k} |V_{kl}| \cdot \mathbb{E} \left[\max_{l \in [K]} \sup_{t \in \mathbb{R}} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right].$$

Define $l_{\min} := \arg \min_{l \in [K]} n_j$ and $n_* := \min_{l \in [K]} n_l$. Then, combining the DKW inequality with a union bound gives that, for any $\eta > 0$,

$$\begin{aligned} \mathbb{P} \left[\max_{l \in [K]} \sup_{t \in \mathbb{R}} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| > \eta \right] &\leq K \cdot \mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{F}_{l_{\min}}^k(t) - \tilde{F}_{l_{\min}}^k(t)| > \eta \right] \\ &\leq 2K e^{-2n_* \eta^2}. \end{aligned}$$

This implies that

$$\begin{aligned} \mathbb{E} \left[\max_{l \in [K]} \sup_{t \in \mathbb{R}} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right] &\leq \int_0^\infty \mathbb{P} \left[\max_{l \in [K]} \sup_{t \in \mathbb{R}} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| > \eta \right] d\eta \\ &\leq 2K \int_0^\infty e^{-2n_* \eta^2} d\eta = K \sqrt{\frac{\pi}{2n_*}}. \end{aligned}$$

Similarly, the DKW inequality also implies that, for any $\eta > 0$,

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| > 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\eta)}{2n_*}} \right] \leq \eta;$$

see Lemma A2. Therefore, setting $\eta = 1/n_*$, we obtain

$$\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| \right] \leq 2 \sum_{l \neq k} |V_{kl}| \left[\sqrt{\frac{\log(2K) + \log(n_*)}{2n_*}} + \frac{1}{n_*} \right].$$

□

Lemma A2. *Under the assumptions of Theorem 2, for any $k \in [K]$ and $\eta > 0$,*

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| > 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\eta)}{2n_*}} \right] \leq \eta.$$

Proof of Lemma A2. It follows from the definitions of $\hat{\Delta}_k(t)$ and $\Delta_k(t)$, and from the DKW inequality, that, for any $\eta > 0$,

$$\begin{aligned} & \mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| > 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\eta)}{2n_*}} \right] \\ & \leq \mathbb{P} \left[\max_{j \in [K]} |\hat{F}_j^k(t) - \tilde{F}_j^k(t)| > \sqrt{\frac{\log(2K) + \log(1/\eta)}{2n_*}} \right] \\ & \leq K \cdot \mathbb{P} \left[|\hat{F}_{l_{\min}}^k(t) - \tilde{F}_{l_{\min}}^k(t)| > \sqrt{\frac{\log(2K) + \log(1/\eta)}{2n_*}} \right] \\ & \leq 2K \cdot \exp \left[-2n_* \frac{\log(2K) + \log(1/\eta)}{2n_*} \right] \\ & = 2K \cdot \exp[-\log(2K) - \log(1/\eta)] = \eta. \end{aligned}$$

□

Proof of Theorem 3. Define the events $\mathcal{A}_1 = \{\hat{\mathcal{I}}_k = \emptyset\}$ and $\mathcal{A}_2 = \{\hat{i}_k = 1\}$. Then,

$$\begin{aligned} & \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k \right] \\ & \leq \mathbb{P}[\mathcal{A}_1] + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \right] \\ & \quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \right]. \end{aligned} \tag{A83}$$

We will now separately bound the three terms on the right-hand-side of (A83). The following notation will be useful for this purpose. For all $i \in [n_k]$, let $U_i \sim \text{Uniform}(0, 1)$ be independent and identically distributed uniform random variables, and denote their order statistics as $U_{(1)} < U_{(2)} < \dots < U_{(n_k)}$.

- The probability of the event \mathcal{A}_1 can be bound from above as:

$$\mathbb{P} \left[\hat{\mathcal{I}}_k = \emptyset \right] \leq \frac{1}{n_*}. \tag{A84}$$

To simplify the notation in the proof of (A84), define

$$d_k := c(n_k) + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2 \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right).$$

Then, under Assumption 4,

$$\begin{aligned}
& 1 - \left(\alpha + \widehat{\Delta}_k(S_{(i)}^k) - \delta(n_k, n_*) \right) \\
&= 1 - \alpha - \Delta_k(S_{(i)}^k) + \delta(n_k, n_*) + \Delta_k(S_{(i)}^k) - \widehat{\Delta}_k(S_{(i)}^k) \\
&\leq 1 - d_k + \delta(n_k, n_*) + \sup_{t \in [0,1]} |\Delta_k(t) - \widehat{\Delta}_k(t)|.
\end{aligned}$$

Further, it follows from Lemma A2 and (18) that, with probability at least $1 - 1/n_*$,

$$\begin{aligned}
& 1 - \left(\alpha + \widehat{\Delta}_k(S_{(i)}^k) - \delta(n_k, n_*) \right) \\
&\leq 1 - d_k + \delta(n_k, n_*) + 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(n_*)}{2n_*}} \\
&\leq 1 - d_k + c(n_k) + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2 \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right) = 1.
\end{aligned}$$

This implies that the set $\widehat{\mathcal{I}}_k$ defined in (15) is non-empty with probability at least $1 - 1/n_*$, because $n_k \in \widehat{\mathcal{I}}_k$ if $1 - [\alpha + \widehat{\Delta}_k(S_{(i)}^k) - \delta(n_k, n_*)] \leq 1$.

- Under \mathcal{A}_2 , the second term on the right-hand-side of (A83) can be written as:

$$\mathbb{P} \left[Y_{n+1} \in \widehat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] = \mathbb{P} \left[\widehat{s}(X_{n+1}, k) \leq S_{(1)}^k \mid Y = k, \mathcal{D} \right] = F_k^k(S_{(1)}^k).$$

Therefore,

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \widehat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \right] \\
&\leq \mathbb{E} \left[F_k^k(S_{(1)}^k) \right] \\
&= V_{kk} \mathbb{E} \left[\widetilde{F}_k^k(S_{(1)}^k) \right] + \sum_{l \neq k} V_{kl} \mathbb{E} \left[\widetilde{F}_l^k(S_{(1)}^k) \right] \\
&= V_{kk} \mathbb{E} \left[U_{(1)} \right] + \sum_{l \neq k} V_{kl} \mathbb{E} \left[\widetilde{F}_l^k(S_{(1)}^k) \right] \\
&= \frac{V_{kk}}{n_k + 1} + \sum_{l \neq k} V_{kl} \mathbb{E} \left[\widetilde{F}_l^k(S_{(1)}^k) \right].
\end{aligned}$$

Next, using Assumption 3, we can write:

$$\begin{aligned}
\mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \widehat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \right] &\leq \frac{V_{kk}}{n_k + 1} + \mathbb{E} \left[\widetilde{F}_k^k(S_{(1)}^k) \right] \sum_{l \neq k} |V_{kl}| \\
&= \frac{V_{kk}}{n_k + 1} + \mathbb{E} \left[U_{(1)} \sum_{l \neq k} |V_{kl}| \right] \quad (\text{A85}) \\
&= \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k + 1}.
\end{aligned}$$

- Under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$, by definition of $\hat{\mathcal{I}}_k$, for any $i \leq \hat{i}_k - 1$,

$$\frac{i}{n_k} < 1 - \left(\alpha + \hat{\Delta}_k(S_{(i)}^k) - \delta(n_k, n_*) \right).$$

Therefore, choosing $i = \hat{i}_k - 1$, we get:

$$\frac{\hat{i}_k}{n_k} < 1 - \left(\alpha + \hat{\Delta}_k(S_{(\hat{i}_k-1)}^k) - \delta(n_k, n_*) \right) + \frac{1}{n_k}.$$

As in the proof of Theorem 2, the probability of coverage conditional on $Y_{n+1} = k$ and on the labeled data in \mathcal{D} can be decomposed as:

$$\begin{aligned} & \mathbb{P} \left[Y \in C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \\ &= \mathbb{P} \left[\tilde{Y} \in C(X, S_{(\hat{i}_k)}^k) \mid \tilde{Y} = k, \mathcal{D} \right] \\ & \quad + \mathbb{P} \left[Y \in C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] - \mathbb{P} \left[\tilde{Y} \in C(X, S_{(\hat{i}_k)}^k) \mid \tilde{Y} = k, \mathcal{D} \right] \\ &= \tilde{F}_k^k(S_{(\hat{i}_k)}^k) + \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq S_{(\hat{i}_k)}^k \mid Y = k, \mathcal{D} \right] - \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq S_{(\hat{i}_k)}^k \mid \tilde{Y} = k, \mathcal{D} \right] \\ &= \tilde{F}_k^k(S_{(\hat{i}_k)}^k) + F_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) = \tilde{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) \\ &= \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) \\ &= \frac{\hat{i}_k}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) \\ &< 1 - \left(\alpha + \hat{\Delta}_k(S_{(\hat{i}_k-1)}^k) - \delta(n_k, n_*) \right) + \frac{1}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) \\ &= 1 - \alpha + \delta(n_k, n_*) + \frac{1}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) - \hat{\Delta}_k(S_{(\hat{i}_k-1)}^k). \end{aligned}$$

To bound the last term above, note that

$$\begin{aligned} \Delta_k(S_{(\hat{i}_k)}^k) - \hat{\Delta}_k(S_{(\hat{i}_k-1)}^k) &= (\Delta_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k-1)}^k)) + (\Delta_k(S_{(\hat{i}_k-1)}^k) - \hat{\Delta}_k(S_{(\hat{i}_k-1)}^k)) \\ &\leq \sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k(t)| + (\Delta_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k-1)}^k)), \end{aligned}$$

where the expected value of the first term above can be bounded using Lemma A1, and the second term is given by

$$\begin{aligned} & \Delta_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k-1)}^k) \\ &= (V_{kk} - 1) \left[\tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k-1)}^k) \right] + \sum_{l \neq k} V_{kl} \left[\tilde{F}_l^k(S_{(\hat{i}_k)}^k) - \tilde{F}_l^k(S_{(\hat{i}_k-1)}^k) \right] \\ &= \sum_{l \neq k} V_{kl} \left[\left(\tilde{F}_l^k(S_{(\hat{i}_k)}^k) - \tilde{F}_l^k(S_{(\hat{i}_k-1)}^k) \right) - \left(\tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k-1)}^k) \right) \right] \\ &\leq 2 \sum_{l \neq k} |V_{kl}| \cdot \max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right|. \end{aligned}$$

Combining the above, we find that, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned} & \mathbb{P} \left[Y \in C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \\ & \leq 1 - \alpha + \delta(n_k, n_*) + \frac{1}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k(t)| \\ & \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right|. \end{aligned}$$

It remains to bound the expectation of the last term. By Assumption 2,

$$\begin{aligned} \max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right| & \leq f_{\max} \max_{2 \leq i \leq n_k} |S_{(i)}^k - S_{(i-1)}^k| \\ & \leq \frac{f_{\max}}{f_{\min}} \max_{2 \leq i \leq n_k} |\tilde{F}_k^k(S_{(i)}^k) - \tilde{F}_k^k(S_{(i-1)}^k)| \\ & \stackrel{d}{=} \frac{f_{\max}}{f_{\min}} \max_{2 \leq i \leq n_k} (U_{(i)} - U_{(i-1)}) \\ & \leq \frac{f_{\max}}{f_{\min}} \max_{1 \leq i \leq n_k+1} D_i, \end{aligned}$$

where $D_1 = U_{(1)}$, $D_i = U_{(i)} - U_{(i-1)}$ for $i = 2, \dots, n_k$, and $D_{n_k+1} = 1 - U_{(n_k)}$. By a standard result on maximum uniform spacing,

$$\mathbb{E} \left[\max_{1 \leq i \leq n_k+1} D_i \right] = \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j}.$$

Therefore, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned} & \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \right] \\ & \leq 1 - \alpha + \delta(n_k, n_*) + \frac{1}{n_k} + \mathbb{E} \left[\tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) \right] + \mathbb{E} \left[\sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k(t)| \right] \\ & \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \mathbb{E} \left[\max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right| \right]. \end{aligned}$$

Thus, applying Lemma A1 and the bound in (A82), it follows that, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \right] \\
& \leq 1 - \alpha + \delta(n_k, n_*) + \frac{1}{n_k} + \mathbb{E} \left[\max_{i \in [n_k]} \left(U_i - \frac{i}{n_k} \right) \right] + \mathbb{E} \left[\sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k(t)| \right] \\
& \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& = 1 - \alpha + \delta(n_k, n_*) + \frac{1}{n_k} + c(n_k) + \mathbb{E} \left[\sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k(t)| \right] \\
& \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& \leq 1 - \alpha + \delta(n_k, n_*) + \frac{1}{n_k} + c(n_k) \\
& \quad + \frac{2 \sum_{l \neq k} |V_{kl}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
& \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& = 1 - \alpha + 2\delta(n_k, n_*) + \frac{1}{n_k} + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j}.
\end{aligned} \tag{A86}$$

Finally, combining (A83) with (A84), (A85), and (A86) leads to the desired result:

$$\begin{aligned}
& \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k \right] \\
& \leq \mathbb{P}[\mathcal{A}_1] + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \right] \\
& \quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \right] \\
& \leq 1 - \alpha + \frac{1}{n_*} + \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k + 1} \\
& \quad + 2\delta(n_k, n_*) + \frac{1}{n_k} + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& \leq 1 - \alpha + \frac{1}{n_*} + 2\delta(n_k, n_*) + \frac{1}{n_k} \left[1 + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \sum_{j=1}^{n_k+1} \frac{1}{j} \right] \\
& \quad + \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k + 1}.
\end{aligned}$$

□

Proof of Proposition 2. Suppose $Y_{n+1} = k$, for some $k \in [K]$. Proceeding exactly as in the proof of Theorem 2, we obtain:

$$\begin{aligned}
& \mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \\
&= 1 - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k)}^k) \\
&= 1 - \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \max\{\hat{\Delta}_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*), -(1 - \alpha)/n_k\} \\
&\quad + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \\
&\quad + \max\{\hat{\Delta}_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*), -(1 - \alpha)/n_k\} - (\Delta_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*)) - \delta(n_k, n_*) \\
&= \left[1 - \frac{\hat{i}_k}{n_k} - \max\{\hat{\Delta}_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*), -(1 - \alpha)/n_k\} \right] \\
&\quad + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) \\
&\quad + \max\{\hat{\Delta}_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*) + (1 - \alpha)/n_k, 0\} \\
&\quad - \left(\Delta_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*) + (1 - \alpha)/n_k \right) - \delta(n_k, n_*) \\
&\leq \left[1 - \frac{\hat{i}_k}{n_k} - \max\{\hat{\Delta}_k(S_{(\hat{i}_k)}^k) - \delta(n_k, n_*), -(1 - \alpha)/n_k\} \right] \\
&\quad + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) + \sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| - \delta(n_k, n_*),
\end{aligned}$$

using the fact that $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta(n_k, n_*) - (1 - \alpha)/n_k$ implies

$$\begin{aligned}
& \left| \max\{\hat{\Delta}_k(t) - \delta(n_k, n_*) + (1 - \alpha)/n_k, 0\} - (\Delta_k(t) - \delta(n_k, n_*) + (1 - \alpha)/n_k) \right| \\
&\leq |\hat{\Delta}_k(t) - \Delta_k(t)|,
\end{aligned}$$

for all $t \in \mathbb{R}$. The proof is then completed by proceeding as in the proof of Theorem 2. \square

A7.3 Adaptive coverage under a bounded label contamination model

Proof of Theorem 4. Suppose $Y_{n+1} = k$, for some $k \in [K]$. By definition of the conformity score function in (1), the event $k \notin \hat{C}^{\text{ci}}(X_{n+1})$ occurs if and only if $\hat{s}(X_{n+1}, k) > \hat{\tau}_k$. We will assume without loss of generality that $\hat{\mathcal{I}}_k^{\text{ci}} \neq \emptyset$ and $\hat{i}_k = \min\{i \in \hat{\mathcal{I}}_k^{\text{ci}}\}$; otherwise, $\hat{\tau}_k = 1$ and the result trivially holds. By proceeding exactly as in the proof of Theorem 2, we have

$$\begin{aligned}
& \mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\leq \alpha + \mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] + c(n_k) - \delta^{\text{ci}}(n_k, n_*).
\end{aligned}$$

The second term on the right-hand-side above can be bound using the DKW inequality, as made precise by Lemma A3. This leads to:

$$\mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right]$$

$$\begin{aligned}
&\leq \alpha + \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
&\quad + 2 \cdot \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + c(n_k) - \delta^{\text{ci}}(n_k, n_*) \\
&= \alpha + 2 \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| - 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \\
&\leq \alpha + 2 \left\{ \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] - \alpha_V \right\} \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}|.
\end{aligned}$$

Finally, taking an expectation with respect to the data used to estimate the confidence region for V cancels the last term on the right-hand-side above, leading to:

$$\mathbb{P} \left[Y \notin C(X, S_{(i_k)}^k) \mid Y = k \right] \leq \alpha.$$

□

Lemma A3. *Under the assumptions of Theorem 4, for any $k \in [K]$,*

$$\begin{aligned}
&\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\leq \mathbb{E} \left[\sup_{t \in \mathbb{R}} \max \left\{ \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t), 0 \right\} \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\leq \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
&\quad + 2 \cdot \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}|.
\end{aligned}$$

