



HAL
open science

A Unified Approach to Fair Online Learning via Blackwell Approachability

Evgenii Chzhen, Christophe Giraud, Gilles Stoltz

► **To cite this version:**

Evgenii Chzhen, Christophe Giraud, Gilles Stoltz. A Unified Approach to Fair Online Learning via Blackwell Approachability. 35th Conference on Neural Information Processing Systems (NeurIPS 2021), 2021, Virtual conference, Australia. hal-03268279v2

HAL Id: hal-03268279

<https://hal.science/hal-03268279v2>

Submitted on 15 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Unified Approach to Fair Online Learning via Blackwell Approachability

Evgenii Chzhen Christophe Giraud Gilles Stoltz

Université Paris-Saclay, CNRS, Laboratoire de mathématiques d’Orsay, 91405, Orsay, France
{evgenii.chzhen, christophe.giraud, gilles.stoltz}@universite-paris-saclay.fr

Abstract

We provide a setting and a general approach to fair online learning with stochastic sensitive and non-sensitive contexts. The setting is a repeated game between the Player and Nature, where at each stage both pick actions based on the contexts. Inspired by the notion of unawareness, we assume that the Player can only access the non-sensitive context before making a decision, while we discuss both cases of Nature accessing the sensitive contexts and Nature unaware of the sensitive contexts. Adapting Blackwell’s approachability theory to handle the case of an unknown contexts’ distribution, we provide a general necessary and sufficient condition for learning objectives to be compatible with some fairness constraints. This condition is instantiated on (group-wise) no-regret and (group-wise) calibration objectives, and on demographic parity as an additional constraint. When the objective is not compatible with the constraint, the provided framework permits to characterise the optimal trade-off between the two.

1 Introduction

Classically, the goal of the decision maker in sequential environment is purely performance driven — she wants to obtain as high reward as if she has had a complete information about the environment. In contrast, algorithmic fairness shifts the attention from the performance-driven behavior by taking into account additional ethical considerations. The latter is often formalized via the notion of fairness constraint [10, 26, 6] on the decision maker’s strategies. The goal of this work is to bring to light Blackwell’s approachability theory as a suitable theoretical formalism for fair online learning under *group fairness* constraints. The appealing feature of this theory is two-fold: first, it gives explicit criteria when learning is possible; second, if this criteria is met, it comes with an explicit strategy.

It is well known that Blackwell’s approachability theory may be used to characterize online learning problems that are tractable and to design strategies to solve them—for instance, for no-regret learning or calibration. Extensive references to such uses may be found in Cesa-Bianchi and Lugosi [5], Perchet [23] and Abernethy et al. [1]. Actually, as noted by the latter two references, no-regret learning, calibration, and approachability imply each other in some sense. The main first achievement of this article is to extend this use to online learning under fairness constraints. This idea, though natural and intuitive, requires some extensions to Blackwell’s approachability theory, like ignoring the target set and having to estimate it.

Related works in fair online learning. Several frameworks have been proposed to tackle various problems of fairness arising in online learning. Blum et al. [4] consider the problem of online prediction with experts and define fairness via (approximate) equality of average payoffs. Hébert-Johnson et al. [16], Gupta et al. [14] consider the problem of group-wise calibration. (In passing, we may note that Gupta et al. [14] consider some techniques with a flavor of approachability.) Bechavod et al. [2] consider the problem of online binary classification with partial feedback and equal opportunity

constraint [15]. We treat the above works as sources of inspiration; they all differ in the specific (sensitive and non-sensitive) information that the Player may or may not access before or after taking an action. We apply the general formalism of approachability theory to give new insights into online learning under fairness constraints, and approach this goal in a unified (and geometric) way. In particular, the generality of this formalism allows to derive (im)possibility results nearly effortlessly. But we also go beyond such a mere compatibility/incompatibility check between the learning objectives and fairness constraints, and note that approachability theory also gives a clear strategy for the study of trade-offs between incompatible learning objectives and fairness constraints, which often arise in batch setup [6].

Outline. We describe our approachability setting in Section 2 and provide some learning objectives (no-regret and calibration) and fairness constraints (group-wise controls, demographic parity, equalized average payoffs) that fit our framework. A slight extension of the classical result of Blackwell [3] is required and discussed in Section 3. We then support the generality of our framework by deriving (im)possibility results for some objective–constraint pairs in Section 4. We also illustrate in Section 5 how this formalism can be used to derive optimal trade-offs (Pareto frontiers) between performance and fairness for incompatible objective–constraint pairs; as an example, we deal with group-wise calibration (studied by [16, 14]) under demographic parity constraint. For the sake of exposition, we deal in Sections 2–5 with stochastic sensitive contexts whose distribution is known; Section 6 explains how to overcome this and develops a theory of approachability relying on ignoring but estimating the target set.

Notation. The Euclidean norm is denoted by $\|\cdot\|$, while the ℓ_1 norm is denoted by $\|\cdot\|_1$. Given a convex closed set $\mathcal{C} \subset \mathbb{R}^d$, we denote by $\text{Proj}_{\mathcal{C}}(\cdot)$ the projection operator onto \mathcal{C} in Euclidean norm.

2 Fair online learning cast as an approachability problem

In this section, we propose a setting for fair online learning based on approachability—a theory introduced by Blackwell [3] (see also the more modern expositions by Perchet [25] or Mertens et al. [22]). More precisely, we consider the following repeated game between a Player and Nature, with stochastic contexts. The existence of these contexts is a (minor) variation on the classical statement of the approachability problem.

The Player and Nature have respective finite action sets \mathcal{A} and \mathcal{B} . The sets of sensitive and non-sensitive contexts are respectively denoted by \mathcal{S} and \mathcal{X} . The set \mathcal{X} is a general Borel set, while \mathcal{S} is a finite set with cardinality denoted by $|\mathcal{S}|$. Typical choices are $\mathcal{S} = \{0, 1\}$ and $\mathcal{X} = \mathbb{R}^m$ for some $m \in \mathbb{N}$. A joint distribution \mathbf{Q} on $\mathcal{X} \times \mathcal{S}$ is fixed and is unknown to the Player. Finally, a (bounded) Borel-measurable vector-valued payoff function $\mathbf{m} : \mathcal{A} \times \mathcal{B} \times \mathcal{X} \times \mathcal{S} \rightarrow \mathbb{R}^d$, as well as a closed target set $\mathcal{C} \subseteq \mathbb{R}^d$, are given and known by the Player.

At each round $t \geq 1$ the pair of non-sensitive and sensitive contexts $(x_t, s_t) \sim \mathbf{Q}$ is generated independently from the past. The Player observes only the non-sensitive context x_t ; while Nature also observes x_t , it may or may not observe the sensitive context s_t . Then, Nature and the Player simultaneously pick (possibly in a randomized fashion) $b_t \in \mathcal{B}$ and $a_t \in \mathcal{A}$, respectively. The Player finally accesses the obtained reward $\mathbf{m}(a_t, b_t, x_t, s_t)$ and the sensitive context s_t , while Nature has a more complete monitoring and may observe a_t and s_t . We introduce an observation operation G to indicate whether Nature observes x_t only—i.e., $G(x_t, s_t) = x_t$, the case of Nature’s unawareness—or whether Nature observes both contexts—i.e., $G(x_t, s_t) = (x_t, s_t)$, the case of Nature’s awareness.

We consider the short-hand notation $\mathbf{m}_t := \mathbf{m}(a_t, b_t, x_t, s_t)$,

$$\bar{\mathbf{m}}_T := \frac{1}{T} \sum_{t=1}^T \mathbf{m}(a_t, b_t, x_t, s_t), \quad \text{and} \quad \bar{\mathbf{c}}_T = \text{Proj}_{\mathcal{C}}(\bar{\mathbf{m}}_T) = \arg \min_{\mathbf{v} \in \mathcal{C}} \|\bar{\mathbf{m}}_T - \mathbf{v}\|$$

for the instantaneous and average payoffs of the player, as well as the Euclidean projection of the latter onto the closed set \mathcal{C} , respectively. The distance of $\bar{\mathbf{m}}_T$ to \mathcal{C} thus equals $d_T := \|\bar{\mathbf{m}}_T - \bar{\mathbf{c}}_T\|$. The game protocol is summarized on the next page.

We recall that the Player does not know the context distribution \mathbf{Q} .

Definition 1. A target set \mathcal{C} is called \mathbf{m} -approachable by the Player under the distribution \mathbf{Q} if there exists a strategy of the Player such that, for all strategies of the Nature, $\bar{\mathbf{m}}_T \rightarrow \mathcal{C}$ a.s.

PROTOCOL 2.1

Parameters: Observation operator G for Nature; distribution \mathbf{Q} on $\mathcal{X} \times \mathcal{S}$

For $t = 1, 2, \dots$

1. Contexts (x_t, s_t) are sampled according to \mathbf{Q} , independently from the past;
2. Simultaneously,
 - Nature observes $G(x_t, s_t)$ and picks $b_t \in \mathcal{B}$;
 - the Player observes x_t and picks an action $a_t \in \mathcal{A}$;
3. The Player observes the reward $\mathbf{m}(a_t, b_t, x_t, s_t)$ and the sensitive context s_t , while Nature observes (a_t, b_t, x_t, s_t) .

Aim: The Player wants to ensure that $\overline{\mathbf{m}}_T \rightarrow \mathcal{C}$ a.s., i.e., $d_T = \|\overline{\mathbf{m}}_T - \overline{\mathbf{c}}_T\| \rightarrow 0$ a.s.

Remark 1 (Awareness for the Player). *We are mostly interested in a Player unaware of the sensitive contexts s_t (Gajane and Pechenizkiy [13]). However, the setting above also covers the case of a Player aware of these contexts: simply consider the lifted non-sensitive contexts $x'_t = (x_t, s_t)$.*

We now describe payoff functions and target sets corresponding to online learning objectives or online fairness constraints. They may be combined together. For instance, vanilla calibration corresponds below to the \mathbf{m}_{cal} -approachability of a set \mathcal{C}_{cal} , demographic parity, to the \mathbf{m}_{DP} -approachability of a set \mathcal{C}_{DP} , so that vanilla calibration under a demographic parity constraint translates into the $(\mathbf{m}_{\text{cal}}, \mathbf{m}_{\text{DP}})$ -approachability of the product set $\mathcal{C}_{\text{cal}} \times \mathcal{C}_{\text{DP}}$. We therefore consider each objective and each constraint as some elementary brick, to be combined with one or several other bricks. We recall that \mathcal{S} is a finite set and will indicate the cases where we only consider $\mathcal{S} = \{0, 1\}$.

We discuss two objectives: no-regret and approximate calibration, as well as three fairness constraints: group-wise (per-group) control, demographic parity, and equal average payoffs.

2.1 Statement of the objectives

For the sake of a more compact exposition, we define the objectives in two forms: global objectives (the vanilla form of objectives) and group-wise objectives. We denote $\gamma_s = \mathbb{P}(s_t = s)$, so that $(\gamma_s)_{s \in \mathcal{S}}$ corresponds to the marginal of \mathbf{Q} on \mathcal{S} .

Objective 1: (Vanilla and group-wise) no-regret. The definition is based on some payoff function r , possibly taking contexts into account: at each round t , the Player obtains the payoff $r(a_t, b_t, x_t, s_t)$. The aim is to get, on average, almost as much payoff as the best constant action, all things equal. The vanilla (average) regret equals

$$R_T = \min_{a \in \mathcal{A}} \frac{1}{T} \sum_{t=1}^T (r(a_t, b_t, x_t, s_t) - r(a, b_t, x_t, s_t)),$$

while the group-wise (average) regret equals

$$R_{\text{gr}, T} = \min_{s \in \mathcal{S}} \min_{a'_s \in \mathcal{A}} \frac{1}{T} \sum_{t=1}^T (r(a_t, b_t, x_t, s_t) - r(a'_s, b_t, x_t, s_t)) \mathbb{I}\{s_t = s\}.$$

The aim is that $\liminf R_T \geq 0$ a.s. (no-regret) and $\liminf R_{\text{gr}, T} \geq 0$ a.s. (group-wise no-regret), respectively. We could replace the $1/T$ factor by a $1/(\gamma_s T)$ factor in the definition of $R_{\text{gr}, T}$, as we will do for the \mathcal{C}_T calibration criterion, but given the wish of a non-negative limit, this is irrelevant.

Denote by $N = |\mathcal{A}|$ the cardinality of \mathcal{A} . No-regret corresponds to the \mathbf{m}_{reg} -approachability of $([0, +\infty))^N$, with the global payoff function $\mathbf{m}_{\text{reg}}(a, b, x, s) = (r(a, b, x, s) - r(a', b, x, s))_{a' \in \mathcal{A}}$. We also duplicate \mathbf{m}_{reg} into the group-wise payoff function

$$\mathbf{m}_{\text{gr-reg}}(a, b, x, s) = (\mathbf{m}_{\text{reg}}(a, b, x, s) \mathbb{I}\{s' = s\})_{s' \in \mathcal{S}}.$$

Group-wise no-regret then corresponds to the $\mathbf{m}_{\text{gr-reg}}$ -approachability of $\mathcal{C}_{\text{gr-reg}} = ([0, +\infty))^{N|\mathcal{S}|}$.

Objective 2: Approximate (vanilla or group-wise) calibration. Online calibration was first solved by Foster and Vohra [12] and Foster [11]; see the monograph by Cesa-Bianchi and Lugosi [5, Section 4.8] for references to other solutions and extensions. For simplicity, we focus on binary outcomes $b_t \in \{0, 1\}$ and ask the Player to provide at each round forecasts a_t in $[0, 1]$, and even in a discretization of $[0, 1]$ based on a fixed number $N \geq 2$ of points:

$$\mathcal{A} = \{a^{(k)} := (k - 1/2)/N, k \in \{1, \dots, N\}\}.$$

Each $x \in [0, 1]$ can be approximated by some $a^{(k)} \in \mathcal{A}$ with $|x - a^{(k)}| \leq 1/(2N)$. At each round, the Player picks $k_t \in \{1, \dots, N\}$ and forecasts $a_t = a^{(k_t)}$. The action set \mathcal{A} can thus be identified with $\{1, \dots, N\}$.

This problem is actually called $1/N$ -calibration or approximate calibration. The global (vanilla) form of the criterion reads

$$C_T = \sum_{k=1}^N \left| \frac{1}{T} \sum_{t=1}^T (a^{(k)} - b_t) \mathbb{I}\{k_t = k\} \right|,$$

while the approximate group-wise calibration criterion is defined as

$$C_{\text{gr},T} = \sum_{s \in \mathcal{S}} \sum_{k=1}^N \left| \frac{1}{\gamma_s T} \sum_{t=1}^T (a^{(k)} - b_t) \mathbb{I}\{k_t = k\} \mathbb{I}\{s_t = s\} \right|.$$

The aim is that $\limsup C_T \leq 1/N$ a.s. or $\limsup C_{\text{gr},T} \leq 1/N$, respectively. Note that unlike vanilla calibration, its group-wise version requires to be calibrated on each sensitive attribute $s \in \mathcal{S}$. In particular, the classical $1/T$ factor is replaced by $1/(\gamma_s T)$, the expected number of appearances of $s_t = s$ for $t = 1, \dots, T$.

Mannor and Stoltz [18] and Abernethy et al. [1] rewrote the problem of approximate calibration as an approachability problem as follows: introduce the global payoff function

$$\mathbf{m}_{\text{cal}}(k, b) = ((a^{(1)} - b) \mathbb{I}\{k = 1\}, \dots, (a^{(N)} - b) \mathbb{I}\{k = N\}),$$

and duplicate it into the group-wise payoff function as follows:

$$\mathbf{m}_{\text{gr-cal}}(k, b, s) = (\mathbf{m}_{\text{cal}}(k, b) \mathbb{I}\{s = s'\} / \gamma_{s'})_{s' \in \mathcal{S}}.$$

The calibration criteria C_T and $C_{\text{gr},T}$ can now be rewritten as the ℓ^1 -norms of the average payoff vectors $\overline{\mathbf{m}}_{\text{cal},T}$ and $\overline{\mathbf{m}}_{\text{gr-cal},T}$. Approximate vanilla calibration thus corresponds to the \mathbf{m}_{cal} -approachability of $\mathcal{C}_{\text{cal}} = \{\mathbf{v} \in \mathbb{R}^N : \|\mathbf{v}\|_1 \leq 1/N\}$, while approximate group-wise calibration corresponds to the $\mathbf{m}_{\text{gr-cal}}$ -approachability of $\mathcal{C}_{\text{gr-cal}} = \{\mathbf{v} \in \mathbb{R}^{N|\mathcal{S}|} : \|\mathbf{v}\|_1 \leq 1/N\}$.

Note that non-sensitive contexts play no role in the calibration objectives, but the Player can (and *must*) leverage these non-sensitive contexts to possibly infer sensitive contexts when handling group-wise calibration.

2.2 Statement of the fairness constraints

Fairness constraint 1: Group-wise objectives. We already considered possibly group-wise objectives above and Section 4 will show that handling them is already a challenge in our setting where the Player is unaware of the sensitive contexts.

Fairness constraint 2: Demographic parity. We will consider it only in the setting of approximate calibration and further restrict our attention to the case of two groups: $\mathcal{S} = \{0, 1\}$. The demographic parity criterion measures the difference between the average forecasts issued for the two groups:

$$D_T = \left| \frac{1}{\gamma_0 T} \sum_{t=1}^T a_t \mathbb{I}\{s_t = 0\} - \frac{1}{\gamma_1 T} \sum_{t=1}^T a_t \mathbb{I}\{s_t = 1\} \right|.$$

Given the discretization used, the wish is that $\limsup D_T \leq 1/N$. Abiding by a demographic parity constraint is equivalent to \mathbf{m}_{DP} -approaching $\mathcal{C}_{\text{DP}} = \{(u, v) \in \mathbb{R}^2 : |u - v| \leq 1/N\}$, where

$$\mathbf{m}_{\text{DP}}(k, s) = (a^{(k)} \mathbb{I}\{s = 0\} / \gamma_0, a^{(k)} \mathbb{I}\{s = 1\} / \gamma_1).$$

Fairness constraint 3: Equalized average payoffs. This criterion is to be combined with a no-regret criterion; in particular, a base payoff function r is considered. We restrict our attention to the case of two groups, $\mathcal{S} = \{0, 1\}$, and measure the difference of average payoffs:

$$P_T = \left| \frac{1}{\gamma_0 T} \sum_{t=1}^T r(a_t, b_t, x_t, s_t) \mathbb{I}\{s_t = 0\} - \frac{1}{\gamma_1 T} \sum_{t=1}^T r(a_t, b_t, x_t, s_t) \mathbb{I}\{s_t = 1\} \right|.$$

Ensuring $\limsup P_T \leq \varepsilon$ corresponds to $\mathbf{m}_{\text{eq-pay}}$ -approaching $\mathcal{C}_{\text{eq-pay}} = \{(u, v) \in \mathbb{R}^2 : |u - v| \leq \varepsilon\}$, where

$$\mathbf{m}_{\text{eq-pay}}(a, b, x, s) = (r(a, b, x, 0) \mathbb{I}\{s = 0\} / \gamma_0, r(a, b, x, 1) \mathbb{I}\{s = 1\} / \gamma_1).$$

Remark 2. Note that in this general form, the equality of average payoffs encompasses the demographic parity constraint. Indeed, the latter is obtained by setting $r(a, b, x, s) = a$ and $\varepsilon = 1/N$.

2.3 Summary table

The table below gives a summary of different criteria and associated pairs of payoff function and target set. We remark that some of the payoff functions depend on the marginals $(\gamma_s)_{s \in \mathcal{S}}$. Meanwhile, Protocol 2.1 assumes the perfect knowledge of the former. In Section 6 we will show how to bypass this issue, transferring all the unknown quantities into the target set and estimating it.

Criterion	Vector payoff function	Closed convex target set
Calibration	$\mathbf{m}_{\text{cal}}(k, b) = ((a^{(k')} - b) \mathbb{I}\{k = k'\})_{k' \in \mathcal{A}}$	$\mathcal{C}_{\text{cal}} = \{\mathbf{v} \in \mathbb{R}^N : \ \mathbf{v}\ _1 \leq 1/N\}$
Group-calibration	$\mathbf{m}_{\text{gr-cal}}(k, b, s) = (\mathbf{m}_{\text{cal}}(k, b) \mathbb{I}\{s = s'\} / \gamma_{s'})_{s' \in \mathcal{S}}$	$\mathcal{C}_{\text{gr-cal}} = \{\mathbf{v} \in \mathbb{R}^{N \mathcal{S} } : \ \mathbf{v}\ _1 \leq 1/N\}$
No-regret	$\mathbf{m}_{\text{reg}}(a, b, x, s) = (r(a, b, x, s) - r(a', b, x, s))_{a' \in \mathcal{A}}$	$\mathcal{C}_{\text{reg}} = ([0, +\infty))^N$
Group-no-regret	$\mathbf{m}_{\text{gr-reg}}(a, b, x, s) = (\mathbf{m}_{\text{reg}}(a, b, x, s) \mathbb{I}\{s' = s\})_{s' \in \mathcal{S}}$	$\mathcal{C}_{\text{gr-reg}} = ([0, +\infty))^{N \mathcal{S} }$
Demographic parity	$\mathbf{m}_{\text{DP}}(k, s) = (a^{(k)} \mathbb{I}\{s = 0\} / \gamma_0, a^{(k)} \mathbb{I}\{s = 1\} / \gamma_1)$	$\mathcal{C}_{\text{DP}} = \{(u, v) \in \mathbb{R}^2 : u - v \leq 1/N\}$
Equalized payoffs	$\mathbf{m}_{\text{eq-pay}}(a, b, x, s) = (r(a, b, x, s') \mathbb{I}\{s = s'\} / \gamma_{s'})_{s' \in \{0, 1\}}$	$\mathcal{C}_{\text{eq-pay}} = \{(u, v) \in \mathbb{R}^2 : u - v \leq \varepsilon\}$

3 Approachability theory adapted

We provide a rather straightforward extension of the approachability theory to deal with Protocol 2.1, namely, with the existence of stochastic contexts, drawn according to an unknown distribution \mathbf{Q} . We want to characterize closed convex sets that are approachable.

