

Swap Agnostic Learning, or Characterizing Omniprediction via Multicalibration*

Parikshit Gopalan
Apple

Michael P. Kim
Cornell University[†]

Omer Reingold
Stanford

Abstract

We introduce and study Swap Agnostic Learning. The problem can be phrased as a game between a *predictor* and an *adversary*: first, the predictor selects a hypothesis h ; then, the adversary plays in response, and for each level set of the predictor $\{x \in \mathcal{X} : h(x) = v\}$ selects a (different) loss-minimizing hypothesis $c_v \in \mathcal{C}$; the predictor wins if h competes with the adaptive adversary's loss. Despite the strength of the adversary, we demonstrate the feasibility of Swap Agnostic Learning for any convex loss.

Somewhat surprisingly, the result follows through an investigation into the connections between Omniprediction [GKR⁺22] and Multicalibration [HKRR18]. Omniprediction is a new notion of optimality for predictors that strengthens classical notions such as agnostic learning. It asks for loss minimization guarantees (relative to a hypothesis class) that apply not just for a specific loss function, but for any loss belonging to a rich family of losses. A recent line of work shows that omniprediction is implied by multicalibration and related multi-group fairness notions [GKR⁺22, GHK⁺23]. This unexpected connection raises the question: *is multi-group fairness necessary for omniprediction?*

Our work gives the first affirmative answer to this question. We establish an *equivalence* between swap variants of omniprediction and multicalibration and swap agnostic learning. Further, swap multicalibration is essentially equivalent to the standard notion of multicalibration, so existing learning algorithms can be used to achieve any of the three notions. Building on this characterization, we paint a complete picture of the relationship between different variants of multi-group fairness, omniprediction, and Outcome Indistinguishability [DKR⁺21]. This inquiry reveals a unified notion of OI that captures all existing notions of omniprediction and multicalibration.

*This manuscript represents the full version of [GKR23]. It subsumes arXiv:2302.06726v1 by the authors, under the title *Characterizing notions of omniprediction via multicalibration*.

[†]This work was completed when the author was at the Miller Institute for Basic Research in Science and the Simons Institute for the Theory of Computing at UC Berkeley.

Contents

1	Introduction	1
2	Our Contributions: Defining and Characterizing Swap Learning	3
2.1	Definitions of Swap Learning and Swap Multicalibration	3
2.2	Characterizing Swap Learning	7
2.3	Discussion and Related Work	10
3	Properties of Swap Multicalibration	12
3.1	Algorithms for Swap Multicalibration	14
4	Swap Agnostic Learning, Swap Omniprediction, and Swap Multicalibration are Equivalent	15
4.1	Statement of Equivalence	17
4.1.1	Preliminary Results	17
4.1.2	Proof of Equivalence	19
4.2	An algorithm for Swap Agnostic Learning	20
5	Characterizing Swap Loss OI via Multicalibration	21
6	Relating Notions of Omniprediction	25

1 Introduction

Since its inception as an extension to Valiant’s PAC framework [Val84, Hau92, KV94], Agnostic Learning has been the central problem of supervised learning theory. Agnostic learning frames the task of supervised learning through loss minimization: given a loss function ℓ , a hypothesis class \mathcal{C} , and $\varepsilon \geq 0$, a predictor h is an agnostic learner if it achieves loss that competes with the minimal achievable within the hypothesis class.

$$\mathbf{E}[\ell(\mathbf{y}, h(\mathbf{x}))] \leq \min_{c \in \mathcal{C}} \mathbf{E}[\ell(\mathbf{y}, c(\mathbf{x}))] + \varepsilon \quad (1)$$

While the agnostic learning paradigm has been remarkably successful, in recent years, researchers have investigated alternative learning paradigms to address concerns of the modern prediction pipeline, including fairness and robustness. The work of [HKRR18] introduced *multicalibration* as a new paradigm for learning fair predictors. Multicalibration asserts fairness as a first-order goal, requiring that predictions appear calibrated even conditioned on membership in one of a potentially-huge collection of subgroups \mathcal{C} of the domain.¹ As a solution concept, multicalibration can be achieved efficiently using a weak learner for \mathcal{C} to identify subgroups where predictions are miscalibrated. In contrast to agnostic learning, multicalibration does not make reference to minimizing any loss.

Yet, it turns out that there are deep connections between multicalibration and loss minimization. This connection was first discovered in the work of [GKR⁺22], who introduced the notion of *omniprediction* as a new solution concept in supervised learning. Intuitively, an omnipredictor is a single predictor that provides an optimality guarantee for many losses *simultaneously*. More formally, for a collection of loss functions \mathcal{L} and a hypothesis class \mathcal{C} , a predictor is an $(\mathcal{L}, \mathcal{C})$ -omnipredictor if *for any loss in the collection* $\ell \in \mathcal{L}$, the predictions achieve loss that competes with the minimal achievable within the class \mathcal{C} . On a technical level, the omniprediction guarantee is made possible by post-processing the predictions after the loss is revealed. Importantly though, the post-processing is a data-free optimization that can be performed efficiently, based only on the loss function and prediction, and not the data distribution. In this way, an omnipredictor can be trained on the dataset once and for all without knowing the specific target loss in advance (beyond the fact that it belongs to \mathcal{L}).

The surprising main result of [GKR⁺22] demonstrates that multicalibrated predictors are omnipredictors for \mathcal{L}_{cvx} , the class of all convex loss functions. In other words, the multicalibration framework is capable of guaranteeing loss minimization in a very strong sense: learning a single multicalibrated predictor gives an agnostic learner, for every target convex loss $\ell \in \mathcal{L}_{\text{cvx}}$. The results of [GKR⁺22] stand in contrast to the conventional wisdom that optimizing predictions for different loss functions requires a separate training procedure for each loss.

On the surface, multicalibration seems like a fundamentally different approach to supervised learning than agnostic learning. The implication of omniprediction from multicalibration, however, suggests a deeper connection between the notions. In fact, follow-up work of [GHK⁺23] gave new constructions of omnipredictors (for different loss classes \mathcal{L}), similarly deriving the guarantees from notions of multi-group fairness. To summarize the state of the art, we have constructions of omnipredictors of various flavors, and all such constructions rely on some variant of multicalibration.

¹This collection of subgroups is suggestively denoted by \mathcal{C} , indicating that (in correspondence with a hypothesis class) subgroup membership can be computed using models of bounded capacity, such as small decision trees, halfspaces, or neural networks.

While multicalibration suffices to guarantee omniprediction, a glaring question remains in the development of this theory: *Is multicalibration necessary for omniprediction?* A closely-related question stems from the observation that multicalibration guarantees omniprediction, despite the fact that the definition of multicalibration does not reference any loss function to minimize. It is natural to ask whether we can arrive at the same solution concept through a more standard supervised learning framing: *Can we characterize multicalibration in the language of loss minimization?* We investigate these questions, exploring the connections between agnostic learning, omniprediction, and multicalibration.

Tight characterizations via swap notions. To understand the relationship between omniprediction and multicalibration, we formulate a new learning task, which we call Swap Agnostic Learning. This *swap* variant of agnostic learning is inspired by the notion of swap regret in the online learning literature [FV98, BM07a], where the learner must achieve vanishing regret, not simply overall, but even conditioned on their decisions. Swap agnostic learning naturally suggests swap notions of omniprediction (and generalizations such as swap loss OI) and multicalibration. At first glance, swap agnostic learning is a considerably more ambitious goal than standard agnostic learning. Nevertheless, our work demonstrates an efficient algorithm for swap agnostic learning for any convex loss function, leveraging only a weak agnostic learner for the hypothesis class.

Our main result about swap agnostic learning actually follows as a corollary of a much broader discovery: *swap agnostic learning and swap variants of omniprediction and multicalibration are all equivalent.* Concretely, we introduce new swap variants of omniprediction and multicalibration (also inspired by swap regret) and show that each of these definitions—despite different technical framings—actually encodes the same solution concept. In other words, once we move to swap variants, multicalibration *is* necessary (and sufficient) for omniprediction, and necessary even for the “weaker” goal of swap agnostic learning. Further, we argue that the original formulation of multicalibration is essentially equivalent to our new swap variant; in particular, the original multicalibration algorithm of [HKRR18] actually produces a swap multicalibrated predictor. Thus, all three swap variants can be achieved efficiently given a weak agnostic learner for the class \mathcal{C} .

While the swap and standard variants of multicalibration are quite similar, for agnostic learning and omniprediction, the distinction is substantial. We prove separations between standard and swap learning for these notions. In combination, our results shed light on the question of whether (standard) omniprediction implies multicalibration, suggesting that the answer is no: (swap) multicalibration is equivalent to the stronger notion of swap omniprediction, which is provably stronger than standard omniprediction.

Our results provide an exact characterization of these swap learning notions, as well as relationships between other learning desiderata explored in recent works. At the extreme, we introduce a swap variant of Loss Outcome Indistinguishability [DKR⁺21, GHK⁺23] and show how it is expressive enough to capture all existing notions in the study of multicalibration and omniprediction. We present a tight characterization of it in terms of swap multicalibration for a suitably defined class \mathcal{C}' which depends on both the loss family \mathcal{L} and the hypothesis class \mathcal{C} . Standard Loss OI (without the swap) was known to be implied by calibrated multiaccuracy for the same class \mathcal{C}' . Our result shows that strengthening to swap loss OI, gives a tight characterization in terms of swap multicalibration.² We describe each of the notions, and our contributions, in detail in the next section.

²At this point, we have introduced multiple notions of swap learning and omniprediction. For the curious reader, the precise relationship between the different notions is depicted in Figure 1.

2 Our Contributions: Defining and Characterizing Swap Learning

In this work, we give new definitions of swap learning notions. A major contribution of this work is defining these notions precisely, within the context of prior work on supervised learning and algorithmic fairness. After defining these notions, we give a tight characterization of them, showing key equivalences and key separations. In this section, we give a technical overview of the definitions and the results.

We begin with the definitions we introduce to study swap notions of learning. We recall prior notions—agnostic learning, omniprediction, multicalibration, and loss OI—and we present our new swap variants of each. Along the way, we include a number of claims about the relationship between the standard and swap variants. Then, with the definitions in place, we give an overview of our results characterizing the notions.

Preliminaries. We work in the *agnostic* learning setting, where we assume a data distribution $(\mathbf{x}, \mathbf{y}) \sim \mathcal{D}$ supported on $\mathcal{X} \times \{0, 1\}$. We take the objects of study in agnostic learning to be real-valued hypotheses $h : \mathcal{X} \rightarrow \mathbb{R}$. Omniprediction and multicalibration study the special case of predictors $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ that map domain elements to probabilities. We denote the Bayes optimal predictor as $p^*(x) = \Pr[\mathbf{y} = 1 | \mathbf{x} = x]$. As is standard in agnostic learning, we make no assumptions about the complexity of $p^* : \mathcal{X} \rightarrow [0, 1]$. While, in principle, predictors and hypotheses may take on continuous values in \mathbb{R} , we restrict our attention to functions supported on finitely-many values. We let $\text{Im}(\tilde{p}) \subseteq [0, 1]$ denote the set of values taken on by $\tilde{p}(x)$.

Given a hypothesis $\tilde{p} : \mathcal{X} \rightarrow \mathbb{R}$, it will be useful to imagine drawing samples from \mathcal{D} in two steps: first, we sample a prediction $\mathbf{v} \in \mathbb{R}$ according to the distribution of $\tilde{p}(\mathbf{x})$ under \mathcal{D} ; then, we draw the random variables (\mathbf{x}, \mathbf{y}) according to the conditional distribution $\mathcal{D} | \tilde{p}(\mathbf{x}) = \mathbf{v}$. Accordingly, let $\mathcal{D}_{\tilde{p}}$ denote the distribution of $\tilde{p}(\mathbf{x})$ where $\mathbf{x} \sim \mathcal{D}$. When the predictor \tilde{p} is clear from context, we use the following notation: for each $v \in \text{Im}(\tilde{p})$, let $\mathcal{D}|_v$ denote the conditional distribution $\mathcal{D} | \tilde{p}(\mathbf{x}) = v$. Note that the distribution $\mathcal{D}|_{\mathbf{v}}$ for randomly drawn $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$ recovers the data distribution \mathcal{D} .

Throughout, we consider loss functions $\ell : \{0, 1\} \times \mathbb{R} \rightarrow \mathbb{R}$, which take a binary label y and a real-valued action t and assign to it a loss value. Two common examples are the squared loss $\ell_2(y, t) = (y - t)^2$ and logistic loss $\ell_{\text{lg}}(y, t) = \log(1 + \exp((1 - 2y)t))$. Of key interest is the loss class \mathcal{L}_{cvx} , which denotes the set of all convex losses satisfying a set of “niceness” conditions, defined formally in Definition 4.1.

2.1 Definitions of Swap Learning and Swap Multicalibration

Swap Agnostic Learning

Swap agnostic learning is defined with respect to a loss function $\ell : \{0, 1\} \times \mathbb{R} \rightarrow \mathbb{R}$ and a hypothesis class $\mathcal{C} \subseteq \{c : \mathcal{X} \rightarrow \mathbb{R}\}$, and can be viewed as a game with two participants:

- **The Predictor** \mathcal{P} selects a hypothesis $h : \mathcal{X} \rightarrow \mathbb{R}$, aiming to minimize the loss ℓ .
- **The Adversary** \mathcal{A} plays in response: for each level set $\{x \in \mathcal{X} : h(x) = v\}$, the adversary may choose a separate loss-minimizing hypothesis $c_v \in \mathcal{C}$.