Further,

$$\begin{aligned}
&\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\leq \frac{2 \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
&\quad + 2 \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right).
\end{aligned}$$

Proof of Lemma A3. To simplify the notation in the following, define

$$\hat{\psi}_k := \min_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right).$$

Note that $\hat{\Delta}_k^{\text{ci}}(t)$ can be equivalently written as:

$$\begin{aligned} \hat{\Delta}_k^{\text{ci}}(t) &:= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| \\ &\quad - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right|. \end{aligned}$$

Therefore, for any $k \in [K]$ and $t \in \mathbb{R}$,

$$\begin{aligned} \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) &= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) - \sum_{l \neq k} V_{kl} \left(\tilde{F}_l^k(t) - \tilde{F}_k^k(t) \right) \\ &\quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \tilde{F}_l^k(t) + \tilde{F}_k^k(t) - \hat{F}_k^k(t) \right) \\ &\quad + \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) \left(\tilde{F}_l^k(t) - \hat{F}_l^k(t) + \hat{F}_k^k(t) - \tilde{F}_k^k(t) \right) \\ &\quad + \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \\ &\quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \tilde{F}_l^k(t) + \tilde{F}_k^k(t) - \hat{F}_k^k(t) \right) \\ &\quad + \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) \left(\tilde{F}_l^k(t) - \hat{F}_l^k(t) + \hat{F}_k^k(t) - \tilde{F}_k^k(t) \right) \\ &\quad + \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} - \hat{\psi}_k \right) \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) + \hat{\psi}_k \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \\ &\quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right|. \end{aligned}$$

Next, note that,

$$\begin{aligned} \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} - \hat{\psi}_k \right) \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) &\leq \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} - \hat{\psi}_k \right) \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &\leq \left[\max_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) - \hat{\psi}_k \right] \cdot \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &= \hat{\zeta}_k \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right|, \end{aligned}$$

because $\hat{V}_{kl}^{\text{upp}} - V_{kl} \geq \hat{\psi}_k$ for all $l \neq k$. Therefore,

$$\begin{aligned} \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) &\leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + |\hat{V}_{kl}^{\text{upp}} - V_{kl}| \right) \\ &\quad + \hat{\zeta}_k \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &\quad + \hat{\psi}_k \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| \\ &\quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right|. \end{aligned}$$

In the event that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$, which occurs with probability $1 - \alpha_V$,

$$0 \leq \hat{V}_{kl}^{\text{upp}} - V_{kl} \leq \hat{\delta}_{kl}^{(V)}, \quad 0 \leq \hat{\psi}_k \leq \hat{\delta}_{kl}^{(V)},$$

which implies $\hat{V}_{kl}^{\text{upp}} - V_{kl} - \hat{\psi}_k \leq \hat{\delta}_{kl}^{(V)} \leq \hat{\delta}_{k*}^{(V)}$. Further, in this event, $\hat{\zeta}_k \leq \hat{\zeta}_k^{\text{upp}}$. Therefore, if $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$,

$$\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) \leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right).$$

Imagine now that it is not the case that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$. Then,

$$\begin{aligned} \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) &= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) - \sum_{l \neq k} V_{kl} \left(\tilde{F}_l^k(t) - \tilde{F}_k^k(t) \right) \\ &\quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \tilde{F}_l^k(t) + \tilde{F}_k^k(t) - \hat{F}_k^k(t) \right) \\ &\quad + \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) \left(\tilde{F}_l^k(t) - \tilde{F}_k^k(t) \right) \\ &\quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right| \\ &\leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| + \sum_{l \neq k} \left| \hat{V}_{kl}^{\text{upp}} - V_{kl} \right|. \end{aligned}$$

In conclusion,

$$\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) \leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)$$

$$+ \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} \left| \hat{V}_{kl}^{\text{upp}} - V_{kl} \right|.$$

From here, proceeding as in the proof of Lemma A1, one easily arrives at:

$$\begin{aligned} & \mathbb{E} \left[\sup_{t \in \mathbb{R}} \max \left\{ \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t), 0 \right\} \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\ & \leq \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\ & \quad + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} \left| \hat{V}_{kl}^{\text{upp}} - V_{kl} \right| \\ & \leq \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\ & \quad + 2 \cdot \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}|. \end{aligned}$$

This completes the first part of the proof.

To prove the second part, recall that

$$\begin{aligned} \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) &= \sum_{l \neq k} \hat{V}_{kl}^{\text{upp}} \left(\hat{F}_l^k(t) - \tilde{F}_l^k(t) + \tilde{F}_k^k(t) - \hat{F}_k^k(t) \right) \\ & \quad + \sum_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) \left(\hat{F}_l^k(t) - \tilde{F}_k^k(t) \right) \\ & \quad - \hat{\delta}_{k*}^{(V)} \left| \sum_{l \neq k} \left(\hat{F}_l^k(t) - \hat{F}_k^k(t) \right) \right| - |\hat{\zeta}_k^{\text{upp}}| \sum_{l \neq k} \left| \hat{F}_l^k(t) - \hat{F}_k^k(t) \right|. \end{aligned}$$

Therefore,

$$\begin{aligned} \left| \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) \right| & \leq \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) + \tilde{F}_k^k(t) - \hat{F}_k^k(t) \right| \\ & \quad + \sum_{l \neq k} \left| \hat{V}_{kl}^{\text{upp}} - V_{kl} \right| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) \\ & \leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \\ & \quad + \sum_{l \neq k} \left| \hat{V}_{kl}^{\text{upp}} - V_{kl} \right| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right). \end{aligned}$$

Then, in the event that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$,

$$\left| \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) \right| \leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right).$$

By contrast, if it is not the case that $V_{kl} \in [\hat{V}_{kl}^{\text{low}}, \hat{V}_{kl}^{\text{upp}}]$ for all $l \neq k$,

$$\begin{aligned}
\left| \hat{\Delta}_k^{\text{ci}}(t) - \Delta_k^{\text{ci}}(t) \right| &\leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \\
&\quad + \sum_{l \neq k} \left| \hat{V}_{kl}^{\text{upp}} - V_{kl} \right| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) \\
&\leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \\
&\quad + \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + |V_{kl}| \right) + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) \\
&\leq 2 \left(\max_{l \in [K]} \left| \hat{F}_l^k(t) - \tilde{F}_l^k(t) \right| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \\
&\quad + 2 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + (K-1) \left(\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right).
\end{aligned}$$

From here, proceeding as in the proof of Lemma A1, one easily arrives at:

$$\begin{aligned}
&\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\leq \frac{2 \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
&\quad + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot 2 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right).
\end{aligned}$$

This completes the second part of the proof. \square

Proof of Theorem 5. We follow the same strategy as in the proof of Theorem 3. Define the events $\mathcal{A}_1 = \{\hat{\mathcal{I}}_k^{\text{ci}} = \emptyset\}$ and $\mathcal{A}_2 = \{\hat{i}_k = 1\}$. Then,

$$\begin{aligned}
&\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\leq \mathbb{P} \left[\mathcal{A}_1 \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_2] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
&\quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right].
\end{aligned} \tag{A87}$$

We will now separately bound the three terms on the right-hand-side of (A87). The following notation will be useful for this purpose. For all $i \in [n_k]$, let $U_i \sim \text{Uniform}(0, 1)$ be independent and identically distributed uniform random variables, and denote their order statistics as $U_{(1)} < U_{(2)} < \dots < U_{(n_k)}$.

- The probability of \mathcal{A}_1 conditional on \hat{V}^{low} and \hat{V}^{upp} can be bound from above as:

$$\mathbb{P} \left[\hat{\mathcal{I}}_k^{\text{ci}} = \emptyset \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \leq \frac{1}{n_*} + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right]. \quad (\text{A88})$$

To simplify the notation in the proof of (A88), define

$$\begin{aligned} d_k^{\text{ci}} := & c(n_k) + \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2\sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right) \\ & + 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right). \end{aligned}$$

Then, under Assumption 5,

$$\begin{aligned} & 1 - \left(\alpha + \hat{\Delta}_k^{\text{ci}}(S_{(i)}^k) - \delta^{\text{ci}}(n_k, n_*) \right) \\ & = 1 - \alpha - \Delta_k(S_{(i)}^k) + \delta^{\text{ci}}(n_k, n_*) + \Delta_k(S_{(i)}^k) - \hat{\Delta}_k^{\text{ci}}(S_{(i)}^k) \\ & \leq 1 - d_k^{\text{ci}} + \delta^{\text{ci}}(n_k, n_*) + \sup_{t \in [0,1]} |\Delta_k(t) - \hat{\Delta}_k^{\text{ci}}(t)|. \end{aligned}$$

Further, it follows from the definition of $\delta^{\text{ci}}(n_k, n_*)$ in (22) and Lemma A4 applied with $\eta = 1/n_*$ that, with probability at least

$$1 - \frac{1}{n_*} - \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right],$$

we have

$$\begin{aligned} & 1 - \left(\alpha + \hat{\Delta}_k^{\text{ci}}(S_{(i)}^k) - \delta^{\text{ci}}(n_k, n_*) \right) \\ & \leq 1 - d_k^{\text{ci}} + \delta^{\text{ci}}(n_k, n_*) + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) + 2 \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \sqrt{\frac{\log(2K) + \log(n_*)}{2n_*}} \\ & \leq 1 - d_k^{\text{ci}} + c(n_k) + \frac{2 \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right)}{\sqrt{n_*}} \left(\frac{1}{\sqrt{n_*}} + 2\sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right) \\ & \quad + 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) = 1. \end{aligned}$$

This implies that the set $\hat{\mathcal{I}}_k^{\text{ci}}$ defined in (15), with $\hat{\Delta}_k(t)$ and $\delta^{\text{ci}}(n_k, n_*)$ instead of $\hat{\Delta}_k^{\text{ci}}(t)$ and $\delta(n_k, n_*)$, is non-empty with probability at least $1 - 1/n_* - \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right]$, because $n_k \in \hat{\mathcal{I}}_k^{\text{ci}}$ if $1 - [\alpha + \hat{\Delta}_k^{\text{ci}}(S_{(i)}^k) - \delta^{\text{ci}}(n_k, n_*)] \leq 1$.

- Under \mathcal{A}_2 , the second term on the right-hand-side of (A83) can be written as:

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] = \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq S_{(1)}^k \mid Y = k, \mathcal{D} \right] = F_k^k(S_{(1)}^k).$$

Therefore, by the independence of \hat{V}^{low} and \hat{V}^{upp} from the other data,

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_2] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \leq \mathbb{E} \left[\tilde{F}_k^k(S_{(1)}^k) \right] + \mathbb{E} \left[F_k^k(S_{(1)}^k) - \tilde{F}_k^k(S_{(1)}^k) \right] \\
& = \mathbb{E} [U_{(1)}] + (V_{kk} - 1) \mathbb{E} \left[\tilde{F}_k^k(S_{(1)}^k) \right] + \sum_{l \neq k} V_{kl} \mathbb{E} \left[\tilde{F}_l^k(S_{(1)}^k) \right] \\
& = V_{kk} \mathbb{E} [U_{(1)}] + \sum_{l \neq k} V_{kl} \mathbb{E} \left[\tilde{F}_l^k(S_{(1)}^k) \right] \\
& = \frac{V_{kk}}{n_k + 1} + \sum_{l \neq k} V_{kl} \mathbb{E} \left[\tilde{F}_l^k(S_{(1)}^k) \right].
\end{aligned}$$

Next, using Assumption 3 as in the proof of Theorem 3, we arrive to:

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_2] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \leq \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k + 1} = \frac{1 + \sum_{l \neq k} (|V_{kl}| - V_{kl})}{n_k + 1} \\
& \leq \frac{1 + 2 \sum_{l \neq k} |V_{kl}|}{n_k + 1}.
\end{aligned} \tag{A89}$$

- Under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$, by definition of $\hat{\mathcal{I}}_k^{\text{ci}}$, for any $i \leq \hat{i}_k - 1$,

$$\frac{i}{n_k} < 1 - \left(\alpha + \hat{\Delta}_k^{\text{ci}}(S_{(i)}^k) - \delta^{\text{ci}}(n_k, n_*) \right).$$

Therefore, choosing $i = \hat{i}_k - 1$, we get:

$$\frac{\hat{i}_k}{n_k} < 1 - \left(\alpha + \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k-1)}^k) - \delta^{\text{ci}}(n_k, n_*) \right) + \frac{1}{n_k}.$$

As in the proof of Theorem 3, the probability of coverage conditional on $Y_{n+1} = k$, on \mathcal{D} , and on the confidence interval $[\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$ can be decomposed as:

$$\begin{aligned}
& \mathbb{P} \left[Y \in C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& = \frac{\hat{i}_k}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) \\
& < 1 - \left(\alpha + \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k-1)}^k) - \delta^{\text{ci}}(n_k, n_*) \right) + \frac{1}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k)}^k) \\
& = 1 - \alpha + \delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) + \Delta_k(S_{(\hat{i}_k-1)}^k).
\end{aligned}$$

To bound the last term above, note that

$$\Delta_k(S_{(\hat{i}_k)}^k) - \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k-1)}^k) = (\Delta_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k-1)}^k)) + (\Delta_k(S_{(\hat{i}_k-1)}^k) - \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k-1)}^k))$$

$$\leq \sup_{t \in \mathbb{R}} |\Delta_k(t) - \widehat{\Delta}_k^{\text{ci}}(t)| + (\Delta_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k-1)}^k)),$$

where the expected value of the first term above can be bounded using Lemma A3, and the second term is given by

$$\begin{aligned} & \Delta_k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k-1)}^k) \\ &= (V_{kk} - 1) \left[\tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k-1)}^k) \right] + \sum_{l \neq k} V_{kl} \left[\tilde{F}_l^k(S_{(\hat{i}_k)}^k) - \tilde{F}_l^k(S_{(\hat{i}_k-1)}^k) \right] \\ &= \sum_{l \neq k} V_{kl} \left[\left(\tilde{F}_l^k(S_{(\hat{i}_k)}^k) - \tilde{F}_l^k(S_{(\hat{i}_k-1)}^k) \right) - \left(\tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k-1)}^k) \right) \right] \\ &\leq 2 \sum_{l \neq k} |V_{kl}| \cdot \max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right|. \end{aligned}$$

Combining the above, we find that, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned} & \mathbb{P} \left[Y \in C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\ &\leq 1 - \alpha + \delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} + \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) + \sup_{t \in \mathbb{R}} |\Delta_k(t) - \widehat{\Delta}_k^{\text{ci}}(t)| \\ &\quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right|. \end{aligned}$$

It remains to bound the expectation of the last term. By Assumption 2,

$$\begin{aligned} \max_{2 \leq i \leq n_k} \max_{l \in [K]} \left| \tilde{F}_l^k(S_{(i)}^k) - \tilde{F}_l^k(S_{(i-1)}^k) \right| &\leq f_{\max} \max_{2 \leq i \leq n_k} |S_{(i)}^k - S_{(i-1)}^k| \\ &\leq \frac{f_{\max}}{f_{\min}} \max_{2 \leq i \leq n_k} |\tilde{F}_k^k(S_{(i)}^k) - \tilde{F}_k^k(S_{(i-1)}^k)| \\ &\stackrel{d}{=} \frac{f_{\max}}{f_{\min}} \max_{2 \leq i \leq n_k} (U_{(i)} - U_{(i-1)}) \\ &\leq \frac{f_{\max}}{f_{\min}} \max_{1 \leq i \leq n_k+1} D_i, \end{aligned}$$

where $D_1 = U_{(1)}$, $D_i = U_{(i)} - U_{(i-1)}$ for $i = 2, \dots, n_k$, and $D_{n_k+1} = 1 - U_{(n_k)}$. By a standard result on maximum uniform spacing,

$$\mathbb{E} \left[\max_{1 \leq i \leq n_k+1} D_i \right] = \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j}.$$

Therefore, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \leq 1 - \alpha + \delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} + c(n_k) + \mathbb{E} \left[\sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k^{\text{ci}}(t)| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& \leq 1 - \alpha + \delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} + c(n_k) \\
& \quad + \frac{2 \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}|}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
& \quad + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + |V_{kl}| \right) + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) \\
& \quad + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& = 1 - \alpha + 2\delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} \\
& \quad + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \cdot \sum_{l \neq k} \left(|\hat{V}_{kl}^{\text{upp}}| + |V_{kl}| \right) - 2\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \\
& \quad + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
& \leq 1 - \alpha + 2\delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} \\
& \quad + 2 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \cdot \left(\mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] - \alpha_V \right) \\
& \quad + (K-1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j}.
\end{aligned} \tag{A90}$$

Above, the second inequality is obtained from Lemma A3 and the bound in (A82).