Pure vs. mixed actions. To conclude the description of the setting, we provide more details on the randomized draws of the (pure) actions a_{t+1} and b_{t+1} of the Player and Nature at round $t + 1$. We denote by h_t the information available to Player at the end of round t , and by H_t the full history of the first t rounds: $h_t = (\mathbf{m}_{t'}, x_{t'}, s_{t'})_{t' \leq t}$ and $H_t = (a_{t'}, b_{t'}, x_{t'}, s_{t'})_{t' \leq t}$. At the beginning of round $t + 1$, the Player thus picks in a h_t -measurable way a measurable family $(\mathbf{p}_{t+1}^x)_{x \in \mathcal{X}}$ of probability distributions over \mathcal{A} (i.e., a collection of distributions such that $x \in \mathcal{X} \mapsto \mathbf{p}_{t+1}^x$ is Borel-measurable), and then draws a_{t+1} independently at random according to the mixed action $\mathbf{p}_{t+1}^{x_{t+1}}$. Similarly, Nature picks in a H_t -measurable way a measurable family $(\mathbf{q}_{t+1}^{G(x, s)})_{(x, s) \in \mathcal{X} \times \mathcal{S}}$ of probability distributions over \mathcal{B} , and uses $\mathbf{q}_{t+1}^{G(x_{t+1}, s_{t+1})}$ to draw b_{t+1} .

Approachability strategy. We adapt the original strategy by Blackwell [3] by assuming the existence of and substituting a sequence of estimates $\hat{\mathbf{Q}}_t$ that are h_t -adapted in place of the unknown distribution \mathbf{Q} . We will assume that this sequence is convergent in the total variation distance in the sense of Assumption 1 below. To state the strategy, we extend linearly \mathbf{m} : for all probability distributions \mathbf{p} over \mathcal{A} and \mathbf{q} over \mathcal{B} , for all $(x, s) \in \mathcal{X} \times \mathcal{S}$,

$$\mathbf{m}(\mathbf{p}, \mathbf{q}, x, s) = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \mathbf{p}(a) \mathbf{q}(b) \mathbf{m}(a, b, x, s).$$

Now, the Player uses an arbitrary measurable family of distributions $(\mathbf{p}_1^x)_{x \in \mathcal{X}}$ for the first round, gets the estimate $\hat{\mathbf{Q}}_1$, and then uses, for rounds $t + 1$, where $t \geq 1$:

$$(\mathbf{p}_{t+1}^x)_{x \in \mathcal{X}} \in \arg \min_{(\mathbf{p}^x)_{x \in \mathcal{X}}} \max_{(\mathbf{q}^{G(x, s)})_{(x, s) \in \mathcal{X} \times \mathcal{S}}} \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x, s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) \right\rangle, \quad (1)$$

where the minimum and maximum are over all measurable families of probability distributions over \mathcal{A} and \mathcal{B} , respectively. The Player then gets access to h_{t+1} and may compute the estimate $\hat{\mathbf{Q}}_{t+1}$ to be used at the next round.

Necessary and sufficient condition for approachability. We were able to work out such a condition under the assumption that \mathbf{Q} can be estimated well enough, e.g., faster than at a $1/\ln^3(T)$ rate in total variation distance. We recall that the total variation distance between two probability distributions \mathbf{Q}_1 and \mathbf{Q}_2 on $\mathcal{X} \times \mathcal{S}$ equals (see, e.g., Devroye [8]):

$$\text{TV}(\mathbf{Q}_1, \mathbf{Q}_2) = \sup_{E \subseteq \mathcal{X} \times \mathcal{S}} |\mathbf{Q}_1(E) - \mathbf{Q}_2(E)| = \frac{1}{2} \int_{\mathcal{X} \times \mathcal{S}} |g_1(x, s) - g_2(x, s)| d\mu(x, s),$$

where the supremum is over all Borel sets E of $\mathcal{X} \times \mathcal{S}$, and where g_1 and g_2 denote densities of \mathbf{Q}_1 and \mathbf{Q}_2 with respect to a common dominating probability distribution μ .

Assumption 1 (fast enough sequential estimation of \mathbf{Q}). *The sequence of (h_t) -adapted estimators $(\hat{\mathbf{Q}}_t)$ used is such that $\sum_{t=1}^{+\infty} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} < +\infty$.*

The above assumption implies both $\frac{1}{T} \sum_{t=1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}$ and $\sum_{t \geq T+1} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}$ converge to zero (with T ; see Appendix A for details). Assumption 1 is trivially satisfied in the case when \mathbf{Q} is known, as it is sufficient to take $\hat{\mathbf{Q}}_t = \mathbf{Q}$. When both \mathcal{X} and \mathcal{S} are finite sets, we may use the empirical frequencies as estimators $\hat{\mathbf{Q}}_t$; they satisfy $\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})] = O(1/t)$; see, e.g., [7, Lemma 3]. The general case of an uncountable \mathcal{X} , e.g., $\mathcal{X} = \mathbb{R}^m$ requires results for density estimation in the L^1 or L^2 norms; such results rely typically on moving averages or kernel estimates and may be found, for instance, in the monographs by Devroye and Györfi [9] and Devroye [8] (see also Tsybakov [28]). Under mild conditions, the estimation takes place at a polynomial rate in total variation distance (e.g., a $T^{-1/5}$ rate in dimension $m = 1$). Note that the needed rate of decrease for $\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]$ in Assumption 1 is extremely slow: a $1/\ln^3(T)$ rate would suffice.

Assumption 2 (boundedness). *We assume that $\|\mathbf{m}\|_{\infty,2} := \max_{a,b \in \mathcal{A} \times \mathcal{B}} \sup_{(x,s) \in \mathcal{X} \times \mathcal{S}} \|\mathbf{m}(a, b, x, s)\| < +\infty$.*

We may now state our main result; the distance of $\bar{\mathbf{m}}_T$ to \mathcal{C} was denoted by d_T in Protocol 2.1.

Theorem 1. *Assume that \mathcal{C} is a closed convex set and that Assumptions 1 (fast enough sequential estimation of \mathbf{Q}) and 2 (bounded reward function) are satisfied, then \mathcal{C} is approachable if and only if*

$$\forall (\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \{0,1\}} \exists (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \text{s.t.} \quad \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \in \mathcal{C}. \quad (2)$$

In this case, the strategy of Eq. (1) achieves the following rates for L^2 and almost-sure convergences:

$$\mathbb{E}[d_T^2] \leq \sqrt{\frac{K}{T}} + 4\|\mathbf{m}\|_{\infty,2} \overbrace{\frac{1}{T} \sum_{t=1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}}^{:= \bar{\Delta}_T} \quad \text{and}$$

$$\mathbb{P}\left(\sup_{t \geq T} d_t \geq \varepsilon\right) \leq \frac{3K}{T\varepsilon^2} + \frac{16\|\mathbf{m}\|_{\infty,2}}{\varepsilon^2} \left(\sqrt{\frac{K}{T-1}} + 2 \left(\sup_{t \geq T} \bar{\Delta}_t \right) \left(\bar{\Delta}_T + \sum_{t \geq T} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} \right) \right)$$

where $K < +\infty$ denotes the maximal distance to \mathcal{C} of an element of the compact set $\mathbf{m}(\mathcal{A}, \mathcal{B}, \mathcal{X}, \mathcal{S})$.

The proof lies in Appendix A. The necessity part of the theorem actually relies on no assumption other than \mathcal{C} being closed; it consists of showing that Nature has a stationary strategy such that there exists $\alpha > 0$ with $d_T \geq \alpha$ in the limit, i.e., the average payoff vectors $\bar{\mathbf{m}}_T$ remain α -away from \mathcal{C} in the limit. This exactly indicates that the underlying fair online learning problem is not tractable: the underlying objectives and underlying fairness constraints cannot be simultaneously satisfied.

4 Working out some objective–constraint pairs: (im)possibility results

In this section we apply Theorem 1 to deal with some examples of objective–constraint pairs described in Sections 2.1 and 2.2. Some of them have been considered before in the literature (sometimes in the batch setup) using various tools [4, 16, 21, 14], as discussed in Section 1.

We keep the original criteria and obtain possibility or impossibility results. This is a first step, meanwhile, Section 5 will explain how to go further and obtain a trade-off, if needed, between the objective and the fairness constraint.

Additional notation. We recall that $\gamma_s = \mathbb{P}(s_t = s)$ and denote by \mathbf{Q}^s the conditional distribution of x_t given $s_t = s$, so that $d\mathbf{Q}(x, s) = \gamma_s d\mathbf{Q}^s(x)$. We denote by $\text{supp}(\mathbf{Q}^s) \subseteq \mathcal{X}$ the support of \mathbf{Q}^s .

Example 1: Vanilla calibration under a demographic parity constraint—achievable. Consider the following payoff function and target set, obtained by simultaneously considering the objective of vanilla calibration and the constraint of demographic parity: $\mathbf{m} = (\mathbf{m}_{\text{cal}}, \mathbf{m}_{\text{DP}})$ and $\mathcal{C} = \mathcal{C}_{\text{cal}} \times \mathcal{C}_{\text{DP}}$.

Defining $\psi(u_1, u_2) := |u_1 - u_2|$, the approachability condition (2) then reads as follows (where we introduce short-hand notation \mathcal{C} and DP):

$$\forall (\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \{0,1\}} \exists (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \text{s.t.} \quad \begin{cases} \mathcal{C} := \left\| \int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{cal}}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}) d\mathbf{Q}(x, s) \right\|_1 \leq \frac{1}{N}; \\ \text{DP} := \psi \left(\int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{DP}}(\mathbf{p}^x, s) d\mathbf{Q}(x, s) \right) \leq \frac{1}{N}. \end{cases} \quad (3)$$

Recalling the notation \mathbf{Q}^0 and \mathbf{Q}^1 for the conditional distributions, we observe that

$$\text{DP} = \left| \int_{\mathcal{X}} \sum_{k=1}^N \mathbf{p}^x(k) a^{(k)} d\mathbf{Q}^0(x) - \int_{\mathcal{X}} \sum_{k=1}^N \mathbf{p}^x(k) a^{(k)} d\mathbf{Q}^1(x) \right|.$$

We now show that the condition in Eq. (3) is satisfied. For any $(\mathbf{q}^{G(x,s)})$, we define the family (\mathbf{p}^x) as the constant family $(\text{dirac}(Q_{\mathcal{A}}))$, where $\text{dirac}(Q_{\mathcal{A}})$ denotes the Dirac mass supported on $Q_{\mathcal{A}}$, the closest point of \mathcal{A} to $Q := \int_{\mathcal{X} \times \{0,1\}} \mathbf{q}^{G(x,s)}(1) d\mathbf{Q}(x, s)$. We have $\text{DP} = 0$ as \mathbf{p}^x does not depend on x . Substituting the expression for \mathbf{m}_{cal} into the definition of \mathcal{C} , we observe that for such a choice of $(\mathbf{p}^x)_{x \in \mathcal{X}}$, we have

$$\mathcal{C} = \left| \int_{\mathcal{X} \times \{0,1\}} (Q_{\mathcal{A}} - \mathbf{q}^{G(x,s)}(1)) d\mathbf{Q}(x, s) \right| \leq \frac{1}{2N} + \underbrace{\left| \int_{\mathcal{X} \times \{0,1\}} (Q - \mathbf{q}^{G(x,s)}(1)) d\mathbf{Q}(x, s) \right|}_{=0},$$

where the inequality holds by taking the effect of discretization in \mathcal{A} into account and by the very definition of Q . The condition of Eq. (3) is thus satisfied. Therefore, under Assumption 1 (the existence of fast enough sequential estimators of \mathbf{Q}) and thanks to Theorem 1, the vanilla calibration and the demographic parity can be achieved simultaneously no matter the monitoring of the Nature.

Example 2: Group-wise no-regret—mixed picture. Let the target set be $\mathcal{C}_{\text{gr-reg}} = ([0, +\infty))^{N|S|}$ and the payoff function be $\mathbf{m}_{\text{gr-reg}}$, i.e., we consider the case of group-wise no-regret under no additional constraint. The approachability condition in Eq. (2) demands that

$$\forall (\mathbf{q}^{G(x,s)}) \exists (\mathbf{p}^x) \quad \text{s.t.} \quad \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}_{\text{gr-reg}}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}) d\mathbf{Q}(x, s) \in ([0, +\infty))^{N|S|}, \quad \text{i.e.,} \quad (4)$$

$$\forall (a', s), \quad \int_{\text{supp}(\mathbf{Q}^s)} \sum_{a \in \mathcal{A}} \mathbf{p}^x(a) \left(\sum_{b \in \mathcal{B}} \mathbf{q}^{G(x,s)}(b) (r(a, b, x, s) - r(a', b, x, s)) \right) d\mathbf{Q}^s(x) \geq 0.$$

No-regret seems a harmless challenge, and it is so when the sensitive context is directly observed by the Player, which we do not assume. (In this case, the Player may simply run several no-regret

algorithms in parallel, one per sensitive group s .) In our context, the direct observation is emulated in some sense when the non-sensitive context x reveals the sensitive context s ; this is the case, for instance, when the supports of the distributions \mathbf{Q}^s are pairwise disjoint. Note, however, that these distributions \mathbf{Q}^s are unknown to the Player and need to be learned. The second part of Proposition 1 shows that in this case, the group-wise no-regret may be controlled. We get a similar control in the case of irrelevant sensitive contexts, i.e., not affecting the payoffs and not used by Nature; see the first part of Proposition 1, which corresponds to the case of vanilla no-regret minimization. In both cases, the group-wise no-regret can be controlled under Assumption 1, thanks to Theorem 1. However, as we show by means of counter-examples, these are the only cases that may be favorably dealt with.

Proposition 1. *The condition of Eq. (4) holds when*

- *the sensitive context is irrelevant, i.e., the payoff function is such that $r(a, b, x, s) = r(a, b, x)$ and Nature's monitoring is $G(x, s) = x$;*
- *for all $s \neq s'$, it holds that $\text{supp}(\mathbf{Q}^s) \cap \text{supp}(\mathbf{Q}^{s'}) = \emptyset$, no matter Nature's monitoring G .*

Otherwise, the condition of Eq. (4) may not hold.

Proof. We mimic the classical proof of no-regret by approachability for the positive results. For the *first* positive result: for any (\mathbf{q}^x) , we define $a^x \in \arg \max_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \mathbf{q}^x(b) r(a, b, x)$ and let $(\mathbf{p}^x) = (\text{dirac}(a^x))$. For the *second* positive result: fix any $(\mathbf{q}^{G(x,s)})$; we define $(\mathbf{p}^x)_{x \in \mathcal{X}}$ point-wise as follows. For all $s \in \mathcal{S}$, all $x \in \text{supp}(\mathbf{Q}^s)$, we set $\mathbf{p}^x = \text{dirac}(a^x)$, where we validly define $a^x \in \arg \max_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} \mathbf{q}^{G(x,s)}(b) r(a, b, x, s)$ on the union of the supports of $(\mathbf{Q}^s)_{s \in \mathcal{S}}$, since they are pair-wise disjoint; we define the a^x arbitrarily elsewhere.

Two counter-examples detailed in Appendix B back up the final part of the proposition: we show that Eq. (4) does not hold. In the first counter-example, the monitoring is $G(x, s) = x$, the payoff function depends on s , and the supports of $(\mathbf{Q}^s)_{s \in \mathcal{S}}$ have non negligible intersection. In the second example, the monitoring is $G(x, s) = (x, s)$, the payoff function does not depend on s , and the supports of $(\mathbf{Q}^s)_{s \in \mathcal{S}}$ have non negligible intersection. \square

Example 3: (Vanilla) no-regret under the equalized average payoffs constraint. For the sake of space we deal with this example in Appendix B, obtaining similar conclusions as that of Blum et al. [4].

5 Group-wise calibration under a demographic parity constraint: trade-off

In this section, we consider the problem of group-wise calibration under the demographic parity constraint; in particular, $\mathcal{S} = \{0, 1\}$. As we will see, except for special cases, the corresponding two error criteria cannot be simultaneously smaller than the desired $1/N$ in the limit. However, a (possibly optimal) trade-off may be set between the calibration error ε and the violation level δ of demographic parity. To that end, we introduce neighborhoods of the original target sets $\mathcal{C}_{\text{gr-cal}}$ and \mathcal{C}_{DP} :

$$\mathcal{C}_{\text{gr-cal}}^\varepsilon = \{\mathbf{v} \in \mathbb{R}^{2N} : \|\mathbf{v}\|_1 \leq \varepsilon\} \quad \text{and} \quad \mathcal{C}_{\text{DP}}^\delta = \{(u, v) \in \mathbb{R}^2 : |u - v| \leq \delta\}.$$

A pair $(\varepsilon, \delta) \in \mathbb{R}_+ \times \mathbb{R}_+$ is said *achievable* when $\mathcal{C}_{\text{gr-cal}}^\varepsilon \times \mathcal{C}_{\text{DP}}^\delta$ is approachable with $\mathbf{m} = (\mathbf{m}_{\text{gr-cal}}, \mathbf{m}_{\text{DP}})$. Theorem 1 provides a characterization of this approachability as well as an associated strategy; in particular, when (ε, δ) is achievable, this strategy ensures that the calibration error C_T and the violation D_T of demographic parity satisfy: $\limsup C_T \leq \varepsilon$ a.s. and $\limsup D_T \leq \delta$ a.s.

The goal of this section is to identify all achievable pairs (ε, δ) . We will do so by determining, for $\delta \geq 0$ of interest, the *smallest* $\varepsilon \geq 0$ such that (ε, δ) is achievable¹; we denote it by $\varepsilon^*(\delta)$. The line $(\delta, \varepsilon^*(\delta))$ is a Pareto frontier.

¹Note that if (ε, δ) is achievable, then (ε', δ') with $\varepsilon' \geq \varepsilon$ and $\delta' \geq \delta$ is also achievable.

Re-parametrization of the problem. Under Assumption 1 (the existence of fast enough sequential estimators of \mathbf{Q}) and thanks to Theorem 1, the $(\mathbf{m}_{\text{gr-cal}}, \mathbf{m}_{\text{DP}})$ -approachability of $\mathcal{C}_{\text{gr-cal}}^\varepsilon \times \mathcal{C}_{\text{DP}}^\delta$ holds if and only if the condition of Eq. (2) is satisfied. The latter can be stated as follows:

$$\forall (\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \{0,1\}} \exists (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \text{s.t.} \quad \begin{cases} \left\| \int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{gr-cal}}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}) d\mathbf{Q}(x, s) \right\|_1 \leq \varepsilon; \\ \psi \left(\int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{DP}}(\mathbf{p}^x, s) d\mathbf{Q}(x, s) \right) \leq \delta, \end{cases} \quad (5)$$

where we recall that $\psi(u_1, u_2) = |u_1 - u_2|$. Now, one can show (see comments after Lemma 3 of Appendix C) that the $\psi(\dots)$ term above is always smaller than $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$. Thus, we can re-parameterize the problem and focus only on $\delta_\tau = \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$, where $\tau \in [0, 1]$.

Computation of the Pareto frontier. The condition of Eq. (5) indicates that

$$\varepsilon^*(\delta_\tau) = \max_{(\mathbf{q}^{G(x,s)})} \min_{(\mathbf{p}^x)} \left\| \int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{gr-cal}}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, s) d\mathbf{Q}(x, s) \right\|_1 \quad (6)$$

s.t. $\psi \left(\int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{DP}}(\mathbf{p}^x, s) d\mathbf{Q}(x, s) \right) \leq \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$.

Propositions 2 and 3 below compute the values (up to the $1/N$ discretization error) of $\varepsilon^*(\delta_\tau)$ in two scenarios, depending on whether Nature observes the sensitive contexts s_t .

Proposition 2 (Nature awareness: $G(x, s) = (x, s)$). *Under Assumption 1 and with the monitoring $G(x, s) = (x, s)$ for Nature, the Pareto frontier $(\varepsilon^*(\delta_\tau), \delta_\tau)_{\tau \in [0,1]}$ of achievable pairs satisfies*

$$\delta_\tau = \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) \quad \text{and} \quad 1 - \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) \leq \varepsilon^*(\delta_\tau) \leq 1 - \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) + \frac{1}{N}.$$

Proposition 3 (Nature unawareness: $G(x, s) = x$). *Under Assumption 1 and with the monitoring $G(x, s) = x$ for Nature, the Pareto frontier $(\varepsilon^*(\delta_\tau), \delta_\tau)_{\tau \in [0,1]}$ of achievable pairs satisfies:*

$$\delta_\tau = \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) \quad \text{and} \quad (1 - \tau) \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) \leq \varepsilon^*(\delta_\tau) \leq (1 - \tau) \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) + \frac{1}{N}.$$

The parameter $\tau \in [0, 1]$ is set by the user.

We observe that in the case when the true label b_t provided by the Nature can be directly influenced by the sensitive attribute s_t , Proposition 2 shows that approximate group-wise calibration with $\varepsilon = 1/N$ is never possible, unless $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) = 1$ (and $\tau = 1$ is picked). The latter case corresponds to the situation when the supports of \mathbf{Q}^0 and \mathbf{Q}^1 are disjoint, hence allowing the Player to infer the sensitive context s from the non-sensitive one x , essentially reducing (up to unknown \mathbf{Q}) the problem to the previously studied setup of Player's awareness [16].

When the true label b_t provided by the Nature is not *directly* influenced by the sensitive attribute s_t (it is influenced by s_t only via x_t), Proposition 3 indicates that calibration is always possible by setting $\tau = 1$, no matter the value of $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$. Interestingly, this proposition also shows that if $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) = 0$, i.e., the x_t and the s_t are independent, then the Player is able to achieve calibration and satisfy the demographic parity constraint simultaneously.