Swap agnostic learning requires that the predictions of h are competitive with the adaptive strategy determined by the adversary $\{c_v : v \in \text{Im}(h)\}$, even though the predictor plays first.

Definition 2.1 (Swap Agnostic Learning). *For a loss function ℓ , hypothesis class \mathcal{C} , and error $\varepsilon \geq 0$, a hypothesis h is a $(\ell, \mathcal{C}, \varepsilon)$ -swap agnostic learner if*

$$\mathbf{E}[\ell(\mathbf{y}, h(\mathbf{x}))] \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_h} \left[\min_{c_v \in \mathcal{C}} \mathbf{E}[\ell(\mathbf{y}, c_v(\mathbf{x}) \mid h(\mathbf{x}) = \mathbf{v})] \right] + \varepsilon. \quad (2)$$

We borrow nomenclature from online learning: a predictor using a swap agnostic learner h has no incentive to “swap” any of their fixed predictions $h(\mathbf{x}) = v$ to predict according to $c_v \in \mathcal{C}$. Contrasting the requirement in (2) to that of agnostic learning in (1), we have switched the order of quantifiers, such that the minimization is taken after the expectation over the choice of $h(\mathbf{x}) = \mathbf{v}$.

Swap agnostic learning strengthens standard agnostic learning, where the predictor only competes against the single best hypothesis $c \in \mathcal{C}$.

Claim 2.2. *If \tilde{p} is a $(\ell, \mathcal{C}, \varepsilon)$ -swap agnostic learner, it is a $(\ell, \mathcal{C}, \varepsilon)$ -agnostic learner.*

Indeed, swap agnostic learning seems to be a much more stringent condition. Provided the class \mathcal{C} contains all constant functions, the adversary can simply imitate the predictor when $h(x) = v$ by choosing $c_v(x) = v$. For such \mathcal{C} , (2) implies that imitation is the best strategy for the adversary (up to ε).

Notation. In this presentation, we use the notation $c_v \in \mathcal{C}$ to emphasize the fact that the minimization occurs after conditioning on the predicted value $\mathbf{v} = v$. Of course, the choice of c_v is still a search over \mathcal{C} , and thus, equivalently could be written as a minimization over $c \in \mathcal{C}$. We use these notations interchangeably.

Swap Omniprediction

We recall the original notion of omniprediction proposed by [GKR⁺22], which requires a predictor to yield an agnostic learner simultaneously for every loss in a collection \mathcal{L} . As in [GKR⁺22, GHK⁺23], we observe that for any loss function ℓ and known distribution on outcomes $\mathbf{y} \sim \text{Ber}(p)$ for $p \in [0, 1]$, there is an optimal action $k_\ell(p)$ that minimizes the expected loss ℓ . Formally, we define the optimal post-processing of predictions by the function $k_\ell : [0, 1] \rightarrow \mathbb{R}$, which specifies the action that minimizes the expected loss.³

$$k_\ell(p) = \arg \min_{t \in \mathbb{R}} \mathbf{E}_{\mathbf{y} \sim \text{Ber}(p)} [\ell(\mathbf{y}, t)]. \quad (3)$$

This observation is particularly powerful because the function k_ℓ is given by a simple, univariate optimization that can be used as a data-free post-processing procedure on top of a predictor \tilde{p} . Concretely, it allows us to define omniprediction, which requires that the function $h_\ell = k_\ell \circ \tilde{p}$ is an $(\ell, \mathcal{C}, \varepsilon)$ -agnostic learner for every $\ell \in \mathcal{L}$.

Definition 2.3 (Omnipredictor, [GKR⁺22]). *For a collection of loss functions \mathcal{L} , a hypothesis class \mathcal{C} , and error $\varepsilon > 0$, a predictor $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ is an $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -omnipredictor if for every $\ell \in \mathcal{L}$,*

$$\mathbf{E}[\ell(\mathbf{y}, k_\ell(\tilde{p}(\mathbf{x})))] \leq \min_{c \in \mathcal{C}} \mathbf{E}[\ell(\mathbf{y}, c(\mathbf{x}))] + \varepsilon. \quad (4)$$

³If there are multiple optima, we break ties arbitrarily.

We propose a strengthened notion of Swap Omniprediction, where the adversary may choose both the loss $\ell_v \in \mathcal{L}$ and hypothesis $c_v \in \mathcal{C}$ based on the prediction $\tilde{p}(x) = v$. For each supported $v \in \text{Im}(\tilde{p})$, we allow an adversarial choice of loss $\ell_v \in \mathcal{L}$; we call the collection $\{\ell_v \in \mathcal{L} : v \in \text{Im}(\tilde{p})\}$ an assignment of loss functions. For each v , we compare the omnipredictor’s action $k_{\ell_v}(v)$ with the optimal choice of hypothesis c_v tailored to the assigned loss ℓ_v over the conditional distribution $\mathcal{D}|\tilde{p}(\mathbf{x}) = v$.

Definition 2.4 (Swap Omnipredictor). *For a collection of loss functions \mathcal{L} , a hypothesis class \mathcal{C} , and error $\varepsilon > 0$, a predictor $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ is a $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -swap omnipredictor if for any assignment of loss functions $\{\ell_v \in \mathcal{L}\}_{v \in \text{Im}(\tilde{p})}$,*

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\mathbf{E}[\ell_{\mathbf{v}}(\mathbf{y}, k_{\ell_{\mathbf{v}}}(\mathbf{v})) \mid \tilde{p}(\mathbf{x}) = \mathbf{v}] \right] \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\min_{c \in \mathcal{C}} \mathbf{E}[\ell_{\mathbf{v}}(\mathbf{y}, c(\mathbf{x})) \mid \tilde{p}(\mathbf{x}) = \mathbf{v}] \right] + \varepsilon. \quad (5)$$

Swap omniprediction gives the adversary considerable power. For instance, the special case where we restrict the adversary’s choice of losses to be constant, $\ell_v = \ell$, realizes swap agnostic learning for ℓ .

Claim 2.5. *If \tilde{p} is a $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -swap omnipredictor, it is a $(\ell, \mathcal{C}, \varepsilon)$ -swap agnostic learner for every $\ell \in \mathcal{L}$, and hence a $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -omnipredictor.*

Analogous to the standard notions (where omniprediction implies agnostic learning), swap omniprediction implies swap agnostic learning for every ℓ ; however, swap omniprediction gives an even stronger guarantee since the adversary’s loss ℓ_v may be chosen in response to the prediction $\tilde{p}(\mathbf{x}) = v$.

Swap Multicalibration

Multicalibration was introduced in the work of [HKRR18] as a notion of algorithmic fairness following [KMR17], but has since seen broad application across many learning applications (e.g., [DKR⁺21, GKR⁺22, KKG⁺22, GJN⁺22]). Informally, multicalibration requires predictions to appear calibrated, not simply overall, but also when we restrict our attention to subgroups within some broad collection \mathcal{C} . The formulation below appears in [GHK⁺23].

Definition 2.6 (Multicalibration, [HKRR18]). *For a hypothesis class \mathcal{C} and $\alpha \geq 0$, a predictor $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ is (\mathcal{C}, α) -multicalibrated if*

$$\max_{c \in \mathcal{C}} \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}[c(\mathbf{x})(\mathbf{y} - \mathbf{v}) \mid \tilde{p}(\mathbf{x}) = \mathbf{v}] \right| \right] \leq \alpha. \quad (6)$$

When $c : \mathcal{X} \rightarrow \{0, 1\}$ is Boolean (and has sufficiently large measure), this definition says that conditioned on $c(\mathbf{x}) = 1$, the calibration violation is small, recovering the definition in [HKRR18].

Swap Multicalibration strengthens multicalibration, extending the pseudorandomness perspective on multicalibration developed in [DKR⁺21, GKR⁺22]. Swap multicalibration requires that for the typical prediction $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$, no hypothesis in $c_v \in \mathcal{C}$ achieves good correlation with the residual labels $\mathbf{y} - \mathbf{v}$ over \mathcal{D} conditioned on $\tilde{p}(\mathbf{x}) = \mathbf{v}$.

Definition 2.7 (Swap Multicalibration). *For a hypothesis class \mathcal{C} and $\alpha \geq 0$, a predictor $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ is (\mathcal{C}, α) -swap multicalibrated if*

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\max_{c \in \mathcal{C}} \left| \mathbf{E}[c(\mathbf{x})(\mathbf{y} - \mathbf{v}) \mid \tilde{p}(\mathbf{x}) = \mathbf{v}] \right| \right] \leq \alpha. \quad (7)$$

Again, the difference between swap and standard multicalibration is in the order of quantifiers. The standard definition requires that for every $c \in \mathcal{C}$, the correlation $|\mathbf{E}_{\mathcal{D}|v}[c(\mathbf{x})(\mathbf{y} - \mathbf{v})]|$ is small in expectation over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$. Swap multicalibration considers the maximum of $|\mathbf{E}_{\mathcal{D}|v}[c(\mathbf{x})(\mathbf{y} - \mathbf{v})]|$ over all $c \in \mathcal{C}$ for each fixing of $\mathbf{v} = v$ and requires this to be small in expectation over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$. It follows that swap multicalibration implies standard multicalibration.

Claim 2.8. *If \tilde{p} is (\mathcal{C}, α) -swap multicalibrated, it is (\mathcal{C}, α) -multicalibrated.*

Swap multicalibration allows an auditor to use a different choice of $c_v \in \mathcal{C}$ to audit our predictor, conditioned on each value \mathbf{v} of the predictions. In this sense, it is reminiscent of our definitions of swap loss minimization. But while swap multicalibration is nominally a stronger notion than standard multicalibration, the complexity of achieving swap multicalibration is essentially the same as multicalibration. We present two arguments to support this view.

- All known algorithms for achieving standard multicalibration [HKRR18, GKR⁺22] already give swap multicalibration. They ensure multicalibration by conditioning on each prediction value, and then using a weak agnostic learner for \mathcal{C} to audit for a violation of the multicalibration condition. This approach naturally yields swap multicalibration. Thus, the same primitive of weak agnostic learning [BLM01] needed for standard multicalibration (and standard agnostic learning) also suffices for swap multicalibration.
- We show that starting from any (\mathcal{C}, α) -multicalibrated predictor \tilde{p} , by suitably discretizing its values, we obtain a predictor that is close to \tilde{p} and (\mathcal{C}, α') -swap multicalibrated, with some degradation in the value of α' .

We formalize both these arguments in Section 3.

Swap Loss Outcome Indistinguishability

The connections between multi-group fairness and omniprediction were extended in [GHK⁺23]. Building on the original Outcome Indistinguishability framework of [DKR⁺21], [GHK⁺23] introduced a related notion they referred to as Loss Outcome Indistinguishability (Loss OI). Intuitively, Loss OI requires that outcomes sampled from the predictive model \tilde{p} are indistinguishable from Nature’s outcomes, according to tests defined by the loss class \mathcal{L} and hypothesis class \mathcal{C} . They showed that $(\mathcal{L}, \mathcal{C})$ -Loss OI implies the standard $(\mathcal{L}, \mathcal{C})$ -omniprediction from [GKR⁺22], and that it is implied by a multi-group fairness notion called calibrated multiaccuracy (for a class of hypotheses \mathcal{C}' that may depend on both \mathcal{L} and \mathcal{C}). For a fixed hypothesis class, calibrated multiaccuracy implies multiaccuracy but is weaker than multicalibration. The Loss OI framework can be used to derive a wide range of omniprediction guarantees starting from multigroup fairness notions of varying strength. In this work, we extend the framework by introducing a swap variant.

Recalling the definition put forth by [DKR⁺21, GHK⁺23], Outcome Indistinguishability requires predictors \tilde{p} to fool a family \mathcal{U} of statistical tests $u : \mathcal{X} \times [0, 1] \times \{0, 1\}$ that take a point $\mathbf{x} \in \mathcal{X}$, a prediction $\tilde{p}(\mathbf{x}) \in [0, 1]$ and a label $\mathbf{y} \in \{0, 1\}$ as their arguments. Formally, we use $(\mathbf{x}, \mathbf{y}^*)$ to denote a sample from the true joint distribution over $\mathcal{X} \times \{0, 1\}$. Then, given a predictor \tilde{p} , we associate it with the random variable with $\mathbf{E}[\tilde{\mathbf{y}}|x] = \tilde{p}(x)$, i.e., where $\tilde{\mathbf{y}}|x \sim \text{Ber}(\tilde{p}(x))$. The variable $\tilde{\mathbf{y}}$ can be viewed as \tilde{p} ’s simulation of Nature’s label \mathbf{y}^* .⁴ The goal is distinguish between the scenarios

⁴In this presentation, we abuse notation and let \mathcal{D} to denote the joint distribution $(\mathbf{x}, \mathbf{y}^*, \tilde{\mathbf{y}})$, where $\mathbf{E}[\mathbf{y}^*|x] = p^*(x)$ and $\mathbf{E}[\tilde{\mathbf{y}}|x] = \tilde{p}(x)$. While the joint distribution of $(\mathbf{y}^*, \tilde{\mathbf{y}})$ is not important to us, for simplicity we assume they are independent given $\mathbf{x} = x$.

where $\mathbf{y} = \mathbf{y}^*$ is generated by *nature* versus where $\mathbf{y} = \tilde{\mathbf{y}}$ is a simulation of nature according to the predictor \tilde{p} .