Combining (A87) with (A88), (A89), and (A90) leads to:

$$\begin{aligned}
& \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \leq \mathbb{P} \left[\mathcal{A}_1 \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_2] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\
& \quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right]
\end{aligned}$$

$$\begin{aligned}
&\leq 1 - \alpha + \frac{1}{n_*} + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] + \frac{1 + 2 \sum_{l \neq k} |V_{kl}|}{n_k + 1} \\
&\quad + 2\delta^{\text{ci}}(n_k, n_*) + \frac{1}{n_k} + 2 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \cdot \left(\mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] - \alpha_V \right) \\
&\quad + (K - 1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k + 1} \sum_{j=1}^{n_k+1} \frac{1}{j} \\
&\leq 1 - \alpha + \frac{1}{n_*} + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] + (K - 1) \left(2\hat{\delta}_{k*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) \\
&\quad + 2\delta^{\text{ci}}(n_k, n_*) + 2 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \cdot \left(\mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] - \alpha_V \right) \\
&\quad + \frac{2}{n_k} \left[1 + \sum_{l \neq k} |V_{kl}| + \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \sum_{j=1}^{n_k+1} \frac{1}{j} \right].
\end{aligned}$$

Finally, taking an expectation with respect to \hat{V}^{low} and \hat{V}^{upp} , we arrive at:

$$\begin{aligned}
&\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{ci}}(X_{n+1}) \mid Y = k \right] \\
&\leq 1 - \alpha + \frac{1}{n_*} + \alpha_V + 2\mathbb{E} \left[\delta^{\text{ci}}(n_k, n_*) \right] + (K - 1) \left(2\mathbb{E} \left[\hat{\delta}_{k*}^{(V)} \right] + \mathbb{E} \left[|\hat{\zeta}_k^{\text{upp}}| \right] \right) \\
&\quad + \frac{2}{n_k} \left[1 + \sum_{l \neq k} |V_{kl}| + \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \sum_{j=1}^{n_k+1} \frac{1}{j} \right] \\
&= 1 - \alpha + \frac{1}{n_*} + \alpha_V + 2c(n_k) + 4\alpha_V \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \\
&\quad + \frac{4 \sum_{l \neq k} \mathbb{E} \left[|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right]}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
&\quad + (K - 1) \left(2\mathbb{E} \left[\hat{\delta}_{k*}^{(V)} \right] + \mathbb{E} \left[|\hat{\zeta}_k^{\text{upp}}| \right] \right) + \frac{2}{n_k} \left[1 + \sum_{l \neq k} |V_{kl}| + \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \sum_{j=1}^{n_k+1} \frac{1}{j} \right] \\
&= 1 - \alpha + \frac{1}{n_*} + \left(1 + 4 \sum_{l \neq k} |\bar{V}_{kl}^{\text{upp}}| \right) \alpha_V + 2c(n_k) \\
&\quad + \frac{4 \sum_{l \neq k} \mathbb{E} \left[|\hat{V}_{kl}^{\text{upp}}| + \hat{\delta}_{kl}^{(V)} \right]}{\sqrt{n_*}} \min \left\{ K \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K) + \log(n_*)}{2}} \right\} \\
&\quad + (K - 1) \left(2\mathbb{E} \left[\hat{\delta}_{k*}^{(V)} \right] + \mathbb{E} \left[|\hat{\zeta}_k^{\text{upp}}| \right] \right) + \frac{2}{n_k} \left[1 + \sum_{l \neq k} |V_{kl}| + \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \sum_{j=1}^{n_k+1} \frac{1}{j} \right].
\end{aligned}$$

□

Lemma A4. *Under the assumptions of Theorem 4, for any $k \in [K]$ and $\eta > 0$,*

$$\begin{aligned} & \mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| > (K-1) \left(2\hat{\delta}_{k^*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) + 2\Upsilon(K, \eta, n_*) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\ & \leq \eta \mathbb{I} \left[V \in [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right], \end{aligned}$$

where $\Upsilon(K, \eta, n_*) := \sqrt{[\log(2K) + \log(1/\eta)]/(2n_*)}$.

Proof of Lemma A4. By proceeding as in the proof of Lemma A3, we obtain that, in the event that $V \in [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}]$, for any $k \in [K]$ and $t \in \mathbb{R}$,

$$|\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| \leq 2 \left(\max_{l \in [K]} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| + (K-1) \left(2\hat{\delta}_{k^*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right).$$

Combined with the DKW inequality, this implies that, for any $\eta > 0$,

$$\begin{aligned} & \mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t)| > (K-1) \left(2\hat{\delta}_{k^*}^{(V)} + |\hat{\zeta}_k^{\text{upp}}| \right) + 2\Upsilon(K, \eta, n_*) \sum_{l \neq k} |\hat{V}_{kl}^{\text{upp}}| \mid \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\ & \leq \mathbb{P} \left[\max_{j \in [K]} |\hat{F}_j^k(t) - \tilde{F}_j^k(t)| > \sqrt{\frac{\log(2K) + \log(1/\eta)}{2n_*}} \mid \mathbb{I} \left[V \in [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \right] \\ & \quad + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] \\ & \leq \eta \mathbb{I} \left[V \in [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right] + \mathbb{I} \left[V \notin [\hat{V}^{\text{low}}, \hat{V}^{\text{upp}}] \right]. \end{aligned}$$

Note that the last inequality above follows exactly as in the proof of Lemma A2. \square

Proof of Proposition 3. This proof combines elements of the proofs of Proposition 2, Theorem 2, and Theorem 4. Suppose $Y_{n+1} = k$, for some $k \in [K]$. Proceeding exactly as in the proof of Theorem 2, we obtain:

$$\begin{aligned} & \mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D}, \hat{V}^{\text{low}}, \hat{V}^{\text{upp}} \right] \\ & = 1 - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k)}^k) \\ & = 1 - \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \max \left\{ \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*), -(1-\alpha)/n_k \right\} \\ & \quad + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) \\ & \quad + \max \left\{ \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*), -(1-\alpha)/n_k \right\} - (\Delta_k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*)) \\ & = 1 - \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \max \left\{ \hat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*), -(1-\alpha)/n_k \right\} \\ & \quad + \hat{F}_k^k(S_{(\hat{i}_k)}^k) - \tilde{F}_k^k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) \end{aligned}$$

$$\begin{aligned}
& + \max \left\{ \widehat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) , -(1 - \alpha)/n_k \right\} \\
& - \max \left\{ \Delta_k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) , -(1 - \alpha)/n_k \right\} \\
= & 1 - \widehat{F}_k^k(S_{(\hat{i}_k)}^k) - \max \left\{ \widehat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) , -(1 - \alpha)/n_k \right\} \\
& + \widehat{F}_k^k(S_{(\hat{i}_k)}^k) - \widetilde{F}_k^k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) \\
& + \max \left\{ \widehat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) , \delta^{\text{ci}}(n_k, n_*) - (1 - \alpha)/n_k \right\} \\
& - \max \left\{ \Delta_k(S_{(\hat{i}_k)}^k) , \delta^{\text{ci}}(n_k, n_*) - (1 - \alpha)/n_k \right\} \\
\leq & 1 - \widehat{F}_k^k(S_{(\hat{i}_k)}^k) - \max \left\{ \widehat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) , -(1 - \alpha)/n_k \right\} \\
& + \widehat{F}_k^k(S_{(\hat{i}_k)}^k) - \widetilde{F}_k^k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) + \max \left\{ \widehat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \Delta_k(S_{(\hat{i}_k)}^k) , 0 \right\} \\
\leq & \left[1 - \frac{\hat{i}_k}{n_k} - \max \left\{ \widehat{\Delta}_k^{\text{ci}}(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) , -(1 - \alpha)/n_k \right\} \right] \\
& + \widehat{F}_k^k(S_{(\hat{i}_k)}^k) - \widetilde{F}_k^k(S_{(\hat{i}_k)}^k) - \delta^{\text{ci}}(n_k, n_*) + \sup_{t \in \mathbb{R}} \max \left\{ \widehat{\Delta}_k^{\text{ci}}(t) - \Delta_k(t) , 0 \right\} ,
\end{aligned}$$

using the assumption that $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta^{\text{ci}}(n_k, n_*) - (1 - \alpha)/n_k$ in the third equality above. The proof is then completed by proceeding as in the proof of Theorem 4, using Lemma A3. □

A7.4 Fitting the label contamination model

Proof of Equation (28). By definition of Q and \tilde{Q} in (27), for any $k, l \in [K]$,

$$\begin{aligned}
\tilde{Q}_{lk} &= \mathbb{P} \left[\hat{f}(X) = k \mid \tilde{Y} = l, \hat{f} \right] \\
&= \sum_{l'=1}^K \mathbb{P} \left[\hat{f}(X) = k, Y = l' \mid \tilde{Y} = l, \hat{f} \right] \\
&= \sum_{l'=1}^K \mathbb{P} \left[Y = l' \mid \tilde{Y} = l, \hat{f} \right] \mathbb{P} \left[\hat{f}(X) = k \mid Y = l', \tilde{Y} = l, \hat{f} \right] \\
&= \sum_{l'=1}^K M_{ll'} \mathbb{P} \left[\hat{f}(X) = k \mid Y = l', \hat{f} \right] \\
&= \sum_{l'=1}^K M_{ll'} Q_{l'k} \\
&= (MQ)_{lk}.
\end{aligned}$$

□

Proof of Equation (A46). The proof of (A46) follows quite directly from Equation (28) and the definitions of $\tilde{\psi}$ and ψ in Equation (A47):

$$\begin{aligned}
\tilde{\psi} &:= \sum_{k=1}^K \mathbb{P} \left[\hat{f}(X) = k, \tilde{Y} = k \mid \hat{f} \right] \\
&= \sum_{k=1}^K \tilde{\rho}_k \tilde{Q}_{kk} = \sum_{k=1}^K \tilde{\rho}_k \sum_{s=1}^K M_{ks} Q_{sk} \\
&= \sum_{k=1}^K \tilde{\rho}_k \left(M_{kk} Q_{kk} + \sum_{s \neq k} M_{ks} Q_{sk} \right) \\
&= \sum_{k=1}^K \tilde{\rho}_k \left(\frac{(1-\epsilon)\rho_k}{(1-\epsilon)\rho_k + \epsilon/K} Q_{kk} + \sum_{s=1}^K \frac{\rho_s \cdot \epsilon/K}{(1-\epsilon)\rho_k + \epsilon/K} Q_{sk} \right) \\
&= (1-\epsilon) \sum_{k=1}^K \rho_k Q_{kk} + \frac{\epsilon}{K} \sum_{k=1}^K \sum_{s=1}^K \rho_s Q_{sk} \\
&= (1-\epsilon) \sum_{k=1}^K \mathbb{P} \left[\hat{f}(X) = k, Y = k \mid \hat{f} \right] + \frac{\epsilon}{K} \sum_{k=1}^K \sum_{s=1}^K \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] \\
&= (1-\epsilon)\psi + \frac{\epsilon}{K}.
\end{aligned}$$

□

Proof of Equation (A56). By the definition of $\hat{\zeta}_k$ in (20),

$$\begin{aligned}
\hat{\zeta}_k &:= \max_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) - \min_{l \neq k} \left(\hat{V}_{kl}^{\text{upp}} - V_{kl} \right) \\
&= \left| \left(\frac{\xi}{K} \cdot \frac{1 + \nu(1 + 2\xi)}{1 + \nu\xi} - \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{1 + \hat{\nu}^{\text{low}}(1 + 2\hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}} \right) \right. \\
&\quad \left. - \left(\frac{\xi}{K} \cdot \frac{1 - \nu}{1 + \nu\xi} - \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{1 - \hat{\nu}^{\text{upp}}}{1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}}} \right) \right| \\
&= \left| \frac{\xi}{K} \cdot \frac{2\nu(1 + \xi)}{1 + \nu\xi} - \frac{\hat{\xi}^{\text{low}}}{K} \left(\frac{1 + \hat{\nu}^{\text{low}}(1 + 2\hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}} - \frac{1 - \hat{\nu}^{\text{upp}}}{1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}}} \right) \right| \\
&= \left| \frac{\xi}{K} \cdot \frac{2\nu(1 + \xi)}{1 + \nu\xi} - \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{(1 + \hat{\xi}^{\text{low}}) (\hat{\nu}^{\text{low}} + \hat{\nu}^{\text{upp}} + 2\hat{\nu}^{\text{low}}\hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})}{(1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}) (1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})} \right|.
\end{aligned}$$

Then, applying the triangle inequality gives us:

$$\hat{\zeta}_k \leq \left| \frac{\xi}{K} \cdot \frac{2\nu(1 + \xi)}{1 + \nu\xi} - \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{2\hat{\nu}^{\text{low}}(1 + \hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}} \right|$$

$$\begin{aligned}
& + \frac{\hat{\xi}^{\text{low}}}{K} \left| \frac{2\hat{\nu}^{\text{low}}(1 + \hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}} - \frac{(1 + \hat{\xi}^{\text{low}})(\hat{\nu}^{\text{low}} + \hat{\nu}^{\text{upp}} + 2\hat{\nu}^{\text{low}}\hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})}{(1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}})(1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})} \right| \\
& = \left| \frac{\xi}{K} \cdot \frac{2\nu(1 + \xi)}{1 + \nu\xi} - \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{2\hat{\nu}^{\text{low}}(1 + \hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}} \right| + \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{(\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}})(1 + \hat{\xi}^{\text{low}})}{(1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}})(1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})}.
\end{aligned}$$

Next, the monotonicity of the function $(\nu, \epsilon) \mapsto (\nu(1 + \epsilon))/((1 + \nu\epsilon))$ in both inputs allows us to conclude that:

$$\begin{aligned}
\hat{\zeta}_k & \leq \left| \frac{\hat{\xi}^{\text{upp}}}{K} \cdot \frac{2\hat{\nu}^{\text{upp}}(1 + \hat{\xi}^{\text{upp}})}{1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{upp}}} - \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{2\hat{\nu}^{\text{low}}(1 + \hat{\xi}^{\text{low}})}{1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}}} \right| \\
& + \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{(\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}})(1 + \hat{\xi}^{\text{low}})}{(1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}})(1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})} \\
& = \frac{2}{K} \cdot \frac{(\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}}) + \left[\hat{\nu}^{\text{upp}} (\hat{\xi}^{\text{upp}})^2 - \hat{\nu}^{\text{low}} (\hat{\xi}^{\text{low}})^2 \right] + \hat{\nu}^{\text{low}}\hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}}\hat{\xi}^{\text{upp}} (\hat{\xi}^{\text{upp}} - \hat{\xi}^{\text{low}})}{(1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{upp}})(1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}})} \\
& + \frac{\hat{\xi}^{\text{low}}}{K} \cdot \frac{(\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}})(1 + \hat{\xi}^{\text{low}})}{(1 + \hat{\nu}^{\text{low}}\hat{\xi}^{\text{low}})(1 + \hat{\nu}^{\text{upp}}\hat{\xi}^{\text{low}})}.
\end{aligned}$$

□

Proof of Equation (A57). We begin by proving the first equation, namely

$$\tilde{\psi} = (1 - \epsilon)\psi + \frac{\epsilon}{K} + \frac{\epsilon\nu}{K}(2\phi - 1).$$

It follows directly from the general estimating equation (28) that

$$\begin{aligned}
\tilde{\psi} & := \sum_{k=1}^K \mathbb{P} \left[\hat{f}(X) = k, \tilde{Y} = k \mid \hat{f} \right] \\
& = \frac{1}{K} \sum_{k=1}^K \tilde{Q}_{kk} \\
& = \frac{1}{K} \sum_{k=1}^K \sum_{s=1}^K M_{ks} Q_{sk} \\
& = \frac{1}{K} \sum_{k=1}^K \left(M_{kk} Q_{kk} + \sum_{s \neq k} M_{ks} Q_{sk} \right) \\
& = \frac{1}{K} \sum_{k=1}^K \left((1 - \epsilon) Q_{kk} + \sum_{s \in \mathcal{B}_k} \frac{\epsilon}{K} (1 + \nu) Q_{sk} + \sum_{s \in \mathcal{B}_k^c} \frac{\epsilon}{K} (1 - \nu) Q_{sk} \right)
\end{aligned}$$

$$\begin{aligned}
&= (1 - \epsilon)\psi + \frac{\epsilon(1 + \nu)}{K^2} \sum_{k=1}^K \sum_{s \in \mathcal{B}_k} Q_{sk} + \frac{\epsilon(1 - \nu)}{K^2} \sum_{k=1}^K \sum_{s \in \mathcal{B}_k^c} Q_{sk} \\
&= (1 - \epsilon)\psi + \frac{\epsilon(1 + \nu)}{K} \sum_{k=1}^K \sum_{s \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] \\
&\quad + \frac{\epsilon(1 - \nu)}{K} \sum_{k=1}^K \sum_{s \in \mathcal{B}_k^c} \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] \\
&= (1 - \epsilon)\psi + \frac{\epsilon}{K} \sum_{k=1}^K \sum_{s=1}^K \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] \\
&\quad + \frac{\epsilon\nu}{K} \sum_{k=1}^K \left(\sum_{s \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] - \sum_{s \in \mathcal{B}_k^c} \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] \right) \\
&= (1 - \epsilon)\psi + \frac{\epsilon}{K} + \frac{\epsilon\nu}{K} \left(2 \sum_{k=1}^K \sum_{s \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = k, Y = s \mid \hat{f} \right] - 1 \right) \\
&= (1 - \epsilon)\psi + \frac{\epsilon}{K} + \frac{\epsilon\nu}{K} (2\phi - 1),
\end{aligned}$$

where

$$\phi := \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = l, Y = k \mid \hat{f} \right].$$

This completes the first part of the proof.