6 Approachability of an unknown target set

A limitation of the calibration problems under demographic parity constraint discussed in Section 4 (Example 1) and Section 5 is that the unknown probabilities γ_0 and γ_1 enter the payoff functions $\mathbf{m}_{\text{gr-cal}}$ and \mathbf{m}_{DP} . We already pointed out this issue in Section 2.3. Even worse, the trade-off claimed in Propositions 2 and 3 relies on the knowledge of the unknown $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$, to set the values of the achievable pair (δ, ε) targeted; that is, the target set is unknown. To bypass the first limitation

we transfer the unknown (γ_0, γ_1) to the target set, which makes the payoff function fully known to the Player. We will then be left with the problem of approaching an unknown target set only. For instance, in the context of Section 5, we can define

$$\widetilde{\mathbf{m}}_{\text{gr-cal}}(k, y, s) = (\mathbf{m}_{\text{cal}}(k, y) \mathbb{I}\{s = s'\})_{s'=0,1} \quad \text{and} \quad \widetilde{\mathbf{m}}_{\text{DP}}(k, s) = (a^{(k)} \mathbb{I}\{s = 0\}, a^{(k)} \mathbb{I}\{s = 1\}),$$

and set $\widetilde{\mathbf{m}} := (\widetilde{\mathbf{m}}_{\text{gr-cal}}, \widetilde{\mathbf{m}}_{\text{DP}})$. Taking into account the definition of \mathbf{m}_{cal} , we note that $\widetilde{\mathbf{m}}$ does not depend on (γ_0, γ_1) . Furthermore, by considering the closed convex target sets

$$\widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon = \left\{ (\mathbf{v}_0, \mathbf{v}_1) \in \mathbb{R}^{2N} : \frac{\|\mathbf{v}_0\|_1}{\gamma_0} + \frac{\|\mathbf{v}_1\|_1}{\gamma_1} \leq \varepsilon \right\}, \quad \widetilde{\mathcal{C}}_{\text{DP}}^\delta = \left\{ (u, v) \in \mathbb{R}^2 : \left| \frac{u}{\gamma_0} - \frac{v}{\gamma_1} \right| \leq \delta \right\},$$

we remark that the $(\widetilde{\mathbf{m}}_{\text{gr-cal}}, \widetilde{\mathbf{m}}_{\text{DP}})$ -approachability of $\widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \widetilde{\mathcal{C}}_{\text{DP}}^\delta$ is equivalent to the $(\mathbf{m}_{\text{gr-cal}}, \mathbf{m}_{\text{DP}})$ -approachability of $\mathcal{C}_{\text{gr-cal}}^\varepsilon \times \mathcal{C}_{\text{DP}}^\delta$. The unknown quantities appear only in the target set $\widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \widetilde{\mathcal{C}}_{\text{DP}}^\delta$ (and δ and ε count as unknown quantities given the trade-off exhibited), while the payoff $\widetilde{\mathbf{m}}$ is known beforehand. Thus, it is sufficient to consider the setup of Protocol 2.1 with an *unknown* target set \mathcal{C} .

Approachability strategy for an unknown target set \mathcal{C} . We still assume that the Player is able to build an h_t -adapted sequence of estimates $\hat{\mathbf{Q}}_t$. Additionally, we assume that for $T_r := 2^r$, with $r \geq 0$, the Player can construct an h_{T_r} -adapted estimate $\hat{\mathcal{C}}_r$ of \mathcal{C} . We discuss this assumption at the end of this section. We define $d(\hat{\mathcal{C}}_r, \mathcal{C}) = \sup_{x \in \hat{\mathcal{C}}_r} d(x, \mathcal{C})$.

Assumption 3. *There exist $B < +\infty$ and a summable non-increasing sequence $(\beta_r)_{r \geq 0}$ such that for all $r \geq 0$, the sets $\hat{\mathcal{C}}_r$ are convex closed, with $\|\mathbf{v} - \text{Proj}_{\hat{\mathcal{C}}_r}(\mathbf{v})\| \leq B$ for all $\mathbf{v} \in \mathbf{m}(\mathcal{A}, \mathcal{B}, \mathcal{X}, \{0, 1\})$,*

$$\mathbb{P}(\mathcal{C} \subset \hat{\mathcal{C}}_r) \geq 1 - 1/(2T_r), \quad \text{and} \quad \max \left\{ \mathbb{E}[d(\hat{\mathcal{C}}_r, \mathcal{C})^2], \mathbb{E}[d(\mathcal{C}, \hat{\mathcal{C}}_r)^2] \right\} \leq \beta_r^2.$$

For all $r \geq 0$ and all $t \in \{T_r, \dots, T_{r+1} - 1\}$, define $\hat{c}_t := \text{Proj}_{\hat{\mathcal{C}}_r}(\overline{\mathbf{m}}_t)$. The idea of the approachability strategy is to use \hat{c}_t in place of $\overline{\mathbf{c}}_t$ in Eq. (1) and update the estimate $\hat{\mathcal{C}}_r$ of the target \mathcal{C} only at the end of rounds $t = T_r$. More precisely, the strategy of the Player is:

$$(\mathbf{p}_{t+1}^x)_{x \in \mathcal{X}} \in \arg \min_{(\mathbf{p}^x)} \max_{(\mathbf{q}^{G(x,s)})} \left\langle \overline{\mathbf{m}}_t - \hat{c}_t, \int \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) \right\rangle. \quad (7)$$

Theorem 2. *Under Assumption 3 and the assumptions of Theorem 1, a convex closed set \mathcal{C} , unknown to the Player, is \mathbf{m} -approachable if and only if Blackwell's condition in Eq. (2) is satisfied. In this case, the strategy of Eq. (7) is an approachability strategy.*

Appendix D.2 provides a proof of Theorem 2. But before we do so, we discuss in Appendix D.1 why and how the target set \mathcal{C} may be estimated by sets $\hat{\mathcal{C}}_r$ satisfying Assumption 3. The construction is idiosyncratic and strongly depends on the problem and exact setting considered (in particular, whether the set of non-sensitive contexts \mathcal{X} is finite or not). We provide an illustration for the target set $\mathcal{C} = \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \widetilde{\mathcal{C}}_{\text{DP}}^\delta$ of Section 5, in the case of a finite set \mathcal{X} .

7 Limitations of the current work and topics for future work

The anonymous reviewers of this article pointed out some limitations to or possible extensions of the current work, which we list now.

We only considered, for the sake of readability, the case of demographic parity with two groups. While the criterion of demographic parity easily extends to more groups, the extension of the trade-offs stated in Section 5 is less clear. It would probably involve the total variation distances $\text{TV}(\mathbf{Q}^s, \mathbf{Q}^{s'})$ between each pair \mathbf{Q}^s and $\mathbf{Q}^{s'}$ of marginal distributions, where $s, s' \in \mathcal{S}$, or the distances $\text{TV}(\mathbf{Q}^s, \mathbf{Q})$.

The complexity of the generic approachability strategy of Section 3 is at least linear in the number of groups (if an approximate solution is used, and even polynomial in this number for an exact solution), see Mannor and Stoltz [18, Sections 3.3 and 3.4]. The convergence rates achieved in Theorem 1

involve total variation distances that would also probably depend in at least a linear fashion on the number of groups, unless some special structure is assumed. Both facts may be an issue for large numbers of groups. More generally, we only provide in this article a generic strategy, that is, a first approach to tackle a given fair online learning problem, but specific strategies may be more efficient and get better regret bounds, in particular for large numbers of groups. (For instance, for group-wise calibration, Gupta et al. [14] base a specific and computationally more efficient strategy on an exponential surrogate loss: this strategy enjoys a sample complexity only logarithmic in the number of groups.) The design of such specific strategies remains largely open.

As mentioned in the introduction, Bechavod et al. [2] consider the objective of online binary classification and an equal-opportunity fairness constraint, under some partial monitoring known as “apple tasting”. In this article, we considered a bandit monitoring for the Player: she observes the reward obtained at each round. Partial monitoring, which was introduced by Rustichini [27], consists of only receiving feedback being a (possibly random) function of the actions played by the Player and Nature. A theory of approachability under partial monitoring was initiated by Perchet [24], who stated a necessary and sufficient condition (see also Mannor et al. [19]); Mannor et al. [20] then exhibited a computationally more efficient strategy, with improved convergence rates, and Kwon and Perchet [17] finally obtained the optimal convergence rates. Therefore, the question to investigate is the extension of the results of this article from a bandit monitoring to a partial monitoring.

Acknowledgments and Disclosure of Funding

Evgenii Chzhen was fully funded by grant ANR-11-LABX-0056-LMH (Labex LMH, part of Programme d’investissements d’avenir). Christophe Giraud received partial support by grant ANR-19-CHIA-0021-01 (“BiSCottE”, Agence Nationale de la Recherche). Gilles Stoltz has no direct funding to acknowledge.

Additional revenues for authors are: Evgenii Chzhen—none; Christophe Giraud—none; Gilles Stoltz—part time employment as an affiliate professor with HEC Paris.

References

- [1] J. Abernethy, P. Bartlett, and E. Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of the 24th Annual Conference on Learning Theory*, volume PMLR 19, pages 27–46, 2011.
- [2] Y. Bechavod, K. Ligett, A. Roth, B. Waggoner, and S. Wu. Equal opportunity in online classification with partial feedback. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [3] D. Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956.
- [4] A. Blum, S. Gunasekar, T. Lykouris, and N. Srebro. On preserving non-discrimination when combining expert advice. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- [5] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [6] A. Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2):153–163, 2017.
- [7] L. Devroye. The equivalence of weak, strong and complete convergence in L_1 for kernel density estimates. *The Annals of Statistics*, 11(3):896–904, 1983.
- [8] L. Devroye. *A Course in Density Estimation*. Birkhäuser, 1987.
- [9] L. Devroye and L. Györfi. *Nonparametric Density Estimation: The L^1 View*. Wiley, 1985.
- [10] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 214–226, 2012.
- [11] D. Foster. A proof of calibration via Blackwell’s approachability theorem. *Games and Economic Behavior*, 29:73–78, 1999.
- [12] D. Foster and R. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
- [13] P. Gajane and M. Pechenizkiy. On formalizing fairness in prediction with machine learning, 2017. Preprint, arXiv:1710.03184.
- [14] V. Gupta, C. Jung, G. Noarov, M. Pai, and A. Roth. Online multivald learning: Means, moments, and prediction intervals, 2021. Preprint, arXiv:2101.01739.
- [15] M. Hardt, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [16] U. Hébert-Johnson, M. Kim, O. Reingold, and G. Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning*, volume PMLR 80, pages 1939–1948, 2018.
- [17] J. Kwon and V. Perchet. Online learning and Blackwell approachability with partial monitoring: Optimal convergence rates. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume PMLR 54, pages 604–613, 2017.

- [18] S. Mannor and G. Stoltz. A geometric proof of calibration. *Mathematics of Operations Research*, 35(4):721–727, 2010.
- [19] S. Mannor, V. Perchet, and G. Stoltz. A primal condition for approachability with partial monitoring. *Journal of Dynamics and Games*, 1(3):447–469, 2013.
- [20] S. Mannor, V. Perchet, and G. Stoltz. Set-valued approachability and online learning with partial monitoring. *Journal of Machine Learning Research*, 15(1):3247–3295, 2014.
- [21] N. Martinez, M. Bertran, and G. Sapiro. Minimax Pareto fairness: A multi objective perspective. In *Proceedings of the 37th International Conference on Machine Learning*, volume PMLR 119, pages 6755–6764, 2020.
- [22] J.-F. Mertens, S. Sorin, and Z. Shmuel. *Repeated Games*. Cambridge University Press, 2015.
- [23] V. Perchet. *Approachability, calibration and regret in games with partial observations*. PhD thesis, Université Pierre et Marie Curie – Paris VI, 2010.
- [24] V. Perchet. Approachability of convex sets in games with partial monitoring. *Journal of Optimization Theory and Applications*, 149(3):665–677, 2011.
- [25] V. Perchet. Approachability, regret and calibration: Implications and equivalences. *Journal of Dynamics & Games*, 1(2):181–254, 2014.
- [26] G. Pleiss, M. Raghavan, F. Wu, J. Kleinberg, and K. Weinberger. On fairness and calibration. In *Advances in Neural Information Processing Systems*, volume 30, pages 5680–5689, 2017.
- [27] A. Rustichini. Minimizing regret: The general case. *Games and Economic Behavior*, 29(1): 224–243, 1999.
- [28] A. Tsybakov. *Introduction to Nonparametric Estimation*. Springer, 2008.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] Fairness results exhibited should only have positive impacts
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Supplementary Material for
“A Unified Approach to Fair Online Learning
via Blackwell Approachability”

by Evgenii Chzhen, Christophe Giraud, Gilles Stoltz

This supplementary material contains all the proofs omitted from the main body. Each section provides the proofs of claims, theorems, or propositions of a section of the main body; more precisely, Appendix A provides proofs for Section 3, Appendix B does so for Section 4, Appendix C, for Section 5, and finally, Appendix D deals with Section 6.

A Proofs for Section 3

We start by proving a claim stated right after Assumption 1: that

$$\sum_{t=1}^{+\infty} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} := C < +\infty \quad \text{entails} \quad \bar{\Delta}_T := \frac{1}{T} \sum_{t=1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} \longrightarrow 0.$$

Indeed,

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} &= \frac{1}{T} \sum_{t=1}^{\lfloor \sqrt{T} \rfloor} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} + \frac{1}{T} \sum_{t=\lfloor \sqrt{T} \rfloor+1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} \\ &\leq \underbrace{\frac{1}{\sqrt{T}} \sum_{t=1}^{\lfloor \sqrt{T} \rfloor} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}}_{\leq C} + \sum_{t=\lfloor \sqrt{T} \rfloor+1}^T \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}, \end{aligned}$$

which converges to 0, as it is the sum of C/\sqrt{T} with a quantity smaller than the remainder of a convergent series.

We recall that for all Borel-measurable functions $f : \mathcal{X} \times \mathcal{S} \rightarrow \mathbb{R}^d$ with $\sup_{(x,s) \in \mathcal{X} \times \mathcal{S}} \|f(x,s)\| \leq M$,

$$\begin{aligned} &\left\| \int_{\mathcal{X} \times \mathcal{S}} f(x,s) d\mathbf{Q}_1(x,s) - \int_{\mathcal{X} \times \mathcal{S}} f(x,s) d\mathbf{Q}_2(x,s) \right\| \\ &\leq \int_{\mathcal{X} \times \mathcal{S}} \|f(x,s)\| |g_1(x,s) - g_2(x,s)| d\mu(x,s) \leq 2M \cdot \text{TV}(\mathbf{Q}_1, \mathbf{Q}_2), \quad (8) \end{aligned}$$

where g_1 and g_2 denote densities of the distributions \mathbf{Q}_1 and \mathbf{Q}_2 with respect to a common dominating measure μ .

We now move to the proof of Theorem 1, which we restate below. It relies on two lemmas stated below in Section A.1. Unless stated otherwise (namely, for matters related to the estimation of \mathbf{Q}), all material is standard and was introduced by Blackwell [3] (see also the more modern expositions by Perchet [25] or Mertens et al. [22]).

Theorem 1. *Assume that \mathcal{C} is a closed convex set and that Assumptions 1 (fast enough sequential estimation of \mathbf{Q}) and 2 (bounded reward function) are satisfied, then \mathcal{C} is approachable if and only if*

$$\forall (\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \{0,1\}} \exists (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \text{s.t.} \quad \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x,s) \in \mathcal{C}. \quad (2)$$

In this case, the strategy of Eq. (1) achieves the following rates for L^2 and almost-sure convergences:

$$\begin{aligned} \mathbb{E}[d_T^2] &\leq \sqrt{\frac{K}{T}} + 4\|\mathbf{m}\|_{\infty,2} \overbrace{\frac{1}{T} \sum_{t=1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}}^{:= \bar{\Delta}_T} \quad \text{and} \\ \mathbb{P}\left(\sup_{t \geq T} d_t \geq \varepsilon\right) &\leq \frac{3K}{T\varepsilon^2} + \frac{16\|\mathbf{m}\|_{\infty,2}}{\varepsilon^2} \left(\sqrt{\frac{K}{T-1}} + 2 \left(\sup_{t \geq T} \bar{\Delta}_t \right) \left(\bar{\Delta}_T + \sum_{t \geq T} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} \right) \right) \end{aligned}$$

where $K < +\infty$ denotes the maximal distance to \mathcal{C} of an element of the compact set $\mathbf{m}(\mathcal{A}, \mathcal{B}, \mathcal{X}, \mathcal{S})$.

Proof of Theorem 1. Part I: Necessity. Assume that the condition in Eq. (2) is not satisfied, then

$$\exists (\mathbf{q}_0^{G(x,s)})_{(x,s) \in \mathcal{X} \times \mathcal{S}} \forall (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \text{s.t.} \quad \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}_0^{G(x,s)}, x, s) d\mathbf{Q}(x,s) \notin \mathcal{C}.$$

Since \mathcal{C} is closed and by continuity of the norm, there exists $\alpha > 0$ such that

$$\forall (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \min_{\mathbf{v} \in \mathcal{C}} \left\| \mathbf{v} - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}_0^{G(x,s)}, x, s) d\mathbf{Q}(x,s) \right\| \geq \alpha. \quad (9)$$

Let Nature play using this distribution $(\mathbf{q}_0^{G(x,s)})_{(x,s)}$ at each stage $t \geq 1$ to draw b_t . Given that the sensitive attributes and contexts (x_t, s_t) are drawn i.i.d., the conditional expectation of the reward of the player at round $t \geq 1$ based on the history $H_{t-1} = (a_{t'}, b_{t'}, x_{t'}, s_{t'})_{t' \leq t-1}$ equals

$$\mathbb{E}[\mathbf{m}(a_t, b_t, x_t, s_t) \mid H_{t-1}] = \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_t^x, \mathbf{q}_0^{G(x,s)}, x, s) d\mathbf{Q}(x, s).$$

Then, for any strategy of the player, it holds by martingale convergence (e.g., by the Hoeffding-Azuma inequality and the Borel–Cantelli lemma, used for each component of \mathbf{m}) that

$$\left\| \frac{1}{T} \sum_{t=1}^T \mathbf{m}(a_t, b_t, x_t, s_t) - \frac{1}{T} \sum_{t=1}^T \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_t^x, \mathbf{q}_0^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \right\| \longrightarrow 0 \quad \text{a.s.}$$

Set $\bar{\mathbf{p}}_T^x := \frac{1}{T} \sum_{t=1}^T \mathbf{p}_t^x$, then the above implies that

$$\left\| \bar{\mathbf{m}}_T - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\bar{\mathbf{p}}_T^x, \mathbf{q}_0^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \right\| \longrightarrow 0 \quad \text{a.s.} \quad (10)$$

By the triangle inequality for the Euclidean norm, Eqs. (9) and (10) entail that

$$\liminf_{T \rightarrow +\infty} d(\bar{\mathbf{m}}_T, \mathcal{C}) = \liminf_{T \rightarrow +\infty} \|\bar{\mathbf{m}}_T - \bar{\mathbf{c}}_T\| = \liminf_{T \rightarrow +\infty} \min_{\mathbf{v} \in \mathcal{C}} \|\mathbf{v} - \bar{\mathbf{m}}_T\| \geq \alpha \quad \text{a.s.}$$

That is, Nature prevents the player from approaching \mathcal{C} (and even: Nature approaches the complement of the α -neighborhood of \mathcal{C}).

Note that in this part we did not use that the target set \mathcal{C} was convex, only that it was a closed set.

Part II: Sufficiency. Recall that we denoted by $d_t := \|\bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t\|_2$ the Euclidean distance of $\bar{\mathbf{m}}_t$ to \mathcal{C} . Observe that by definition of the projections $\bar{\mathbf{c}}_{t+1}$ and $\bar{\mathbf{c}}_t$ and by expanding the square norm,

$$\begin{aligned} d_{t+1}^2 &\leq \|\bar{\mathbf{m}}_{t+1} - \bar{\mathbf{c}}_t\|^2 = \left\| \frac{t}{t+1} (\bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t) + \frac{1}{t+1} (\mathbf{m}_{t+1} - \bar{\mathbf{c}}_t) \right\|^2 \\ &= \left(\frac{t}{t+1} \right)^2 d_t^2 + \frac{\|\mathbf{m}_{t+1} - \bar{\mathbf{c}}_t\|^2}{(t+1)^2} + \frac{2t}{(t+1)^2} \langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \mathbf{m}_{t+1} - \bar{\mathbf{c}}_t \rangle. \end{aligned} \quad (11)$$

Moreover, we have, by definition of $(\mathbf{p}_{t+1}^x)_{x \in \mathcal{X}}$ in Eq. (1) as the argmin of a maximum,

$$\begin{aligned} &\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \mathbf{m}_{t+1} - \bar{\mathbf{c}}_t \rangle \\ &= \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \mathbf{m}_{t+1} - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) \right\rangle \\ &\quad + \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) - \bar{\mathbf{c}}_t \right\rangle \\ &\leq \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \mathbf{m}_{t+1} - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) \right\rangle \\ &\quad + \min_{(\mathbf{p}^x)_x} \max_{(\mathbf{q}^{G(x,s)})_{G(x,s)}} \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) - \bar{\mathbf{c}}_t \right\rangle. \end{aligned} \quad (12)$$

Furthermore, the Cauchy-Schwarz inequality, followed by an application of the bound of Eq. (8), indicates that for all $(\mathbf{p}^x)_x$ and all $(\mathbf{q}^{G(x,s)})_{(x,s)}$,

$$\begin{aligned} &\left| \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \right\rangle \right| \\ &\leq d_t \cdot \left\| \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\hat{\mathbf{Q}}_t(x, s) - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \right\| \\ &\leq 2d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \cdot \|\mathbf{m}\|_{\infty, 2}. \end{aligned}$$

Hence, using twice this bound in Eq. (12) and introducing

$$Z_{t+1} := \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \mathbf{m}_{t+1} - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \right\rangle, \quad (13)$$

we obtain

$$\begin{aligned} \langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \mathbf{m}_{t+1} - \bar{\mathbf{c}}_t \rangle &\leq Z_{t+1} + 4d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \cdot \|\mathbf{m}\|_{\infty,2} \\ &\quad + \min_{(\mathbf{p}^x)_x} \max_{(\mathbf{q}^{G(x,s)})_{G(x,s)}} \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) - \bar{\mathbf{c}}_t \right\rangle. \end{aligned} \quad (14)$$

We recall that the Euclidean projection \mathbf{c} of a vector \mathbf{n} onto a closed convex set $\mathcal{C} \subset \mathbb{R}^d$ satisfies:

$$\forall \mathbf{c}' \in \mathcal{C}, \quad \langle \mathbf{n} - \mathbf{c}, \mathbf{c}' - \mathbf{c} \rangle \leq 0.$$

Thus, thanks to von Neumann's minmax theorem (for the equality) and the Blackwell's condition in Eq. (2) together with the above-recalled property of the projection,

$$\begin{aligned} &\min_{(\mathbf{p}^x)_x} \max_{(\mathbf{q}^{G(x,s)})_{G(x,s)}} \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) - \bar{\mathbf{c}}_t \right\rangle \\ &= \max_{(\mathbf{q}^{G(x,s)})_{G(x,s)}} \min_{(\mathbf{p}^x)_x} \left\langle \bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) - \bar{\mathbf{c}}_t \right\rangle \leq 0. \end{aligned} \quad (15)$$

Hence, combining Eqs. (14) and (15) with Eq. (11), and bounding $\|\mathbf{m}_{t+1} - \bar{\mathbf{c}}_t\|^2$ by K (by definition of K), we have obtained so far

$$d_{t+1}^2 \leq \left(\frac{t}{t+1} \right)^2 d_t^2 + \frac{K}{(t+1)^2} + \frac{2t}{(t+1)^2} \left(Z_{t+1} + 4d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \cdot \|\mathbf{m}\|_{\infty,2} \right). \quad (16)$$

The $4d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \cdot \|\mathbf{m}\|_{\infty,2}$ is the sole difference to the standard proof of approachability. We deal with it by adapting the conclusions of the original proof.