Formally, we require than for every $u \in \mathcal{U}$,

$$\mathbf{E}_{\mathcal{D}}[u(\mathbf{x}, \tilde{p}(\mathbf{x}), \mathbf{y}^*)] \approx_{\varepsilon} \mathbf{E}_{\mathcal{D}}[u(\mathbf{x}, \tilde{p}(\mathbf{x}), \tilde{\mathbf{y}})].$$

Loss OI specializes this to a specific family of tests arising in the analysis of omnipredictors.

Definition 2.9 (Loss OI, [GHK+23]). *For a collection of loss functions \mathcal{L} , hypothesis class \mathcal{C} , and $\varepsilon \geq 0$, define the family of tests $\mathcal{U}(\mathcal{L}, \mathcal{C}) = \{u_{\ell, c}\}_{\ell \in \mathcal{L}, c \in \mathcal{C}}$ where*

$$u_{\ell, c}(x, v, y) = \ell(y, k_{\ell}(v)) - \ell(y, c(x)). \quad (8)$$

A predictor $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ is $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -loss OI if for every $u \in \mathcal{U}(\mathcal{L}, \mathcal{C})$, it holds that

$$\left| \mathbf{E}_{(\mathbf{x}, \mathbf{y}^*) \sim \mathcal{D}} [u(\mathbf{x}, \tilde{p}(\mathbf{x}), \mathbf{y}^*)] - \mathbf{E}_{(\mathbf{x}, \tilde{\mathbf{y}}) \sim \mathcal{D}} [u(\mathbf{x}, \tilde{p}(\mathbf{x}), \tilde{\mathbf{y}})] \right| \leq \varepsilon. \quad (9)$$

By design, Loss OI implies omniprediction.

Claim 2.10 (Proposition 4.5, [GHK+23]). *If the predictor \tilde{p} is $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -loss OI, then it is an $(\mathcal{L}, \mathcal{C}, \varepsilon)$ -omnipredictor.*

Indeed, if the expected value of u is nonpositive for all $u \in \mathcal{U}(\mathcal{L}, \mathcal{C})$, then \tilde{p} must achieve loss competitive with all $c \in \mathcal{C}$. The argument leverages the fact that u must be nonpositive when $\tilde{\mathbf{y}} \sim \text{Ber}(\tilde{p}(\mathbf{x}))$ —after all, in this world \tilde{p} is the Bayes optimal. By indistinguishability, \tilde{p} must also be optimal in the world where outcomes are drawn as \mathbf{y}^* . The converse, however, is not always true.

Here, we introduce Swap Loss OI. In this swap variant, we allow the choice of distinguisher to depend on the predicted value.

Definition 2.11 (Swap Loss OI). *For a collection of loss functions \mathcal{L} , hypothesis class \mathcal{C} and $\varepsilon \geq 0$, for an assignment of loss functions $\{\ell_v \in \mathcal{L}\}_{v \in \text{Im}(\tilde{p})}$ and hypotheses $\{h_v \in \mathcal{H}\}_{v \in \text{Im}(\tilde{p})}$, denote $u_v = u_{\ell_v, c_v} \in \mathcal{U}(\mathcal{L}, \mathcal{C})$. A predictor \tilde{p} is $(\mathcal{L}, \mathcal{C}, \alpha)$ -swap loss OI if for all such assignments,*

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left| \mathbf{E}_{\mathcal{D}_{\mathbf{v}}} [u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*) - u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \tilde{\mathbf{y}})] \right| \leq \alpha.$$

Swap Loss OI naturally strengthens the standard variant of Loss OI: in Loss OI, ℓ, c are fixed and cannot depend on \mathbf{v} . In fact, Swap Loss OI also strengthens swap omniprediction.

Lemma 2.12. *If the predictor \tilde{p} satisfies $(\mathcal{L}, \mathcal{C}, \alpha)$ -swap loss OI, then*

- *it is $(\mathcal{L}, \mathcal{C}, \alpha)$ -loss OI.*
- *it is an $(\mathcal{L}, \mathcal{C}, \alpha)$ -swap omnipredictor.*

We prove this lemma in Section 5.

2.2 Characterizing Swap Learning

With the definitions of swap learning notions in place, we can state our main results.

Equivalence of swap notions. Our first main result (Theorem 4.3) shows the following equivalence between swap notions of multicalibration, omniprediction and agnostic learning. Specifically, the theorem shows that multicalibration is intimately related to omniprediction for the set of convex loss functions \mathcal{L}_{cvx} .⁵

Theorem 1 (Informal statement of Theorem 4.3). *For any hypothesis class \mathcal{C} and predictor $\tilde{p} : \mathcal{X} \rightarrow \{0, 1\}$, the following notions are equivalent:*

- \tilde{p} is \mathcal{C} -swap multicalibrated.
- \tilde{p} is an $(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -swap omnipredictor.
- \tilde{p} is an (ℓ_2, \mathcal{C}) -swap agnostic learner.

In other words, once we move to the swap variants of these notions, minimizing the squared error (ℓ_2) suffices to minimize *every* convex loss function. By the equivalence to multicalibration, we immediately obtain an efficient algorithm for achieving swap agnostic learning for any convex loss, using only an oracle for standard agnostic learning (see Section 3).

Swap Loss OI is multicalibration over augmented class. Our second main result (Theorem 5.1) gives a characterization of Swap Loss OI in terms of swap multicalibration. As we discussed, swap loss OI is the strongest notion of omniprediction to date: all existing notions of omniprediction can be implemented via swap loss OI. We show that this indistinguishability notion can actually be defined as a variant of multicalibration. In order to state the result, we recall the notion of discrete derivatives of loss functions used in [GHK⁺23].

Definition 2.13. *For a loss function $\ell : \{0, 1\} \times \mathbb{R} \rightarrow \mathbb{R}$, define $\partial\ell : \mathbb{R} \rightarrow \mathbb{R}$ as*

$$\partial\ell(t) = \ell(1, t) - \ell(0, t).$$

For a family of loss functions \mathcal{L} and a hypothesis class \mathcal{C} , let the classes of discrete derivatives be given as $\partial\mathcal{L} = \{\partial\ell\}_{\ell \in \mathcal{L}}$ and $\partial\mathcal{L} \circ \mathcal{C} = \{\partial\ell \circ c\}_{\ell \in \mathcal{L}, c \in \mathcal{C}}$.

We prove the equivalence of $(\mathcal{L}, \mathcal{C})$ -swap loss OI (for loss classes \mathcal{L} that satisfy the same niceness property as above and contain the squared loss) and $\partial\mathcal{L} \circ \mathcal{C}$ -swap multicalibration.

Theorem 2 (Informal statement of Theorem 5.1). *Let \mathcal{L} be a family of “nice” loss functions containing the squared loss ℓ_2 . For any hypothesis class \mathcal{C} , the following notions are equivalent:*

- $(\mathcal{L}, \mathcal{C})$ -swap loss OI.
- $(\partial\mathcal{L} \circ \mathcal{C})$ -swap multicalibration.

Interestingly, the same class $\partial\mathcal{L} \circ \mathcal{C}$ appears in the work of [GHK⁺23], where they show that $(\mathcal{L}, \mathcal{C})$ -loss OI is implied by calibration and $\partial\mathcal{L} \circ \mathcal{C}$ -multiaccuracy. We show that the stronger notion of $\partial\mathcal{L} \circ \mathcal{C}$ -multicalibration in fact gives an equivalence to $(\mathcal{L}, \mathcal{C})$ -swap loss OI. Along the way, we prove that swap loss OI for just the squared loss implies calibration.

⁵For technical reasons, the loss class \mathcal{L}_{cvx} consists of all convex loss functions that also satisfy a few “niceness” properties, stated formally in Definition 4.1).

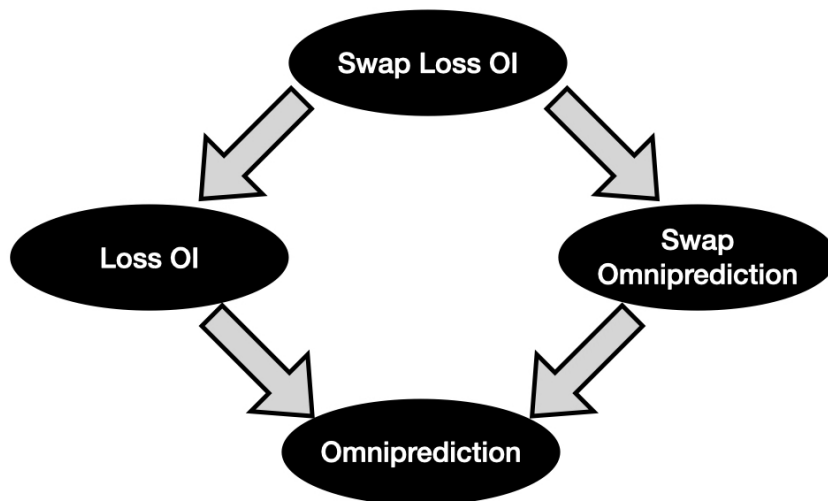


Figure 1: Relation between notions of omniprediction

Separating notions of omniprediction. Rounding out our main results, we paint a complete picture of the relationships between various omniprediction notions by establishing implications and separations between them. The implications are illustrated in Figure 1. That loss OI implies omniprediction (without swap) was proved in [GHK⁺23]; the remaining implications are proved in Section 6. We complement these implications with the following separation result.

Theorem 3. *$(\mathcal{L}, \mathcal{C})$ -Loss OI and $(\mathcal{L}, \mathcal{C})$ -swap omniprediction are incomparable.*

The separation result is proved using the ℓ_2 and ℓ_4 loss, both of which are convex. The theorem statement follows as a corollary of the lemmas proved in Section 6. Our results show that Figure 1 cannot be simplified further for generic choices of \mathcal{L} and \mathcal{C} (although other implications could hold for specific loss and hypothesis classes). In all, we conclude the following statements:

- Swap loss OI is more powerful than loss OI, and swap omniprediction is more powerful than standard omniprediction.
- Swap loss OI is the most powerful of the four notions, whereas standard omniprediction is the least powerful.

Summary of contributions. Our work adds a new and significant connection in the theory of agnostic learning and multi-group fairness. In particular, by borrowing the notion of swapping from online learning, we uncover surprisingly-powerful, but feasible solution concepts for supervised learning. Namely, we can achieve the swap agnostic learning notions assuming only sample access to the data source and a weak agnostic learner for the collection \mathcal{C} . The equivalence we draw between swap agnostic learning for the squared error and multicalibration with swap omniprediction *for all convex losses* highlights the power of simple goals like squared error minimization and calibration.

Our work gives a complete characterization of the notions we study at the “upper end” of strength (i.e., swap variants). A fascinating outstanding question to address in future research

is whether there is a similar characterization of *standard* omniprediction in terms of multi-group fairness notions.

2.3 Discussion and Related Work

Theorem 1 shows the first *reverse* connection between an omniprediction notion and multi-group fairness. It tells us that multicalibration, while it does not explicitly mention any loss function, actually *requires* loss minimization in a deep sense. Note that convexity plays a key role in the statement; indeed, an example from [GKR⁺22] shows that \mathcal{C} -swap multicalibration does not give (even standard) omniprediction for all non-convex losses.

The equivalence to swap agnostic ℓ_2 -loss minimization shows that the problem of finding a multicalibrated predictor can be formulated as a loss minimization problem, using the language of swap agnostic learning. One can also show that (ℓ_2, \mathcal{C}) -agnostic learning (without the swap) results in a predictor that guarantees \mathcal{C} -multiaccuracy. A similar result appears in [GHK⁺23, Theorem 5.6]. In a sense, our results characterizing multicalibration can be viewed as a lifting of the earlier result on multiaccuracy to a more powerful (swap) model. This “hierarchy” of implications bears similarity to results from the outcome indistinguishability literature, where multicalibration corresponds to indistinguishability in the prediction-access model, whereas multiaccuracy corresponds to indistinguishability in the weaker no-access model [DKR⁺21].

Multicalibration as no-swap-regret relative to \mathcal{C} . Our notions of swap loss minimization and swap omniprediction are inspired by notions of swap regret minimization in online learning [FV98, FV99, BM07a]. These classic results of [FV98, FV99] relate calibration and swap regret, whereas [BM07a] show a generic way to convert a low external regret algorithm to ones with low internal and swap regret. The study of online learning and regret minimization is extensive, with deep connections to calibration and equilibria in game theory, we refer the readers to [CBL06, BM07b] for comprehensive surveys.