Let us now prove the second equation, namely

$$\tilde{\phi} = \phi - \epsilon(1 - \nu) \left(\phi - \frac{1}{2} \right).$$

It follows directly from the general estimating equation (28) that

$$\begin{aligned}
\tilde{\phi} &:= \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \mathbb{P} \left[\hat{f}(X) = l, \tilde{Y} = k \mid \hat{f} \right] \\
&= \frac{1}{K} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \tilde{Q}_{kl} \\
&= \frac{1}{K} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \sum_{s=1}^K M_{ks} Q_{sl} \\
&= \frac{1}{K} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \left(M_{kk} Q_{kl} + \sum_{s \neq k} M_{ks} Q_{sl} \right)
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{K} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \left((1-\epsilon)Q_{kl} + \sum_{s \in \mathcal{B}_k} \frac{\epsilon}{K} (1+\nu) Q_{sl} + \sum_{s \in \mathcal{B}_k^c} \frac{\epsilon}{K} (1-\nu) Q_{sk} \right) \\
&= (1-\epsilon) \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \frac{Q_{kl}}{K} + \frac{\epsilon(1+\nu)}{K} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \sum_{s \in \mathcal{B}_k} \frac{Q_{sl}}{K} + \frac{\epsilon(1-\nu)}{K} \sum_{k=1}^K \sum_{l \in \mathcal{B}_k} \sum_{s \in \mathcal{B}_k^c} \frac{Q_{sl}}{K} \\
&= (1-\epsilon)\phi + \frac{\epsilon(1+\nu)}{K} \frac{K}{2} \sum_{k=1}^K \sum_{s \in \mathcal{B}_k} \frac{Q_{sk}}{K} + \frac{\epsilon(1-\nu)}{K} \frac{K}{2} \sum_{k=1}^K \sum_{s \in \mathcal{B}_k^c} \frac{Q_{sk}}{K} \\
&= (1-\epsilon)\phi + \frac{\epsilon(1+\nu)}{2}\phi + \frac{\epsilon(1-\nu)}{2}(1-\phi).
\end{aligned}$$

This concludes the second part of the proof. \square

A7.5 Extensions of preliminary theoretical results

Proof of Theorem A1. The proof strategy is similar to that of Theorem 1. By definition of the conformity score function in (1), $Y_{n+1} \in \hat{C}(X_{n+1})$ if and only if $\hat{s}(X_{n+1}, Y_{n+1}) \leq \hat{\tau}_{Y_{n+1}}$. Therefore,

$$\begin{aligned}
&\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \right] \\
&= \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \right] + \left(\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \right] - \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \right] \right).
\end{aligned}$$

The proof is then completed by noting that the second term on the right-hand-side above can be written as:

$$\begin{aligned}
&\mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \right] - \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \right] \\
&= \sum_{k=1}^K \left(\rho_k \cdot \mathbb{P} \left[Y_{n+1} \in \hat{C}(X_{n+1}) \mid Y_{n+1} = k \right] - \tilde{\rho}_k \cdot \mathbb{P} \left[\tilde{Y}_{n+1} \in \hat{C}(X_{n+1}) \mid \tilde{Y}_{n+1} = k \right] \right) \\
&= \sum_{k=1}^K \left(\rho_k \cdot \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq \hat{\tau} \mid Y_{n+1} = k \right] - \tilde{\rho}_k \cdot \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq \hat{\tau} \mid \tilde{Y}_{n+1} = k \right] \right) \\
&= \mathbb{E} \left[\sum_{k=1}^K \left[\rho_k F_k^k(\hat{\tau}) - \tilde{\rho}_k \tilde{F}_k^k(\hat{\tau}) \right] \right] = \mathbb{E} [\Delta(\hat{\tau})].
\end{aligned}$$

\square

Proof of Corollary 3. By Theorem A1, it suffices to prove $\mathbb{E}[\Delta(t)] \geq 0$ for all $t \in \mathbb{R}$. To establish that, note that combining (A68) with Proposition 1 gives:

$$\Delta(t) = \sum_{k=1}^K \left[\rho_k F_k^k(t) - \tilde{\rho}_k \tilde{F}_k^k(t) \right] = \sum_{k=1}^K \left[\rho_k F_k^k(t) - \sum_{l=1}^K M_{kl} \tilde{\rho}_l F_l^k(t) \right].$$

To simplify the notation in the following, define

$$W_{kl} := \mathbb{P}[\tilde{Y} = k \mid X, Y = l] = M_{kl} \frac{\tilde{\rho}_k}{\rho_l}.$$

Then,

$$\begin{aligned} \Delta(t) &= \sum_{l=1}^K \rho_l F_l^l(t) - \sum_{l=1}^K \rho_l \sum_{k=1}^K W_{kl} F_l^k(t) \\ &= \sum_{l=1}^K \rho_l F_l^l(t) - \sum_{l=1}^K \rho_l \left[W_{ll} F_l^l(t) + \sum_{k \neq l} W_{kl} F_l^k(t) \right] \\ &\geq \sum_{l=1}^K \rho_l F_l^l(t) - \sum_{l=1}^K \rho_l \left[W_{ll} F_l^l(t) + \left(\sum_{k \neq l} W_{kl} \right) \max_{k \neq l} F_l^k(t) \right] \\ &= \sum_{l=1}^K \rho_l F_l^l(t) - \sum_{l=1}^K \rho_l \left[W_{ll} F_l^l(t) + \left(\sum_{k=1}^K W_{kl} \right) \max_{k \neq l} F_l^k(t) - W_{ll} \max_{k \neq l} F_l^k(t) \right] \\ &= \sum_{l=1}^K \rho_l F_l^l(t) - \sum_{l=1}^K \rho_l \left[W_{ll} F_l^l(t) + \max_{k \neq l} F_l^k(t) - W_{ll} \max_{k \neq l} F_l^k(t) \right] \\ &= \sum_{l=1}^K \rho_l (1 - W_{ll}) \left[F_l^l(t) - \max_{k \neq l} F_l^k(t) \right] \geq 0, \end{aligned}$$

where the last inequality is given by (A71). This implies that $\Delta(\hat{\tau}) \geq 0$ almost-surely. \square

A7.6 Prediction sets with marginal coverage

Proof of Theorem A2. The proof strategy is similar to those of Theorem A1 and Theorem 2. By definition of the conformity score function in (1), the event $Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1})$ occurs if and only if $\hat{s}(X_{n+1}, Y_{n+1}) \leq \hat{\tau}^{\text{marg}}$. Further, note that

$$\tilde{F}^k(t) = \sum_{k=1}^K \tilde{\rho}_k \tilde{F}_k^k(t).$$

Therefore, the probability of a miscoverage event conditional on the data in \mathcal{D} can be decomposed as:

$$\begin{aligned} &\mathbb{P} \left[Y_{n+1} \notin \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \\ &= \mathbb{P} \left[\tilde{Y}_{n+1} \notin \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \\ &\quad + \left(\mathbb{P} \left[Y_{n+1} \notin \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] - \mathbb{P} \left[\tilde{Y}_{n+1} \notin \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \right) \\ &= \sum_{k=1}^K \tilde{\rho}_k \left[1 - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) \right] \end{aligned}$$

$$\begin{aligned}
& + \sum_{k=1}^K \rho_k \mathbb{P} \left[\hat{s}(X_{n+1}, k) > S_{(\hat{i}^{\text{marg}})}^k \mid \mathcal{D}, Y_{n+1} = k \right] \\
& - \sum_{k=1}^K \tilde{\rho}_k \mathbb{P} \left[\hat{s}(X_{n+1}, k) > S_{(\hat{i}^{\text{marg}})}^k \mid \mathcal{D}, \tilde{Y}_{n+1} = k \right] \\
& = \sum_{k=1}^K \tilde{\rho}_k \left[1 - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) \right] + \sum_{k=1}^K \left[\rho_k - \rho_k F_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{\rho}_k + \tilde{\rho}_k \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) \right] \\
& = 1 - \sum_{k=1}^K \tilde{\rho}_k \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}})}) \\
& = 1 - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}})}) \\
& = 1 - \hat{F}(S_{(\hat{i}^{\text{marg}})}) - \hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) \\
& \quad + \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] + \left[\hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}})}) \right] \\
& = \left\{ 1 - \frac{\hat{i}^{\text{marg}}}{n_{\text{cal}}} - \hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) + \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right\} - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \\
& \quad + \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] + \left[\hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}})}) \right] \\
& \leq \sup_{i \in \hat{\mathcal{I}}^{\text{marg}}} \left\{ 1 - \frac{i}{n_{\text{cal}}} - \hat{\Delta}(S_{(i)}) + \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right\} - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \\
& \quad + \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] + \sup_{t \in \mathbb{R}} \left[\hat{\Delta}(t) - \Delta(t) \right].
\end{aligned}$$

By definition of $\hat{\mathcal{I}}^{\text{marg}}$, for all $i \in \hat{\mathcal{I}}^{\text{marg}}$,

$$1 - \frac{i}{n_{\text{cal}}} - \hat{\Delta}(S_{(i)}) + \delta^{\text{marg}}(n_{\text{cal}}, n_*) \leq \alpha,$$

which implies a.s.

$$\sup_{i \in \hat{\mathcal{I}}^{\text{marg}}} \left\{ 1 - \frac{i}{n_{\text{cal}}} - \hat{\Delta}(S_{(i)}) + \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right\} \leq \alpha.$$

Therefore,

$$\begin{aligned}
& \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] \\
& \leq \alpha + \mathbb{E} \left[\sup_{t \in \mathbb{R}} \left[\hat{\Delta}(t) - \Delta(t) \right] \right] + \mathbb{E} \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] - \delta^{\text{marg}}(n_{\text{cal}}, n_*).
\end{aligned}$$

The first expected value on the right-hand-side above can be bound using the DKW inequality, as made precise by Lemma A5. This leads to:

$$\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] \leq \alpha + \mathbb{E} \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] - \delta^{\text{marg}}(n_{\text{cal}}, n_*)$$

$$\begin{aligned}
& + \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \\
& \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\}.
\end{aligned}$$

Finally, the remaining expected value can be bound as the corresponding term in the proof of Theorem 2. Let $U_1, \dots, U_{n_{\text{cal}}}$ be i.i.d. uniform random variables on $[0, 1]$, and denote their order statistics as $U_{(1)}, \dots, U_{(n_{\text{cal}})}$. Then,

$$\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \stackrel{d}{=} \frac{\hat{i}^{\text{marg}}}{n_{\text{cal}}} - U_{(\hat{i}^{\text{marg}})}.$$

Therefore, it follows from (17) that

$$\mathbb{E} \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] \leq \mathbb{E} \left[\sup_{i \in [n_{\text{cal}}]} \left\{ \frac{i}{n_{\text{cal}}} - U_{(i)} \right\} \right] = c(n_{\text{cal}}). \quad (\text{A91})$$

$$\begin{aligned}
\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] & \leq \alpha + c(n_{\text{cal}}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \\
& + \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \\
& \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\} \\
& = \alpha.
\end{aligned}$$

□

Lemma A5. *Under the assumptions of Theorem A2,*

$$\begin{aligned}
\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)| \right] & \leq \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \\
& \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\}.
\end{aligned}$$

Proof of Lemma A5. Note that, for any $k \in [K]$ and $t \in \mathbb{R}$,

$$\begin{aligned}
\hat{\Delta}(t) - \Delta(t) & = \sum_{k=1}^K (\rho_k V_{kk} - \tilde{\rho}_k) \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] + \sum_{k=1}^K \rho_k \sum_{l \neq k} V_{kl} \left[\hat{F}_l^k(t) - \tilde{F}_l^k(t) \right] \\
& = \sum_{k=1}^K \rho_k (V_{kk} - 1) \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] + \sum_{k=1}^K \rho_k \sum_{l \neq k} V_{kl} \left[\hat{F}_l^k(t) - \tilde{F}_l^k(t) \right]
\end{aligned}$$

$$\begin{aligned}
& + \sum_{k=1}^K (\rho_k - \tilde{\rho}_k) \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] \\
= & \sum_{k=1}^K \rho_k \sum_{l \neq k} V_{kl} \left\{ \left[\hat{F}_l^k(t) - \tilde{F}_l^k(t) \right] - \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] \right\} \\
& + \sum_{k=1}^K (\rho_k - \tilde{\rho}_k) \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right],
\end{aligned}$$

where in the last equality we used the fact that V has row sums equal to 1 because it is the inverse of M , which has row sums equal to 1. Therefore,

$$\begin{aligned}
|\hat{\Delta}(t) - \Delta(t)| & \leq 2 \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} |V_{kl}| \right) \left(\max_{l \neq k} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right) \\
& \quad + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \left(\max_{k \in [K]} |\hat{F}_k^k(t) - \tilde{F}_k^k(t)| \right) \\
& \leq 2 \left(\max_{k \in [K]} \sum_{l \neq k} |V_{kl}| \right) \left(\max_{k, l \in [K]: l \neq k} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right) \sum_{k=1}^K \rho_k \\
& \quad + \left(\max_{k \in [K]} |\hat{F}_k^k(t) - \tilde{F}_k^k(t)| \right) \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \\
& = 2 \left(\max_{k \in [K]} \sum_{l \neq k} |V_{kl}| \right) \left(\max_{k, l \in [K]: l \neq k} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right) \\
& \quad + \left(\sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \left(\max_{k \in [K]} |\hat{F}_k^k(t) - \tilde{F}_k^k(t)| \right) \\
& \leq \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \left(\max_{k, l \in [K]} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right),
\end{aligned}$$

and thus

$$\mathbb{E} \left[|\hat{\Delta}(t) - \Delta(t)| \right] \leq \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \mathbb{E} \left[\max_{k, l \in [K]} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right].$$

As in the proof of Lemma A1, define

$$l_{\min} := \arg \min_{l \in [K]} n_l. \qquad n_* := \min_{l \in [K]} n_l.$$

Then, combining the DKW inequality with a union bound leads to:

$$\mathbb{E} \left[\max_{k, l \in [K]} \sup_{t \in \mathbb{R}} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| \right] \leq K^2 \sqrt{\frac{\pi}{2n_*}}.$$

Similarly, the DKW inequality also implies that, for any $\eta > 0$,

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)| > \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \sqrt{\frac{\log(2K^2) + \log(1/\eta)}{2n_*}} \right] \leq \eta.$$

see Lemma A6. Therefore, setting $\eta = 1/n_*$, we obtain

$$\mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)| \right] \leq \left[2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right] \left[\sqrt{\frac{\log(2K^2) + \log(n_*)}{2n_*}} + \frac{1}{n_*} \right].$$

This concludes the proof that

$$\begin{aligned} \mathbb{E} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)| \right] &\leq \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \\ &\quad \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\}. \end{aligned}$$

□

Lemma A6. *Under the assumptions of Theorem A2, for any $\eta > 0$,*

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)| > \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \sqrt{\frac{\log(2K^2) + \log(1/\eta)}{2n_*}} \right] \leq \eta.$$

Proof of Lemma A6. It follows from the definitions of $\hat{\Delta}(t)$ and $\Delta(t)$, and from the DKW inequality, that, for any $\eta > 0$ and $k_0 \in [K]$,

$$\begin{aligned} &\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)| > \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k| \right) \sqrt{\frac{\log(2K^2) + \log(1/\eta)}{2n_*}} \right] \\ &\leq \mathbb{P} \left[\max_{k, l \in [K]} |\hat{F}_l^k(t) - \tilde{F}_l^k(t)| > \sqrt{\frac{\log(2K^2) + \log(1/\eta)}{2n_*}} \right] \\ &\leq K^2 \cdot \mathbb{P} \left[|\hat{F}_{l_{\min}}^{k_0}(t) - \tilde{F}_{l_{\min}}^{k_0}(t)| > \sqrt{\frac{\log(2K^2) + \log(1/\eta)}{2n_*}} \right] \\ &\leq 2K^2 \cdot \exp \left[-2n_* \frac{\log(2K^2) + \log(1/\eta)}{2n_*} \right] \\ &= 2K^2 \cdot \exp \left[-\log(2K^2) - \log(1/\eta) \right] \\ &= \eta. \end{aligned}$$

□

Proof of Theorem A3. The proof combines the strategies from the proof of Theorem 3 and the proof of Theorem A2. Define the events $\mathcal{A}_1 = \{\hat{\mathcal{I}}^{\text{marg}} = \emptyset\}$ and $\mathcal{A}_2 = \{\hat{i}^{\text{marg}} = 1\}$. Then,

$$\begin{aligned} \mathbb{P}\left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1})\right] &\leq \mathbb{P}[\mathcal{A}_1] + \mathbb{E}\left[\mathbb{P}\left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D}\right] \mathbb{I}[\mathcal{A}_2]\right] \\ &\quad + \mathbb{E}\left[\mathbb{P}\left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D}\right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c]\right]. \end{aligned} \quad (\text{A92})$$

We will now separately bound the three terms on the right-hand-side of (A92). The following notation will be useful for this purpose. For all $i \in [n_{\text{cal}}]$, let $U_i \sim \text{Uniform}(0, 1)$ be independent and identically distributed uniform random variables, and denote their order statistics as $U_{(1)} < U_{(2)} < \dots < U_{(n_{\text{cal}})}$.