Before we do so, we note that the Z_{t+1} introduced in Eq. (13) form a martingale difference sequence with respect to the history H_t : indeed, $\bar{\mathbf{m}}_t$ and $\bar{\mathbf{c}}_t$ are H_t -measurable and so are the $(\mathbf{p}_{t+1}^x)_x$ and the $(\mathbf{q}_{t+1}^{G(x,s)})_{x,s}$; since in addition (x_{t+1}, s_{t+1}) is drawn independently from everything according to \mathbf{Q} and a_{t+1} and b_{t+1} are drawn independently at random according to $\mathbf{p}_{t+1}^{x_t}$ and $\mathbf{q}^{G(x_t, s_t)}$, we have

$$\mathbb{E}[\mathbf{m}_{t+1} | H_t] = \mathbb{E}[\mathbf{m}(a_{t+1}, b_{t+1}, x_{t+1}, s_{t+1}) | H_t] = \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) d\mathbf{Q}(x, s),$$

so that $\mathbb{E}[Z_{t+1} | H_t] = 0$.

Part II: Sufficiency—convergence in L^2 . In particular, taking expectations in Eq. (16) and applying the tower rule (for the first inequality) and applying the Cauchy-Schwarz inequality (for the second inequality), we have

$$\begin{aligned} \mathbb{E}[d_{t+1}^2] &\leq \left(\frac{t}{t+1} \right)^2 \mathbb{E}[d_t^2] + \frac{K}{(t+1)^2} + \frac{8t\|\mathbf{m}\|_{\infty,2}}{(t+1)^2} \mathbb{E}[d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q})] \\ &\leq \left(\frac{t}{t+1} \right)^2 \mathbb{E}[d_t^2] + \frac{K}{(t+1)^2} + \frac{8t\|\mathbf{m}\|_{\infty,2}}{(t+1)^2} \sqrt{\mathbb{E}[d_t^2]} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}. \end{aligned} \quad (17)$$

Applying Lemma 1 below, we get

$$\sqrt{\mathbb{E}[d_T^2]} \leq B_T := \sqrt{\frac{K}{T}} + 4\|\mathbf{m}\|_{\infty,2} \underbrace{\frac{1}{T} \sum_{t=1}^{T-1} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}}_{:= \bar{\Delta}_T}. \quad (18)$$

By (a consequence of) Assumption 1, the second term in the right-hand side converges to zero, and we obtain convergence in L^2 .

Part II: Sufficiency—almost-sure convergence. We define

$$S_T := d_T^2 + \mathbb{E} \left[\sum_{t \geq T} \left(\frac{K}{(t+1)^2} + \frac{8t \|\mathbf{m}\|_{\infty,2}}{(t+1)^2} d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \right) \middle| H_T \right],$$

and note that $(S_T)_{T \geq 1}$ is a non-negative super-martingale with respect to the filtration induced by $(H_T)_{T \geq 1}$; indeed, the recursion of Eq. (16) entails, together with $(t/(t+1))^2 \leq 1$ and $\mathbb{E}[Z_{T+1} | H_T] = 0$:

$$\begin{aligned} \mathbb{E}[S_{T+1} | H_T] &= \mathbb{E}[d_{T+1}^2 | H_T] + \mathbb{E} \left[\sum_{t \geq T+1} \left(\frac{K}{(t+1)^2} + \frac{8t \|\mathbf{m}\|_{\infty,2}}{(t+1)^2} d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \right) \middle| H_T \right] \\ &\leq d_T^2 + \mathbb{E} \left[\sum_{t \geq T} \left(\frac{K}{(t+1)^2} + \frac{8t \|\mathbf{m}\|_{\infty,2}}{(t+1)^2} d_t \cdot \text{TV}(\hat{\mathbf{Q}}_t, \mathbf{Q}) \right) \middle| H_T \right] = S_T. \end{aligned}$$

We may thus use $d_T^2 \leq S_T$ and apply Doob's maximal inequality for non-negative super-martingales (Lemma 2):

$$\mathbb{P} \left(\sup_{T' \geq T} d_{T'} \geq \varepsilon \right) = \mathbb{P} \left(\sup_{T' \geq T} d_{T'}^2 \geq \varepsilon^2 \right) \leq \mathbb{P} \left(\sup_{T' \geq T} S_{T'} \geq \varepsilon^2 \right) \leq \frac{\mathbb{E}[S_T]}{\varepsilon^2}.$$

The proof is concluded by upper bounding $\mathbb{E}[S_T]$. The tower rule, the Cauchy-Schwarz inequality, and the bound $t/(t+1)^2 \leq 1/(t+1) \leq 1/t$ yield

$$\mathbb{E}[S_T] \leq \mathbb{E}[d_T^2] + \sum_{t \geq T} \frac{K}{(t+1)^2} + 8 \|\mathbf{m}\|_{\infty,2} \sum_{t \geq T} \sqrt{\mathbb{E}[d_t^2]} \frac{\sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}}{t}.$$

We substitute the bound from Eq. (18), keeping in mind that the total variation distance is always smaller than 1:

$$\mathbb{E}[S_T] \leq \mathbb{E}[d_T^2] + \frac{K}{T} + 8 \|\mathbf{m}\|_{\infty,2} \sum_{t \geq T} \frac{1}{t} \sqrt{\frac{K}{t}} + 32 \|\mathbf{m}\|_{\infty,2}^2 \sum_{t \geq T} \bar{\Delta}_t \frac{\sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]}}{t}.$$

Eq. (18) also implies, together with $(a+b)^2 \leq 2a^2 + 2b^2$, that $\mathbb{E}[d_T^2] \leq 2K/T + 32 \|\mathbf{m}\|_{\infty,2}^2 (\bar{\Delta}_T)^2$. All in all, we get the final bound

$$\mathbb{E}[S_T] \leq \frac{3K}{T} + \frac{16 \|\mathbf{m}\|_{\infty,2} \sqrt{K}}{\sqrt{T-1}} + 32 \|\mathbf{m}\|_{\infty,2}^2 \left(\sup_{t \geq T+1} \bar{\Delta}_t \right) \left(\bar{\Delta}_T + \sum_{t \geq T} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\hat{\mathbf{Q}}_t, \mathbf{Q})]} \right).$$

□

A.1 Two lemmas used in the proof of Theorem 1

The following lemma is an ad-hoc and new, but elementary, tool to deal with the additional term appearing in Eq. (17) compared to the original proof of approachability.

Lemma 1. *Let $t^* \geq 0$, and consider two non-negative sequences $(d_t)_{t \geq t^*}$ and $(\delta_t)_{t \geq t^*}$ fulfilling, for $t \geq t^*$, the recursive inequality*

$$d_{t+1}^2 \leq \left(\frac{t}{t+1} \right)^2 d_t^2 + \frac{K}{(t+1)^2} + \frac{2t}{(t+1)^2} \delta_t d_t. \quad (19)$$

Then, for all $t \geq t^* + 1$,

$$d_t \leq \frac{\sqrt{K(t-t^*)}}{t} + \frac{1}{t} \sum_{t'=t^*}^{t-1} \delta_{t'} + \frac{t^* d_{t^*}}{t}.$$

In particular, if $(d_t)_{t \geq 1}$ and $(\delta_t)_{t \geq 1}$ are two non-negative sequences fulfilling the recursive inequality (19) for $t \geq 1$, and if $d_1 \leq \sqrt{K}$, then, for all $t \geq 1$,

$$d_t \leq \sqrt{\frac{K}{t}} + \frac{1}{t} \sum_{t'=1}^{t-1} \delta_{t'}.$$

Proof. Second part of the lemma. Let us first check that the second part of the lemma follows from the first part. Setting $d_0 = \delta_0 = 0$, the sequences $(d_t)_{t \geq 0}$ and $(\delta_t)_{t \geq 0}$ fulfill Eq. (19) for $t \geq t^* = 0$, hence

$$d_t \leq \sqrt{\frac{K}{t}} + \frac{1}{t} \sum_{t'=0}^{t-1} \delta_{t'} = \sqrt{\frac{K}{t}} + \frac{1}{t} \sum_{t'=1}^{t-1} \delta_{t'},$$

where the equality in the right-hand side comes from $\delta_0 = 0$.

First part of the lemma. Set $U_t = t d_t$ and $\Delta_t^* = \delta_{t^*} + \dots + \delta_t$ with the convention that $\Delta_t^* = 0$ for all $t < t^*$. It is equivalent to prove that for all $t \geq t^* + 1$, we have

$$U_t \leq \sqrt{K(t - t^*)} + \Delta_{t-1}^* + U_{t^*}. \quad (20)$$

We observe that Eq. (20) trivially holds for $t = t^*$. Assume that Eq. (20) holds for $t \geq t^*$. By assumption, we have $U_{t+1}^2 \leq U_t^2 + K + 2\delta_t U_t \leq (U_t + \delta_t)^2 + K$. Substituting Eq. (20) together with the fact that $U_t \geq 0$ and $\delta_t \geq 0$, we get

$$\begin{aligned} U_{t+1}^2 &\leq (U_t + \delta_t)^2 + K \leq (\sqrt{K(t - t^*)} + \Delta_t^* + U_{t^*})^2 + K \\ &= K(t + 1 - t^*) + (\Delta_t^* + U_{t^*})^2 + 2\sqrt{K(t - t^*)} (\Delta_t^* + U_{t^*}) \\ &\leq (\sqrt{K(t + 1 - t^*)} + \Delta_t^* + U_{t^*})^2. \end{aligned}$$

We have proved that $U_{t+1} \leq \sqrt{K(t + 1 - t^*)} + \Delta_t^* + U_{t^*}$, and we conclude by induction. \square

Two maximal inequalities for martingales are called Doob's inequality. We use the less famous one, for non-negative super-martingales.

Lemma 2 (One of Doob's maximal inequalities). *Let $(S_n)_{n \geq 1}$ be a non-negative super-martingale, then*

$$\mathbb{P}\left(\sup_{m \geq n} S_m \geq \eta\right) \leq \frac{\mathbb{E}[S_n]}{\eta}.$$

B Proofs for Section 4

We first detail the two counter-examples alluded at in the proof of Proposition 1, relative to Example 2 on group-wise no-regret. We then discuss Example 3 on vanilla no-regret under the equalized average payoffs constraint.

B.1 Counter-examples for group-wise no-regret

First counter-example. We take $\mathcal{S} = \mathcal{A} = \mathcal{B} = \{0, 1\}$ and let \mathcal{X} be an arbitrary finite set. The monitoring is assumed to be $G(x, s) = x$. Finally, we consider the specific payoff function

$$\forall (a, b, x) \in \mathcal{A} \times \mathcal{B} \times \mathcal{X}, \quad r(a, b, x, 0) = a^2 \quad \text{and} \quad r(a, b, x, 1) = (a - 1)^2.$$

The integral conditions in Eq. (4) read: for all $a' \in \{0, 1\}$,

$$\int_{\mathcal{X}} \sum_{a \in \{0, 1\}} \mathbf{p}^x(a) a^2 d\mathbf{Q}^0(x) \geq (a')^2 \quad \text{and} \quad \int_{\mathcal{X}} \sum_{a \in \{0, 1\}} \mathbf{p}^x(a) (a - 1)^2 d\mathbf{Q}^1(x) \geq (a' - 1)^2,$$

or equivalently, simply

$$\int_{\text{supp}(\mathbf{Q}^0)} \mathbf{p}^x(1) d\mathbf{Q}^0(x) \geq 1 \quad \text{and} \quad \int_{\text{supp}(\mathbf{Q}^1)} \mathbf{p}^x(0) d\mathbf{Q}^1(x) \geq 1.$$

As $\mathbf{p}^x(1) \in [0, 1]$, the fact that the first integral above is larger than 1 entails that $\mathbf{p}^x(1) = 1$ on $\text{supp}(\mathbf{Q}^0)$. Similarly, $\mathbf{p}^x(0) = 1$ on $\text{supp}(\mathbf{Q}^1)$. As we also have $\mathbf{p}^x(0) + \mathbf{p}^x(1) = 1$ for all $x \in \mathcal{X}$, we see that the condition in Eq. (4) cannot hold as soon as $\text{supp}(\mathbf{Q}^0) \cap \text{supp}(\mathbf{Q}^1) \neq \emptyset$.

Second counter-example. Again, we take $\mathcal{S} = \mathcal{A} = \mathcal{B} = \{0, 1\}$ and let \mathcal{X} be an arbitrary finite set but assume this time that Nature's monitoring is $G(x, s) = (x, s)$. Another difference is that we consider a payoff function not depending on s :

$$\forall (a, b, x, s) \in \mathcal{A} \times \mathcal{B} \times \mathcal{X} \times \mathcal{S}, \quad r(a, b, x, s) = \mathbb{I}\{a = b\} = 1 - (a - b)^2.$$

Nature picks the following difficult family of distributions: $\mathbf{q}^{(x, 0)} = (1, 0)^\top$ and $\mathbf{q}^{(x, 1)} = (0, 1)^\top$ for all $x \in \mathcal{X}$, so that $\mathbf{q}^{(x, s)}(b) = 1$ if and only if $b = s$. The integral conditions in Eq. (4) therefore read: for all $s \in \{0, 1\}$,

$$\min_{a' \in \{0, 1\}} \int_{\mathcal{X}} \sum_{a \in \{0, 1\}} \mathbf{p}^x(a) (r(a, s, x, s) - r(a', s, x, s)) d\mathbf{Q}^s(x) = \int_{\mathcal{X}} \mathbf{p}^x(s) d\mathbf{Q}^s(x) - 1 \geq 0.$$

From here we conclude similarly to the previous counter-example.

B.2 Vanilla no-regret under the equalized average payoffs constraint

Blum et al. [4, Section 4] study online regret minimization under a constraint of equal average payoffs, that is, they discuss the $(\mathbf{m}_{\text{reg}}, \mathbf{m}_{\text{eq-pay}})$ -approachability of $\mathcal{C}_{\text{reg}} \times \mathcal{C}_{\text{eq-pay}}$, with the notation of Section 2.

Their setting is different from the setting considered in this article, as the latter relies on the no-regret based on a fixed base payoff function r , while the former considers prediction with expert advice, that may be assimilated to an adversarially chosen sequence (r_t) of payoff functions.

Yet, we mimic the spirit of their results, which is two-fold.

First, we show an impossibility result for the simultaneous satisfaction of the vanilla no-regret objective and the constraint of equal average payoffs, i.e., for the $(\mathbf{m}_{\text{reg}}, \mathbf{m}_{\text{eq-pay}})$ -approachability of $\mathcal{C}_{\text{reg}} \times \mathcal{C}_{\text{eq-pay}}$. We do so for an example of binary online classification. This corresponds to Theorem 4 of Blum et al. [4, Section 4].

Second, we provide a positive result for the mentioned approachability problem, in the case of a Player aware of the sensitive contexts, i.e., following Remark 1, when the Player accesses the contexts $x'_t = (x_t, s_t)$. This corresponds to Theorem 3 of Blum et al. [4, Section 4].

Before we do so, we first instantiate the approachability condition of Eq. (2) with the vector payoff function $\mathbf{m} = (\mathbf{m}_{\text{reg}}, \mathbf{m}_{\text{eq-pay}})$ and the target set $\mathcal{C} = \mathcal{C}_{\text{reg}} \times \mathcal{C}_{\text{eq-pay}}$; it reads: $\forall (\mathbf{q}^{G(x,s)})_{(x,s)} \exists (\mathbf{p}^x)_x$ such that

$$\int_{\mathcal{X} \times \{0,1\}} r(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \geq \max_{a' \in \mathcal{A}} \int_{\mathcal{X} \times \{0,1\}} r(a', \mathbf{q}^{G(x,s)}, x, s) d\mathbf{Q}(x, s) \quad (21)$$

$$\text{and} \quad \left| \int_{\mathcal{X}} r(\mathbf{p}^x, \mathbf{q}^{G(x,0)}, x, 0) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} r(\mathbf{p}^x, \mathbf{q}^{G(x,1)}, x, 1) d\mathbf{Q}^1(x) \right| \leq \varepsilon. \quad (22)$$

Second, we also introduce some additional notation.

Additional notation and reminder on total variation distance. Recall that we denoted by \mathbf{Q}^0 and \mathbf{Q}^1 the two marginals of \mathbf{Q} on \mathcal{X} . We fix some measure μ which dominates both \mathbf{Q}^0 and \mathbf{Q}^1 , e.g., $\mu = \mathbf{Q}^0 + \mathbf{Q}^1$, and denote by g_0 and g_1 densities of \mathbf{Q}^0 and \mathbf{Q}^1 with respect to μ . We introduce the following three sets (defined up to μ -neglectable events):

$$\begin{aligned} \mathcal{X}_0 &= \{x \in \mathcal{X} : g_0(x) > g_1(x)\}, \\ \mathcal{X}_1 &= \{x \in \mathcal{X} : g_1(x) > g_0(x)\}, \\ \mathcal{X}_= &= \{x \in \mathcal{X} : g_1(x) = g_0(x)\}. \end{aligned}$$

Using the above defined sets and densities, we remind that the total variation distance between \mathbf{Q}^0 and \mathbf{Q}^1 can be expressed in the following equivalent ways (see, e.g., Devroye [8] or Tsybakov [28, Lemma 2.1]):

$$\begin{aligned} \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) &= \frac{1}{2} \int_{\mathcal{X}} |g_0(x) - g_1(x)| d\mu(x) \\ &= \int_{\mathcal{X}_1} (g_1(x) - g_0(x)) d\mu(x) = \int_{\mathcal{X}_0} (g_0(x) - g_1(x)) d\mu(x) \\ &= 1 - \int_{\mathcal{X}} \min\{g_0(x), g_1(x)\} d\mu(x). \end{aligned}$$

We may now describe the impossibility example.

Impossibility example for online classification. Binary classification corresponds to the sets of actions $\mathcal{A} = \mathcal{B} = \{0, 1\}$ and to the payoff function $r(a, b, x, s) = \mathbb{I}\{a = b\}$. In particular, for all distributions \mathbf{q} and \mathbf{q} , for all contexts (x, s) ,

$$r(\mathbf{p}, \mathbf{q}, x, s) = \mathbf{p}(0) \mathbf{q}(0) + \mathbf{p}(1) \mathbf{q}(1) := \langle \mathbf{p}, \mathbf{q} \rangle.$$

We focus our attention on the monitoring $G(x, s) = x$, which gives less freedom to Nature. Our impossibility result holds in particular in the case of the more complete monitoring $G(x, s) = (x, s)$.

We will have Nature pick distributions $(\mathbf{q}^x)_{x \in \mathcal{X}}$ such that $\mathbf{q}^x(0) > 1/2$ for all $x \in \mathcal{X}$; the maximum in the right-hand side of Eq. (21) is then achieved for $a' = 0$. Because of this and with the notion introduced, the regret criterion of Eq. (21) may be rewritten as

$$\int_{\mathcal{X} \times \{0,1\}} \underbrace{(\langle \mathbf{p}^x, \mathbf{q}^x \rangle - \mathbf{q}^x(0))}_{\leq 0} d\mathbf{Q}(x, s) \geq 0.$$

The inequality $\langle \mathbf{p}^x, \mathbf{q}^x \rangle - \mathbf{q}^x(0) \leq 0$ holds because $\mathbf{q}^x(0) > \mathbf{q}^x(1)$ by the constraint $\mathbf{q}^x(0) > 1/2$; this inequality is strict unless $\mathbf{p}^x(1) = 0$. Therefore, the regret constraint imposes $\langle \mathbf{p}^x, \mathbf{q}^x \rangle = \mathbf{q}^x(0)$ and $\mathbf{p}^x(1) = 0$ on the support of \mathbf{Q} (which is the union of the supports of \mathbf{Q}^0 and \mathbf{Q}^1).

The constraint of equal average payoffs relies on the following difference, which we rewrite based on the equality just proved:

$$\int_{\mathcal{X}} \langle \mathbf{p}^x, \mathbf{q}^x \rangle d\mathbf{Q}^0(x) - \int_{\mathcal{X}} \langle \mathbf{p}^x, \mathbf{q}^x \rangle d\mathbf{Q}^1(x) = \int_{\mathcal{X}} \mathbf{q}^x(0) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} \mathbf{q}^x(0) d\mathbf{Q}^1(x).$$

We let Nature pick the distributions (\mathbf{q}^x) defined by

$$\mathbf{q}^x(0) = \begin{cases} 1 & \text{for } x \in \text{supp}(\mathbf{Q}^0), \\ 1/2 + \varepsilon & \text{for } x \in \text{supp}(\mathbf{Q}^1). \end{cases}$$

We also replace $d\mathbf{Q}^0$ and $d\mathbf{Q}^1$ by $g_0 d\mu$ and $g_1 d\mu$, respectively. The difference in average payoffs thus rewrites, given the various expressions of the total variation distance recalled above:

$$\begin{aligned} \int_{\mathcal{X}} \mathbf{q}^x(0) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} \mathbf{q}^x(0) d\mathbf{Q}^1(x) &= \int_{\mathcal{X}} \mathbf{q}^x(0) (g_0(x) - g_1(x)) d\mu(x) \\ &= \underbrace{\int_{\mathcal{X}_0} (g_0(x) - g_1(x)) d\mu(x)}_{=\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)} + \left(\frac{1}{2} + \varepsilon\right) \underbrace{\int_{\mathcal{X}_1 \cup \mathcal{X}_=} (g_0(x) - g_1(x)) d\mu(x)}_{=-\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)} \\ &= \left(\frac{1}{2} - \varepsilon\right) \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1). \end{aligned}$$

All in all, the equal average payoffs constraint of Eq. (22), and thus, the approachability condition of Eq. (2), hold if and only if

$$\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) \leq \frac{\varepsilon}{1/2 - \varepsilon},$$

i.e., if the distributions \mathbf{Q}^0 and \mathbf{Q}^1 are close enough.