Concretely in our language, Foster and Vohra [FV98] characterize calibration as (ℓ_2, V) -swap agnostic learning, where $V = \{v \in [0, 1]\}$ is the set of all constant predictors. Since they work in the online setting, they phrase their result in terms of swap regret relative to strategies in V (the terminology of swap regret comes from [BM07a]). Since we work in the batch setting, we can rephrase their result in terms of swap agnostic learning, which is a batch analogue of swap regret. [FV98, Theorem 3] characterizes the calibration score of a predictor \tilde{p} in terms of the regret it suffers relative to all constant predictions in V measured under the squared loss (which measures how much it gains by predicting $w \in [0, 1]$ in place of $v \in \text{Im}(\tilde{p})$). Theorem 1 gives a similar characterization for \mathcal{C} -multicalibrated predictors, where the no-regret guarantee now applies to strategies $c \in \mathcal{C}$. A \mathcal{C} -multicalibrated predictor cannot reduce its squared loss by predicting $c \in \mathcal{C}$ in place of $v \in \text{Im}(\tilde{p})$, moreover this connection goes both ways.

Recent work has established rich connections between online learning and regret minimization on one hand, and multigroup fairness notions on the other: multicalibration in the online setting has been considered in the work of [GJN⁺22], while [BL20, GJKR18] relate multigroup fairness to questions in online learning. Our work contributes a new facet to these connections: it introduces new notions motivated by the regret minimization literature to agnostic loss minimization and omniprediction, and shows that these new notions have tight characterizations in terms of multi-group fairness notions.

Loss minimization notion	Multi-group fairness guarantee	Reference
$(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -swap loss OI	$\partial\mathcal{L}_{\text{cvx}} \circ \mathcal{C}$ -swap multicalibration	Theorem 5.1 (\Leftrightarrow)
$(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -swap omniprediction	\mathcal{C} -swap multicalibration	Theorem 4.3 (\Leftrightarrow)
$(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -loss OI	Calibration + $\partial\mathcal{L}_{\text{cvx}} \circ \mathcal{C}$ - multiaccuracy	[GHK ⁺ 23] (\Leftarrow)
$(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -Omniprediction	\mathcal{C} - multicalibration	[GKR ⁺ 22] (\Leftarrow)

Table 1: Computational guarantees needed for $(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -loss minimization

Comparing notions for \mathcal{L}_{cvx} . Convex losses play a key role in many of the proofs and equivalences that we establish. In particular, for each notion of omniprediction / OI we consider, instantiated with $\mathcal{L} = \mathcal{L}_{\text{cvx}}$ (for arbitrary \mathcal{C}), there is a corresponding multi-group fairness notion that guarantees the notion. We list the simplest known multi-group guarantee for achieving each notion in Table 1. Observe that the implications for the swap notions go both ways—the multi-group guarantee actually characterizes the notion of loss minimization. For the non-swap versions of omniprediction and loss-OI, a tight characterization in terms of some multigroup notion is unknown. Concretely, we leave open the question of whether there is a multi-group characterization of $(\mathcal{L}_{\text{cvx}}, \mathcal{C})$ -omniprediction, which has been alluded to since the original omniprediction work [GKR⁺22, GHK⁺23].

Multi-group fairness. Notions of multi-group fairness were introduced in the work of [HKRR18, KGZ19] and [KNRW17], following [KMR17]. The flagship notion of multicalibration has been extended to several other settings including multiclass predictions [DKR⁺22], real-valued labels [GJN⁺22, JLP⁺21], online learning [GJN⁺22] and importance weights [GRSW22]. Alternate definitions and extensions of the standard notions of multicalibration have been proposed in [GKR⁺22, GKSZ22, DDZ23, DLLT23]. Multicalibration has also proved to have unexpected connections to many other domains, including computational indistinguishability [DKR⁺21], domain adaptation [KKG⁺22], and boosting [GKR⁺22, GHK⁺23].

Independent concurrent work. Independent work of Dwork, Lee, Lin and Tankala [DLLT23] defines the notion of *strict multicalibration*. They connect this notion to the Szemerédi regularity lemma and its weaker version due to Frieze and Kannan. Their notions bears important similarities to swap multicalibration, but it is different. Like swap multicalibration, strict multicalibration involves switching the order of expectations and max. But they require a statistical closeness guarantee, whereas we only require a *first order guarantee*, that $c(\mathbf{x})$ be uncorrelated with $\mathbf{y} - \mathbf{v}$ conditioned on \mathbf{v} . For Boolean c , these notions are essentially equivalent. In the setting of real-valued functions \mathcal{C} , however, they seem rather different. Swap multicalibration is closed under linear/convex combinations of \mathcal{C} , whereas strict multicalibration might not be. In the real-valued setting, strict multicalibration seems to require a stronger learning primitive than weak agnostic learner for \mathcal{C} , which suffices for swap multicalibration. The analogous statistical distance notion of multiaccuracy is at least as powerful as the notion of hypothesis OI for arbitrary loss families \mathcal{L} , which is known to require weak learning for $\partial\mathcal{L} \circ \mathcal{C}$ [GHK⁺23, Theorem 4.9].

The independent work of Globus-Harris, Harrison, Kearns, Roth and Sorrell [GHHK⁺23] relates multicalibration to real-valued boosting to minimize ℓ_2 loss. The implication that (ℓ_2, \mathcal{C}) swap-agnostic learning implies multicalibration follows from Part(1) of Theorem 3.2 in their paper.

They prove that a violation of multicalibration leads to a better strategy for the adversary in the (ℓ_2, \mathcal{C}) -swap minimization game. They do not consider the notion of swap multicalibration, so their result is not a tight characterization unlike our equivalence of Items (1) and (3) in Theorem 4.3.

Organization of Manuscript

The remainder of the manuscript is structured as follows. In Section 3, we give a deeper overview of swap multicalibration, arguing that it should be thought of as essentially equivalent to standard multicalibration. In Section 4, we prove the main equivalence of the three notions of swap learning. In Section 5, we give the characterization of swap loss OI in terms of swap multicalibration over the augmented class. Finally, in Section 6, we prove separation results, establishing the remaining non-implications depicted in Figure 1.

3 Properties of Swap Multicalibration

In this section, we formalize the intuitive claim that standard multicalibration as defined by [HKRR18] is essentially equivalent to swap multicalibration. We start, however, with a concrete difference between the standard definition of multicalibration and the swap variant. Starting from the definitions, we denote the multicalibration error of \tilde{p} with respect to \mathcal{C} under \mathcal{D} as

$$\text{MCE}_{\mathcal{D}}(f, \mathcal{C}) = \max_{c \in \mathcal{C}} \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [c(\mathbf{x})(\mathbf{y} - \mathbf{v})] \right| \right]$$

and the swap multicalibration error of \tilde{p} with respect to \mathcal{C} under \mathcal{D} as

$$\text{sMCE}_{\mathcal{D}}(\tilde{p}, \mathcal{C}) = \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [c(\mathbf{x})(\mathbf{y} - \mathbf{v})] \right| \right].$$

To contrast the notions, consider the following notions of “bad” prediction intervals for a given \mathcal{C} and predictor \tilde{p} .

Definition 3.1. Fix a hypothesis class \mathcal{C} . Let $\beta \in [0, 1]$. For each $c \in \mathcal{C}$, let the β -bad prediction intervals for c be given as the $v \in [0, 1]$ with multicalibration violation of at least β .

$$\text{Bad}_{c,\beta}(\tilde{p}) = \left\{ v \in [0, 1] : \left| \mathbf{E}_{\mathcal{D}|v} [c(\mathbf{x})(\mathbf{y} - v)] \right| \geq \beta \right\}$$

Further, let the global β -bad prediction intervals for \mathcal{C} be the union over $c \in \mathcal{C}$.

$$\text{Bad}_{\mathcal{C},\beta}(\tilde{p}) = \bigcup_{c \in \mathcal{C}} \text{Bad}_{c,\beta}(\tilde{p}) = \left\{ v \in [0, 1] : \exists c \in \mathcal{C} \text{ s.t. } \left| \mathbf{E}_{\mathcal{D}|v} [c(\mathbf{x})(\mathbf{y} - v)] \right| \geq \beta \right\}.$$

Under standard multicalibration, we can bound the probability of $\text{Bad}_{c,\beta}(\tilde{p})$ under $\mathcal{D}_{\tilde{p}}$ for each $c \in \mathcal{C}$ by an application of Markov’s inequality. Note that the set $\text{Bad}_{c,\beta}(\tilde{p})$ depends on our choice of $c \in \mathcal{C}$. That is, different prediction intervals might be bad for different choices of c . In principle, one might hope to take a union bound over all $c \in \mathcal{C}$ to bound the probability of landing in a globally-bad set $\text{Bad}_{\mathcal{C},\beta}(\tilde{p})$. But for large (or infinite) hypothesis classes \mathcal{C} , the resulting bound becomes vacuous quickly.

In contrast, swap multicalibration allows us to define a global set $\text{Bad}_{\mathcal{C},\beta}(\tilde{p})$ of bad prediction intervals $v \subseteq \text{Im}(\tilde{p})$ with small probability under $\mathcal{D}_{\tilde{p}}$. That is, with good probability under $\mathcal{D}_{\tilde{p}}$, the correlation between *every* $c \in \mathcal{C}$ and the residue $\mathbf{y} - v$ conditioned on $\tilde{p}(\mathbf{x}) = v$ will be small. This difference is summarized in the following lemma.

Lemma 3.2. *If \tilde{p} is (\mathcal{C}, α) -multicalibrated, then for every $c \in \mathcal{C}$,*

$$\Pr_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} [\mathbf{v} \in \text{Bad}_{c,\beta}(\tilde{p})] \leq \frac{\alpha}{\beta}.$$

If \tilde{p} is (\mathcal{C}, α) -swap multicalibrated, then

$$\Pr_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} [\mathbf{v} \in \text{Bad}_{\mathcal{C},\beta}(\tilde{p})] \leq \frac{\alpha}{\beta}.$$

Standard versus swap multicalibration. In definition, swap multicalibration is a stronger notion than standard multicalibration. That said, we show that, both analytically and algorithmically, the two notions are tightly connected. In particular, any method (existing and hypothetical) for achieving (standard) \mathcal{C} -multicalibration can be used to achieve \mathcal{C} -swap multicalibration without significant deterioration in the accuracy parameter.

First, we show that any predictor \tilde{p} that guarantees standard multicalibration can be converted into a predictor \bar{p} that guarantees swap multicalibration, using a simple bucketing procedure. This bucketing causes some degradation in the multicalibration error parameter, but the class \mathcal{C} remains unaffected.

Definition 3.3. *Let $\delta \in [0, 1]$ so that $m = 1/\delta \in \mathbb{Z}$. Define $B_j = [(j-1)\delta, j\delta]$ for $j \in [m-1]$ and $B_m = [1-\delta, 1]$. Define the predictor \bar{p}_δ where for every x such that $\tilde{p}(x) \in B_j$, $\bar{p}_\delta(x) = j\delta$.*

Claim 3.4. *Let $\tilde{p} : \mathcal{X} \rightarrow [0, 1]$ be a (\mathcal{C}, α) -multicalibrated. For any $\delta \in [0, 1]$ where $1/\delta \in \mathbb{Z}$, the predictor \bar{p}_δ is $(\mathcal{C}, 2\sqrt{\alpha/\delta} + \delta)$ -swap multicalibrated, and $\max_{x \in \mathcal{X}} |\tilde{p}(x) - \bar{p}_\delta(x)| \leq \delta$.*

While the bound of the theorem is valid for all δ , the swap multicalibration guarantee is only meaningful when $\delta \geq \Omega(\alpha)$. Thus, by discretizing a (\mathcal{C}, α) -multicalibrated predictor, we obtain a $(\mathcal{C}, O(\alpha^{1/3}))$ -swap multicalibrated predictor. Note that the resulting predictor \bar{p} is actually close to the original \tilde{p} in ℓ_∞ ; in other words, any (standard) multicalibrated predictor is actually close to a swap multicalibrated predictor.