- The probability of the event \mathcal{A}_1 can be bound from above as:

$$\mathbb{P}\left[\hat{\mathcal{I}}_k = \emptyset\right] \leq \frac{1}{n_*}. \quad (\text{A93})$$

To simplify the notation in the proof of (A93), define

$$d^{\text{marg}} := c(n_{\text{cal}}) + \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|\right) \frac{1 + 2\sqrt{\log(2K^2) + \log(n_*)}}{\sqrt{n_*}}.$$

Then, under Assumption 6,

$$\begin{aligned} 1 - \left(\alpha + \hat{\Delta}(S_{(i)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*)\right) &= 1 - \alpha - \delta^{\text{marg}}(S_{(i)}) + \delta^{\text{marg}}(n_{\text{cal}}, n_*) + \delta^{\text{marg}}(S_{(i)}) - \hat{\Delta}(S_{(i)}) \\ &\leq 1 - d^{\text{marg}} + \delta^{\text{marg}}(n_{\text{cal}}, n_*) + \sup_{t \in [0,1]} |\delta^{\text{marg}}(t) - \hat{\Delta}(t)|. \end{aligned}$$

Further, it follows from Lemma A6 and (A76) that, with probability at least $1 - 1/n_*$,

$$\begin{aligned} 1 - \left(\alpha + \hat{\Delta}(S_{(i)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*)\right) &\leq 1 - d^{\text{marg}} + \delta^{\text{marg}}(n_{\text{cal}}, n_*) \\ &\quad + \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|\right) \sqrt{\frac{\log(2K^2) + \log(n_*)}{2n_*}} \\ &\leq 1 - d^{\text{marg}} + c(n_{\text{cal}}) \\ &\quad + \left(2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|\right) \frac{1 + 2\sqrt{\log(2K^2) + \log(n_*)}}{\sqrt{n_*}} \\ &= 1. \end{aligned}$$

This implies that the set $\hat{\mathcal{I}}^{\text{marg}}$ defined in (A73) is non-empty with probability at least $1 - 1/n_*$, because $n_{\text{cal}} \in \hat{\mathcal{I}}^{\text{marg}}$ if $1 - [\alpha + \hat{\Delta}(S_{(i)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*)] \leq 1$.

- Under \mathcal{A}_2 , the second term on the right-hand-side of (A92) can be written as:

$$\begin{aligned} \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] &= \mathbb{P} \left[\hat{s}(X_{n+1}, Y_{n+1}) \leq S_{(1)} \mid \mathcal{D} \right] \\ &= F(S_{(1)}) \\ &= \sum_{k=1}^K \rho_k F_k^k(S_{(1)}) \\ &= \sum_{k=1}^K \rho_k \sum_{l=1}^K V_{kl} \tilde{F}_l^k(S_{(1)}). \end{aligned}$$

Then, using Assumption 3, we obtain:

$$\begin{aligned} \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] &\leq \sum_{k=1}^K \rho_k \left(V_{kk} + \sum_{l \neq k} |V_{kl}| \right) \tilde{F}_k^k(S_{(1)}) \\ &= \sum_{k=1}^K \frac{\rho_k}{\tilde{\rho}_k} \left(V_{kk} + \sum_{l \neq k} |V_{kl}| \right) \tilde{\rho}_k \tilde{F}_k^k(S_{(1)}) \\ &\leq \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \right] \sum_{k=1}^K \tilde{\rho}_k \tilde{F}_k^k(S_{(1)}) \\ &= \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \right] \tilde{F}(S_{(1)}). \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \right] &\leq \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \right] \mathbb{E} \left[\tilde{F}(S_{(1)}) \right] \\ &= \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \right] \mathbb{E} [U_{(1)}] \quad (\text{A94}) \\ &= \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \right] \frac{1}{n_{\text{cal}} + 1}. \end{aligned}$$

- Under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$, by definition of $\hat{\mathcal{I}}^{\text{marg}}$, for any $i \leq \hat{i}^{\text{marg}} - 1$,

$$\frac{i}{n_{\text{cal}}} < 1 - \left(\alpha + \hat{\Delta}(S_{(i)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right).$$

Therefore, choosing $i = \hat{i}^{\text{marg}} - 1$, we get:

$$\frac{\hat{i}^{\text{marg}}}{n_{\text{cal}}} < 1 - \left(\alpha + \widehat{\Delta}(S_{(\hat{i}^{\text{marg}}-1)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right) + \frac{1}{n_{\text{cal}}}.$$

As in the proofs of Theorems 3 and A2, the probability of coverage conditional on the labeled data in \mathcal{D} can be written as:

$$\begin{aligned} & \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \\ &= \tilde{F}(S_{(\hat{i}^{\text{marg}})}) + \Delta(S_{(\hat{i}^{\text{marg}})}) \\ &= \hat{F}(S_{(\hat{i}^{\text{marg}})}) + \Delta(S_{(\hat{i}^{\text{marg}})}) + \left[\tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \hat{F}(S_{(\hat{i}^{\text{marg}})}) \right] \\ &= \frac{\hat{i}^{\text{marg}}}{n_{\text{cal}}} + \Delta(S_{(\hat{i}^{\text{marg}})}) + \left[\tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \hat{F}(S_{(\hat{i}^{\text{marg}})}) \right] \\ &< 1 - \left(\alpha + \widehat{\Delta}(S_{(\hat{i}^{\text{marg}}-1)}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right) + \frac{1}{n_{\text{cal}}} \\ &\quad + \Delta(S_{(\hat{i}^{\text{marg}})}) + \left[\tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \hat{F}(S_{(\hat{i}^{\text{marg}})}) \right] \\ &= 1 - \alpha + \delta^{\text{marg}}(n_{\text{cal}}, n_*) + \frac{1}{n_{\text{cal}}} \\ &\quad + \left[\tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \hat{F}(S_{(\hat{i}^{\text{marg}})}) \right] + \left[\Delta(S_{(\hat{i}^{\text{marg}})}) - \widehat{\Delta}(S_{(\hat{i}^{\text{marg}}-1)}) \right]. \end{aligned}$$

The last term above can be bound by proceeding as in the proof of Theorem 3:

$$\begin{aligned} & \Delta(S_{(\hat{i}^{\text{marg}})}) - \widehat{\Delta}(S_{(\hat{i}^{\text{marg}}-1)}) \\ & \leq \sup_{t \in \mathbb{R}} |\Delta(t) - \widehat{\Delta}(t)| + (\Delta(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}}-1)})), \end{aligned}$$

where the expected value of the first term above can be bounded using Lemma A5, and the second term is given by

$$\begin{aligned} & \Delta(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}}-1)}) \\ &= \sum_{k=1}^K (\rho_k V_{kk} - \tilde{\rho}_k) \left[\tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\ &\quad + \sum_{k=1}^K \rho_k \sum_{l \neq k} V_{kl} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\ &= \sum_{k=1}^K \tilde{\rho}_k \cdot \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} \left[\tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\ &\quad + \sum_{k=1}^K \rho_k \sum_{l \neq k} V_{kl} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\ &\leq \left(\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} \right) \sum_{k=1}^K \tilde{\rho}_k \left[\tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \end{aligned}$$

$$\begin{aligned}
& + \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \max_{l \neq k} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
& = \left(\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} \right) \left[\tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
& \quad + \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \max_{l \neq k} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
& \leq \left(\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} \right) \max_{2 \leq i \leq n_{\text{cal}}} \left[\tilde{F}(S_{(i)}) - \tilde{F}(S_{(i-1)}) \right] \\
& \quad + \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \max_{l \neq k} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right],
\end{aligned}$$

where the first inequality follows from the fact that, by definition of the order statistics, $\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) \geq \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)})$ for all $l \in [K]$. By a standard result on maximum uniform spacing,

$$\begin{aligned}
\mathbb{E} \left[\max_{2 \leq i \leq n_{\text{cal}}} \left[\tilde{F}(S_{(i)}) - \tilde{F}(S_{(i-1)}) \right] \right] & = \mathbb{E} \left[\max_{1 \leq i \leq n_{\text{cal}}+1} D_i \right] \\
& = \frac{1}{n_{\text{cal}} + 1} \sum_{j=1}^{n_{\text{cal}}+1} \frac{1}{j},
\end{aligned}$$

where $D_1 = U_{(1)}$, $D_i = U_{(i)} - U_{(i-1)}$ for $i = 2, \dots, n_{\text{cal}}$, and $D_{n_{\text{cal}}+1} = 1 - U_{(n_{\text{cal}})}$. Therefore,

$$\begin{aligned}
& \mathbb{E} \left[\Delta(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
& \leq \left(\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} \right) \frac{1}{n_{\text{cal}} + 1} \sum_{j=1}^{n_{\text{cal}}+1} \frac{1}{j} \\
& \quad + \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \mathbb{E} \left[\max_{l \neq k} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \right].
\end{aligned}$$

By Assumption 2, for any $k \in [K]$,

$$\begin{aligned}
& \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \mathbb{E} \left[\max_{l \neq k} \left[\tilde{F}_l^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_l^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \right] \\
& \leq \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \mathbb{E} \left[f_{\max} \cdot \left[S_{(\hat{i}^{\text{marg}})} - S_{(\hat{i}^{\text{marg}}-1)} \right] \right] \\
& \leq \sum_{k=1}^K \rho_k \left(\sum_{l \neq k} V_{kl} \right) \mathbb{E} \left[\frac{f_{\max}}{f_{\min}} \cdot \left[\tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \right]
\end{aligned}$$

$$\begin{aligned}
&= \frac{f_{\max}}{f_{\min}} \sum_{k=1}^K \tilde{\rho}_k \cdot \frac{\rho_k}{\tilde{\rho}_k} \left(\sum_{l \neq k} V_{kl} \right) \mathbb{E} \left[\tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
&\leq \frac{f_{\max}}{f_{\min}} \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \sum_{k=1}^K \tilde{\rho}_k \mathbb{E} \left[\tilde{F}_k^k(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}_k^k(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
&= \frac{f_{\max}}{f_{\min}} \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \mathbb{E} \left[\tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
&\leq \frac{f_{\max}}{f_{\min}} \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \mathbb{E} \left[\max_{2 \leq i \leq n_{\text{cal}}} |\tilde{F}(S_{(i)}) - \tilde{F}(S_{(i-1)})| \right] \\
&= \frac{f_{\max}}{f_{\min}} \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \mathbb{E} \left[\max_{2 \leq i \leq n_{\text{cal}}} D_i \right] \\
&= \frac{f_{\max}}{f_{\min}} \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \frac{1}{n_{\text{cal}} + 1} \sum_{j=1}^{n_{\text{cal}}+1} \frac{1}{j}.
\end{aligned}$$

Therefore, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned}
&\mathbb{E} \left[\Delta(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}}-1)}) \right] \\
&\leq \left(\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} \right) \frac{1}{n_{\text{cal}} + 1} \sum_{j=1}^{n_{\text{cal}}+1} \frac{1}{j} \\
&\quad + \frac{f_{\max}}{f_{\min}} \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \frac{1}{n_{\text{cal}} + 1} \sum_{j=1}^{n_{\text{cal}}+1} \frac{1}{j} \\
&\leq \frac{\sum_{j=1}^{n_{\text{cal}}+1} 1/j}{n_{\text{cal}} + 1} \cdot \left[\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} + \frac{f_{\max}}{f_{\min}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right].
\end{aligned}$$

At this point, we have proved that

$$\begin{aligned}
& \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \mathbb{I} [\mathcal{A}_1^c \cap \mathcal{A}_2^c] \right] \\
& \leq 1 - \alpha + \delta^{\text{marg}}(n_{\text{cal}}, n_*) + \frac{1}{n_{\text{cal}}} + \mathbb{E} \left[\max_{i \in [n_{\text{cal}}]} |\tilde{F}(S_{(i)}) - \hat{F}(S_{(i)})| \right] \\
& \quad + \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \\
& \quad \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\} \\
& \quad + \frac{\sum_{j=1}^{n_{\text{cal}}+1} 1/j}{n_{\text{cal}} + 1} \cdot \left[\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} + \frac{f_{\max}}{f_{\min}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \\
& = 1 - \alpha + \delta^{\text{marg}}(n_{\text{cal}}, n_*) + \frac{1}{n_{\text{cal}}} + c(n_{\text{cal}}) \tag{A95} \\
& \quad + \frac{2 \max_{k \in [K]} \sum_{l \neq k} |V_{kl}| + \sum_{k=1}^K |\rho_k - \tilde{\rho}_k|}{\sqrt{n_*}} \\
& \quad \cdot \min \left\{ K^2 \sqrt{\frac{\pi}{2}}, \frac{1}{\sqrt{n_*}} + \sqrt{\frac{\log(2K^2) + \log(n_*)}{2}} \right\} \\
& \quad + \frac{\sum_{j=1}^{n_{\text{cal}}+1} 1/j}{n_{\text{cal}} + 1} \cdot \left[\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} + \frac{f_{\max}}{f_{\min}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right] \\
& = 1 - \alpha + 2\delta^{\text{marg}}(n_{\text{cal}}, n_*) + \frac{1}{n_{\text{cal}}} + \\
& \quad + \frac{\sum_{j=1}^{n_{\text{cal}}+1} 1/j}{n_{\text{cal}} + 1} \cdot \left[\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} + \frac{f_{\max}}{f_{\min}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right],
\end{aligned}$$

where the inequality above follows from Lemma A5 and the bound in (A82).

Finally, combining (A92) with (A93), (A94), and (A95) leads to the desired result:

$$\begin{aligned}
& \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \right] \\
& \leq \mathbb{P} [\mathcal{A}_1] + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \mathbb{I} [\mathcal{A}_2] \right] \\
& \quad + \mathbb{E} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \mathbb{I} [\mathcal{A}_1^c \cap \mathcal{A}_2^c] \right] \\
& \leq 1 - \alpha + 2\delta^{\text{marg}}(n_{\text{cal}}, n_*) + \frac{1}{n_{\text{cal}}} + \frac{1}{n_*} + \left[\max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l=1}^K |V_{kl}| \right) \right] \frac{1}{n_{\text{cal}}} \\
& \quad + \frac{\sum_{j=1}^{n_{\text{cal}}+1} 1/j}{n_{\text{cal}}} \cdot \left[\max_{k \in [K]} \frac{\rho_k V_{kk} - \tilde{\rho}_k}{\tilde{\rho}_k} + \frac{f_{\max}}{f_{\min}} \cdot \max_{k \in [K]} \left(\frac{\rho_k}{\tilde{\rho}_k} \sum_{l \neq k} V_{kl} \right) \right].
\end{aligned}$$

□

Proof of Theorem A3. This proof combines elements of the proofs of Proposition 2 and Theorem A2. Proceeding exactly as in the proof of Theorem A2, we obtain:

$$\begin{aligned}
& \mathbb{P} \left[Y_{n+1} \notin \hat{C}^{\text{marg}}(X_{n+1}) \mid \mathcal{D} \right] \\
&= 1 - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) - \Delta(S_{(\hat{i}^{\text{marg}})}) \\
&= 1 - \hat{F}(S_{(\hat{i}^{\text{marg}})}) - \max \left\{ \hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) , 0 \right\} - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \\
&\quad + \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] \\
&\quad + \max \left\{ \hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) , 0 \right\} - \left[\Delta(S_{(\hat{i}^{\text{marg}})}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \right] \\
&\leq \left[1 - \frac{\hat{i}^{\text{marg}}}{n_{\text{cal}}} - \max \left\{ \hat{\Delta}(S_{(\hat{i}^{\text{marg}})}) - \delta^{\text{marg}}(n_{\text{cal}}, n_*) , 0 \right\} \right] - \delta^{\text{marg}}(n_{\text{cal}}, n_*) \\
&\quad + \left[\hat{F}(S_{(\hat{i}^{\text{marg}})}) - \tilde{F}(S_{(\hat{i}^{\text{marg}})}) \right] + \sup_{t \in \mathbb{R}} |\hat{\Delta}(t) - \Delta(t)|,
\end{aligned}$$

using the fact that $\inf_{t \in \mathbb{R}} \Delta(t) \geq \delta(n_{\text{cal}}, n_*)$ implies

$$|\max\{\hat{\Delta}(t) - \delta(n_{\text{cal}}, n_*), 0\} - (\Delta(t) - \delta(n_{\text{cal}}, n_*))| \leq |\hat{\Delta}(t) - \Delta(t)|$$

for all $t \in \mathbb{R}$. The proof is then completed by proceeding as in the proof of Theorem A2. □

A7.7 Prediction sets with calibration-conditional coverage

Proof of Theorem A4. The proof is similar to that of Theorem 2. Suppose $Y_{n+1} = k$, for some $k \in [K]$. We assume without loss of generality that $\sum_{l=1}^K |V_{kl}| \neq 0$; otherwise, there is no label contamination and the result is trivially true. By definition of the conformity score function in (1), the event $k \notin \hat{C}(X_{n+1})$ occurs if and only if $\hat{s}(X_{n+1}, k) > \hat{\tau}_k$. We will assume without loss of generality that $\hat{\mathcal{I}}_k \neq \emptyset$ and $\hat{i}_k = \min\{i \in \hat{\mathcal{I}}_k\}$; otherwise, $\hat{\tau}_k = 1$ and the result trivially holds. As in the proof of Theorem 2, the probability of miscoverage conditional on $Y_{n+1} = k$ and on the labeled data in \mathcal{D} can be bounded from above as:

$$\begin{aligned}
& \mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \\
&\leq \alpha - \delta^{\text{cc}}(n_k, n_*, \gamma) + \sup_{t \in \mathbb{R}} [\hat{\Delta}_k(t) - \Delta_k(t)] + \sup_{t \in \mathbb{R}} [\hat{F}_k^k(t) - \tilde{F}_k^k(t)].
\end{aligned}$$

We know from Lemma A2 that, for any $\gamma_1 > 0$,

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| > 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_1)}{2n_*}} \right] \leq \gamma_1.$$

Similarly, it follows directly from the DKW inequality that, for any $\gamma_2 > 0$,

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] > \sqrt{\frac{\log(1/\gamma_2)}{2n_k}} \right] \leq \gamma_2.$$

Therefore, for any $\gamma, \gamma_1, \gamma_2 > 0$ such that $\gamma = \gamma_1 + \gamma_2$, with probability at least $1 - \gamma$,

$$\begin{aligned} & \mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \\ & \leq \alpha - \delta^{\text{cc}}(n_k, n_*, \gamma) + 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_1)}{2n_*}} + \sqrt{\frac{\log(1/\gamma_2)}{2n_k}}. \end{aligned}$$

Finally, setting

$$\gamma_1 = \frac{\gamma}{2} \cdot \frac{\sum_{l \neq k} |V_{kl}|}{\sum_{l=1}^K |V_{kl}|}, \quad \gamma_2 = \gamma \left(1 - \frac{1}{2} \cdot \frac{\sum_{l \neq k} |V_{kl}|}{\sum_{l=1}^K |V_{kl}|} \right),$$

gives the desired result; that is, with probability at least $1 - \gamma$,

$$\mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \leq \alpha.$$

□

Proof of Theorem A5. The proof follows an approach similar to that of the proof of Theorem 3. Define the events $\mathcal{A}_1 = \{\hat{\mathcal{I}}_k = \emptyset\}$ and $\mathcal{A}_2 = \{\hat{i}_k = 1\}$. Then,

$$\begin{aligned} & \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid \mathcal{D}, Y = k \right] \\ & \leq \mathbb{I}[\mathcal{A}_1] + \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \\ & \quad + \mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c]. \end{aligned} \tag{A96}$$

We will now separately bound the three terms on the right-hand-side of (A96). The following notation will be useful for this purpose. For all $i \in [n_k]$, let $U_i \sim \text{Uniform}(0, 1)$ be independent and identically distributed uniform random variables, and denote their order statistics as $U_{(1)} < U_{(2)} < \dots < U_{(n_k)}$.