This is typically not the case, and having such a small distance between \mathbf{Q}^0 and \mathbf{Q}^1 should be considered a degenerate case. The limit case $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) = 0$ indeed corresponds to the case when the sensitive attributes s_t are independent of the non-sensitive contexts x_t .

Positive result for a Player aware of the s_t and a fair-in-isolation payoff function r . The positive result will be exhibited in the same spirit as the one of Theorem 3 of Blum et al. [4, Section 4]. This spirit is interesting but somewhat limited, as it relies on a (heavy) fair-in-isolation assumption. The latter indeed indicates that for all sequences of contexts and observations, the average loss achieved by a given expert is the same among sensitive groups. This “for all sequences” requirement is particularly demanding. (A question not answered in Blum et al. [4] is the existence of experts that are fair in isolation, for general reward functions r_t or general loss functions ℓ_t , and metrics \mathcal{M} , using their notation.) See comments after Eq. (23) below for the adaptation of this assumption in our context.

A second ingredient for the positive result is that the Player accesses the sensitive contexts s_t . Following Remark 1, this translates into our setting by considering that the Player accesses the contexts $x'_t = (x_t, s_t)$; hence, the distributions picked by the Player will be indexed by (x, s) in this example. Nature’s monitoring is $G(x, s) = (x, s)$ as well.

Given $(\mathbf{q}^{(x,s)})$, to fulfill the no-regret condition

$$\int_{\mathcal{X} \times \{0,1\}} r(\mathbf{p}^{(x,s)}, \mathbf{q}^{(x,s)}, x, s) d\mathbf{Q}(x, s) \geq \max_{a' \in \mathcal{A}} \int_{\mathcal{X} \times \{0,1\}} r(a', \mathbf{q}^{(x,s)}, x, s) d\mathbf{Q}(x, s),$$

the Player may pick $\mathbf{p}^{(x,s)}$ based only on s :

$$\mathbf{p}^{(x,s)} = \text{dirac}(a^s), \quad \text{where} \quad a^s \in \arg \max_{a \in \mathcal{A}} \int_{\mathcal{X}} r(a, \mathbf{q}^{(x,s)}, x, s) d\mathbf{Q}^s(x).$$

This corresponds to using separate no-regret algorithms in the construction of Theorem 3 of Blum et al. [4], one algorithm per sensitive context. The no-regret algorithms based on approachability used here actually have a regret converging to zero in the limit (they do not just approach the set of non-negative numbers) and thus share the same “not worse but not better” property with respect to the best action a as the one used in Theorem 3 of Blum et al. [4].

The constraint of equal average payoffs requires that the following difference is smaller than ε in absolute values:

$$\begin{aligned} \int_{\mathcal{X}} r(\mathbf{p}^{(x,0)}, \mathbf{q}^{(x,0)}, x, 0) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} r(\mathbf{p}^{(x,1)}, \mathbf{q}^{(x,1)}, x, 1) d\mathbf{Q}^1(x) \\ = \max_{a \in \mathcal{A}} \int_{\mathcal{X}} r(a, \mathbf{q}^{(x,0)}, x, 0) d\mathbf{Q}^0(x) - \max_{a \in \mathcal{A}} \int_{\mathcal{X}} r(a, \mathbf{q}^{(x,1)}, x, 1) d\mathbf{Q}^1(x). \end{aligned} \quad (23)$$

This constraint is automatically taken care of by the fair-in-isolation assumption: its analogue in our context (keeping in mind that the actions $a \in \mathcal{A}$ play here the role of the experts in Blum et al. [4]) is to require that for all distributions $(\mathbf{q}^{(x,s)})$ picked by the opponent,

$$\forall a \in \mathcal{A}, \quad \left| \int_{\mathcal{X}} r(a, \mathbf{q}^{(x,0)}, x, 0) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} r(a, \mathbf{q}^{(x,1)}, x, 1) d\mathbf{Q}^1(x) \right| \leq \varepsilon.$$

This basically corresponds to an assumption on the effective range of the payoff function r and is therefore a heavy assumption, of limited interest.

C Proofs for Section 5

We recalled various expressions of the total variation distance in Section B.2. We keep the notation defined therein. The following inequality will be used repeatedly in our proofs.

Lemma 3. *For all Borel-measurable functions $f : \mathcal{X} \rightarrow [0, 1]$,*

$$\left| \int_{\mathcal{X}} f \, d\mathbf{Q}^0 - \int_{\mathcal{X}} f \, d\mathbf{Q}^1 \right| = \left| \int_{\mathcal{X}} f (g_0 - g_1) \, d\mu \right| \leq \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1).$$

Proof. The proof heavily relies on the fact that f takes values in $[0, 1]$. Since, by definitions of \mathcal{X}_0 , \mathcal{X}_1 , and $\mathcal{X}_=$,

$$\int_{\mathcal{X}} f (g_0 - g_1) \, d\mu = \int_{\mathcal{X}_0} \underbrace{f (g_0 - g_1)}_{>0} \, d\mu + \int_{\mathcal{X}_1} \underbrace{f (g_0 - g_1)}_{<0} \, d\mu,$$

we have

$$-\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) = - \int_{\mathcal{X}_1} (g_1 - g_0) \, d\mu \leq \int_{\mathcal{X}} f (g_0 - g_1) \, d\mu \leq \int_{\mathcal{X}_0} (g_0 - g_1) \, d\mu = \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1),$$

which concludes the proof. \square

An application of Lemma 3 is the bound $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$ on the $\psi(\dots)$ quantity of Eq. (5). Indeed,

$$\begin{aligned} \text{DP} &:= \psi \left(\int_{\mathcal{X} \times \{0,1\}} \mathbf{m}_{\text{DP}}(\mathbf{p}^x, s) \, d\mathbf{Q}(x, s) \right) \\ &= \left| \int_{\mathcal{X}} \sum_{k=1}^N \mathbf{p}^x(k) a^{(k)} \, d\mathbf{Q}^0(x) - \int_{\mathcal{X}} \sum_{k=1}^N \mathbf{p}^x(k) a^{(k)} \, d\mathbf{Q}^1(x) \right| \\ &= \left| \int_{\mathcal{X}} A(\mathbf{p}^x) \, d\mathbf{Q}^0(x) - \int_{\mathcal{X}} A(\mathbf{p}^x) \, d\mathbf{Q}^1(x) \right| \leq \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1), \end{aligned}$$

where we introduced

$$A(\mathbf{p}^x) := \sum_{k=1}^N \mathbf{p}^x(k) a^{(k)} \in [0, 1].$$

Before providing the proofs of Propositions 2 and 3 we introduce some additional short-hand notation. We set

$$t^* := \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1).$$

The objective function of the maxmin problem in (6), relative to group-wise calibration, is denoted by and equals

$$\begin{aligned} \text{GC} &:= \left\| \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}_{\text{gr-cal}}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}) \, d\mathbf{Q}(x, s) \right\|_1 \\ &= \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^x(k) (a^{(k)} - \mathbf{q}^{G(x,0)}(1)) g_0(x) \, d\mu(x) \right| \\ &\quad + \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^x(k) (a^{(k)} - \mathbf{q}^{G(x,1)}(1)) g_1(x) \, d\mu(x) \right|. \end{aligned}$$

The problem of Eq. (6) can now be written as

$$\varepsilon^*(\delta_\tau) = \max_{(\mathbf{q}^{G(x,s)})} \min_{(\mathbf{p}^x)} \{ \text{GC} : \text{DP} \leq \tau t^* \}. \quad (24)$$

The proof technique for each of Propositions 2 and 3 consists of two steps. First, by setting some convenient family $(\mathbf{q}^{G(x,s)})_{(x,s)}$, we obtain a lower bound on $\varepsilon^*(\delta_\tau)$. Second, by exhibiting some

convenient family $(\mathbf{p}^x)_x$, possibly based on the knowledge of $(\mathbf{q}^{G(x,s)})_{(x,s)}$, an upper bound on $\varepsilon^*(\delta_\tau)$ is derived.

The definitions of these families will be based on a rounding operator $p \in [0, 1] \mapsto \Pi_{\mathcal{A}} \in \mathcal{A}$, that maps a number $p \in [0, 1]$ to the closest element in the grid \mathcal{A} . Note that by definition of \mathcal{A} and $\Pi_{\mathcal{A}}$, it holds that $|p - \Pi_{\mathcal{A}}(p)| \leq 1/(2N)$ for all $p \in [0, 1]$. We are finally in the position of proving Propositions 2 and 3, and start with the former.

Proof of Proposition 2. Fix some $\tau \in [0, 1]$. Recall that $G(x, s) = (x, s)$, so that families of distributions picked by Nature are “truly” indexed by (x, s) .

For the lower bound on $\varepsilon^*(\delta_\tau)$, we consider, for all $x \in \mathcal{X}$,

$$\mathbf{q}^{(x,0)} = \text{dirac}(1) \quad \text{and} \quad \mathbf{q}^{(x,0)} = \text{dirac}(0).$$

Then, for all choices $(\mathbf{p}^x)_x$, using that $\sum_{k=1}^N \mathbf{p}^x(k) = 1$ for each $x \in \mathcal{X}$:

$$\begin{aligned} \text{GC} &= \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^x(k) (a^{(k)} - 1) g_0(x) \, d\mu(x) \right| + \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^x(k) a^{(k)} g_1(x) \, d\mu(x) \right| \\ &= \sum_{k=1}^N \int_{\mathcal{X}} \mathbf{p}^x(k) (1 - a^{(k)}) g_0(x) \, d\mu(x) + \sum_{k=1}^N \int_{\mathcal{X}} \mathbf{p}^x(k) a^{(k)} g_1(x) \, d\mu(x) \\ &= \underbrace{\int_{\mathcal{X}} g_0(x) \, d\mu(x)}_{=1} + \underbrace{\int_{\mathcal{X}} A(\mathbf{p}^x) (g_1(x) - g_0(x)) \, d\mu(x)}_{\text{absolute value equals DP}} \geq 1 - \text{DP}, \end{aligned}$$

so that the rewriting of Eq. (24) entails $\varepsilon^*(\delta_\tau) \geq 1 - \tau t^*$, as claimed.

To derive an upper bound on $\varepsilon^*(\delta_\tau)$, we consider, for each $(\mathbf{q}^{(x,s)})_{(x,s) \in \mathcal{X} \times \mathcal{S}}$ and each $x \in \mathcal{X}$,

$$\mathbf{p}^{\tau,x} = (1 - \tau) \cdot \text{dirac}(\Pi_{\mathcal{A}}(1/2)) + \tau \cdot \text{dirac}(f(x)),$$

$$\text{where} \quad f(x) = \begin{cases} \Pi_{\mathcal{A}}(\mathbf{q}^{(x,1)}(1)) & \text{if } x \in \mathcal{X}_1 \cup \mathcal{X}_=; \\ \Pi_{\mathcal{A}}(\mathbf{q}^{(x,0)}(1)) & \text{if } x \in \mathcal{X}_0. \end{cases} \quad (25)$$

Note that for this strategy of the Player, $\text{DP} \leq \tau t^*$; indeed, $A(\mathbf{p}^{\tau,x}) = (1 - \tau) \cdot \Pi_{\mathcal{A}}(1/2) + \tau \cdot f(x)$, so that

$$\begin{aligned} \text{DP} &= \left| \int_{\mathcal{X}} A(\mathbf{p}^{\tau,x}) \, d\mathbf{Q}^0(x) - \int_{\mathcal{X}} A(\mathbf{p}^{\tau,x}) \, d\mathbf{Q}^1(x) \right| = \tau \cdot \left| \int_{\mathcal{X}} f(x) \, d\mathbf{Q}^0(x) - \int_{\mathcal{X}} f(x) \, d\mathbf{Q}^1(x) \right| \\ &\leq \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1), \end{aligned}$$

where we applied Lemma 3 for the final inequality. Moreover, the choice of Eq. (25) ensures that $\text{GC} \leq 1 - \tau t^* + 1/N$, as we will prove below. This will lead to $\varepsilon^*(\delta_\tau) \leq 1 - \tau t^* + 1/N$ and will conclude the proof. Indeed,

$$\begin{aligned} \text{GC} &= \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^{\tau,x}(k) (a^{(k)} - \mathbf{q}^{(x,0)}(1)) g_0(x) \, d\mu(x) \right| \\ &\quad + \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^{\tau,x}(k) (a^{(k)} - \mathbf{q}^{(x,1)}(1)) g_1(x) \, d\mu(x) \right| \\ &= (1 - \tau) \left| \int_{\mathcal{X}} (\Pi_{\mathcal{A}}(1/2) - \mathbf{q}^{(x,0)}(1)) g_0(x) \, d\mu(x) \right| \\ &\quad + (1 - \tau) \left| \int_{\mathcal{X}} (\Pi_{\mathcal{A}}(1/2) - \mathbf{q}^{(x,1)}(1)) g_1(x) \, d\mu(x) \right| \\ &\quad + \tau \left| \int_{\mathcal{X}} (f(x) - \mathbf{q}^{(x,0)}(1)) g_0(x) \, d\mu(x) \right| + \tau \left| \int_{\mathcal{X}} (f(x) - \mathbf{q}^{(x,1)}(1)) g_1(x) \, d\mu(x) \right|. \end{aligned}$$

We replace $f(x)$ by its specific values and take care of all rounding operators $\Pi_{\mathcal{A}}$ by adding a $2 \times 1/(2N) = 1/N$ term after application of triangle inequalities:

$$\begin{aligned}
\text{GC} &\leq \frac{1}{N} + (1-\tau) \left| \int_{\mathcal{X}} \overbrace{(1/2 - \mathbf{q}^{(x,0)}(1))}^{\in[-1/2, 1/2]} g_0(x) \, d\mu(x) \right| + (1-\tau) \left| \int_{\mathcal{X}} \overbrace{(1/2 - \mathbf{q}^{(x,1)}(1))}^{\in[-1/2, 1/2]} g_1(x) \, d\mu(x) \right| \\
&\quad + \tau \left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} \overbrace{(\mathbf{q}^{(x,1)}(1) - \mathbf{q}^{(x,0)}(1))}^{\in[-1, 1]} g_0(x) \, d\mu(x) \right| + \tau \left| \int_{\mathcal{X}_0} \overbrace{(\mathbf{q}^{(x,0)}(1) - \mathbf{q}^{(x,1)}(1))}^{\in[-1, 1]} g_1(x) \, d\mu(x) \right| \\
&\leq \frac{1}{N} + \frac{1-\tau}{2} + \frac{1-\tau}{2} + \tau \int_{\mathcal{X}_1 \cup \mathcal{X}_=} g_0(x) \, d\mu(x) + \tau \int_{\mathcal{X}_0} g_1(x) \, d\mu(x) \\
&= 1 - \tau + \tau \underbrace{\int_{\mathcal{X}} \min\{g_0(x), g_1(x)\} \, d\mu(x)}_{=1-t^*} + \frac{1}{N} = 1 - \tau t^* + \frac{1}{N},
\end{aligned}$$

where we used one of the expressions of $t^* = \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$ in the last equality. \square

Proof of Proposition 3. Fix some $\tau \in [0, 1]$. Recall that $G(x, s) = x$, so that families of distributions picked by Nature are only indexed by x and may not depend on s .

For the lower bound on $\varepsilon^*(\delta_\tau)$, we consider $(\mathbf{q}^x)_{x \in \mathcal{X}}$ defined as

$$\mathbf{q}^x = \begin{cases} \text{dirac}(1) & \text{if } x \in \mathcal{X}_1 \cup \mathcal{X}_=; \\ \text{dirac}(0) & \text{if } x \in \mathcal{X}_0. \end{cases}$$

Then, for all choices $(\mathbf{p}^x)_x$, using the notation $A(\mathbf{p}^x)$ and the fact that $\sum_{k=1}^N \mathbf{p}^x(k) = 1$ for each $x \in \mathcal{X}$:

$$\begin{aligned}
\text{GC} &= \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^x(k) (a^{(k)} - \mathbf{q}^x(1)) g_0(x) \, d\mu(x) \right| + \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^x(k) (a^{(k)} - \mathbf{q}^x(1)) g_1(x) \, d\mu(x) \right| \\
&= \sum_{k=1}^N \left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} \mathbf{p}^x(k) (a^{(k)} - 1) g_0(x) \, d\mu(x) + \int_{\mathcal{X}_0} \mathbf{p}^x(k) a^{(k)} g_0(x) \, d\mu(x) \right| \\
&\quad + \sum_{k=1}^N \left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} \mathbf{p}^x(k) (a^{(k)} - 1) g_1(x) \, d\mu(x) + \int_{\mathcal{X}_0} \mathbf{p}^x(k) a^{(k)} g_1(x) \, d\mu(x) \right| \\
&\geq \left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} (A(\mathbf{p}^x) - 1) g_0(x) \, d\mu(x) + \int_{\mathcal{X}_0} A(\mathbf{p}^x) g_0(x) \, d\mu(x) \right| \\
&\quad + \left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} (A(\mathbf{p}^x) - 1) g_1(x) \, d\mu(x) + \int_{\mathcal{X}_0} A(\mathbf{p}^x) g_1(x) \, d\mu(x) \right| \\
&= \left| \int_{\mathcal{X}} A(\mathbf{p}^x) g_0(x) \, d\mu(x) - \int_{\mathcal{X}_1 \cup \mathcal{X}_=} g_0(x) \, d\mu(x) \right| \\
&\quad + \left| \int_{\mathcal{X}} A(\mathbf{p}^x) g_1(x) \, d\mu(x) - \int_{\mathcal{X}_1 \cup \mathcal{X}_=} g_1(x) \, d\mu(x) \right|, \\
&\geq \underbrace{\left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} (g_1(x) - g_0(x)) \, d\mu(x) \right|}_{=t^*} - \underbrace{\left| \int_{\mathcal{X}_1 \cup \mathcal{X}_=} A(\mathbf{p}^x) (g_1(x) - g_0(x)) \, d\mu(x) \right|}_{\text{absolute value equals DP}} \geq t^* - \text{DP}.
\end{aligned}$$

where all the inequalities follows from the triangle inequality. Eq. (24) entails $\varepsilon^*(\delta_\tau) \geq (1 - \tau)t^*$, as claimed.

To derive an upper bound on $\varepsilon^*(\delta_\tau)$, we consider, for each $(\mathbf{q}^x)_{x \in \mathcal{X}}$ and each $x \in \mathcal{X}$,

$$\mathbf{p}^{\tau,x} = (1 - \tau) \cdot \text{dirac}(\Pi_{\mathcal{A}}(Q)) + \tau \cdot \text{dirac}(\Pi_{\mathcal{A}}(\mathbf{q}^x(1)))$$

$$\text{where } Q = \int_{\mathcal{X}} \mathbf{q}^u(1) g_0(u) d\mu(u). \quad (26)$$

Note that for this strategy of the Player, $\text{DP} \leq \tau t^*$; indeed, $A(\mathbf{p}^{\tau,x}) = (1 - \tau) \cdot \Pi_{\mathcal{A}}(Q) + \tau \cdot \Pi_{\mathcal{A}}(\mathbf{q}^x(1))$, so that

$$\begin{aligned} \text{DP} &= \left| \int_{\mathcal{X}} A(\mathbf{p}^{\tau,x}) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} A(\mathbf{p}^{\tau,x}) d\mathbf{Q}^1(x) \right| \\ &= \tau \cdot \left| \int_{\mathcal{X}} \Pi_{\mathcal{A}}(\mathbf{q}^x(1)) d\mathbf{Q}^0(x) - \int_{\mathcal{X}} \Pi_{\mathcal{A}}(\mathbf{q}^x(1)) d\mathbf{Q}^1(x) \right| \leq \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1), \end{aligned}$$

where we applied Lemma 3 for the final inequality. Moreover, the choice of Eq. (26) ensures that $\text{GC} \leq (1 - \tau)t^* + 1/N$, as we will prove below. This will lead to $\varepsilon^*(\delta_\tau) \leq (1 - \tau)t^* + 1/N$ and will conclude the proof. Indeed,

$$\begin{aligned} \text{GC} &= \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^{\tau,x}(k) (a^{(k)} - \mathbf{q}^x(1)) g_0(x) d\mu(x) \right| + \sum_{k=1}^N \left| \int_{\mathcal{X}} \mathbf{p}^{\tau,x}(k) (a^{(k)} - \mathbf{q}^x(1)) g_1(x) d\mu(x) \right| \\ &= (1 - \tau) \left| \int_{\mathcal{X}} (\Pi_{\mathcal{A}}(Q) - \mathbf{q}^x(1)) g_0(x) d\mu(x) \right| + (1 - \tau) \left| \int_{\mathcal{X}} (\Pi_{\mathcal{A}}(Q) - \mathbf{q}^x(1)) g_1(x) d\mu(x) \right| \\ &\quad + \tau \left| \int_{\mathcal{X}} \underbrace{(\Pi_{\mathcal{A}}(\mathbf{q}^x(1)) - \mathbf{q}^x(1))}_{\leq 1/(2N)} g_0(x) d\mu(x) \right| + \tau \left| \int_{\mathcal{X}} \underbrace{(\Pi_{\mathcal{A}}(\mathbf{q}^x(1)) - \mathbf{q}^x(1))}_{\leq 1/(2N)} g_1(x) d\mu(x) \right|. \end{aligned}$$

Taking into account the rounding errors, i.e., replacing the two occurrences of $\Pi_{\mathcal{A}}(Q)$ by Q by adding twice a $(1 - \tau)/(2N)$ term, we get

$$\begin{aligned} \text{GC} &\leq (1 - \tau) \left| \int_{\mathcal{X}} (Q - \mathbf{q}^x(1)) g_0(x) d\mu(x) \right| + (1 - \tau) \left| \int_{\mathcal{X}} (Q - \mathbf{q}^x(1)) g_1(x) d\mu(x) \right| + \frac{1}{N} \\ &= (1 - \tau) \underbrace{\left| Q - \int_{\mathcal{X}} \mathbf{q}^x(1) g_0(x) d\mu(x) \right|}_{=0} + (1 - \tau) \left| Q - \int_{\mathcal{X}} \mathbf{q}^x(1) g_1(x) d\mu(x) \right| + \frac{1}{N} \\ &= (1 - \tau) \left| \int_{\mathcal{X}} \mathbf{q}^x(1) g_0(x) d\mu(x) - \int_{\mathcal{X}} \mathbf{q}^x(1) g_1(x) d\mu(x) \right| + \frac{1}{N} \leq (1 - \tau)t^* + 1/N, \end{aligned}$$

where the last equality holds by definition of Q as an integral and we applied Lemma 3 for the final inequality. \square

D Proofs for Section 6

In this section, we go over the results alluded at in Section 6. We first illustrate that Assumption 3 (which indicates that the target set \mathcal{C} should be estimated in some way) is realistic. We do so in Section D.1 by dealing with the most involved situation discussed in this article, namely, the example discussed at the beginning of Section 6. We then provide in Section D.2 a more complete statement of Theorem 2, with convergence rates, and prove it.