Proof of Claim 3.4. The ℓ_∞ bound is immediate from the definition of \bar{p} . We bound the swap multicalibration error of \bar{p} . We have $\bar{p}(\mathbf{x}) = j\delta$ iff $\tilde{p}(\mathbf{x}) \in B_j$, so that $|\bar{p}(\mathbf{x}) - j\delta| \leq \delta$ holds conditioned on this event. So

$$\begin{aligned} \text{sMCE}_{\mathcal{D}}(\bar{p}, \mathcal{C}) &= \sum_{j \in [m]} \Pr[\bar{p}(\mathbf{x}) = j\delta] \max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}}[c(\mathbf{x})(\mathbf{y} - j\delta) | \bar{p}(\mathbf{x}) = j\delta] \right| \\ &= \sum_{j \in [m]} \Pr[\tilde{p}(\mathbf{x}) \in B_j] \max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}}[c(\mathbf{x})(\mathbf{y} - j\delta) | \tilde{p}(\mathbf{x}) \in B_j] \right| \\ &\leq \sum_{j \in [m]} \Pr[\tilde{p}(\mathbf{x}) \in B_j] \left(\delta + \max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}}[c(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x})) | \tilde{p}(\mathbf{x}) \in B_j] \right| \right) \\ &\leq \delta + \sum_{j \in [m]} \Pr[\tilde{p}(\mathbf{x}) \in B_j] \max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}}[c(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x})) | \tilde{p}(\mathbf{x}) \in B_j] \right| \end{aligned} \quad (10)$$

Let us fix a bucket B_j and a particular $c \in \mathcal{C}$. For $\beta \geq \alpha$ to be specified later we have

$$\begin{aligned}
|\mathbf{E}[c(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x})) | \tilde{p}(\mathbf{x}) \in B_j]| &\leq \Pr[c(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x})) \geq \beta | \tilde{p}(\mathbf{x}) \in B_j] + \beta \Pr[c(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x})) \leq \beta | \tilde{p}(\mathbf{x}) \in B_j] \\
&\leq \frac{\Pr[\tilde{p}(\mathbf{x}) \in \text{Bad}_{c,\beta}(\tilde{p}) \cap B_j]}{\Pr[\tilde{p}(\mathbf{x}) \in B_j]} + \beta \\
&\leq \frac{\Pr[\tilde{p}(\mathbf{x}) \in \text{Bad}_{c,\beta}(\tilde{p})]}{\Pr[\tilde{p}(\mathbf{x}) \in B_j]} + \beta \\
&\leq \frac{\alpha/\beta}{\Pr[\tilde{p}(\mathbf{x}) \in B_j]} + \beta.
\end{aligned}$$

Since this bound holds for every c , it holds for the max over $c \in \mathcal{C}$ conditioned on $\tilde{p}(\mathbf{x}) \in B_j$. Hence

$$\begin{aligned}
\sum_{j \in [m]} \Pr[\tilde{p}(\mathbf{x}) \in B_j] \max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}}[c(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x})) | \tilde{p}(\mathbf{x}) \in B_j] \right| &\leq \sum_{j \in [m]} \Pr[\tilde{p}(\mathbf{x}) \in B_j] \left(\frac{\alpha/\beta}{\Pr[\tilde{p}(\mathbf{x}) \in B_j]} + \beta \right) \\
&\leq \frac{\alpha}{\beta \delta} + \beta,
\end{aligned}$$

where we use $m = 1/\delta$. Plugging this back into Equation (10) gives

$$\text{sMCE}_{\mathcal{D}}(\bar{p}, \mathcal{C}) = \mathbf{E}_{\mathbf{v} \sim \bar{p}_{\mathcal{D}}} \left[\max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}}[c(\mathbf{x})(\mathbf{y} - \mathbf{v})] \right| \right] \leq \frac{\alpha}{\beta \delta} + \beta + \delta.$$

Taking $\beta = \sqrt{\alpha/\delta}$ gives the desired claim. ■

While this generic transformation suffers a polynomial loss in the accuracy parameter, such a loss seems not to be necessary algorithmically. As we describe next, known algorithms for achieving multicalibration [HKRR18, GKR+22] actually guarantee swap multicalibration without any modification.

3.1 Algorithms for Swap Multicalibration

In Lemma 3.6, we show that swap multicalibration is feasible, assuming an oracle for the agnostic learning problem on \mathcal{C} . First, we give a brief overview of the prior work on algorithms for standard multicalibration.

The work introducing multicalibration [HKRR18] gives a boosting-style algorithm for learning multicalibrated predictors that has come to be known as MCBoost. The algorithm is an iterative procedure: starting with a trivial predictor, the MCBoost searches for a supported value $v \in \text{Im}(\tilde{p})$ and “subgroup” $c_v \in \mathcal{C}$ that violate the multicalibration condition. Note that some care has to be taken to ensure that the predictor \tilde{p} stays supported on finitely many values, and that each of these values maintains significant measure in the data distribution $\mathcal{D}_{\tilde{p}}$. In this pseudocode, we ignore these issues; [HKRR18] handles them in full detail.

Importantly, the search over \mathcal{C} for condition (11) can be reduced to weak agnostic learning, defined formally below. In each iteration, we pass the learner samples drawn from the data distribution, but labeled according to $\mathbf{z} = \mathbf{y} - v$ when $\tilde{p}(\mathbf{x}) = v$.

Algorithm 1 MCBBoost [HKRR18]

Parameters: hypothesis class \mathcal{C} and $\alpha > 0$

Given: Dataset S sampled from \mathcal{D}

Initialize: $\tilde{p}(x) \leftarrow 1/2$.

Repeat:

if $\exists v \in \text{Im}(\tilde{p})$ and $c_v \in \mathcal{C}$ such that

$$\mathbf{E}[c_v(\mathbf{x}) \cdot (\mathbf{y} - v) \mid \tilde{p}(\mathbf{x}) = v] > \text{poly}(\alpha) \quad (11)$$

update $\tilde{p}(x) \leftarrow \tilde{p}(x) + \eta c_v(x) \cdot \mathbf{1}[\tilde{p}(x) = v]$

Return: \tilde{p}

Definition 3.5 (Weak agnostic learning). *Suppose \mathcal{D} is a data distribution supported on $\mathcal{X} \times [-1, 1]$. For a hypothesis class \mathcal{C} , a weak agnostic learner WAL solves the following promise problem: for some accuracy parameter $\alpha > 0$, if there exists some $c \in \mathcal{C}$ such that*

$$\mathbf{E}_{(\mathbf{x}, \mathbf{z}) \sim \mathcal{D}} [c(\mathbf{x}) \cdot \mathbf{z}] \geq \alpha$$

then WAL_α returns some $h : \mathcal{X} \rightarrow \mathbb{R}$ such that⁶

$$\mathbf{E}_{(\mathbf{x}, \mathbf{z}) \sim \mathcal{D}} [h(\mathbf{x}) \cdot \mathbf{z}] \geq \text{poly}(\alpha).$$

Weak agnostic learning is a basic supervised learning primitive used in boosting algorithms. Through the connection to boosting, weak agnostic learning is polynomial-time equivalent to agnostic learning, and inherits its statistical efficiency (scaling with the VC-dimension of \mathcal{C}), and also its worst-case computational hardness [Pie12].

MCBoost reduces the problem of swap multicalibration to the standard weak agnostic learning task. It is not hard to see that while MCBBoost was designed to guarantee multicalibration, it actually guarantees swap multicalibration, using an oracle for weak agnostic learning. By this reduction, we establish the feasibility of (\mathcal{C}, α) -swap multicalibration for any learnable \mathcal{C} .

Lemma 3.6. *Suppose the collection \mathcal{C} has a weak agnostic learner, WAL. For any $\alpha > 0$, MCBBoost makes at most $\text{poly}(1/\alpha)$ calls to WAL_α , and returns a (\mathcal{C}, α) -swap multicalibrated predictor.*

Proof. The iteration complexity of MCBBoost is inverse polynomially related to the size of the violations we discover in (11). A standard potential argument and the formal reduction of search to weak agnostic learning can be found in [HKRR18].

By the termination condition, we can see that \tilde{p} must actually be (\mathcal{C}, α) -swap multicalibrated. In particular, when the algorithm terminates, then for all $v \in \text{Im}(\tilde{p})$, we have that

$$\max_{c_v \in \mathcal{C}} \mathbf{E}[c_v(\mathbf{x}) \cdot (\mathbf{y} - v) \mid \tilde{p}(\mathbf{x}) = v] \leq \text{poly}(\alpha) \leq \alpha.$$

Therefore, averaging over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$, we obtain the guarantee. ■

⁶We are informal about the polynomial factor in the guarantee of the weak agnostic learner. The smaller the exponent, the stronger the learning guarantee; i.e., we want WAL_α to return a hypothesis with correlation with \mathbf{z} as close to $\Omega(\alpha)$ as possible.

4 Swap Agnostic Learning, Swap Omniprediction, and Swap Multicalibration are Equivalent

Our main result is an equivalence between swap agnostic learning, swap omniprediction, and swap multicalibration. Concretely, this equivalence shows that swap agnostic learning for the squared error is sufficient to guarantee swap omniprediction for all (nice) convex loss functions. We begin with some preliminaries, then formally state and prove the equivalence. We conclude the section by showing how to use the existing framework for learning multicalibrated predictors to achieve swap agnostic learning for any convex loss.

Nice loss functions. For a loss function $\ell : \{0, 1\} \times \mathbb{R} \rightarrow \mathbb{R}$, we extend ℓ linearly to allow the first argument to take values in the range $p \in [0, 1]$ as:

$$\ell(p, t) = \mathbf{E}_{\mathbf{y} \sim \text{Ber}(p)} [\ell(\mathbf{y}, t)] = p \cdot \ell(1, t) + (1 - p) \cdot \ell(0, t).$$

We say the loss function is *convex* if for $y \in \{0, 1\}$, $\ell(y, t)$ is a convex function of t . By linearity, this convexity property holds for $\ell(p, t)$ for all p .

As in [GHK⁺23], for a loss ℓ , we define the partial difference function $\partial\ell : \mathbb{R} \rightarrow \mathbb{R}$ as

$$\partial\ell(t) = \ell(1, t) - \ell(0, t).$$

We define a class of “nice” loss functions, which obey a minimal set of boundedness and Lipschitzness conditions.

Definition 4.1. For a constant $B > 0$, a loss function is *B-nice* if there exists an interval $I_\ell \subseteq \mathbb{R}$ such that the following conditions hold:

1. (*Optimality*) If $\Pi_\ell : \mathbb{R} \rightarrow I_\ell$ denotes projection onto the interval I_ℓ , then $\ell(p, \Pi_\ell(t)) \leq \ell(p, t)$ for all $t \in \mathbb{R}$ and $p \in [0, 1]$.
2. (*Lipschitzness*) For $y \in \{0, 1\}$ and $t \in I_\ell$, $\ell(y, t)$ is 1-Lipschitz as a function of t .
3. (*Bounded difference*) For $t \in I_\ell$, $|\partial\ell(t)| \leq B$.

The class of all *B-nice* loss functions by $\mathcal{L}(B)$. The subset of *B-nice* convex loss functions is denoted by $\mathcal{L}_{\text{cvx}}(B)$.

Bounded loss functions generalize the idea of loss functions defined over a fixed interval of \mathbb{R} (by optimality) and of bounded output range (by bounded difference). By optimality of nice loss functions, we may assume $k_\ell : [0, 1] \rightarrow I_\ell$ since we do not increase the loss by projection onto I_ℓ . Indeed, for convex losses ℓ , the natural choice for I_ℓ is to take the interval $[k_\ell(0), k_\ell(1)]$. The bounded difference condition implies that ℓ is Lipschitz in its first argument, a property that will be useful.

Lemma 4.2. For every $\ell \in \mathcal{L}(B)$ and $t_0 \in I_\ell$, the function $\ell(p, t_0)$ is *B-Lipschitz* w.r.t. p .

Concept classes. For a concept class of functions $\mathcal{C} : \{c : \mathcal{X} \rightarrow \mathbb{R}\}$, we assume that \mathcal{C} is closed under negation, and it contains the constant functions 0 and 1. Denoting $\|\mathcal{C}\|_\infty = \max |c(x)|$ over $c \in \mathcal{C}, x \in \mathcal{X}$, we say that \mathcal{C} is bounded if $\|\mathcal{C}\|_\infty \leq 1$. For $W \in \mathbb{R}^+$, let $\text{Lin}(\mathcal{C}, W)$ be all functions that can be expressed as a (W -sparse) linear combination of base concepts from \mathcal{C} ,

$$c_w(x) = \sum_{c \in \mathcal{C}} w_c \cdot c(x), \quad \sum_{c \in \mathcal{C}} |w_c| \leq W.$$

Note that for bounded \mathcal{C} , the norm of linear combinations scales gracefully with the sparsity, $\|\text{Lin}(\mathcal{C}, W)\|_\infty \leq W$. We define $\text{Lin}(\mathcal{C})$ to be the set of all linear combinations with no restriction on the weights of the coefficients.

4.1 Statement of Equivalence

With these notions in place, we can formally state the main result.

Theorem 4.3. *Let \tilde{p} be a predictor, \mathcal{C} be a bounded hypothesis class, and $\mathcal{L}_{\text{cvx}}(B)$ be the class of B -nice convex loss functions. The following properties are equivalent.⁷*

1. \tilde{p} is (\mathcal{C}, α_1) -swap multicalibrated.
2. \tilde{p} is an $(\mathcal{L}_{\text{cvx}}(B), \text{Lin}(\mathcal{C}, W), O((W + B)\alpha_2))$ -swap omnipredictor, for all $W \geq 1, B \geq 0$.
3. \tilde{p} is an $(\ell_2, \text{Lin}(\mathcal{C}, 2), \alpha_3)$ -swap agnostic learner.

4.1.1 Preliminary Results

In preparation for proving the theorem, we establish some preliminary results. We define a function $\alpha : \text{Im}(\tilde{p}) \rightarrow [-1, 1]$ which measures the maximum correlation between $c \in \mathcal{C}$ and $\mathbf{y} - v$, conditioned on a prediction value $v \in \text{Im}(\tilde{p})$. Let

$$\alpha(v) = \left| \max_{c \in \mathcal{C}} \mathbf{E}_{\mathcal{D}|v} [c_v(\mathbf{x})(\mathbf{y} - v)] \right|.$$

Using this notation, (\mathcal{C}, α_0) -swap multicalibration can be written as

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} [\alpha(\mathbf{v})] \leq \alpha_0.$$

We observe that swap multicalibration is closed under bounded linear combinations of \mathcal{C} , like with standard multicalibration.