- The probability of the event \mathcal{A}_1 can be bound from above as:

$$\mathbb{P} \left[\hat{\mathcal{I}}_k = \emptyset \right] \leq \gamma_2. \tag{A97}$$

To simplify the notation in the proof of (A97), define

$$d_k := \sqrt{\frac{\log(1/\gamma_1)}{2n_k}} + 4 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_2)}{2n_*}}.$$

Then, under Assumption 7,

$$\begin{aligned}
& 1 - \left(\alpha + \widehat{\Delta}_k(S_{(i)}^k) - \delta^{\text{cc}}(n_k, n_*, \gamma) \right) \\
&= 1 - \alpha - \Delta_k(S_{(i)}^k) + \delta^{\text{cc}}(n_k, n_*, \gamma) + \Delta_k(S_{(i)}^k) - \widehat{\Delta}_k(S_{(i)}^k) \\
&\leq 1 - d_k + \delta^{\text{cc}}(n_k, n_*, \gamma) + \sup_{t \in [0,1]} |\Delta_k(t) - \widehat{\Delta}_k(t)|.
\end{aligned}$$

Further, it follows from Lemma A2 and (A80) that, with probability at least $1 - \gamma_2$,

$$\begin{aligned}
& 1 - \left(\alpha + \widehat{\Delta}_k(S_{(i)}^k) - \delta^{\text{cc}}(n_k, n_*, \gamma) \right) \\
&\leq 1 - d_k + \delta^{\text{cc}}(n_k, n_*, \gamma) + 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_2)}{2n_*}} \\
&\leq 1 - d_k + \sqrt{\frac{\log(1/\gamma_1)}{2n_k}} + 4 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\gamma_2)}{2n_*}} \\
&= 1.
\end{aligned}$$

This implies that the set $\widehat{\mathcal{I}}_k$ defined in (15) is non-empty with probability at least $1 - \gamma_2$, because $n_k \in \widehat{\mathcal{I}}_k$ if $1 - [\alpha + \widehat{\Delta}_k(S_{(i)}^k) - \delta^{\text{cc}}(n_k, n_*, \gamma)] \leq 1$.

- Under \mathcal{A}_2 , the second term on the right-hand-side of (A96) can be written as:

$$\begin{aligned}
\mathbb{P} \left[Y_{n+1} \in \widehat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] &= \mathbb{P} \left[\hat{s}(X_{n+1}, k) \leq S_{(1)}^k \mid Y = k, \mathcal{D} \right] \\
&= F_k^k(S_{(1)}^k).
\end{aligned}$$

Therefore, using Assumption 3, we can write that, for any $a > 0$,

$$\begin{aligned}
& \mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \widehat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \geq a \right] \\
&\leq \mathbb{P} \left[F_k^k(S_{(1)}^k) \geq a \right] \\
&\leq \mathbb{P} \left[\left(V_{kk} + \sum_{l \neq k} |V_{kl}| \right) \tilde{F}_k^k(S_{(1)}^k) \geq a \right] \\
&= \mathbb{P} \left[U_{(1)} \geq \frac{a}{V_{kk} + \sum_{l \neq k} |V_{kl}|} \right] \\
&= \left(1 - \frac{a}{V_{kk} + \sum_{l \neq k} |V_{kl}|} \right)^{n_k} \\
&\leq \frac{1}{n_k} \cdot \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{a}.
\end{aligned}$$

In particular, choosing $a = (V_{kk} + \sum_{l \neq k} |V_{kl}|)/(n_k \bar{\gamma})$, for any $\bar{\gamma} \in (0, 1)$, gives us:

$$\mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \widehat{C}^{\text{cc}}(X_{n+1}) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_2] \leq \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k \bar{\gamma}} \right] \geq 1 - \bar{\gamma}. \quad (\text{A98})$$

- By proceeding as in the proof of Theorem 3, using Assumption 2 we obtain that, under $\mathcal{A}_1^c \cap \mathcal{A}_2^c$,

$$\begin{aligned}
& \mathbb{P} \left[Y \in C(X, S_{(i_k)}^k) \mid Y = k, \mathcal{D} \right] \\
& < 1 - \alpha + \delta^{\text{cc}}(n_k, n_*, \gamma) + \frac{1}{n_k} + \sup_{t \in \mathbb{R}} \left[\tilde{F}_k^k(t) - \hat{F}_k^k(t) \right] + \sup_{t \in \mathbb{R}} |\Delta_k(t) - \hat{\Delta}_k(t)| \\
& + 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \max_{2 \leq i \leq n_k} \left| \tilde{F}_k^k(S_{(i)}^k) - \tilde{F}_k^k(S_{(i-1)}^k) \right|.
\end{aligned} \tag{A99}$$

Now, recall that for any $\bar{\gamma} \in (0, 1)$, by the DKW inequality,

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right] > \sqrt{\frac{\log(1/\bar{\gamma})}{2n_k}} \right] \leq \bar{\gamma}.$$

Similarly, we know from Lemma A2 that, for any $\bar{\gamma} \in (0, 1)$,

$$\mathbb{P} \left[\sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| > 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(1/\bar{\gamma})}{2n_*}} \right] \leq \bar{\gamma}.$$

It remains to bound the last term on the right-hand-side of (A99). Note that, for any $\bar{\gamma} \in (0, 1)$,

$$\mathbb{P} \left[\max_{2 \leq i \leq n_k} \left| \tilde{F}_k^k(S_{(i)}^k) - \tilde{F}_k^k(S_{(i-1)}^k) \right| > \bar{\gamma} \right] = \mathbb{P} \left[\max_{1 \leq i \leq n_k+1} D_i > \bar{\gamma} \right],$$

where $U_i \sim \text{Uniform}(0, 1)$ are independent and identically distributed uniform random variables for all $i \in [n_k]$, and $D_1 = U_{(1)}$, $D_i = U_{(i)} - U_{(i-1)}$ for $i = 2, \dots, n_k$, and $D_{n_k+1} = 1 - U_{(n_k)}$.

Further, by combining the Markov inequality with a standard result on the asymptotic behavior of the maximum uniform spacing, we obtain that, for any $a > 0$,

$$\mathbb{P} \left[\max_{1 \leq i \leq n_k+1} D_i > a \right] \leq \frac{1}{a} \cdot \mathbb{E} \left[\max_{1 \leq i \leq n_k+1} D_i \right] = \frac{1}{a} \cdot \frac{1}{n_k} \sum_{j=1}^{n_k+1} \frac{1}{j}.$$

Therefore, setting $a = (\sum_{j=1}^{n_k+1} 1/j)/(n_k \bar{\gamma})$ for any $\bar{\gamma} > 0$, we obtain that

$$\mathbb{P} \left[\max_{1 \leq i \leq n_k+1} D_i > \frac{1}{\bar{\gamma}} \cdot \frac{1}{n_k} \sum_{j=1}^{n_k+1} \frac{1}{j} \right] \leq \bar{\gamma}.$$

With a union bound, this implies that, for any $\bar{\gamma} \in (0, 1)$,

$$\mathbb{P} \left[\mathbb{P} \left[Y \in C(X, S_{(i_k)}^k) \mid Y = k, \mathcal{D} \right] \mathbb{I}[\mathcal{A}_1^c \cap \mathcal{A}_2^c] > 1 - \alpha + \omega(n_k, n_*, \gamma, \bar{\gamma}) \right] \leq \bar{\gamma}, \tag{A100}$$

where

$$\begin{aligned}
& \omega(n_k, n_*, \gamma, \bar{\gamma}) \\
&= \delta^{\text{cc}}(n_k, n_*, \gamma) + \frac{1}{n_k} + \sqrt{\frac{\log(3/\bar{\gamma})}{2n_k}} + 2 \sum_{l \neq k} |V_{kl}| \sqrt{\frac{\log(2K) + \log(3/\bar{\gamma})}{2n_*}} \\
&+ 2 \sum_{l \neq k} |V_{kl}| \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{n_k} \cdot \left[\log(n_k + 1) + \frac{3}{\bar{\gamma}} \sum_{j=1}^{n_k+1} \frac{1}{j} \right].
\end{aligned}$$

Finally, combining (A96) with (A97), (A98), and (A100) leads to

$$\begin{aligned}
\mathbb{P} \left[\mathbb{P} \left[Y_{n+1} \in \hat{C}^{\text{cc}}(X_{n+1}) \mid \mathcal{D}, Y = k \right] \geq 1 - \alpha + \omega(n_k, n_*, \gamma, \bar{\gamma}) + \frac{V_{kk} + \sum_{l \neq k} |V_{kl}|}{n_k \bar{\gamma}} \right] \\
\leq \gamma_2 + 2\bar{\gamma}.
\end{aligned}$$

Thus, the desired result is obtained by setting

$$\bar{\gamma} = \frac{\gamma - \gamma_2}{2} = \frac{\gamma}{2} \left(1 - \frac{1}{2} \cdot \frac{\sum_{l \neq k} |V_{kl}|}{\sum_{l=1}^K |V_{kl}|} \right).$$

□

Proof of Theorem A4. This proof combines elements of the proofs of Proposition 2 and Theorem A4. Suppose $Y_{n+1} = k$, for some $k \in [K]$. Proceeding exactly as in the proof of Theorem A4, we obtain:

$$\begin{aligned}
& \mathbb{P} \left[Y \notin C(X, S_{(\hat{i}_k)}^k) \mid Y = k, \mathcal{D} \right] \\
& \leq \left[1 - \frac{\hat{i}_k}{n_k} - \max \left\{ \hat{\Delta}_k(S_{(\hat{i}_k)}) - \delta^{\text{cc}}(n_k, n_*), -\sqrt{\frac{\log(1/\gamma)}{2n_k}} \right\} \right] - \delta^{\text{cc}}(n_k, n_*) \\
& \quad + \left[\hat{F}_k^k(S_{(\hat{i}_k)}) - \tilde{F}_k^k(S_{(\hat{i}_k)}) \right] \\
& \quad + \max \left\{ \hat{\Delta}_k(S_{(\hat{i}_k)}) - \delta^{\text{cc}}(n_k, n_*), -\sqrt{\frac{\log(1/\gamma)}{2n_k}} \right\} - \left[\Delta_k(S_{(\hat{i}_k)}) - \delta^{\text{cc}}(n_k, n_*) \right] \\
& = \left[1 - \frac{\hat{i}_k}{n_k} - \max \left\{ \hat{\Delta}_k(S_{(\hat{i}_k)}) - \delta^{\text{cc}}(n_k, n_*), -\sqrt{\frac{\log(1/\gamma)}{2n_k}} \right\} \right] - \delta^{\text{cc}}(n_k, n_*) \\
& \quad + \left[\hat{F}_k^k(S_{(\hat{i}_k)}) - \tilde{F}_k^k(S_{(\hat{i}_k)}) \right] \\
& \quad + \max \left\{ \hat{\Delta}_k(S_{(\hat{i}_k)}) + \sqrt{\frac{\log(1/\gamma)}{2n_k}} - \delta^{\text{cc}}(n_k, n_*), 0 \right\}
\end{aligned}$$

$$\begin{aligned}
& - \left[\Delta_k(S_{(\hat{i}_k)}) + \sqrt{\frac{\log(1/\gamma)}{2n_k}} - \delta^{\text{cc}}(n_k, n_*) \right] \\
& \leq \alpha - \delta^{\text{cc}}(n_k, n_*, \gamma) + \sup_{t \in \mathbb{R}} |\hat{\Delta}_k(t) - \Delta_k(t)| + \sup_{t \in \mathbb{R}} \left[\hat{F}_k^k(t) - \tilde{F}_k^k(t) \right].
\end{aligned}$$

using the fact that $\inf_{t \in \mathbb{R}} \Delta_k(t) \geq \delta^{\text{cc}}(n_k, n_*) - \sqrt{\log(1/\gamma)/(2n_k)}$ implies

$$\begin{aligned}
& \left| \max \left\{ \hat{\Delta}_k(S_{(\hat{i}_k)}) + \sqrt{\frac{\log(1/\gamma)}{2n_k}} - \delta^{\text{cc}}(n_k, n_*), 0 \right\} - \left[\Delta_k(t) + \sqrt{\frac{\log(1/\gamma)}{2n_k}} - \delta^{\text{cc}}(n_k, n_*) \right] \right| \\
& \leq |\hat{\Delta}_k(t) - \Delta_k(t)|
\end{aligned}$$

for all $t \in \mathbb{R}$. □

Comparison to worst-case coverage bounds

Proof of Corollary 2. The expression for $\Delta_k(t)$ in (12) can be equivalently rewritten as

$$\Delta_k(t) = \sum_{l \neq k} V_{kl} \left(\tilde{F}_l^k(t) - \tilde{F}_k^k(t) \right),$$

which implies

$$-\sum_{l \neq k} |V_{kl}| \leq \Delta_k(t) \leq \sum_{l \neq k} |V_{kl}|, \quad \forall t \in \mathbb{R}.$$

The proof is then completed by applying Theorem 1. □

A8 Supplementary numerical results

A8.1 Simulations under a known label contamination model

A8.1.1 Additional views and performance metrics




Figure A5: Performances of different conformal methods on simulated data with randomly contaminated labels, as a function of the number of calibration samples. The reported empirical coverage and average size of the prediction sets are stratified based on the true label of the test points. The strength parameter of the label contamination process is $\epsilon = 0.1$. Other details are as in Figure 2.




Figure A6: Performances of different conformal methods on simulated data with varying numbers of calibration samples, as a function of the label contamination strength. Other details are as in Figure 2.

A8.1.2 The effect of the number of classes

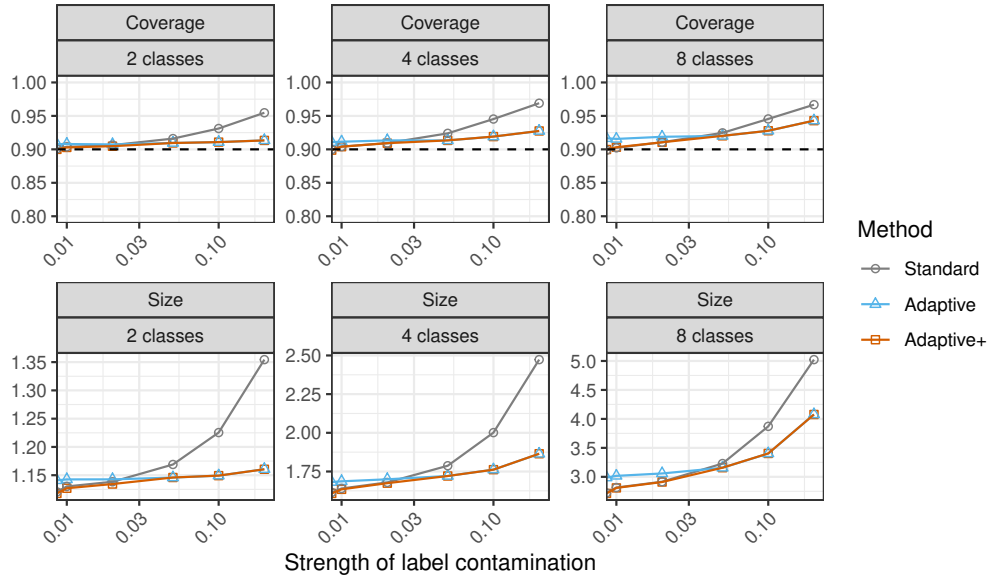


Figure A7: Performances of different conformal methods on simulated classification data with different numbers of labels, as a function of the contamination strength. The strength parameter of the label contamination process is $\epsilon = 0.1$. Other details are as in Figure A6.