D.1 Assumption 3 is realistic

The beginning of Section 6 explained why and how performing an optimal trade-off between accuracy in group-calibration and unfairness in terms of demographic parity amounts to studying the $(\widetilde{\mathbf{m}}_{\text{gr-cal}}, \widetilde{\mathbf{m}}_{\text{DP}})$ -approachability of $\mathcal{C} = \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \widetilde{\mathcal{C}}_{\text{DP}}^\delta$, where $(\widetilde{\mathbf{m}}_{\text{gr-cal}}, \widetilde{\mathbf{m}}_{\text{DP}})$ is a known vector payoff function and where $\mathcal{C} := \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \widetilde{\mathcal{C}}_{\text{DP}}^\delta$ is unknown:

$$\widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon = \left\{ (\mathbf{v}_0, \mathbf{v}_1) \in \mathbb{R}^{2N} : \frac{\|\mathbf{v}_0\|_1}{\gamma_0} + \frac{\|\mathbf{v}_1\|_1}{\gamma_1} \leq \varepsilon \right\}, \quad \widetilde{\mathcal{C}}_{\text{DP}}^\delta = \left\{ (u, v) \in \mathbb{R}^2 : \left| \frac{u}{\gamma_0} - \frac{v}{\gamma_1} \right| \leq \delta \right\},$$

with $\varepsilon = (1 - \tau) \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$ and $\delta = \tau \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$ for some *known* $\tau \in [0, 1]$ but an unknown $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$, and with unknown probabilities γ_0, γ_1 . The parameter τ controls the desired trade-off between the calibration error and the discrepancy in demographic parity and thus is left as a parameter of user's choice.

We recall that the strategy of the Player proceeds in phases: at each time $T_r := 2^r$ for $r \geq 1$, the Player updates the estimate $\hat{\mathcal{C}}_r$ of \mathcal{C} . The focus of this section is to provide a sequence of estimates $\hat{\mathcal{C}}_r$ of \mathcal{C} fulfilling Assumption 3. The latter is a key requirement for the existence of an approachability strategy stated in Theorem 2. We must therefore prove that it is a realistic assumption.

The four requirements of Assumption 3. For the convenience of the reader, we restate the various requirements of Assumption 3, giving them nicknames, to be able to refer to them easily in the sequel: for all $r \geq 0$, the sets $\hat{\mathcal{C}}_r$

- (CC) are convex closed;
- (Proj-dist) satisfy $\|\mathbf{v} - \text{Proj}_{\hat{\mathcal{C}}_r}(\mathbf{v})\| \leq B$, for all $\mathbf{v} \in \mathbf{m}(\mathcal{A}, \mathcal{B}, \mathcal{X}, \{0, 1\})$;
- (Super-set) satisfy $\mathbb{P}(\mathcal{C} \subset \hat{\mathcal{C}}_r) \geq 1 - 1/(2T_r)$;
- (L2-Hausdorff) satisfy $\max\left\{ \mathbb{E}[d(\hat{\mathcal{C}}_r, \mathcal{C})^2], \mathbb{E}[d(\mathcal{C}, \hat{\mathcal{C}}_r)^2] \right\} \leq \beta_r^2$.

The constant $B < +\infty$ is independent of r and the sequence $(\beta_r)_{r \geq 0}$ is summable and non-increasing. The vector payoff function \mathbf{m} above refers to $(\widetilde{\mathbf{m}}_{\text{gr-cal}}, \widetilde{\mathbf{m}}_{\text{DP}})$.

(Proj-dist) requires that the distance of a possible vector payoff to sets $\hat{\mathcal{C}}_r$ are uniformly controlled. (Super-set) requires that the $\hat{\mathcal{C}}_r$ are, with high probability, super-sets of \mathcal{C} . Finally, (L2-Hausdorff) requires that some L^2 criterion of Hausdorff distance between sets is controlled. We will go over each of these requirements but first deepen our reduction scheme.

In the sequel, and as in the main body of the paper, we focus on the case where $\gamma_0 > 0$ and $\gamma_1 > 0$, i.e., there are two effective values for the sensitive contexts.

But first, a further reduction. Recall that the average vector payoff, described in Section 6, is equal to

$$\frac{1}{T} \sum_{t=1}^T (\widetilde{\mathbf{m}}_{\text{gr-cal}}(a_t, b_t, x_t, s_t), \widetilde{\mathbf{m}}_{\text{DP}}(a_t, b_t, x_t, s_t)),$$

where the first $2N$ components always lie in the interval $[-1, 1]$, while the last two ones lie in the interval $[0, 1]$. Therefore, in the definition of $\widetilde{\mathcal{C}}_{\text{DP}}^\delta$, we may restrict our attention to $(u, v) \in [0, 1]^2$ and use rather the alternative definition

$$\widetilde{\mathcal{C}}_{\text{DP}}^\delta := \left\{ (u, v) \in [0, 1]^2 : \left| \frac{u}{\gamma_0} - \frac{v}{\gamma_1} \right| \leq \delta \right\} = \left\{ (u, v) \in [0, 1]^2 : |\gamma_1 u - \gamma_0 v| \leq \gamma_0 \gamma_1 \delta \right\}.$$

As for $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$, given we are studying a calibration problem, we note that the ε of interest lie in $[0, 1]$ (with 0 included). Vectors $(\mathbf{v}_0, \mathbf{v}_1)$ of $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$ satisfy in particular that $\|\mathbf{v}_0\|_1 + \|\mathbf{v}_1\|_1 \leq \varepsilon$, which shows that $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \subseteq B_{\mathbb{R}^{2N}}^{\ell_1}$, where $B_{\mathbb{R}^{2N}}^{\ell_1} = \{\mathbf{v} \in \mathbb{R}^{2N} : \|\mathbf{v}\|_1 \leq 1\}$ is the unit ℓ_1 ball in \mathbb{R}^{2N} . Therefore,

$$\begin{aligned}\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon &= \left\{ (\mathbf{v}_0, \mathbf{v}_1) \in B_{\mathbb{R}^{2N}}^{\ell_1} : \frac{\|\mathbf{v}_0\|_1}{\gamma_0} + \frac{\|\mathbf{v}_1\|_1}{\gamma_1} \leq \varepsilon \right\} \\ &= \left\{ (\mathbf{v}_0, \mathbf{v}_1) \in B_{\mathbb{R}^{2N}}^{\ell_1} : \gamma_1 \|\mathbf{v}_0\|_1 + \gamma_0 \|\mathbf{v}_1\|_1 \leq \gamma_0 \gamma_1 \varepsilon \right\}.\end{aligned}$$

Plug-in estimation of \mathcal{C} . We consider estimators $\hat{\gamma}_{0,t}, \hat{\gamma}_{1,t} \in [0, 1]$ of γ_0, γ_1 and an estimator $\widehat{M}_t \in [0, 1]$ of $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$, based on the first t i.i.d. samples from \mathbf{Q} , see the end of the section for examples. We substitute them in the definitions of the target sets. We actually perform a careful such substitution by considering possibly data-dependent parameters $\alpha_1(t) \in (0, 1]$ and $\alpha_2(t) \in (0, 1]$, to be specified by the analysis, that will provide the needed upper confidence bounds (i.e., super-set condition). More precisely, we define estimators of $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$ and $\tilde{\mathcal{C}}_{\text{DP}}^\delta$ by

$$\begin{aligned}\widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon(t) &= \left\{ (\mathbf{v}_0, \mathbf{v}_1) \in B_{\mathbb{R}^{2N}}^{\ell_1} : \hat{\gamma}_{1,t} \|\mathbf{v}_0\|_1 + \hat{\gamma}_{0,t} \|\mathbf{v}_1\|_1 \leq \hat{\gamma}_{0,t} \hat{\gamma}_{1,t} \hat{\varepsilon}_t + \alpha_1(t) + 4\alpha_2(t) \right\}, \\ \widehat{\mathcal{C}}_{\text{DP}}^\delta(t) &= \left\{ (u, v) \in [0, 1] : |\hat{\gamma}_{1,t} u - \hat{\gamma}_{0,t} v| \leq \hat{\gamma}_{0,t} \hat{\gamma}_{1,t} \hat{\delta}_t + \alpha_1(t) + 4\alpha_2(t) \right\},\end{aligned}$$

where $\hat{\varepsilon}_t = (1 - \tau) \widehat{M}_t$ and $\hat{\delta}_t = \tau \widehat{M}_t$. We then set $\hat{\mathcal{C}}_r := \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon(T_r) \times \widehat{\mathcal{C}}_{\text{DP}}^\delta(T_r)$.

Requirements (CC) and (Proj-dist) hold. We observe that both $\widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon(t)$ and $\widehat{\mathcal{C}}_{\text{DP}}^\delta(t)$ are convex, closed, and bounded. The boundedness of these sets and the fact that $\mathbf{m} = (\widetilde{\mathbf{m}}_{\text{gr-cal}}, \widetilde{\mathbf{m}}_{\text{DP}})$ is bounded as well ensure the (Proj-dist) property.

Choice of $\alpha_1(T_r)$ and $\alpha_2(T_r)$, part 1. We introduce the following sets, indicating that some confidence bounds around the introduced estimators hold, of widths smaller than the introduced parameters $\alpha_1(T_r)$ and $\alpha_2(T_r)$. These sets need only to be considered at times T_r , where $r \geq 1$:

$$\Omega_{T_r}^{\alpha_1, \alpha_2} := \left\{ \left| \widehat{M}_{T_r} - \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) \right| \leq \alpha_1(T_r) \quad \text{and} \quad \forall s \in \{0, 1\}, |\hat{\gamma}_{s, T_r} - \gamma_s| \leq \alpha_2(T_r) \right\}.$$

We assume in the sequel that we could pick all $\alpha_1(T_r)$ and $\alpha_2(T_r)$ such that for all $r \geq 1$,

$$\mathbb{P}(\Omega_{T_r}^{\alpha_1, \alpha_2}) \geq 1 - \frac{1}{2T_r}, \quad (27)$$

and explain, in the final part of this section, how this can be ensured.

Requirement (Super-set) holds. It follows from the assumption above on the probability of $\Omega_{T_r}^{\alpha_1, \alpha_2}$ and from the following lemma.

Lemma 4. *On the event $\Omega_{T_r}^{\alpha_1, \alpha_2}$ defined above, it holds that $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \subseteq \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon(T_r)$ and $\tilde{\mathcal{C}}_{\text{DP}}^\delta \subseteq \widehat{\mathcal{C}}_{\text{DP}}^\delta(T_r)$, thus*

$$\mathcal{C} = \left(\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \tilde{\mathcal{C}}_{\text{DP}}^\delta \right) \subseteq \hat{\mathcal{C}}_r = \left(\widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon(T_r) \times \widehat{\mathcal{C}}_{\text{DP}}^\delta(T_r) \right).$$

Proof. For brevity, we drop the dependencies in T_r in the notation.

Part I: $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \subset \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$. We fix some $(\mathbf{v}_0, \mathbf{v}_1) \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$. By assumption,

$$\gamma_1 \|\mathbf{v}_0\|_1 + \gamma_0 \|\mathbf{v}_1\|_1 \leq (1 - \tau) \gamma_0 \gamma_1 \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1). \quad (28)$$

Furthermore, since $\|\mathbf{v}_0\|_1 + \|\mathbf{v}_1\|_1 \leq 1$, it holds on $\Omega^{\alpha_1, \alpha_2}$ that

$$\begin{aligned}\gamma_1 \|\mathbf{v}_0\|_1 + \gamma_0 \|\mathbf{v}_1\|_1 &\geq \hat{\gamma}_1 \|\mathbf{v}_0\|_1 + \hat{\gamma}_0 \|\mathbf{v}_1\|_1 - 2\alpha_2, \\ \gamma_0 \gamma_1 \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) &\leq \hat{\gamma}_0 \hat{\gamma}_1 \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1) + \alpha_2 \leq \dots \leq \hat{\gamma}_0 \hat{\gamma}_1 \cdot \widehat{M} + \alpha_1 + 2\alpha_2.\end{aligned} \quad (29)$$

Thus, in view of Eq. (28) and the definition of $\widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$, it holds that $(\mathbf{v}_0, \mathbf{v}_1) \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$.

Part II: $\tilde{\mathcal{C}}_{\text{DP}}^\delta \subset \hat{\mathcal{C}}_{\text{DP}}^\delta$. We fix some $(u, v) \in \tilde{\mathcal{C}}_{\text{DP}}^\delta$. By assumption, $u, v \in [0, 1]$ and

$$|\gamma_1 u - \gamma_0 v| \leq \tau \gamma_0 \gamma_1 \cdot \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1). \quad (30)$$

Furthermore, on $\Omega^{\alpha_1, \alpha_2}$,

$$|\gamma_1 u - \gamma_0 v| = |\hat{\gamma}_1 u - \hat{\gamma}_0 v + (\gamma_1 - \hat{\gamma}_1)u - (\gamma_0 - \hat{\gamma}_0)v| \geq |\hat{\gamma}_1 u - \hat{\gamma}_0 v| - 2\alpha_2.$$

In view of Eq. (30) and the second bound of Eq. (29), we conclude that $(u, v) \in \hat{\mathcal{C}}_{\text{DP}}^\delta$ on $\Omega^{\alpha_1, \alpha_2}$. \square

Requirement (L2-Hausdorff) holds. We bound separately the two expectations appearing in (L2-Hausdorff). As in the proof above, we omit the dependencies in T_r in the notation.

Part I: bound on $\mathbb{E}[d(\mathcal{C}, \hat{\mathcal{C}}_r)^2]$. By definition of d , given that we are dealing with Euclidean projections onto a product set and are bounding square Euclidean distances, we have the decomposition:

$$\begin{aligned} \mathbb{E}[d(\mathcal{C}, \hat{\mathcal{C}}_r)^2] &= \mathbb{E}\left[\sup_{\mathbf{x} \in \mathcal{C}} d(\mathbf{x}, \hat{\mathcal{C}}_r)^2\right] = \mathbb{E}\left[\sup_{\mathbf{x} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \tilde{\mathcal{C}}_{\text{DP}}^\delta} d(\mathbf{x}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \tilde{\mathcal{C}}_{\text{DP}}^\delta)^2\right] \\ &= \mathbb{E}\left[\sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2\right] + \mathbb{E}\left[\sup_{(u, v) \in \tilde{\mathcal{C}}_{\text{DP}}^\delta} d((u, v), \tilde{\mathcal{C}}_{\text{DP}}^\delta)^2\right]. \end{aligned} \quad (31)$$

We start with the first term in the right-hand side of (31). As $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$ always contains the null vector and $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \subseteq B_{\mathbb{R}^{2N}}^{\ell_1}$,

$$\sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 \leq \sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} \|\mathbf{v}\|^2 \leq \sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} \|\mathbf{v}\|_1 \leq 1. \quad (32)$$

In addition, Lemma 4 ensures that on $\Omega^{\alpha_1, \alpha_2}$ we have $\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \subset \hat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$, and hence $d(\tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon, \hat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon) = 0$, on $\Omega^{\alpha_1, \alpha_2}$. Thus, we can write

$$\begin{aligned} \mathbb{E}\left[\sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2\right] &= \mathbb{E}\left[\left(\mathbb{I}\{\Omega^{\alpha_1, \alpha_2}\} + (1 - \mathbb{I}\{\Omega^{\alpha_1, \alpha_2}\})\right) \sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2\right] \\ &\leq 0 + 1 - \mathbb{P}(\Omega^{\alpha_1, \alpha_2}) \leq \frac{1}{2T_r}, \end{aligned}$$

where the inequality comes from the assumption made in Eq. (27) combined with Eq. (32).

A bound $1/T_r$ on the second term of Eq. (31) follows similarly, using that

$$\sup_{(u, v) \in \tilde{\mathcal{C}}_{\text{DP}}^\delta} d((u, v), \tilde{\mathcal{C}}_{\text{DP}}^\delta)^2 \leq \sup_{(u, v) \in \tilde{\mathcal{C}}_{\text{DP}}^\delta} u^2 + v^2 \leq 2.$$

Hence, $\mathbb{E}[d(\mathcal{C}, \hat{\mathcal{C}}_r)^2] \leq 3/(2T_r)$.

Part II: bound on $\mathbb{E}[d(\hat{\mathcal{C}}_r, \mathcal{C})^2]$. We start in a similar manner:

$$\begin{aligned} \mathbb{E}[d(\hat{\mathcal{C}}_r, \mathcal{C})^2] &= \mathbb{E}\left[\sup_{\mathbf{x} \in \hat{\mathcal{C}}_r} d(\mathbf{x}, \mathcal{C})^2\right] = \mathbb{E}\left[\sup_{\mathbf{x} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \tilde{\mathcal{C}}_{\text{DP}}^\delta} d(\mathbf{x}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon \times \tilde{\mathcal{C}}_{\text{DP}}^\delta)^2\right] \\ &= \underbrace{\mathbb{E}\left[\sup_{\mathbf{v} \in \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2\right]}_{\leq 1 \text{ a.s.}} + \underbrace{\mathbb{E}\left[\sup_{(u, v) \in \tilde{\mathcal{C}}_{\text{DP}}^\delta} d((u, v), \tilde{\mathcal{C}}_{\text{DP}}^\delta)^2\right]}_{\leq 2 \text{ a.s.}}. \end{aligned} \quad (33)$$

As in Part I, we start with the first term in the right-side of (33) and we split the expectation into two parts

$$\mathbb{E} \left[\sup_{\mathbf{v} \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 \right] \leq \mathbb{E} \left[\mathbb{I}\{\Omega^{\alpha_1, \alpha_2}\} \sup_{\mathbf{v} \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 \right] + \frac{1}{2T_r}. \quad (34)$$

Let us upper-bound the right-hand side expectation. We introduce some local short-hand notation. Given two real numbers a, b , we denote by $a \vee b$ and $a \wedge b$ the maximum and minimum between a and b , respectively. We set $\alpha := \alpha_1 \vee \alpha_2$ and now show that

$$\text{on } \Omega^{\alpha_1, \alpha_2}, \quad \sup_{\mathbf{v} \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 \leq \alpha^{2/3} \left(\frac{81}{(\gamma_0 \gamma_1)^2} \vee \frac{10}{\gamma_0 \wedge \gamma_1} \right). \quad (35)$$

We fix some $\mathbf{v} = (\mathbf{v}_0, \mathbf{v}_1) \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$ and set $\mathbf{v}' = (\mathbf{v}'_0, \mathbf{v}'_1) := \lambda \mathbf{v}$ with

$$\lambda := 1 \wedge \left(\left(\frac{\gamma_0 \gamma_1 \varepsilon}{\gamma_0 \gamma_1 \varepsilon + 8\alpha} \right) \left(\frac{\hat{\gamma}_1}{\gamma_1} \wedge \frac{\hat{\gamma}_0}{\gamma_0} \right) \right).$$

The fact that $\mathbf{v} \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$ entails that on $\Omega^{\alpha_1, \alpha_2}$,

$$\hat{\gamma}_1 \|\mathbf{v}_0\|_1 + \hat{\gamma}_0 \|\mathbf{v}_1\|_1 \leq \hat{\gamma}_0 \hat{\gamma}_1 \varepsilon + 5\alpha \leq ((\gamma_0 + \alpha) \wedge 1) ((\gamma_1 + \alpha) \wedge 1) ((\varepsilon + \alpha) \wedge 1) + 5\alpha \leq \gamma_0 \gamma_1 \varepsilon + 8\alpha. \quad (36)$$

Here, and in what follows, we repeatedly use that $\gamma_0, \gamma_1, \varepsilon$ and their estimates all lie in $[0, 1]$. Furthermore, for the above-defined \mathbf{v}' , we can write on $\Omega^{\alpha_1, \alpha_2}$, by definition of λ ,

$$\gamma_1 \|\mathbf{v}'_0\|_1 + \gamma_0 \|\mathbf{v}'_1\|_1 \leq \lambda \left(\frac{\gamma_1}{\hat{\gamma}_1} \vee \frac{\gamma_0}{\hat{\gamma}_0} \right) \underbrace{(\hat{\gamma}_1 \|\mathbf{v}_0\|_1 + \hat{\gamma}_0 \|\mathbf{v}_1\|_1)}_{\leq \gamma_0 \gamma_1 \varepsilon + 8\alpha} \leq \gamma_0 \gamma_1 \varepsilon,$$

implying that $\mathbf{v}' \in \widehat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon$. Thus, $d(\mathbf{v}, \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon) \leq \|\mathbf{v} - \mathbf{v}'\| = (1 - \lambda) \|\mathbf{v}\|$ on $\Omega^{\alpha_1, \alpha_2}$. Since $\|\mathbf{v}\|_1 \leq 1$, we have $\|\mathbf{v}\| \leq \sqrt{\|\mathbf{v}\|_1} \leq 1$. All in all, we obtained the following upper bound on $\Omega^{\alpha_1, \alpha_2}$:

$$d(\mathbf{v}, \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 \leq \|\mathbf{v} - \mathbf{v}'\|^2 \leq (1 - \lambda)^2 \wedge \|\mathbf{v}\|_1.$$

We now bound separately each term to obtain the bound (35). First, on $\Omega^{\alpha_1, \alpha_2}$, we have

$$\begin{aligned} 1 \geq \lambda &\geq \left(\frac{\gamma_0 \gamma_1 \varepsilon}{\gamma_0 \gamma_1 \varepsilon + 8\alpha} \right) \left(\frac{\gamma_1 - \alpha}{\gamma_1} \wedge \frac{\gamma_0 - \alpha}{\gamma_0} \right) \\ &\geq \frac{\gamma_0 \gamma_1 \varepsilon - \alpha \varepsilon (\gamma_0 \vee \gamma_1)}{\gamma_0 \gamma_1 \varepsilon + 8\alpha} \geq 1 - \frac{9\alpha}{\gamma_0 \gamma_1 \varepsilon + 8\alpha} \geq 1 - \frac{9\alpha}{\gamma_0 \gamma_1 \varepsilon}, \end{aligned}$$

and thus,

$$(1 - \lambda)^2 \leq \frac{81\alpha^2}{(\gamma_0 \gamma_1 \varepsilon)^2}.$$

Second, for $\|\mathbf{v}\|_1$, we start from (36) and write

$$(\gamma_1 \wedge \gamma_0 - \alpha) \|\mathbf{v}\|_1 \leq (\hat{\gamma}_1 \wedge \hat{\gamma}_0) \|\mathbf{v}\|_1 \leq \hat{\gamma}_1 \|\mathbf{v}_0\|_1 + \hat{\gamma}_0 \|\mathbf{v}_1\|_1 \leq \gamma_0 \gamma_1 \varepsilon + 8\alpha,$$

from which we get $(\gamma_1 \wedge \gamma_0) \|\mathbf{v}\|_1 \leq \gamma_0 \gamma_1 \varepsilon + 9\alpha$, which in turn yields

$$\|\mathbf{v}\|_1 \leq \varepsilon + \frac{9\alpha}{\gamma_0 \wedge \gamma_1}.$$

The bound $(1 - \lambda)^2$ is convenient to use when $\varepsilon \geq \alpha^{2/3}$, while the bound on $\|\mathbf{v}\|_1$ will be used when $\varepsilon \leq \alpha^{2/3}$. When combining them by distinguishing these two cases, the \wedge symbol needs to be replaced by a \vee symbol, so, on $\Omega^{\alpha_1, \alpha_2}$

$$\begin{aligned} d(\mathbf{v}, \widetilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 &\leq (1 - \lambda)^2 \wedge \|\mathbf{v}\|_1 \leq \frac{81\alpha^2}{(\gamma_0 \gamma_1 \varepsilon)^2} \wedge \left(\varepsilon + \frac{9\alpha}{\gamma_0 \wedge \gamma_1} \right) \\ &\leq \frac{81\alpha^{2/3}}{(\gamma_0 \gamma_1)^2} \vee \left(\alpha^{2/3} + \frac{9\alpha}{\gamma_0 \wedge \gamma_1} \right), \end{aligned}$$

which entails the claimed bound (35), via $\alpha \leq \alpha^{2/3} \leq 1$. From (34) and (35), we get

$$\mathbb{E} \left[\sup_{\mathbf{v} \in \hat{\mathcal{C}}_{\text{gr-cal}}^\varepsilon} d(\mathbf{v}, \tilde{\mathcal{C}}_{\text{gr-cal}}^\varepsilon)^2 \right] \leq \frac{1}{2T_r} + C_{\gamma_0, \gamma_1}^{\text{gr-cal}} \mathbb{E}[\alpha^{2/3}],$$

for some constant $C_{\gamma_0, \gamma_1}^{\text{gr-cal}}$ only depending on γ_0 and γ_1 .