Claim 4.4. *For every $h \in \text{Lin}(\mathcal{C}, W)$ and $v \in \text{Im}(\tilde{p})$, we have*

$$\max_{h \in \text{Lin}(\mathcal{C}, W)} \left| \mathbf{E}_{\mathcal{D}} [h(\mathbf{x})(\mathbf{y} - v) | \tilde{p}(\mathbf{x}) = v] \right| \leq W\alpha(v).$$

Let $p_v^* = \mathbf{E}[\mathbf{y} | \tilde{p}(\mathbf{x}) = v]$. Then

$$|p_v^* - v| \leq \alpha(v). \tag{12}$$

⁷When we say these conditions are equivalent, we mean that they imply each other with parameters α_i that are polynomially related. The relations we derive result in at most a quadratic loss in parameters.

Proof. Suppose that $h \in \text{Lin}(\mathcal{C}, W)$ of the form $h(x) = \sum_{c \in \mathcal{C}} w_c \cdot c(x)$. From Claim 2.8, we know that the multicalibration violation for $c \in \mathcal{C}$ is bounded by $\alpha(v)$ for every $v \in \text{Im}(\tilde{p})$.

$$\begin{aligned} |\mathbf{E}[h(\mathbf{x})(\mathbf{y} - v) \mid \tilde{p}(\mathbf{x}) = v]| &= \left| \mathbf{E} \left[\sum_{c \in \mathcal{C}} w_c \cdot c(\mathbf{x})(\mathbf{y} - v) \mid \tilde{p}(\mathbf{x}) = v \right] \right| \\ &\leq \left(\sum_{c \in \mathcal{C}} |w_c| \right) \cdot \max_{c \in \mathcal{C}} |\mathbf{E}[c(\mathbf{x})(\mathbf{y} - v) \mid \tilde{p}(\mathbf{x}) = v]| \\ &\leq W \cdot \alpha(v) \end{aligned}$$

The inequalities follow by Holder's inequality and the assumed bound on the weight of W for $h \in \text{Lin}(\mathcal{C}, W)$.

Equation (12) follows since $1 \in \mathcal{C}$. ■

Claim 4.5. For $h \in \text{Lin}(\mathcal{C}, w)$, $v \in \text{Im}(\tilde{p})$ and $y \in \{0, 1\}$, define the following conditional expectations:

$$\begin{aligned} \mu(h : v) &= \mathbf{E}[h(\mathbf{x}) \mid \tilde{p}(\mathbf{x}) = v] \\ \mu(h : v, y) &= \mathbf{E}[h(\mathbf{x}) \mid (\tilde{p}(\mathbf{x}), \mathbf{y}) = (v, y)]. \end{aligned}$$

Then for each $y \in \{0, 1\}$

$$\Pr[\mathbf{y} = y \mid \tilde{p}(\mathbf{x}) = v] |\mu(h : v, y) - \mu(h : v)| \leq (W + 1)\alpha(v). \quad (13)$$

Proof. Recall that $\text{Cov}[\mathbf{y}, \mathbf{z}] = \mathbf{E}[\mathbf{y}\mathbf{z}] - \mathbf{E}[\mathbf{y}]\mathbf{E}[\mathbf{z}]$. For any $h \in \text{Lin}(\mathcal{C}, W)$ we have

$$\begin{aligned} |\text{Cov}[\mathbf{y}, h(\mathbf{x}) \mid \tilde{p}(\mathbf{x}) = v]| &= |\mathbf{E}[h(\mathbf{x})(\mathbf{y} - \mathbf{E}[\mathbf{y}]) \mid \tilde{p}(\mathbf{x}) = v]| \\ &= |\mathbf{E}[h(\mathbf{x})(\mathbf{y} - v) \mid \tilde{p}(\mathbf{x}) = v]| + |\mathbf{E}[(v - \mathbf{y}) \mid \tilde{p}(\mathbf{x}) = v]| \\ &\leq (W + 1)\alpha(v) \end{aligned}$$

where we use the fact that $h \in \text{Lin}(\mathcal{C}, W)$ and $1 \in \mathcal{C}$. Since $\mathbf{y} \in \{0, 1\}$, this implies the claimed bounds by standard properties of covariance (see [GKR⁺22, Corollary 5.1]). ■

Next we show the following lemma, which shows that one can replace $h(\mathbf{x})$ by the constant $\Pi_\ell(\mu(h : v))$ without a large increase in the loss.

Lemma 4.6. For all $h \in \text{Lin}(\mathcal{C}, W)$, $v \in \text{Im}(\tilde{p})$ and loss $\ell \in \mathcal{L}_{\text{cvx}}(B)$, we have

$$\mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, \Pi_\ell(\mu(h : v)))] \leq \mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, h(\mathbf{x}))] + 2(W + 1)\alpha(v). \quad (14)$$

Proof. For any $y \in \{0, 1\}$,

$$\begin{aligned} \mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, h(\mathbf{x})) \mid (\tilde{p}(\mathbf{x}), \mathbf{y}) = (v, y)] &= \mathbf{E}_{\mathcal{D}|_v} [\ell(y, h(\mathbf{x})) \mid (\tilde{p}(\mathbf{x}), \mathbf{y}) = (v, y)] \\ &\geq \ell(y, \mathbf{E}[h(\mathbf{x}) \mid (\tilde{p}(\mathbf{x}), \mathbf{y}) = (v, y)]) \end{aligned} \quad (15)$$

$$\begin{aligned} &= \ell(y, \mu(h : v, y)) \\ &\geq \ell(y, \Pi_\ell(\mu(h : v, y))). \end{aligned} \quad (16)$$

where Equation (15) uses Jensen's inequality, and Equation (16) uses the optimality of projection for nice loss functions. Further, by the 1-Lipschitzness of ℓ on I_ℓ , and of Π_ℓ on \mathbb{R}

$$\begin{aligned} \ell(y, \Pi_\ell(\mu(h : v, y))) - \ell(y, \Pi_\ell(\mu(h : v))) &\leq |\Pi_\ell(\mu(h : v, y)) - \Pi_\ell(\mu(h : v))| \\ &\leq |\mu(h : v, y) - \mu(h : v)| \end{aligned} \quad (17)$$

Hence we have

$$\begin{aligned} &\mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, \Pi_\ell(\mu(h : v))) - \mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, h(\mathbf{x}))]] \\ &= \sum_{y \in \{0,1\}} \Pr[\mathbf{y} = y | \tilde{p}(\mathbf{x}) = v] (\ell(y, \Pi_\ell(\mu(h : v))) - \mathbf{E}[\ell(y, h(\mathbf{x})) | (\tilde{p}(\mathbf{x}), \mathbf{y}) = (v, y)]) \\ &\leq \sum_{y \in \{0,1\}} \Pr[\mathbf{y} = y | \tilde{p}(\mathbf{x}) = v] (\ell(y, \Pi_\ell(\mu(h : v))) - \ell(y, \Pi_\ell(\mu(h : v, y)))) \quad (\text{By Equation (16)}) \\ &\leq \sum_{y \in \{0,1\}} \Pr[\mathbf{y} = y | \tilde{p}(\mathbf{x}) = v] |\mu(h : v, y) - \mu(h : v)| \quad (\text{by Equation (17)}) \\ &\leq 2(W + 1)\alpha(v). \quad (\text{By Equation (13)}) \end{aligned}$$

■

Next we compare $\Pi_\ell(\mu(h : v))$ with $k_\ell(v)$. It is clear that the latter is better for minimizing loss when $\mathbf{y} \sim \text{Ber}(v)$, by definition. We need to compare the losses when $\mathbf{y} \sim \text{Ber}(p_v^*)$. But p_v^* and v are at most $\alpha(v)$ apart by Equation (12). Hence, by using Lipschitzness, one can infer that $k_\ell(v)$ is better than $\Pi_\ell(\mu(h : v))$ and hence $h(\mathbf{x})$. This is formalized in the following lemma and its proof.

Lemma 4.7. *For all $v \in \text{Im}(\tilde{p})$, $\ell \in \mathcal{L}_{\text{cvx}}(B)$ and $h \in \text{Lin}(\mathcal{C}, W)$, we have*

$$\mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, k_\ell(\tilde{p}(\mathbf{x})))] \leq \mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, h(\mathbf{x}))] + 2(W + B + 1)\alpha(v). \quad (18)$$

Proof. By the definition of k_ℓ , $k_\ell(v)$ minimizes expected loss when $\mathbf{y} \sim \text{Ber}(v)$, so

$$\ell(v, k_\ell(v)) \leq \ell(v, \Pi_\ell(\mu(h : v))) \quad (19)$$

On the other hand,

$$\mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, t)] = \ell(p_v^*, t), \text{ where } p_v^* = \mathbf{E}_{\mathcal{D}|_v} [\mathbf{y}].$$

Thus our goal is compare the losses $\ell(p_v^*, t)$ for $t = k_\ell(v)$ and $t = \Pi_\ell(\mu(h : v))$. Hence, applying Lemma 4.2 gives

$$\begin{aligned} \ell(p_v^*, k_\ell(v)) &\leq \ell(v, k_\ell(v)) + \alpha(v)B \\ -\ell(p_v^*, \Pi_\ell(\mu(h : v))) &\leq -\ell(v, \Pi_\ell(\mu(h : v))) + \alpha(v)B \end{aligned}$$

Subtracting these inequalities and then using Equation (19) gives

$$\ell(p_v^*, k_\ell(v)) - \ell(p_v^*, \Pi_\ell(\mu(h : v))) \leq \ell(v, k_\ell(v)) - \ell(v, \Pi_\ell(\mu(h : v))) + 2B\alpha(v) \quad (20)$$

$$\leq 2\alpha(v)B. \quad (21)$$

We can now write

$$\begin{aligned}
\mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, k_\ell(v))] &= \ell(p_v^*, k_\ell(v)) \\
&\leq \ell(p_v^*, \Pi_\ell(\mu(h : v))) + 2\alpha(v)B && \text{(by Equation (21))} \\
&= \mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, \Pi_\ell(\mu(h : v)))] + 2\alpha(v)B \\
&\leq \mathbf{E}_{\mathcal{D}|_v} [\ell(\mathbf{y}, h(\mathbf{x}))] + 2(W + 1)\alpha(v) + 2B\alpha(v). && \text{(by Equation (14))}
\end{aligned}$$

■

4.1.2 Proof of Equivalence

We now complete the proof of Theorem 4.3.

Proof of Theorem 4.3. (1) \implies (2) Fix a (\mathcal{C}, α_1) -swap multicalibrated predictor \tilde{p} . Fix a choice of loss functions $\{\ell_v \in \mathcal{L}_{\text{cvx}}\}_{v \in \text{Im}(f)}$ and hypotheses $\{h_v \in \mathcal{H}\}_{v \in \text{Im}(f)}$. For each v , we apply Equation (18) with the loss $\ell = \ell_v$, hypothesis $h = h_v$ to get

$$\mathbf{E}_{\mathcal{D}|_v} [\ell_v(\mathbf{y}, k_{\ell_v}(v))] \leq \mathbf{E}_{\mathcal{D}|_v} [\ell_v(\mathbf{y}, h_v(\mathbf{x}))] + 2(W + B + 1)\alpha(v).$$

We now take expectations over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$, and use the $\mathbf{E}[\alpha(\mathbf{v})] \leq \alpha_1$ to derive the desired implication.

(2) \implies (3) with $\alpha_3 = 7\alpha_2$ because ℓ_2 is a 1/2-nice loss function, so we plug in $B = 1/2$ and $W = 2$ into claim (2).