A8.1.3 The effect of the classifier




Figure A8: Performances of different conformal methods on simulated data using different machine learning classifiers, as a function of the number of calibration samples. The strength parameter of the label contamination process is $\epsilon = 0.1$. Other details are as in Figure 2.

A8.1.4 The effect of the data distribution




Figure A9: Performances of different conformal methods on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. All methods guarantee 90% label-conditional coverage. The data are simulated from a logistic model with random parameters. Other details are as in Figure 2.




Figure A10: Performances of different conformal methods on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. The data are simulated from a heteroscedastic decision-tree model. Other details are as in Figure 2.

A8.1.5 The effect of the label contamination process




Figure A11: Performances of different conformal methods on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. The contamination process has a block-like structure. Other details are as in Figure A9.

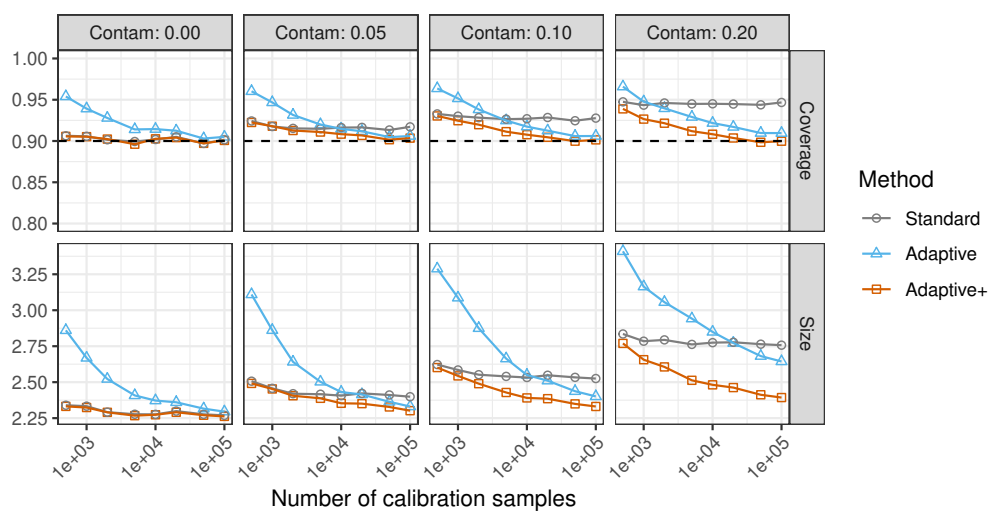


Figure A12: Performances of different conformal methods on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. The label contamination process has a random heterogeneous structure. Other details are as in Figure A9.

A8.1.6 Prediction sets with marginal coverage




Figure A13: Performances of different conformal prediction methods with marginal coverage on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. The dashed horizontal line indicates the 90% nominal marginal coverage level. Other details are as in Figure 2.




Figure A14: Performances of different conformal prediction methods with marginal coverage on simulated data with varying numbers of calibration samples, as a function of the label contamination strength. The dashed horizontal line indicates the 90% nominal marginal coverage level. Other details are as in Figure A6.




Figure A15: Performances of different conformal methods on simulated data with label contamination. The label contamination process follows a classical randomized response model. The Adaptive+ method is applied with and without perfect knowledge of the contaminated label frequencies. Note that there is no significant difference in performance here, because it is easy to estimate $\tilde{\rho}$ accurately from the available data. The nominal marginal coverage level is 90%. Other details are as in Figure A9.




Figure A16: Performances of different conformal methods on simulated data with label contamination. The label contamination process has a random heterogeneous structure. The Adaptive+ method is applied with and without perfect knowledge of the contaminated label frequencies. Note that there is no significant difference in performance here, because it is easy to estimate $\tilde{\rho}$ accurately from the available data. The nominal marginal coverage level is 90%. Other details are as in Figure A15.

A8.1.7 Prediction sets with calibration-conditional coverage




Figure A17: Performances of different conformal methods on simulated data with random label contamination of varying strength, as a function of the number of calibration samples. All methods guarantee 90% calibration and label-conditional coverage with probability at least 90% over the calibration data. Other details are as in Figure A13.

A8.2 Simulations under a bounded label contamination model

A8.2.1 Randomized response model

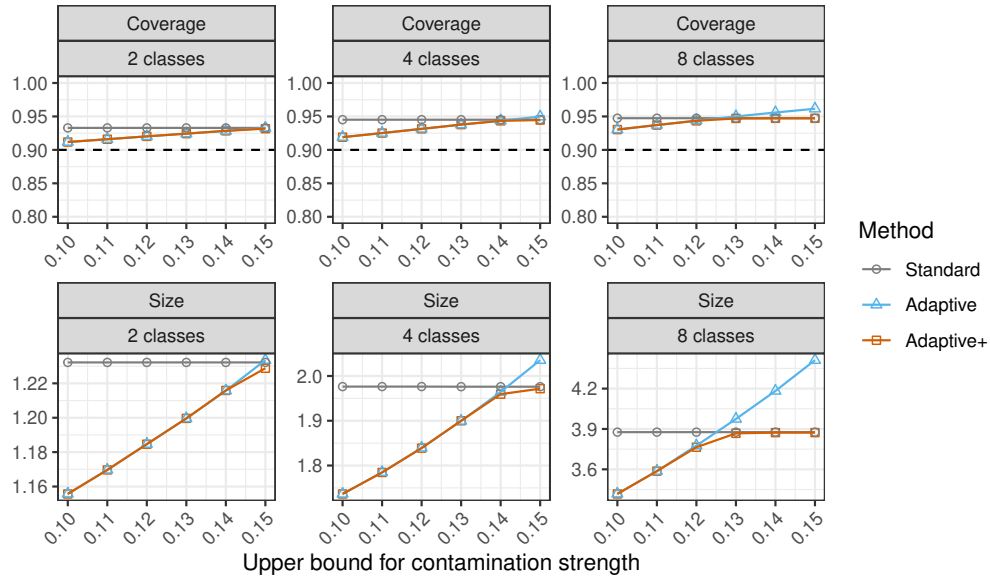


Figure A18: Performances of different conformal methods on simulated data with random label contamination based on a randomized response model with unknown contamination strength parameter ϵ . The results are shown as a function of the known upper confidence bound for ϵ , whose true value is $\epsilon = 0.1$. Other details are as in Figure 3.

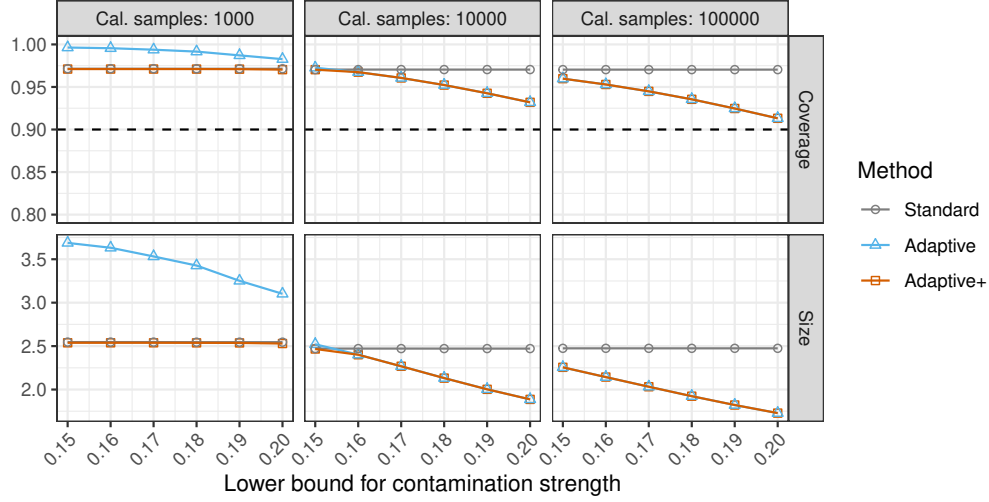


Figure A19: Performances of different conformal methods on simulated data with random label contamination, as a function of the known lower confidence bound for the contamination strength parameter $\epsilon = 0.2$. The results are stratified based on the numbers of calibration samples. Other details are as in Figure 3.




Figure A20: Performances of different conformal methods on simulated data with random label contamination, as a function of the known lower confidence bound for the contamination strength parameter $\epsilon = 0.1$. The results are stratified based on the numbers of calibration samples. Other details are as in Figure 3.

A8.2.2 Two-level randomized response model




Figure A21: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model with partly unknown parameters. The Adaptive+ method is implemented by applying the specialized version of Algorithm 2 described in Section A3.2, based on a 99% confidence interval $[\hat{\epsilon}^{\text{upp}}, \epsilon]$ for $\epsilon = 0.2$ and a degenerate interval $[\nu, \nu]$ for $\nu \in \{0, 0.25, 0.5, 0.75, 1\}$. The results are shown as a function of $\hat{\epsilon}^{\text{upp}}$. The calibration set size is 10,000. Other details are as in Figure 3.




Figure A22: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The Adaptive+ method is applied based on a 99% confidence interval $[\hat{\epsilon}^{\text{upp}}, \epsilon]$ for $\epsilon = 0.2$ and a degenerate interval $[\nu, \nu]$ for $\nu \in \{0, 0.25, 0.5, 0.75, 1\}$. The results are shown as a function of $\hat{\epsilon}^{\text{upp}}$. The calibration set size is 100,000. Other details are as in Figure A21.

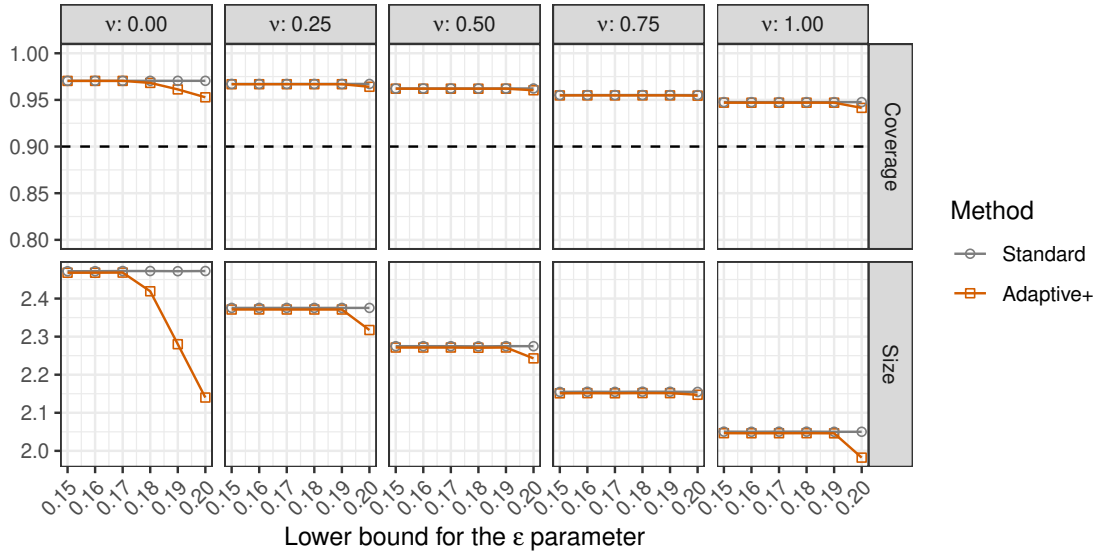


Figure A23: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The Adaptive+ method is applied based on a 99% confidence interval $[\hat{\epsilon}^{\text{upp}}, \epsilon]$ for $\epsilon = 0.2$ and an interval $[\max\{0, \nu - 0.02\}, \min\{\nu + 0.02, 1\}]$ for $\nu \in \{0, 0.25, 0.5, 0.75, 1\}$. The results are shown as a function of $\hat{\epsilon}^{\text{upp}}$. The calibration set size is 10,000. Other details are as in Figure A21.




Figure A24: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The Adaptive+ method is applied based on a 99% confidence interval $[\hat{\epsilon}^{\text{upp}}, \epsilon]$ for $\epsilon = 0.2$ and an interval $[\max\{0, \nu - 0.02\}, \min\{\nu + 0.02, 1\}]$ for $\nu \in \{0, 0.25, 0.5, 0.75, 1\}$. The results are shown as a function of $\hat{\epsilon}^{\text{upp}}$. The calibration set size is 100,000. Other details are as in Figure A23.

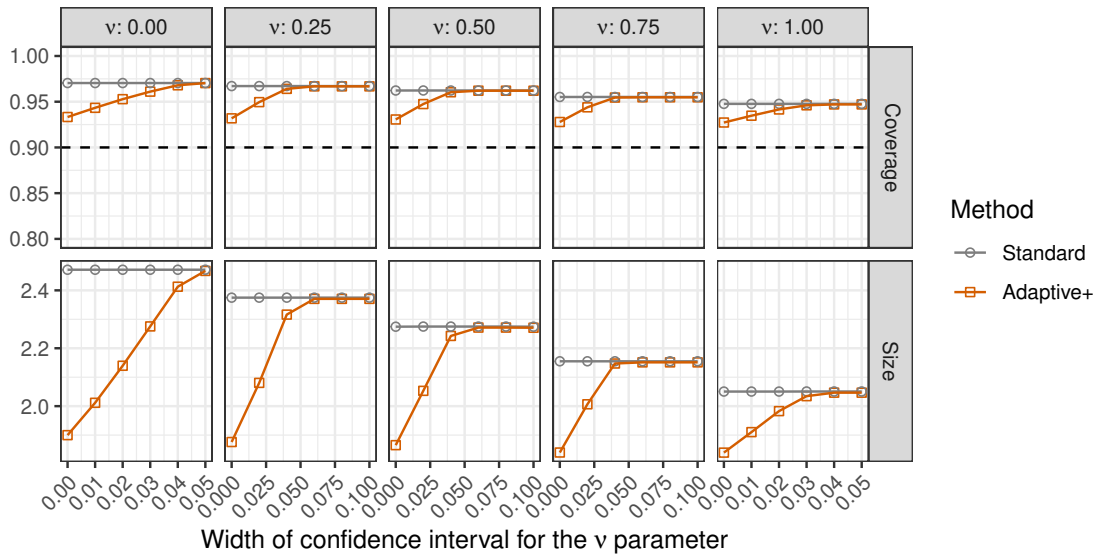


Figure A25: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The Adaptive+ method is implemented by applying the specialized version of Algorithm 2 described in Section A3.2, based on a 99% symmetric confidence interval $[\hat{\nu}^{\text{low}}, \hat{\nu}^{\text{upp}}]$ for $\nu \in \{0, 0.25, 0.5, 0.75, 1\}$ and a degenerate interval $[\epsilon, \epsilon]$ for $\epsilon = 0.2$. The results are shown as a function of $\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}}$. The calibration set size is 10,000. Other details are as in Figure A21.




Figure A26: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model with partly unknown parameters. The calibration set size is 100,000. Other details are as in Figure A25.




Figure A27: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The Adaptive+ method is implemented by applying the specialized version of Algorithm 2 described in Section A3.2, based on a 99% symmetric confidence interval $[\hat{\nu}^{\text{low}}, \hat{\nu}^{\text{upp}}]$ for $\nu \in \{0, 0.25, 0.5, 0.75, 1\}$ and a degenerate interval $[\epsilon - 0.02, \epsilon]$ for $\epsilon = 0.2$. The results are shown as a function of $\hat{\nu}^{\text{upp}} - \hat{\nu}^{\text{low}}$. The calibration set size is 10,000. Other details are as in Figure A21.




Figure A28: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model with partly unknown parameters. The calibration set size is 100,000. Other details are as in Figure A27.

A8.3 Robustness to model estimation

A8.3.1 Randomized response model

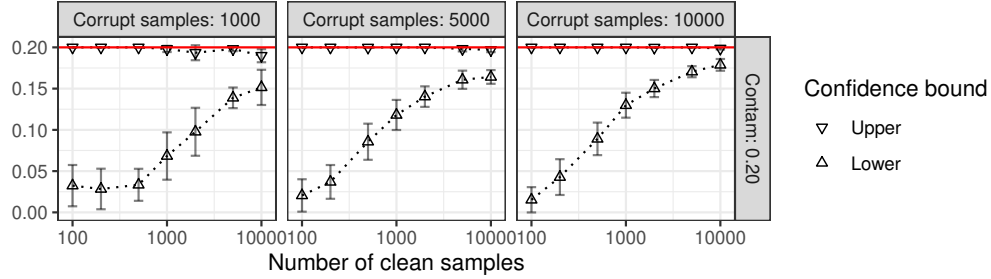


Figure A29: Bootstrap confidence intervals for the parameter ϵ of a randomized response model describing the label contamination process. The results are shown as a function of the number of clean and contaminated samples used to fit the model. These confidence intervals are utilized by the *Adaptive+ (CI)* method in Figure 4.