The second term of the decomposition (33) of $\mathbb{E}[d(\hat{\mathcal{C}}_r, \mathcal{C})^2]$ can be handled similarly, leading to the existence of a constant C_{γ_0, γ_1} , only depending on γ_0 and γ_1 , such that

$$\mathbb{E} \left[d(\hat{\mathcal{C}}_r, \mathcal{C})^2 \right] \leq \frac{3}{2T_r} + C_{\gamma_0, \gamma_1} \mathbb{E}[\alpha^{2/3}],$$

where the expectation in the right-hand side is due to the fact that α_1, α_2 might be data-dependent.

Combining Part I and Part II. The bound of Part II contains an additional term compared to the one of Part I. We have thus have proved so far (writing again the dependencies on T_r):

$$\max \left\{ \mathbb{E} [d(\hat{\mathcal{C}}_r, \mathcal{C})^2], \mathbb{E} [d(\mathcal{C}, \hat{\mathcal{C}}_r)^2] \right\} \leq \frac{3}{2T_r} + C_{\gamma_0, \gamma_1} \mathbb{E}[\alpha(T_r)^{2/3}]. \quad (37)$$

To get the desired property (L2-Hausdorff), we only need to make sure that the right hand side of (37) can be upper bounded by β_r^2 where (β_r) is non-increasing and summable. Recall that our proof also relied on the assumption (27). We now illustrate that indeed, $\alpha_1(T_r) \leq 1$ and $\alpha_2(T_r) \leq 1$ may be set in a way such that all these facts hold. For the sake of simplicity, we provide the illustration for the case of finite set \mathcal{X} .

Choice of $\alpha_1(T_r)$ and $\alpha_2(T_r)$, part 2: illustration for finite sets \mathcal{X} . Based on the T -sample $(x_t, s_t)_{1 \leq t \leq T}$ with distribution \mathbf{Q} , we denote by

$$N_{s,T} = \sum_{t=1}^T \mathbb{I}\{s_t = s\}$$

the number of occurrences of the value $s \in \{0, 1\}$ of the sensitive context, and consider the empirical frequencies $\hat{\gamma}_{0,T} = N_{0,T}/T$ and $\hat{\gamma}_{1,T} = N_{1,T}/T$ to estimate the frequencies γ_0 and γ_1 of the sensitive contexts.

The choice of \widehat{M}_T , and hence, the one of $\alpha_1(T)$, depend heavily on the possibly additional assumptions on the marginal distributions \mathbf{Q}^0 and \mathbf{Q}^1 . We illustrate such a choice for the case where \mathcal{X} is a finite set. In that case, we may consider the empirical distributions $\hat{\mathbf{Q}}_T^0$ and $\hat{\mathbf{Q}}_T^1$ for these marginals: for each $s \in \{0, 1\}$, $\hat{\mathbf{Q}}_T^s$ is some arbitrary distribution over \mathcal{X} (say, the uniform distribution) when $N_{s,T} = 0$, and otherwise, for each $x \in \mathcal{X}$,

$$\hat{\mathbf{Q}}_T^s(x) = \frac{1}{N_{s,T}} \sum_{t=1}^T \mathbb{I}\{x_t = x, s_t = s\}.$$

Then, we consider the plug-in estimate $\widehat{M}_T := \text{TV}(\hat{\mathbf{Q}}_T^0, \hat{\mathbf{Q}}_T^1)$ of $\text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)$.

Proof of (27), part I. We set

$$\alpha_2(T) = 1 \wedge \sqrt{\frac{\log(8T)}{2T}}$$

and note that by Hoeffding's inequality (and the fact that we only have two classes and that probabilities sum up to 1), for those T for which $\alpha_2(T) < 1$,

$$\begin{aligned} \mathbb{P} \left(\forall s \in \{0, 1\}, |\hat{\gamma}_{s,T} - \gamma_s| > \alpha_2(T) \right) &= \mathbb{P} \left(|\hat{\gamma}_{0,T} - \gamma_0| > \alpha_2(T) \right) \\ &\leq 2 \exp(-2T \alpha_2(T)^2) = \frac{1}{4T}. \end{aligned} \quad (38)$$

For T such that $\alpha_2(T) = 1$, the probability above is null, as $|\hat{\gamma}_{s,T} - \gamma_s| \leq 1$ a.s., and therefore, the final $1/(4T)$ bound holds in particular.

Proof of (27), part II. We set $\theta(0) = 1$ and $\theta(n) := \sqrt{\frac{|\mathcal{X}| + \log(8T)}{2n}}$ for $n \geq 1$, and define

$$\alpha_1(T) := 1 \wedge (\theta(N_{0,T}) + \theta(N_{1,T})).$$

We now prove that

$$\mathbb{P}\left(|\widehat{M}_T - \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)| > \alpha_1(T)\right) \leq \frac{1}{4T}. \quad (39)$$

The property (27) then follows from the bounds (38) and (39) at $T = T_r$.

Using that $|\widehat{M}_T - \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)| \leq 1$ a.s. (for the first inequality in the display below) and the triangle inequality

$$|\widehat{M}_T - \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)| \leq \text{TV}(\mathbf{Q}^0, \widehat{\mathbf{Q}}_T^0) + \text{TV}(\mathbf{Q}^1, \widehat{\mathbf{Q}}_T^1)$$

(for the second inequality in the display below), we have

$$\begin{aligned} \mathbb{P}\left(|\widehat{M}_T - \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)| > \alpha_1(T)\right) &= \mathbb{P}\left(|\widehat{M}_T - \text{TV}(\mathbf{Q}^0, \mathbf{Q}^1)| > \theta(N_{0,T}) + \theta(N_{1,T})\right) \\ &\leq \mathbb{P}\left(\text{TV}(\mathbf{Q}^0, \widehat{\mathbf{Q}}_T^0) + \text{TV}(\mathbf{Q}^1, \widehat{\mathbf{Q}}_T^1) > \theta(N_{0,T}) + \theta(N_{1,T})\right) \\ &\leq \sum_{s \in \{0,1\}} \mathbb{P}\left(\text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}_T^s) > \theta(N_{s,T})\right). \end{aligned} \quad (40)$$

The conclusion (39) follows from showing that for each $s \in \{0, 1\}$,

$$\mathbb{P}\left(\text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}_T^s) > \theta(N_{s,T})\right) \leq \frac{1}{8T}.$$

A useful auxiliary result to that end is the following. Denote by $\widehat{\mathbf{P}}^n$ the empirical frequencies of some probability distribution \mathbf{P} on \mathcal{X} based on a sample of deterministic size $n \geq 1$. Hoeffding's inequality and a union bound over the $\leq 2^{|\mathcal{X}|}$ subsets of \mathcal{X} ensure that for all $\theta \geq 0$,

$$\mathbb{P}\left(\text{TV}(\mathbf{P}, \widehat{\mathbf{P}}^n) \geq \theta\right) = \mathbb{P}\left(\max_{A \subset \mathcal{X}} (\mathbf{P}(A) - \widehat{\mathbf{P}}^n(A)) \geq \theta\right) \leq 2^{|\mathcal{X}|} \exp(-2n\theta^2). \quad (41)$$

In our case, note however that the estimators $\widehat{\mathbf{Q}}_T^s$ at time T are built on a random number $N_{s,T}$ of samples. We therefore decompose the probability of interest according to the values of $N_{s,T}$: for each $s \in \{0, 1\}$,

$$\begin{aligned} \mathbb{P}\left(\text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}_T^s) > \theta(N_{s,T})\right) &= \sum_{n=0}^T \mathbb{P}\left(N_{s,T} = n \text{ and } \text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}_T^s) > \theta(n)\right) \\ &= \sum_{n=1}^T \mathbb{P}\left(N_{s,T} = n \text{ and } \text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}_T^s) > \theta(n)\right) \\ &= \sum_{n=1}^T \mathbb{P}(N_{s,T} = n) \mathbb{P}\left(\text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}^{s,n}) > \theta(n)\right), \end{aligned}$$

where the second equality follows from the choice $\theta(0) = 1$ and the fact that a total variation is always smaller than 1, and where the third equality follows by conditional independence with $\widehat{\mathbf{Q}}^{s,n}$ denoting the empirical distribution based on a \mathbf{Q}^s -sample of size n . Substituting the bound (41) and the definition of the $\theta(n)$, we get

$$\begin{aligned} \mathbb{P}\left(\text{TV}(\mathbf{Q}^s, \widehat{\mathbf{Q}}_T^s) > \theta(N_{s,T})\right) &\leq 2^{|\mathcal{X}|} \sum_{n=1}^T \mathbb{P}(N_{s,T} = n) \exp(-2n\theta(n)^2) \\ &= 2^{|\mathcal{X}|} \sum_{n=1}^T \mathbb{P}(N_{s,T} = n) \exp(-|\mathcal{X}| - \log(8T)) \\ &\leq \underbrace{(2/e)^{|\mathcal{X}|}}_{\leq 1} \frac{1}{8T} \underbrace{\sum_{n=1}^T \mathbb{P}(N_{s,T} = n)}_{\leq 1} \leq \frac{1}{8T}, \end{aligned}$$

which is exactly what remained to be proven.

Control of the right-hand side of (37). It involves $\mathbb{E}[\alpha(T_r)^{2/3}]$, where

$$\begin{aligned}\alpha(T_r) &= \alpha_1(T_r) \vee \alpha_2(T_r) = \alpha_1(T_r) = 1 \wedge \left(\sqrt{\frac{|\mathcal{X}| + \log(8T)}{2N_{0,T}}} + \sqrt{\frac{|\mathcal{X}| + \log(8T)}{2N_{1,T}}} \right) \\ &\leq \sum_{s \in \{0,1\}} 1 \wedge \sqrt{\frac{|\mathcal{X}| + \log(8T)}{2N_{s,T}}}.\end{aligned}$$

Now, note that $N_{s,T}$ follows the binomial distribution with parameters γ_s and T . Thus, for each $s \in \{0, 1\}$,

$$\begin{aligned}\mathbb{E} \left[\left(\frac{|\mathcal{X}| + \log(8T)}{2N_{s,T}} \right)^{1/3} \wedge 1 \right] &\leq \mathbb{P}(N_{s,T} \leq T\gamma_s/2) + \left(\frac{|\mathcal{X}| + \log(8T)}{\gamma_s T} \right)^{1/3} \\ &\leq \exp(-\gamma_s^2 T/2) + \left(\frac{|\mathcal{X}| + \log(8T)}{\gamma_s T} \right)^{1/3},\end{aligned}$$

where we applied Hoeffding's inequality to get the last bound. Therefore, recalling that $T_r = 2^r$, the bound of Eq. (37) may be further bounded as: for all $r \geq 2$,

$$\frac{3}{2T_r} + C_{\gamma_0, \gamma_1} \mathbb{E}[\alpha(T_r)^{2/3}] \leq C'_{\gamma_0, \gamma_1} \left(\frac{r}{2^r} \right)^{1/3} =: \beta_r^2,$$

for some constant C'_{γ_0, γ_1} depending only on γ_0 and γ_1 . We observe that $\beta_r = \sqrt{C'_{\gamma_0, \gamma_1}} (r 2^{-r})^{1/6}$ is non-increasing for $r \geq 2$ and summable, as required.

We emphasize that, while exact values of $\alpha_1(T_r)$ and $\alpha_2(T_r)$ are needed for the construction of the set-estimate $\hat{\mathcal{C}}_r$, the knowledge of β_r is not required by the algorithm (its choice is required for the sake of the theoretical analysis only).

D.2 Proof of Theorem 2

We actually prove a more complete and more precise version of Theorem 2.

Theorem 3 (contains Theorem 2). *Under Assumption 3 and the assumptions of Theorem 1, a convex closed set \mathcal{C} , unknown to the Player, is \mathbf{m} -approachable if and only if Blackwell's condition in Eq. (2) is satisfied. In this case, the strategy of Eq. (7) is an approachability strategy. It achieves the following rates for L^2 convergence: for all $r \geq 1$ and all $t \in [T_r, T_{r+1} - 1]$,*

$$\sqrt{\mathbb{E}[d_t^2]} \leq \frac{\sqrt{6B^2 + 8B\|\mathbf{m}\|_{\infty,2}}}{(\sqrt{2}-1)\sqrt{t}} + \frac{4\|\mathbf{m}\|_{\infty,2}}{t} \sum_{t'=1}^{t-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} + \frac{4}{t} \sum_{r'=0}^r T_{r'} \beta_{r'}.$$

It also achieves the following rates for almost-sure convergence: for all $r \geq 1$,

$$\mathbb{P} \left(\sup_{t \geq T_r} d_t \geq 2\varepsilon \right) \leq \frac{\Xi_r}{\varepsilon^2} + \frac{1}{\varepsilon^2} \sum_{r' \geq r} \beta_{r'}^2,$$

where Ξ_r is defined in Eq. (60), page 40, and converges to 0.

Comments after Assumption 1 explain why the middle term in the L^2 bound vanishes. Assumption 3 indicates that the series $(\beta_r)_{r \geq 1}$ is summable, hence the following sequence of Cesaro averages built on it also vanishes:

$$\bar{\beta}_r := \frac{1}{T_{r+1}} \sum_{r'=0}^r T_{r'} \beta_{r'} \rightarrow 0. \quad (42)$$

Finally, the series $(\beta_r^2)_{r \geq 1}$ is also summable, hence its associated sequence of remainder sums also vanishes:

$$\sum_{r' \geq r} \beta_{r'}^2 \rightarrow 0.$$

We now move to the proof. We simply note at this stage that the condition $\|\mathbf{v} - \text{Proj}_{\hat{\mathcal{C}}_r}(\mathbf{v})\| \leq B$ for all $\mathbf{v} \in \mathbf{m}(\mathcal{A}, \mathcal{B}, \mathcal{X}, \{0, 1\})$ of Assumption 3 also holds, by convexity, for all \mathbf{v} in the convex hull of $\mathbf{m}(\mathcal{A}, \mathcal{B}, \mathcal{X}, \{0, 1\})$.

Proof. The proof is required only for the sufficiency, since the necessity was proven in Theorem 1.

Recall that, from the perspective of the Player, the game proceeds in phases lasting from $T_r := 2^r$ to $T_{r+1} - 1 := 2^{r+1} - 1$. For each time $t \in [T_r, T_{r+1} - 1]$, the Player uses $\hat{\mathcal{C}}_r$ as an estimate of the true target set \mathcal{C} , and updates to $\hat{\mathcal{C}}_{r+1}$ only at $t = T_{r+1}$. The initial stage of the proof is split into two parts: first, we closely follow the proof of Theorem 1 and analyze the game for $t \in [T_r, T_{r+1} - 1]$; then, we handle the case of transition from $\hat{\mathcal{C}}_r$ to $\hat{\mathcal{C}}_{r+1}$.

We introduce the following short-hand notation:

$$\hat{d}_t := \|\bar{\mathbf{m}}_t - \hat{\mathbf{c}}_t\| \quad \text{and} \quad \Omega_r = \{\mathcal{C} \subset \hat{\mathcal{C}}_r\}.$$

Note that unlike the quantity of interest $d_t = \|\bar{\mathbf{m}}_t - \bar{\mathbf{c}}_t\|$, which is equal to the distance from the average payoff $\bar{\mathbf{m}}_t$ along the trajectory to the *true* target set $\bar{\mathbf{c}}_t$, the distance \hat{d}_t is with respect to the currently used estimate $\hat{\mathcal{C}}_r$. The key insight of the proof is hidden in the fact that, if Ω_r occurs, then the approachability condition, which is met by \mathcal{C} , is also met by the super-set estimate $\hat{\mathcal{C}}_r$.

Convergence in L^2 . Let us start with the following observation, which relates d_t to \hat{d}_t , based on Assumption 3. We have for $t \in [T_r, T_{r+1} - 1]$,

$$\begin{aligned} d_t &= \|\bar{\mathbf{m}}_t - \text{Proj}_{\mathcal{C}} \bar{\mathbf{m}}_t\| \leq \|\bar{\mathbf{m}}_t - \text{Proj}_{\mathcal{C}} \text{Proj}_{\hat{\mathcal{C}}_r} \bar{\mathbf{m}}_t\| \\ &\leq \|\bar{\mathbf{m}}_t - \text{Proj}_{\hat{\mathcal{C}}_r} \bar{\mathbf{m}}_t\| + \|\text{Proj}_{\hat{\mathcal{C}}_r} \bar{\mathbf{m}}_t - \text{Proj}_{\mathcal{C}} \text{Proj}_{\hat{\mathcal{C}}_r} \bar{\mathbf{m}}_t\| \\ &\leq \hat{d}_t + d(\hat{\mathcal{C}}_r, \mathcal{C}). \end{aligned} \quad (43)$$

Hence, according to the fourth item of Assumption 3 and the L^2 -triangular inequality, we have

$$\sqrt{\mathbb{E}[d_t^2]} \leq \sqrt{\mathbb{E}[\hat{d}_t^2]} + \beta_r. \quad (44)$$

Since $\beta_r \rightarrow 0$ according to Assumption 3, the latter implies that, if $\mathbb{E}[\hat{d}_t^2] \rightarrow 0$, then $\mathbb{E}[d_t^2] \rightarrow 0$.

As already mentioned, to prove the L^2 -convergence, we consider two cases. In the first case, we study the evolution of the game within one phase, that is for $t \in [T_r, T_{r+1} - 2]$ – the case where we project onto $\hat{\mathcal{C}}_r$. The second case is when $t = T_{r+1} - 1$, that is, when in the next round we are going to update the estimate $\hat{\mathcal{C}}_r$.