It remains to prove that (3) \implies (1). We show the contrapositive, that if \tilde{p} is not α_1 multicalibrated, then $f = \tilde{p}$ is not a $(\ell_2, \text{Lin}(\mathcal{C}, 2), \alpha_3)$ -swap agnostic learner. By the definition of multicalibration, for every $v \in \text{Im}(\tilde{p})$, there exist c_v such that

$$\begin{aligned}
\mathbf{E}_{\mathcal{D}|_v} [c_v(\mathbf{x})(\mathbf{y} - \tilde{p}(\mathbf{x}))] &= \alpha(v) \\
\mathbf{E}_{\mathcal{D}_{\tilde{p}}} [\alpha(\mathbf{v})] &\geq \alpha_1.
\end{aligned}$$

By negating c_v if needed, we may assume $\alpha(v) \geq 0$ for all v . We now define the updated hypothesis h' where

$$h'(x) = v + \alpha(v)c_v(x) \text{ for } x \in \tilde{p}^{-1}(v)$$

A standard calculation (included in the Appendix) shows that

$$\mathbf{E}_{\mathcal{D}|_v} [(\mathbf{y} - v)^2] - \mathbf{E}_{\mathcal{D}|_v} [(\mathbf{y} - h'(\mathbf{x}))^2] \geq \alpha(v)^2.$$

Taking expectation over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$, we have

$$\mathbf{E}_{\mathcal{D}_{\tilde{p}}} \left(\mathbf{E}_{\mathcal{D}|_v} [(\mathbf{y} - v)^2] - \mathbf{E}_{\mathcal{D}|_v} [(\mathbf{y} - h'(\mathbf{x}))^2] \right) \geq \mathbf{E}_{\mathcal{D}_{\tilde{p}}} [\alpha(\mathbf{v})^2] \geq \mathbf{E}_{\mathcal{D}_{\tilde{p}}} [\alpha(\mathbf{v})]^2 \geq \alpha_1^2.$$

It remains to show that $v + \alpha_v c_v(x) \in \text{Lin}(\mathcal{C}, 2)$. Note that

$$\alpha_v = \mathbf{E}_{\mathcal{D}|_v} [c(\mathbf{x})(\mathbf{y} - v)] \leq \max |c(\mathbf{x})| \max(|y - \tilde{p}(x)|) \leq 1$$

since $c(\mathbf{x}), y - v \in [-1, 1]$. Hence $h'(x) = w_1 \cdot 1 + w_2 c(v)$ where $|w_1| + |w_2| \leq 2$. This contradicts the definition of an $(\ell_2, \text{Lin}(\mathcal{C}, 2), \alpha_3)$ -swap agnostic learner if $\alpha_3 < \alpha_1^2$. ■

Algorithm 2 Swap Agnostic Learning via MCBost

Parameters: loss ℓ , hypothesis class \mathcal{C} , and $\varepsilon > 0$, let $\alpha = \text{poly}(\varepsilon)$

Given: Dataset S sampled from \mathcal{D}

Run:

$\tilde{p} \leftarrow \text{MCBoost}_{\mathcal{C}, \alpha}(S)$

$h_\ell \leftarrow k_\ell \circ \tilde{p}$

Return: h_ℓ

4.2 An algorithm for Swap Agnostic Learning

The equivalence from Theorem 4.3 suggests an immediate strategy for obtaining a swap agnostic learner. First, learn a swap multicalibrated predictor; then, return the predictor, post-processed to an optimal hypothesis according to ℓ . Importantly, MCBost reduces the problem of swap multicalibration (and thus, swap agnostic learning) to a standard agnostic learning task. In all, we can combine the MCBost algorithm for a class \mathcal{C} with a specific loss ℓ to obtain a (ℓ, \mathcal{C}) -swap agnostic learner.

Corollary 4.8 (Informal). *For any (nice) convex loss ℓ , hypothesis class \mathcal{C} , and $\varepsilon > 0$, Algorithm 2 returns a $(\ell, \mathcal{C}, \varepsilon)$ -swap agnostic learner from a sample of $m \leq \text{VC}(\mathcal{C}) \cdot \text{poly}(1/\varepsilon)$ data points drawn from \mathcal{D} , after making $\leq \text{poly}(1/\varepsilon)$ calls to weak agnostic learner for \mathcal{C} .*

5 Characterizing Swap Loss OI via Multicalibration

We show that $(\mathcal{L}, \mathcal{C})$ -swap loss OI and $(\partial\mathcal{L} \circ \mathcal{C})$ -swap multicalibration are equivalent for nice loss functions.

Theorem 5.1. *Let $\mathcal{L} \subseteq \mathcal{L}(B)$ be a family of B -nice loss functions such that $\ell_2 \in \mathcal{L}$. Then $(\partial\mathcal{L} \circ \mathcal{C}, \alpha_1)$ -swap multicalibration and $(\mathcal{L}, \mathcal{C}, \alpha_2)$ -swap loss OI are equivalent.⁸*

In preparation for this, we first give the proof of Lemma 2.12.

Proof of Lemma 2.12. The proof of Part (1) follows the proof of [GHK⁺23, Proposition 4.5], showing that loss OI implies omniprediction. By the definition of k_{ℓ_v} , for every $x \in \mathcal{X}$ such that $\tilde{p}(x) = v$

$$\begin{aligned} \mathbf{E}_{\tilde{\mathbf{y}} \sim \text{Ber}(v)} u_v(x, v, \tilde{\mathbf{y}}) &= \mathbf{E}_{\tilde{\mathbf{y}} \sim \text{Ber}(v)} [\ell_v(\tilde{\mathbf{y}}, k_{\ell_v}(v)) - \ell_v(\tilde{\mathbf{y}}, c_v(x))] \\ &= \ell_v(v, k_{\ell_v}(v)) - \ell_v(v, c_v(x)) \\ &\leq 0 \end{aligned}$$

Hence this also holds in expectation under $\mathcal{D}|_v$, which only considers points where $\tilde{p}(\mathbf{x}) = v$:

$$\mathbf{E}_{\mathcal{D}|_v} [u_v(\mathbf{x}, v, \tilde{\mathbf{y}})] \leq 0.$$

⁸Here equivalence means that there are reductions in either direction that lose a multiplicative factor of $(B+1)$ in the error.

Since \tilde{p} satisfies swap loss OI, we deduce that

$$\mathbf{E}_{\mathcal{D}|v} [u_{\mathbf{v}}(\mathbf{x}, v, \mathbf{y}^*)] \leq \alpha(v)$$

Taking expectations over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$ and using the definition of u_v , we get

$$\begin{aligned} \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k_{\ell_{\mathbf{v}}}(\mathbf{v})) - \ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x}))] &= \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}} [u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*)] \\ &\leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} [\alpha(\mathbf{v})] \leq \alpha \end{aligned}$$

Rearranging the outer inequality gives

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k_{\ell_{\mathbf{v}}}(\mathbf{v}))] \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x}))] + \alpha.$$

Part (2) is implied by taking $\ell_v = \ell$ for every v . ■

Proof of Lemma 2.12. The proof of Part (1) follows the proof of [GHK⁺23, Proposition 4.5], showing that loss OI implies omniprediction. By the definition of k_{ℓ_v} , for every $x \in \mathcal{X}$ such that $\tilde{p}(x) = v$

$$\begin{aligned} \mathbf{E}_{\tilde{\mathbf{y}} \sim \text{Ber}(v)} u_v(x, v, \tilde{\mathbf{y}}) &= \mathbf{E}_{\tilde{\mathbf{y}} \sim \text{Ber}(v)} [\ell_v(\tilde{\mathbf{y}}, k_{\ell_v}(v)) - \ell_v(\tilde{\mathbf{y}}, c_v(x))] \\ &= \ell_v(v, k_{\ell_v}(v)) - \ell_v(v, c_v(x)) \\ &\leq 0 \end{aligned}$$

Hence this also holds in expectation under $\mathcal{D}|v$, which only considers points where $\tilde{p}(\mathbf{x}) = v$:

$$\mathbf{E}_{\mathcal{D}|v} [u_{\mathbf{v}}(\mathbf{x}, v, \tilde{\mathbf{y}})] \leq 0.$$

Since \tilde{p} satisfies swap loss OI, we deduce that

$$\mathbf{E}_{\mathcal{D}|v} [u_{\mathbf{v}}(\mathbf{x}, v, \mathbf{y}^*)] \leq \alpha(v)$$

Taking expectations over $\mathbf{v} \sim \mathcal{D}_{\tilde{p}}$ and using the definition of u_v , we get

$$\begin{aligned} \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k_{\ell_{\mathbf{v}}}(\mathbf{v})) - \ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x}))] &= \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}} [u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*)] \\ &\leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} [\alpha(\mathbf{v})] \leq \alpha \end{aligned}$$

Rearranging the outer inequality gives

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k_{\ell_{\mathbf{v}}}(\mathbf{v}))] \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x}))] + \alpha.$$

Part (2) is implied by taking $\ell_v = \ell$ for every v . ■

We use the following simple claim from [GHK⁺23].

Claim 5.2 (Lemma 4.8, [GHK⁺23]). For random variables $\mathbf{y}_1, \mathbf{y}_2 \in \{0, 1\}$ and $t \in \mathbb{R}$,

$$\mathbf{E}[\ell(\mathbf{y}_1, t) - \ell(\mathbf{y}_2, t)] = \mathbf{E}[(\mathbf{y}_1 - \mathbf{y}_2)\partial\ell(t)]. \quad (22)$$

We record two corollaries of this claim. These can respectively be seen as strengthenings of the two parts of Theorem [GHK⁺23, Theorem 4.9], which respectively characterized hypothesis OI in terms of multiaccuracy and decision OI in terms of calibration. We generalize these to the swap setting.

Corollary 5.3. For every choice of $\{\ell_v, c_v\}_{v \in \text{Im}(\tilde{p})}$, we have

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, c_{\mathbf{v}}(\mathbf{x}))] \right| \right] = \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [(\mathbf{y}^* - \tilde{\mathbf{y}})\partial\ell_{\mathbf{v}} \circ c_{\mathbf{v}}(\mathbf{x})] \right| \right]. \quad (23)$$

Hence if \tilde{p} is $(\partial\mathcal{L} \circ \mathcal{C}, \alpha)$ -swap multicalibrated, then

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, c_{\mathbf{v}}(\mathbf{x}))] \right| \right] \leq \alpha.$$

Proof. Equation (23) is derived by applying Equation (22) to the LHS. Assuming that \tilde{p} is $(\partial\mathcal{L} \circ \mathcal{C}, \alpha)$ -swap multicalibrated, we have

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [(\mathbf{y}^* - \tilde{\mathbf{y}})\partial\ell_{\mathbf{v}} \circ c_{\mathbf{v}}(\mathbf{x})] \right| \right] \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \max_{c' \in \partial\mathcal{L} \circ \mathcal{C}} \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [(\mathbf{y}^* - \tilde{\mathbf{y}})c'(\mathbf{x})] \right| \right] \leq \alpha. \quad \blacksquare$$

Corollary 5.4. Let $\{\ell_v\}_{v \in \text{Im}(f)}$ be a collection of loss B -nice loss functions. Let $k(v) = k_{\ell_v}(v)$. If \tilde{p} is α -calibrated then

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k(\mathbf{v})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k(\mathbf{v}))] \right| \right] \leq B\alpha. \quad (24)$$

Proof. We have

$$\begin{aligned} \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k(\mathbf{v})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k(\mathbf{v}))] \right| \right] &= \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [(\mathbf{y}^* - \mathbf{v})\partial\ell_{\mathbf{v}}(k(\mathbf{v}))] \right| \right] \\ &= \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \partial\ell_{\mathbf{v}}(k(\mathbf{v})) \right| \left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [\mathbf{y}^* - \mathbf{v}] \right| \right] \\ &\leq B \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}_{|\mathbf{v}}} [\mathbf{y}^* - \mathbf{v}] \right| \right] \\ &\leq B\alpha. \end{aligned}$$

where we use the fact that $k(v) \in I_{\ell}$, and so $|\partial\ell_v(k(v))| \leq B$. \blacksquare

Finally, we show the following key technical lemma which explains why the ℓ_2 loss has a special role.

Lemma 5.5. If \tilde{p} is $(\{\ell_2\}, \mathcal{C}, \alpha)$ -swap OI, then it is α -calibrated.

Proof. Observe that $\ell_2(y, v) = (y - v)^2/2$ so $k_{\ell_2}(v) = v$. Hence,

$$\begin{aligned} u_{\ell_2,0}(x, v, y) &= \ell_2(y, k_{\ell_2}(v)) - \ell_2(y, 0) \\ &= ((y - v)^2 - y^2)/2 \\ &= -vy + v^2/2. \end{aligned} \tag{25}$$

Recall that $\{0, 1\} \subset \mathcal{C}$. The implication of swap loss OI when we take $c_v = 0$ for all v is that

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [u_{\ell_2,0}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*) - u_{\ell_2,0}(\mathbf{x}, \mathbf{v}, \tilde{\mathbf{y}})] \right| \right] \leq \alpha.$$

We can simplify the LHS using Equation (25) to derive

$$\begin{aligned} \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [(-\mathbf{v}\mathbf{y}^* + \mathbf{v}^2/2) - (-\mathbf{v}\tilde{\mathbf{y}} + \mathbf{v}^2/2)] \right| \right] &= \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\mathbf{v}(\mathbf{y}^* - \tilde{\mathbf{y}})] \right| \right] \\ &= \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\mathbf{v} \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\tilde{\mathbf{y}} - \mathbf{y}^*] \right| \right] \leq \alpha. \end{aligned} \tag{26}$$

Considering the case where $c_v = 1$ for all v gives

$$\begin{aligned} u_{\ell_2,1}(x, v, y) &= \ell_2(y, k_{\ell_2}(v)) - \ell_2(y, 1) \\ &= ((y - v)^2 - (1 - y)^2)/2 \\ &= (1 - v)y + (v^2 - 1)/2. \end{aligned}$$

We derive the following implication of swap loss OI by taking $c_v = 0$ for all v :

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [u_{\ell_2,1}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*) - u_{\ell_2,1}(\mathbf{x}, \mathbf{v}, \tilde{\mathbf{y}})] \right| \right] = \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[(1 - \mathbf{v}) \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\tilde{\mathbf{y}} - \mathbf{y}^*] \right| \right] \leq \alpha \tag{27}$$

Adding the bounds from Equations (26) and (27) we get

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\mathbf{v} - \mathbf{y}^*] \right| \right] = \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\tilde{\mathbf{y}} - \mathbf{y}^*] \right| \right] \leq \alpha$$

■

We can now complete the proof of Theorem 5.1.