Figure A30: Performances of different conformal prediction methods, as a function of the numbers of clean and contaminated samples used to fit a randomized response label contamination model. The fixed upper bound for $\epsilon = 0.2$ utilized by the *Adaptive+ (CI)* method, namely Algorithm 2, is equal to $\bar{\epsilon} = 0.25$. Other details are as in Figure 4.




Figure A31: Bootstrap confidence intervals for the parameter ϵ of a randomized response model describing the label contamination process. The results are shown as a function of the number of clean and contaminated samples used to fit the model. These confidence intervals are utilized by the *Adaptive+ (CI)* method in Figure A30.




Figure A32: Performances of different conformal methods on simulated data with random label contamination following a randomized response model. The results are shown as a function of the number of clean samples used to fit the noise parameter ϵ . The results are stratified based on the true value of ϵ . Other details are as in Figure 4.




Figure A33: Performances of different conformal methods on simulated data with random label contamination following a randomized response model. The results are shown as a function of the number of clean samples used to fit the noise parameter ϵ . The number of possible classes is $K = 4$. Other details are as in Figure A32.




Figure A34: Performances of different conformal methods on simulated data with random label contamination following a randomized response model. The results are shown as a function of the number of clean samples used to fit the noise parameter ϵ . The number of possible classes is $K = 8$. Other details are as in Figure A32.

A8.3.2 Two-level randomized response model

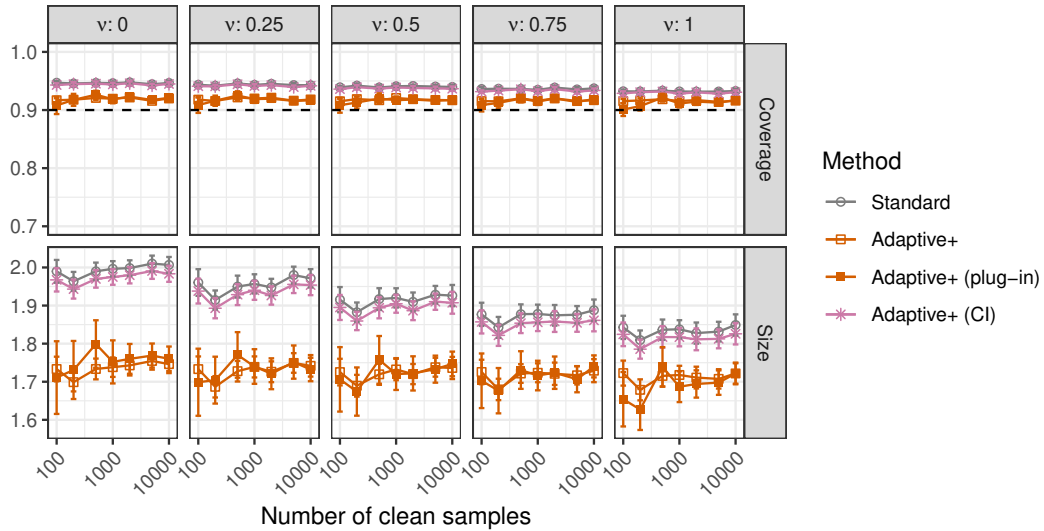


Figure A35: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The results are shown as a function of the number of clean samples used to fit the parameters ϵ and ν , for different true values of ν . The true parameter ϵ is equal to 0.1. The number of possible classes is $K = 4$. Other details are as in Figure A33.

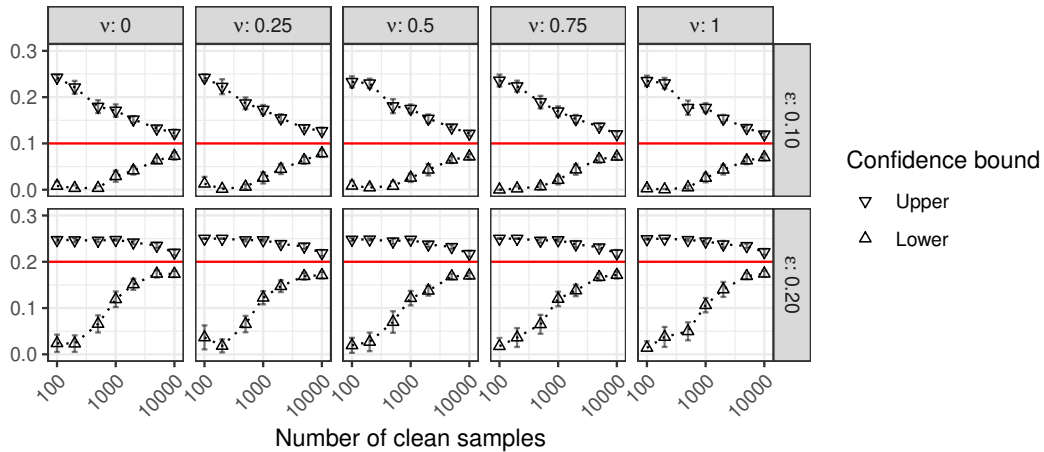


Figure A36: Bootstrap confidence intervals for the parameter ϵ in a two-level randomized response model for the label contamination process. The results are shown as a function of the number of clean samples used to fit the model, stratified based on the true values of the model parameters ϵ and ν . The number of contaminated training samples is 10,000. These confidence intervals correspond to those utilized by the *Adaptive+ (CI)* method in Figure A35.




Figure A37: Bootstrap confidence intervals for the parameter ν in a two-level randomized response model for the label contamination process. Other details are as in Figure A36.

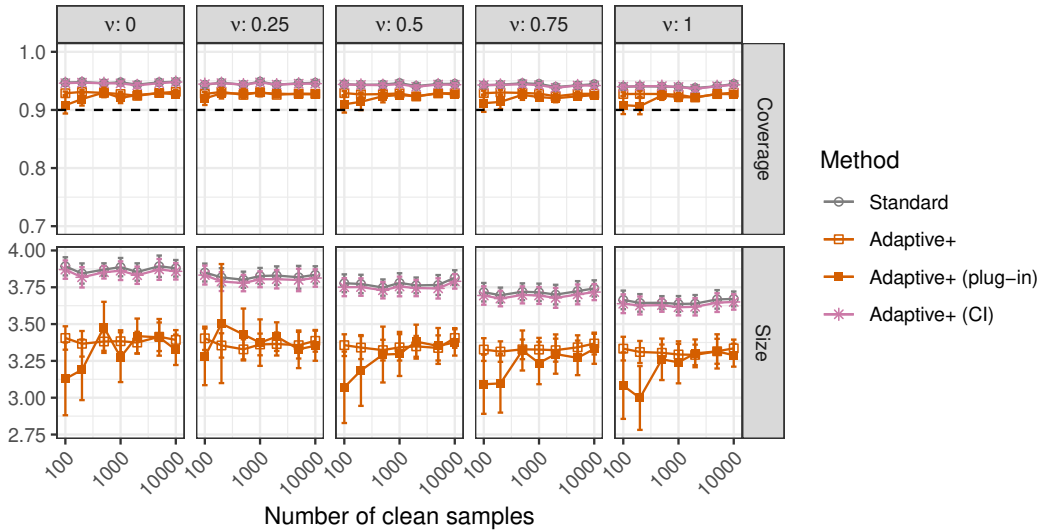


Figure A38: Performances of different conformal methods on simulated data with random label contamination from a two-level randomized response model. The results are shown as a function of the number of clean samples used to fit the parameters ϵ and ν , for different true values of ν . The true parameter ϵ is equal to 0.1. The number of possible classes is $K = 8$. Other details are as in Figure A35.

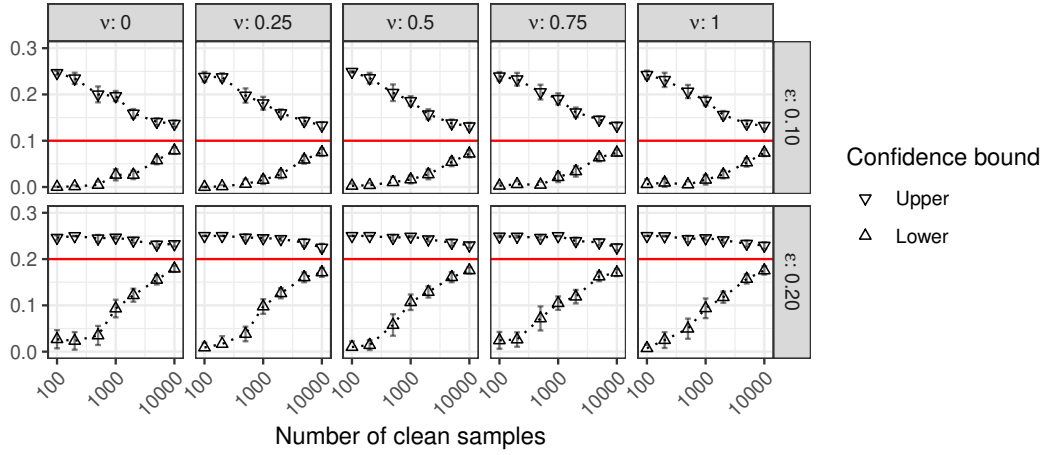


Figure A39: Bootstrap confidence intervals for the parameter ϵ in a two-level randomized response model for the label contamination process. The results are shown as a function of the number of clean samples used to fit the model, stratified based on the true values of the model parameters ϵ and ν . The number of contaminated training samples is 10,000. The number of possible labels is $K = 8$. These confidence intervals correspond to those utilized by the *Adaptive+* (CI) method in Figure A38.

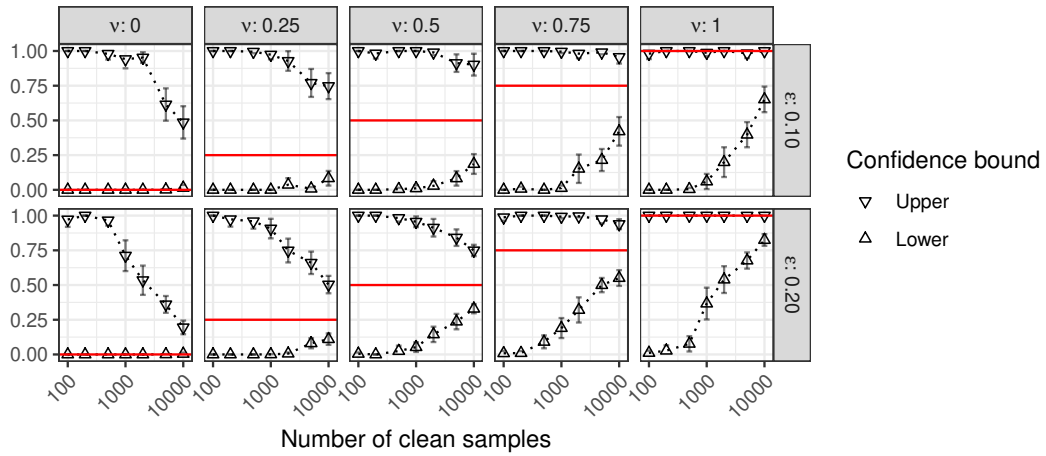


Figure A40: Bootstrap confidence intervals for the parameter ν in a two-level randomized response model for the label contamination process. The number of possible labels is $K = 8$. Other details are as in Figure A39.

A8.4 Robustness to model mis-specification

This section demonstrates the robustness of our adaptive methods to other types of mis-specifications in the label contamination process, going beyond the issue of estimating the noise parameter ϵ of the randomized response model described in Section A3.1. For this purpose, we generate synthetic data with $K = 4$ possible labels as explained in Section 4.1, but using a label contamination process with a block-like structure. That is, the contamination process is described by a transition matrix $T \in [0, 1]^{K \times K}$, defined as $T_{kl} = \mathbb{P}[\tilde{Y} = k \mid X, Y = l]$ for all $l, k \in [K]$, given by $T = (1 - \epsilon)I_K + \epsilon/K \cdot B_{K,2}$, where $B_{K,2}$ is a block-diagonal matrix with $K/2$ constant blocks equal to J_2 —the 2×2 matrix of ones. Then, we apply the optimistic version of Algorithm 1 under the mis-specified assumption that $T = (1 - \epsilon)I_K + \epsilon/K \cdot J_K$, separately using either the correct value of ϵ or a plug-in empirical estimate obtained from an independent model-fitting data set as in Section A3.1.4. Consistently with the previous section, we refer to the former method as *Adaptive+* and to the latter as *Adaptive+ (plug-in)*. Figure A41 compares the performances of our *Adaptive+* and *Adaptive+ (plug-in)* to that of the standard conformal inference approach, as a function of the true ϵ and of the number of clean model-fitting data points used to estimate this parameter. The number of calibration data points here is set equal to 10,000. The results show that both methods are quite robust to model mis-specification and can generally produce more informative prediction sets compared to the standard benchmark.




Figure A41: Performances of different conformal methods on simulated data, as a function of the number of clean samples used to estimate a mis-specified contamination model. The true model matrix M has a block-diagonal structure, and the results are stratified based on the contamination strength. Other details are as in Figure 2.

Finally, Figure A42 demonstrates that our methods also enjoy similar robustness under different label contamination processes, focusing specifically on a transition matrix $T = (1 - \epsilon)I_K + \epsilon/K \cdot U_K$, where U_K is a matrix of i.i.d. uniform random numbers on $[0, 1]$, standardized to have its columns sum to one.




Figure A42: Performances of different conformal methods on simulated data, as a function of the number of clean samples used to estimate a mis-specified contamination model. The contamination model matrix M has a random structure. Other details are as in Figure A41.

A8.5 Demonstrations with CIFAR-10 image data

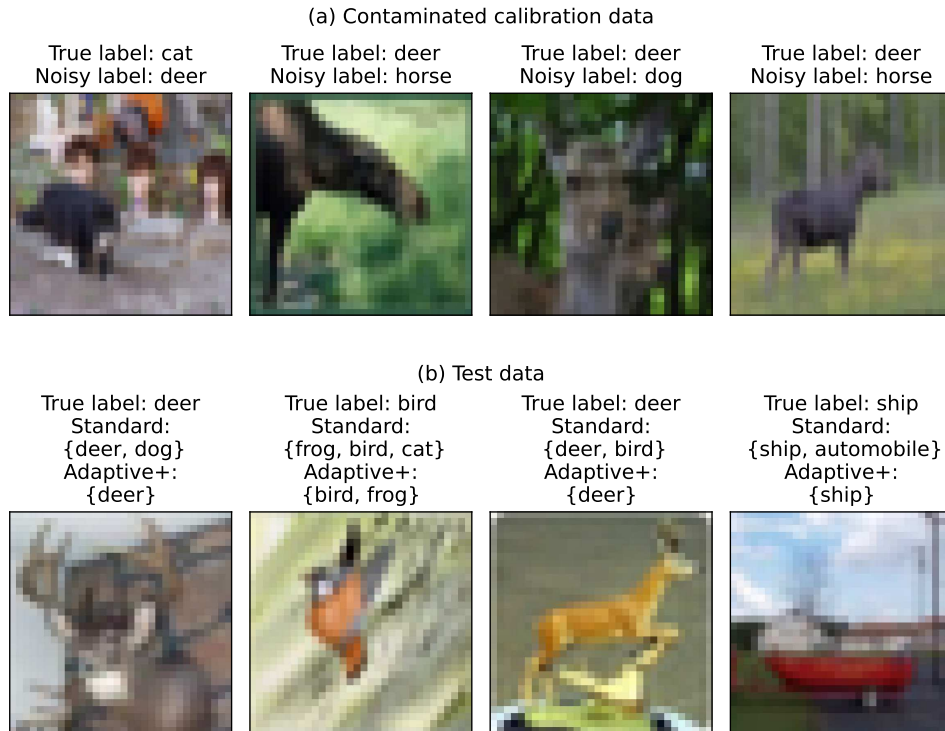


Figure A43: Demonstration of some CIFAR-10 images with contaminated labels (a), and of conformal prediction sets (b) obtained with and without correcting for the presence of label contamination in the calibration data, as in the experiments of Figure 1.




Figure A44: Performances of different conformal methods on CIFAR-10 image data with noisy human-assigned labels, as a function of calibration set size. The horizontal dashed line indicates the nominal 90% label-conditional coverage level. Other details are as in Figure 1.

A9 Additional details on numerical experiments

A9.1 Experiments with heteroscedastic decision-tree model

We set $d = 50$ and generate each sample of features $X \in \mathbb{R}^d$ independently as follows: $X_1 = +1$ w.p. $3/4$, and $X_1 = -1$ w.p. $1/4$; $X_2 = +1$ w.p. $3/4$, and $X_2 = -2$ w.p. $1/4$; $X_3 = +1$ w.p. $1/4$, and $X_3 = -2$ w.p. $1/2$; X_4 is uniformly distributed on $\{1, \dots, 4\}$; and $X_j \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$ for all $j \in \{5, \dots, d\}$. The labels Y belong to one of $K = 4$ possible classes, and their conditional distribution given $X = x$ is given by the decision tree shown in Figure A45, which only depends on the first four features.




Figure A45: A toy heteroscedastic decision-tree model for $P_{Y|X}$ used to generate synthetic in some numerical experiments.