Case $T_r \leq t \leq T_{r+1} - 2$: Defining

$$Z_{t+1} := \left\langle \bar{\mathbf{m}}_t - \hat{\mathbf{c}}_t, \mathbf{m}_{t+1} - \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) \, d\mathbf{Q}(x, s) \right\rangle, \quad (45)$$

$$\text{and } B_t := \left\langle \bar{\mathbf{m}}_t - \hat{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}_{t+1}^x, \mathbf{q}_{t+1}^{G(x,s)}, x, s) \, d\mathbf{Q}(x, s) - \hat{\mathbf{c}}_t \right\rangle, \quad (46)$$

we can write

$$\begin{aligned} \hat{d}_{t+1}^2 &\leq \|\bar{\mathbf{m}}_{t+1} - \hat{\mathbf{c}}_t\|^2 \\ &\leq \frac{t^2}{(t+1)^2} \hat{d}_t^2 + \frac{1}{(t+1)^2} \|\mathbf{m}_{t+1} - \hat{\mathbf{c}}_t\|^2 + \frac{2t}{(t+1)^2} (Z_{t+1} + B_t). \end{aligned} \quad (47)$$

As in the proof of Theorem 1, the main non-standard analysis is connected with the treatment of B_t . Observe that thanks to Assumption 3, we always have $|B_t| \leq B(B + 2\|\mathbf{m}\|_{\infty,2})$, hence

$$B_t \leq B_t \mathbb{I}\{\Omega_r\} + B(B + 2\|\mathbf{m}\|_{\infty,2}) \mathbb{I}\{\Omega_r^c\}. \quad (48)$$

Furthermore, similarly as for Eq. (14), we have on Ω_r

$$\begin{aligned} B_t &\leq 4\|\mathbf{m}\|_{\infty,2} \hat{d}_t \cdot \text{TV}(\mathbf{Q}, \hat{\mathbf{Q}}_t) \\ &\quad + \min_{(\mathbf{p}^x)_{x \in \mathcal{X}}} \max_{(\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \mathcal{S}}} \left\langle \bar{\mathbf{m}}_t - \hat{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) \, d\mathbf{Q}(x, s) - \hat{\mathbf{c}}_t \right\rangle. \end{aligned} \quad (49)$$

Since, by definition of Ω_r , we have the inclusion $\mathcal{C} \subset \hat{\mathcal{C}}_r$ on Ω_r , Blackwell's condition (2) implies that, on Ω_r ,

$$\forall (\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \{0,1\}} \exists (\mathbf{p}^x)_{x \in \mathcal{X}} \quad \text{s.t.} \quad \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) \, d\mathbf{Q}(x, s) \in \mathcal{C} \subset \hat{\mathcal{C}}_r.$$

The first item of Assumption 3 requires $\hat{\mathcal{C}}_r$ to be closed convex almost surely. Hence, using the property of Euclidean projection onto a convex closed set, in conjunction with von Neumann's minmax theorem, we conclude that, on Ω_r , it holds that

$$\min_{(\mathbf{p}^x)_{x \in \mathcal{X}}} \max_{(\mathbf{q}^{G(x,s)})_{(x,s) \in \mathcal{X} \times \mathcal{S}}} \left\langle \bar{\mathbf{m}}_t - \hat{\mathbf{c}}_t, \int_{\mathcal{X} \times \mathcal{S}} \mathbf{m}(\mathbf{p}^x, \mathbf{q}^{G(x,s)}, x, s) \, d\mathbf{Q}(x, s) - \hat{\mathbf{c}}_t \right\rangle \leq 0.$$

The above inequality, combined with Eqs. (47)–(49), yields

$$\begin{aligned} \hat{d}_{t+1}^2 &\leq \frac{t^2}{(t+1)^2} \hat{d}_t^2 + \frac{1}{(t+1)^2} \|\mathbf{m}_{t+1} - \hat{\mathbf{c}}_t\|^2 \\ &\quad + \frac{2t}{(t+1)^2} \left(Z_{t+1} + 4\|\mathbf{m}\|_{\infty,2} \hat{d}_t \cdot \text{TV}(\mathbf{Q}, \hat{\mathbf{Q}}_t) \mathbb{I}\{\Omega_r\} + B(B+2\|\mathbf{m}\|_{\infty,2}) \mathbb{I}\{\Omega_r^c\} \right). \end{aligned} \quad (50)$$

Since $(Z_t)_{t \geq 1}$ is martingale difference (by the same arguments as for the proof of Theorem 1), taking expectations from both sides of the above inequality, in conjunction with the condition on $\mathbb{P}(\Omega_r)$ of Assumption 3 and the Cauchy-Schwartz inequality, yields

$$\begin{aligned} \mathbb{E}[\hat{d}_{t+1}^2] &\leq \frac{t^2}{(t+1)^2} \mathbb{E}[\hat{d}_t^2] + \frac{B^2}{(t+1)^2} \\ &\quad + \frac{2t}{(t+1)^2} \left(4\|\mathbf{m}\|_{\infty,2} \sqrt{\mathbb{E}[\hat{d}_t^2]} \cdot \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_t)]} + \frac{B(B+2\|\mathbf{m}\|_{\infty,2})}{2T_r} \right). \end{aligned}$$

We deduce from the above that, for all $t \in [T_r, T_{r+1} - 2]$, since $t/(2T_r) \leq 1$,

$$\begin{aligned} \mathbb{E}[\hat{d}_{t+1}^2] &\leq \frac{t^2}{(t+1)^2} \mathbb{E}[\hat{d}_t^2] + \frac{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}{(t+1)^2} \\ &\quad + \frac{2t}{(t+1)^2} \left(4\|\mathbf{m}\|_{\infty,2} \sqrt{\mathbb{E}[\hat{d}_t^2]} \cdot \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_t)]} \right). \end{aligned} \quad (51)$$

Applying Lemma 1 with $t^* = T_r$, $K = 3B^2 + 4B\|\mathbf{m}\|_{\infty,2}$, and $\delta_t = 4\|\mathbf{m}\|_{\infty,2} \cdot \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_t)]}$, we obtain that, for all $t \in [T_r, T_{r+1} - 1]$,

$$\begin{aligned} \sqrt{\mathbb{E}[\hat{d}_t^2]} &\leq \frac{\sqrt{(3B^2 + 4B\|\mathbf{m}\|_{\infty,2})(t - T_r)}}{t} + \frac{4}{t} \sum_{t'=T_r}^{t-1} \|\mathbf{m}\|_{\infty,2} \cdot \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} \\ &\quad + \frac{T_r}{t} \sqrt{\mathbb{E}[\hat{d}_{T_r}^2]}. \end{aligned} \quad (52)$$

Case $t = T_{r+1} - 1$: In this case, when passing from t to $t+1$, the Player updates the estimate of the target set \mathcal{C} , which incurs additional price. In particular, the established recursion in Eq (52) does not hold, since by definition $\hat{d}_{T_{r+1}} = \|\bar{\mathbf{m}}_{T_{r+1}} - \hat{\mathbf{c}}_{T_{r+1}}\|$, where $\hat{\mathbf{c}}_{T_{r+1}}$ is the projection onto $\hat{\mathcal{C}}_{r+1}$. However, note that the argument of the first case still holds if we fix the set onto which we project. More formally, the inequality (51) still holds at $t = T_{r+1} - 1$, if we replace $\hat{d}_{T_{r+1}}$ in the left-hand side by

$$\tilde{d}_{T_{r+1}} := \|\bar{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{\mathcal{C}}_r}(\bar{\mathbf{m}}_{T_{r+1}})\|.$$

Hence $\sqrt{\mathbb{E}[\tilde{d}_{T_{r+1}}^2]}$ is smaller than the right-hand side of Eq. (52) with $t = T_{r+1}$. Applying the same argument as in (43), and applying Minkowski's inequality, we get

$$\begin{aligned} \sqrt{\mathbb{E}[\tilde{d}_{T_{r+1}}^2]} &= \sqrt{\mathbb{E}[\|\bar{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{\mathcal{C}}_{r+1}}(\bar{\mathbf{m}}_{T_{r+1}})\|^2]} \\ &\leq \sqrt{\mathbb{E}[\|\bar{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{\mathcal{C}}_r}(\bar{\mathbf{m}}_{T_{r+1}})\|^2]} + \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_r, \hat{\mathcal{C}}_{r+1})^2]} \\ &= \sqrt{\mathbb{E}[\tilde{d}_{T_{r+1}}^2]} + \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_r, \hat{\mathcal{C}}_{r+1})^2]}. \end{aligned}$$

Recalling that the bound in Eq. (52) holds for $\sqrt{\mathbb{E}[\hat{d}_{T_{r+1}}^2]}$, and using the above derived relation, we get for all $r \geq 0$

$$\begin{aligned} \sqrt{\mathbb{E}[\hat{d}_{T_{r+1}}^2]} &\leq \frac{\sqrt{(3B^2 + 4B\|\mathbf{m}\|_{\infty,2})(T_{r+1} - T_r)}}{T_{r+1}} + \frac{4}{T_{r+1}} \sum_{t'=T_r}^{T_{r+1}-1} \|\mathbf{m}\|_{\infty,2} \cdot \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} \\ &\quad + \frac{T_r}{T_{r+1}} \sqrt{\mathbb{E}[\hat{d}_{T_r}^2]} + \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_r, \hat{\mathcal{C}}_{r+1})^2]}. \end{aligned} \quad (53)$$

Multiplying Eq. (53) by T_{r+1} on both sides and rearranging, we deduce that for all $r \geq 0$

$$\begin{aligned} \left(T_{r+1} \sqrt{\mathbb{E}[\hat{d}_{T_{r+1}}^2]} - T_r \sqrt{\mathbb{E}[\hat{d}_{T_r}^2]} \right) &\leq \sqrt{(3B^2 + 4B\|\mathbf{m}\|_{\infty,2})(T_{r+1} - T_r)} \\ &\quad + 4 \sum_{t'=T_r}^{T_{r+1}-1} \|\mathbf{m}\|_{\infty,2} \cdot \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} \\ &\quad + T_{r+1} \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_r, \hat{\mathcal{C}}_{r+1})^2]}. \end{aligned}$$

Summing up the above inequalities over $r \geq 0$, and using the fact that, by Assumption 3, $\hat{d}_1 \leq B$, we obtain, with the convention $T_{-1} = 0$,

$$\begin{aligned} \sqrt{\mathbb{E}[\hat{d}_{T_r}^2]} &\leq \sqrt{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}} \frac{1}{T_r} \sum_{r'=0}^r \sqrt{T_{r'} - T_{r'-1}} \\ &\quad + 4\|\mathbf{m}\|_{\infty,2} \frac{1}{T_r} \sum_{t'=1}^{T_r-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} + \frac{1}{T_r} \sum_{r'=1}^r T_{r'} \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_{r'-1}, \hat{\mathcal{C}}_{r'})^2]}. \end{aligned} \quad (54)$$

To conclude the convergence in L^2 , we observe that $d(\hat{\mathcal{C}}_{r'-1}, \hat{\mathcal{C}}_{r'}) \leq d(\hat{\mathcal{C}}_{r'-1}, \mathcal{C}) + d(\mathcal{C}, \hat{\mathcal{C}}_{r'})$ and hence, the triangle inequality for L^2 -norms and Assumption 3 yield

$$\sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_{r'-1}, \hat{\mathcal{C}}_{r'})^2]} \leq \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_{r'-1}, \mathcal{C})^2]} + \sqrt{\mathbb{E}[d(\mathcal{C}, \hat{\mathcal{C}}_{r'})^2]} \leq \beta_{r'-1} + \beta_{r'} \leq 2\beta_{r'-1}. \quad (55)$$

Substituting the above bound in Eq. (54) (and reindexing, using that $T_{r'} = 2T_{r'-1}$), we get for all $r \geq 1$

$$\begin{aligned} \sqrt{\mathbb{E}[\hat{d}_{T_r}^2]} &\leq \frac{\sqrt{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}}{T_r} \sum_{r'=0}^r \sqrt{T_{r'} - T_{r'-1}} + 4\|\mathbf{m}\|_{\infty,2} \frac{1}{T_r} \sum_{t'=1}^{T_r-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} \\ &\quad + \frac{4}{T_r} \sum_{r'=0}^{r-1} T_{r'} \beta_{r'} \\ &\leq \frac{\sqrt{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}}{(\sqrt{2}-1)\sqrt{T_r}} + 4\|\mathbf{m}\|_{\infty,2} \frac{1}{T_r} \sum_{t'=1}^{T_r-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} + \frac{4}{T_r} \sum_{r'=0}^{r-1} T_{r'} \beta_{r'}. \end{aligned} \quad (56)$$

The $(\sqrt{2}-1)\sqrt{T_r}$ factor in the denominator of the first term of the final bound was obtained as follows:

$$\sum_{r'=0}^r \sqrt{T_{r'} - T_{r'-1}} = 1 + \sum_{r'=1}^r \sqrt{2^{r'-1}} = 1 + \frac{\sqrt{2^r} - 1}{\sqrt{2} - 1} \leq \frac{\sqrt{2^r}}{\sqrt{2} - 1} = \frac{\sqrt{T_r}}{\sqrt{2} - 1}.$$

Combining the first inequality of the two inequalities of Eqs. (56) with Eq. (52), we get for all $r \geq 1$ and $t \in [T_r, T_{r+1} - 1]$,

$$\begin{aligned} \sqrt{\mathbb{E}[\hat{d}_t^2]} &\leq \frac{\sqrt{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}}{t} \left(\sqrt{t - T_r} + \sum_{r'=0}^r \sqrt{T_{r'} - T_{r'-1}} \right) \\ &\quad + \frac{4\|\mathbf{m}\|_{\infty,2}}{t} \sum_{t'=1}^{t-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} + \frac{4}{t} \sum_{r'=0}^{r-1} T_{r'} \beta_{r'} \\ &\leq \frac{\sqrt{6B^2 + 8B\|\mathbf{m}\|_{\infty,2}}}{(\sqrt{2}-1)\sqrt{t}} + \frac{4\|\mathbf{m}\|_{\infty,2}}{t} \sum_{t'=1}^{t-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} + \frac{4}{t} \sum_{r'=0}^{r-1} T_{r'} \beta_{r'}, \end{aligned} \quad (57)$$

where the last inequality follows from

$$\sqrt{t - T_r} + \sum_{r'=0}^r \sqrt{T_{r'} - T_{r'-1}} \leq \sum_{r'=0}^{r+1} \sqrt{T_{r'} - T_{r'-1}} \leq \frac{\sqrt{T_{r+1}}}{\sqrt{2} - 1} = \frac{\sqrt{2T_r}}{\sqrt{2} - 1}.$$

Combining inequality (57) with (44), i.e., adding β_r to the bound above, and using $T_r/t \leq 1$, we conclude the stated bound for the L^2 convergence.

Almost-sure convergence. We observe that, according to (43), by union bounds, Markov's inequality, and the third item of Assumption 3, we have

$$\begin{aligned} \mathbb{P}\left[\sup_{t \geq T_r} d_t \geq 2\varepsilon\right] &\leq \mathbb{P}\left[\sup_{t \geq T_r} \hat{d}_t \geq \varepsilon\right] + \mathbb{P}\left[\sup_{r' \geq r} d(\hat{C}_{r'}, \mathcal{C}) \geq \varepsilon\right] \\ &\leq \mathbb{P}\left[\sup_{t \geq T_r} \hat{d}_t \geq \varepsilon\right] + \frac{1}{\varepsilon^2} \sum_{r' \geq r} \beta_{r'}^2. \end{aligned}$$

In what follows, we bound $\mathbb{P}\left[\sup_{t \geq T_r} \hat{d}_t \geq \varepsilon\right]$ by Ξ_r/ε^2 , where Ξ_r is defined in Eq. (60).

As in Theorem 1, we introduce a super-martingale S_t bounding \hat{d}_t^2 and whose expectation vanishes; however, the analysis is more involved here due to additional difficulties connected to handling the switches between regimes. More precisely, let us define, for $t \in [T_r, T_r - 1]$,

$$\begin{aligned} V_t &= \frac{B^2}{(t+1)^2} + \frac{2t}{(t+1)^2} \left(4\|\mathbf{m}\|_{\infty,2} \hat{d}_t \cdot \text{TV}(\mathbf{Q}, \hat{\mathbf{Q}}_t) + B(B + 2\|\mathbf{m}\|_{\infty,2}) \mathbb{I}\{\Omega_r^c\}\right) \\ &\quad + 2B d(\hat{C}_r, \hat{C}_{r+1}) \mathbb{I}\{t = T_{r+1} - 1\}. \end{aligned}$$

Using the above defined V_t , we additionally introduce the process

$$S_T = \hat{d}_T^2 + \sum_{t \geq T} \mathbb{E}[V_t | H_T]. \quad (58)$$

We observe that, by Assumption 3 and the triangle inequality,

$$\begin{aligned} \hat{d}_{T_{r+1}}^2 - \tilde{d}_{T_{r+1}}^2 &= (\hat{d}_{T_{r+1}} - \tilde{d}_{T_{r+1}}) \overbrace{(\hat{d}_{T_{r+1}} + \tilde{d}_{T_{r+1}})}^{\leq 2B} \\ &\leq 2B \left(\|\overline{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{C}_{r+1}}(\overline{\mathbf{m}}_{T_{r+1}})\| - \|\overline{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{C}_r}(\overline{\mathbf{m}}_{T_{r+1}})\| \right) \\ &\leq 2B \left(\|\overline{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{C}_{r+1}}(\text{Proj}_{\hat{C}_r}(\overline{\mathbf{m}}_{T_{r+1}}))\| - \|\overline{\mathbf{m}}_{T_{r+1}} - \text{Proj}_{\hat{C}_r}(\overline{\mathbf{m}}_{T_{r+1}})\| \right) \\ &\leq 2B \|\text{Proj}_{\hat{C}_r}(\overline{\mathbf{m}}_{T_{r+1}}) - \text{Proj}_{\hat{C}_{r+1}}(\text{Proj}_{\hat{C}_r}(\overline{\mathbf{m}}_{T_{r+1}}))\| \\ &\leq 2B \cdot d(\hat{C}_r, \hat{C}_{r+1}). \end{aligned}$$

Thus, in view of Eq. (50), and recalling that the right-hand side of Eq. (50) bounds rather $\tilde{d}_{T_{r+1}}$ at $t = T_{r+1} - 1$, the following recursive relation holds for any $t \in [T_r, T_{r+1} - 1]$:

$$\hat{d}_{t+1}^2 \leq \hat{d}_t^2 + V_t + \frac{2t}{(t+1)^2} Z_{t+1}.$$

Recalling that $\mathbb{E}[Z_{t+1} | H_t] = 0$, we deduce

$$\mathbb{E}[S_{T+1} | H_T] = \mathbb{E}[\hat{d}_{T+1}^2 | H_T] + \sum_{t \geq T+1} \mathbb{E}[V_t | H_T] \leq \hat{d}_T^2 + \sum_{t \geq T} \mathbb{E}[V_t | H_T] = S_T,$$

which means that $(S_T)_{T \geq 1}$ is a super-martingale.

Since, by definition of S_T , it holds that $\hat{d}_T^2 \leq S_T$, Doob's maximal inequality for non-negative super-martingales (Lemma 2) gives

$$\mathbb{P}\left(\sup_{t \geq T_r} \hat{d}_t \geq \varepsilon\right) \leq \mathbb{P}\left(\sup_{t \geq T_r} S_t \geq \varepsilon^2\right) \leq \frac{\mathbb{E}[S_{T_r}]}{\varepsilon^2}.$$

It only remains to bound $\mathbb{E}[S_{T_r}]$ by Ξ_r .

Note that by the Cauchy-Schwarz inequality and the bound of Eq. (55),

$$\mathbb{E}[d(\hat{\mathcal{C}}_{r'}, \hat{\mathcal{C}}_{r'+1})] \leq \sqrt{\mathbb{E}[d(\hat{\mathcal{C}}_{r'}, \hat{\mathcal{C}}_{r'+1})^2]} \leq 2\beta_{r'}.$$

Thanks to this inequality, to $t\mathbb{P}(\Omega_r^c) \leq t/(2T_r) \leq 1$ for $t \in [T_r, T_{r+1} - 1]$, and other manipulations that are standard by now, the expectation of the sum appearing in the definition (58) of the supermartingale S_T can be bounded as

$$\begin{aligned} \sum_{t \geq T_r} \mathbb{E}[V_t] &\leq \sum_{t \geq T_r} \frac{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}{(t+1)^2} + \sum_{t \geq T_r} \frac{8t\|\mathbf{m}\|_{\infty,2}}{(t+1)^2} \sqrt{\mathbb{E}[\hat{d}_t^2]} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_t)]} \\ &\quad + 2B \sum_{r' \geq r} \mathbb{E}[d(\hat{\mathcal{C}}_{r'}, \hat{\mathcal{C}}_{r'+1})] \\ &\leq \frac{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}{T_r} + 4B \sum_{r' \geq r} \beta_{r'} + \sum_{t \geq T_r} \frac{8\|\mathbf{m}\|_{\infty,2}}{t} \sqrt{\mathbb{E}[\hat{d}_t^2]} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_t)]}. \end{aligned} \tag{59}$$

To bound the right hand side of the above inequality, we observe that for $t \geq T_r$, by Eq. (57), we have

$$\begin{aligned} \sqrt{\mathbb{E}[\hat{d}_t^2]} &\leq \frac{\sqrt{6B^2 + 8B\|\mathbf{m}\|_{\infty,2}}}{(\sqrt{2}-1)\sqrt{t}} + 4\|\mathbf{m}\|_{\infty,2} \overbrace{\max_{t \geq T_r} \frac{1}{t} \sum_{t'=1}^{t-1} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]}}^{=: \bar{\Delta}_{T_r}^*} \\ &\quad + 4 \underbrace{\max_{r'' \geq r} \frac{1}{T_{r''}} \sum_{r'=0}^{r''-1} T_{r'} \beta_{r'}}_{=: \beta_r^*}. \end{aligned}$$

Substituting the above bound into Eq. (59), using $\sum_{t \geq T_r} t^{-3/2} \leq 2/\sqrt{T_r - 1}$ and $\text{TV}(\mathbf{Q}, \hat{\mathbf{Q}}_{t'}) \leq 1$, we obtain

$$\begin{aligned} \sum_{t \geq T_r} \mathbb{E}[V_t] &\leq \frac{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}{T_r} + 16\|\mathbf{m}\|_{\infty,2} \frac{\sqrt{6B^2 + 8B\|\mathbf{m}\|_{\infty,2}}}{(\sqrt{2}-1)\sqrt{T_r-1}} + 4B \sum_{r' \geq r} \beta_{r'} \\ &\quad + 32\|\mathbf{m}\|_{\infty,2} (\|\mathbf{m}\|_{\infty,2} \bar{\Delta}_{T_r}^* + \beta_r^*) \sum_{t \geq T_r} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]}. \end{aligned}$$

Finally, we take into account the definition of the super martingale S_T in Eq. (58) and the upper bound of Eq. (56), which we square, using that $(x+y+z)^2 \leq 2x^2 + 2(y+z)^2$. Doing so, and performing some crude boundings for the sake of readability, we get the final bound $\mathbb{E}[S_{T_r}] \leq \Xi_r$, where

$$\begin{aligned} \Xi_r &:= (1 + 2(\sqrt{2}-1)^{-2}) \frac{3B^2 + 4B\|\mathbf{m}\|_{\infty,2}}{T_r} + 16\|\mathbf{m}\|_{\infty,2} \frac{\sqrt{6B^2 + 8B\|\mathbf{m}\|_{\infty,2}}}{(\sqrt{2}-1)\sqrt{T_r-1}} + 4B \sum_{r' \geq r} \beta_{r'} \\ &\quad + 32(\|\mathbf{m}\|_{\infty,2} \bar{\Delta}_{T_r}^* + \beta_r^*) \left(\|\mathbf{m}\|_{\infty,2} \bar{\Delta}_{T_r}^* + \beta_r^* + \|\mathbf{m}\|_{\infty,2} \sum_{t \geq T_r} \frac{1}{t} \sqrt{\mathbb{E}[\text{TV}^2(\mathbf{Q}, \hat{\mathbf{Q}}_{t'})]} \right). \end{aligned} \tag{60}$$

As indicated in Eq. (42), the Cesaro averages $\bar{\beta}_r$, which are positive, tend to 0; therefore, we also have $\beta_r^* \rightarrow 0$. For similar reasons, and as already noted for Theorem 1, the term $\bar{\Delta}_{T_r}^*$ also vanishes under Assumption 1. The latter also implies that the final term in Eq. (60) vanishes. Other terms clearly vanish or were already discussed for the L^2 -convergence. All in all, $\Xi_r \rightarrow 0$, as claimed. \square