Proof of Theorem 5.1. We first show the forward implication, that swap multicalibration implies swap loss OI.

Since $\ell_2 \in \mathcal{L}$ and $1 \in \mathcal{C}$, we have $\partial \ell_2 \circ 1 = 1 \in \partial \mathcal{L} \circ \mathcal{C}$. This implies that \tilde{p} is α -multicalibrated, since

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [1(\mathbf{y} - \mathbf{v})] \right| \right] \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\max_{c \in \mathcal{C}} \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [c(\mathbf{x})(\mathbf{y} - \mathbf{v})] \right| \right] \leq \alpha.$$

Consider any collection of losses $\{\ell_v\}_{v \in \text{Im}(\tilde{p})}$. Applying Corollary 5.4, we have

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k(\mathbf{v})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k(\mathbf{v}))] \right| \right] \leq B\alpha.$$

On the other hand, by Corollary 5.3, we have for every choice of $\{\ell_v, c_v\}_{v \in \text{Im}(\tilde{p})}$,

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, c_{\mathbf{v}}(\mathbf{x}))] \right| \right] \leq \alpha.$$

Hence for any choice of $\{u_v\}_{v \in \text{Im}(\tilde{p})}$ we can bound

$$\begin{aligned} & \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*) - u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \tilde{\mathbf{y}})] \right| \\ & \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k(\mathbf{v})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k(\mathbf{v}))] \right| + \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c_{\mathbf{v}}(\mathbf{x})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, c_{\mathbf{v}}(\mathbf{x}))] \right| \right] \\ & \leq (B+1)\alpha \end{aligned}$$

which shows that \tilde{p} satisfies swap loss OI with $\alpha_2 = (D+1)\alpha_1$.

Next we show the reverse implication: if \tilde{p} satisfies $(\mathcal{L}, \mathcal{C}, \alpha_2)$ -swap loss OI, then it satisfies $(\partial\mathcal{L} \circ \mathcal{C}, \alpha_1)$ -swap multicalibration. The first step is to observe that by lemma 5.5, since $\ell_2 \in \mathcal{L}$, the predictor \tilde{p} is α_2 calibrated. Since any $\ell \in \mathcal{L}$ is B -nice, we have

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k(\mathbf{v})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k(\mathbf{v}))] \right| \right] = \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [(\mathbf{y}^* - \tilde{\mathbf{y}})k(\mathbf{v})] \right| \right] \leq B\alpha_2.$$

For any $\{\ell_v, c_v\}_{v \in \text{Im}(f)}$, since

$$u_v(x, v, y) = \ell_v(y, k_\ell(v)) + \ell_v(y, c_v(x))$$

we can write

$$\begin{aligned} & \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c(\mathbf{x})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, c(\mathbf{x}))] \right| \right] \\ & \leq \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \mathbf{y}^*) - u_{\mathbf{v}}(\mathbf{x}, \mathbf{v}, \tilde{\mathbf{y}})] \right| + \left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, k(\mathbf{v})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, k(\mathbf{v}))] \right| \right] \\ & \leq (B+1)\alpha_2. \end{aligned}$$

But by Equation (23), the LHS can be written as

$$\mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\ell_{\mathbf{v}}(\mathbf{y}^*, c(\mathbf{x})) - \ell_{\mathbf{v}}(\tilde{\mathbf{y}}, c(\mathbf{x}))] \right| \right] = \mathbf{E}_{\mathbf{v} \sim \mathcal{D}_{\tilde{p}}} \left[\left| \mathbf{E}_{\mathcal{D}|\mathbf{v}} [\partial\ell_{\mathbf{v}} \circ c_{\mathbf{v}}(\mathbf{x})(\mathbf{y}^* - \mathbf{v})] \right| \right]$$

This shows that \tilde{p} is $(\partial\mathcal{L} \circ \mathcal{C}, (B+1)\alpha_2)$ -swap multicalibrated. ■

6 Relating Notions of Omniprediction

In this work, we have discussed the four different notions of omniprediction defined to date.

- 00) Omniprediction, as originally defined by [GKR⁺22].
- 01) Loss OI, from [GHK⁺23].
- 10) Swap omniprediction.

11) Swap Loss OI.

In order to compare them, we can ask which of these notions implies the other for any fixed choice of loss class \mathcal{L} and hypothesis class \mathcal{C} .

- Loss OI implies omniprediction by [GHK⁺23, Proposition 4.5].
- Swap omniprediction implies omniprediction by Claim 2.5.
- Swap loss OI implies both loss OI and swap multicalibration by Lemma 2.12.

These relationships were summarized in Figure 1. Further, this picture captures all the implications that hold for all $(\mathcal{L}, \mathcal{C})$. Next, we show that for any implication not drawn in the diagram, there exists some (natural) choice of $(\mathcal{L}, \mathcal{C})$, where the implication does not hold. In particular, we prove that neither loss OI nor swap omniprediction implies the other for all $(\mathcal{L}, \mathcal{C})$. This separates these notions from swap loss OI, since swap loss OI implies both these notions.⁹ By similar reasoning, it separates omniprediction from both these loss OI and swap omniprediction, since omniprediction is implied by either of them.

Swap omniprediction does not imply loss OI. We prove this non-implication using a counterexample used in [GHK⁺23]. In particular, they show that omniprediction does not imply loss OI [GHK⁺23, Theorem 4.6], and the same example in fact shows that swap omniprediction does not imply loss OI. In their example, we have \mathcal{D} on $\{\pm 1\}^3 \times [0, 1]$ where the marginal on $\{\pm 1\}^3$ is uniform, and $p^*(x) = (1 + x_1x_2x_3)/2$, whereas $\tilde{p}(x) = 1/2$ for all x . We take $\mathcal{C} = \{1, x_1, x_2, x_3\}$. Since $\tilde{p} = 1/2$ is constant, it is easy to check that $\tilde{p} - p^* = -x_1x_2x_3/2$ is uncorrelated with \mathcal{C} . Hence \tilde{p} satisfies swap multicalibration (which is the same as multicalibration or even multiaccuracy in this setting where \tilde{p} is constant). Hence by Theorem 4.3, \tilde{p} is an $(\mathcal{L}_{\text{cvx}}(1), \text{Lin}_{\mathcal{C}}, 0)$ -swap omnipredictor. [GHK⁺23, Theorem 4.6] prove that \tilde{p} is not loss OI for the ℓ_4 loss. Hence we have the following result.

Lemma 6.1. *The predictor \tilde{p} is $(\mathcal{C}, 0)$ -swap multicalibrated and hence it is a $(\{\ell_4\}, \text{Lin}(\mathcal{C}), 0)$ -swap omnipredictor. But it is not $(\{\ell_4\}, \text{Lin}(\mathcal{C}, 1), \varepsilon)$ -loss OI for $\varepsilon < 4/9$.*

We remark that the construction extends to all ℓ_p losses for even $p > 2$. Hence even for convex losses, the notions of swap omniprediction and loss-OI seem incomparable.

Loss OI does not imply swap omniprediction. Next we construct an example showing that loss OI need not imply swap omniprediction. We consider the set of all GLM losses defined below, which contain common losses including the squared loss and the logistic loss.

Definition 6.2. *Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a convex, differentiable function such that $[0, 1] \subseteq \text{Im}(g')$. Define its matching loss to be $\ell_g = g(t) - yt$. Define $\mathcal{L}_{\text{GLM}} = \{\ell_g\}$ be the set of all such loss functions.*

[GHK⁺23] shows a general decomposition result that reduces achieving loss OI to a calibration condition and a multiaccuracy condition. Whereas arbitrary losses might require multiaccuracy for the more powerful class $\partial\mathcal{L} \circ \mathcal{C}$, for \mathcal{L}_{GLM} we have $\partial\mathcal{L}_{\text{GLM}} \circ \mathcal{C} = \mathcal{C}$. This is formalized in the following result.

⁹For instance if loss OI implied swap loss OI, it would also imply swap omniprediction, which our claim shows it does not.

$x = (x_1, x_2)$	$p^*(x)$	$\tilde{p}(x)$
$(-1, -1)$	0	$\frac{1}{8}$
$(+1, -1)$	$\frac{1}{4}$	$\frac{1}{8}$
$(-1, +1)$	1	$\frac{7}{8}$
$(+1, +1)$	$\frac{3}{4}$	$\frac{7}{8}$

Table 2: Separating loss-OI and swap-resilient omniprediction

Lemma 6.3 (Theorem 5.3, [GHK⁺23]). *If \tilde{p} is ε_1 -calibrated and $(\mathcal{C}, \varepsilon_2)$ -multiaccurate, then it is $(\mathcal{L}_{\text{GLM}}, \text{Lin}(\mathcal{C}, W), \delta)$ -loss OI for $\delta = \varepsilon_1 + W\varepsilon_2$.*

In light of the above result, it suffices to find a predictor that is calibrated and multiaccurate (and hence satisfies loss OI), but not multicalibrated, hence not swap multicalibrated. By Theorem 4.3 it is not an $(\{\ell_2\}, \text{Lin}_{\mathcal{C}}, \delta)$ -swap omnipredictor for δ less than some constant.

We show a separation between loss OI and swap omniprediction using the predictors $p^*, \tilde{p} : \{\pm 1\}^2 \rightarrow [0, 1]$ defined in Table 6.

Lemma 6.4. *Consider the distribution \mathcal{D} on $\{\pm 1\}^2 \times \{0, 1\}$ where the marginal on $\{\pm 1\}^2$ is uniform and $\mathbf{E}[y|x] = p^*(x)$. Let $\mathcal{C} = \{1, x_1, x_2\}$.*

1. $\tilde{p} \in \text{Lin}(\mathcal{C}, 1)$. Moreover, it minimizes the squared error over all hypotheses from $\text{Lin}(\mathcal{C})$.
2. \tilde{p} is perfectly calibrated and $(\mathcal{C}, 0)$ -multiaccurate. So it is $(\mathcal{L}_{\text{GLM}}, \text{Lin}(\mathcal{C}), 0)$ -loss OI.
3. \tilde{p} is not (\mathcal{C}, α) -multicalibrated for $\alpha < 1/8$. It is not $(\ell_2, \text{Lin}(\mathcal{C}), \delta)$ -swap agnostic learner for $\delta < 1/64$.

Proof. We compute Fourier expansions for the two predictors:

$$p^*(x) = \frac{1}{8}(4 + 3x_2 - x_1x_2) \quad (28)$$

$$\tilde{p}(x) = \frac{1}{8}(4 + 3x_2) \quad (29)$$

This shows that $\tilde{p} \in \text{Lin}(\mathcal{C})$, and moreover that it is the optimal approximation to p^* in $\text{Lin}(\mathcal{C})$, as it is the projection of p^* onto $\text{Lin}(\mathcal{C})$. This shows that \tilde{p} is an $(\ell_2, \text{Lin}(\mathcal{C}), 0)$ -agnostic learner.

It is easy to check that \tilde{p} is perfectly calibrated. It is $(\mathcal{C}, 0)$ -multiaccurate, since it is the projection of p^* onto $\text{Lin}(\mathcal{C})$, so $\tilde{p} - p^*$ is orthogonal to $\text{Lin}(\mathcal{C})$. Hence we can apply Lemma 6.3 to conclude that it is $(\mathcal{L}_{\text{GLM}}, \text{Lin}(\mathcal{C}), 0)$ -loss OI, where \mathcal{L}_{GLM} which contains the squared loss.

To show that \tilde{p} is not swap-agnostic, we observe that conditioning on the value of $\tilde{p}(x) = (4 + 3x_2)/8$ is equivalent to conditioning on $x_2 \in \{\pm 1\}$. For each value of x_2 , the restriction of p^* which is now linear in x_1 belongs to $\text{Lin}(\mathcal{C})$. Indeed if we condition on $\tilde{p}(x) = 1/8$ so that $x_2 = -1$, we have

$$p^*(x) = \frac{1}{2} - \frac{3}{8} + \frac{1}{8}x_1 = \frac{1 + x_1}{8}.$$

Conditioned on $\tilde{p}(x) = 7/8$ so that $x_2 = 1$, we have

$$p^*(x) = \frac{1}{2} + \frac{3}{8} - \frac{1}{8}x_1 = \frac{7 - x_1}{8}.$$

Hence we have

$$\mathbf{E}_{v \sim \mathcal{D}_{\tilde{p}}} \left[\left[\min_{h \in \text{Lin}(\mathcal{C})} \mathbf{E}[(\mathbf{y} - h(\mathbf{x}))^2 | f(\mathbf{x}) = v] \right] \right] = \mathbf{E}[(y - p^*(x))^2] = \text{Var}[\mathbf{y}],$$

whereas the variance decomposition of squared loss gives

$$\begin{aligned} \mathbf{E}[(\mathbf{y} - \tilde{p}(\mathbf{x}))^2] &= \mathbf{E}[(\mathbf{y} - p^*(\mathbf{x}))^2] + \mathbf{E}[(p^*(\mathbf{x}) - \tilde{p}(\mathbf{x}))^2] \\ &= \text{Var}[\mathbf{y}] + \frac{1}{64} \mathbf{E}[(x_1 x_2)^2] \\ &= \text{Var}[\mathbf{y}] + \frac{1}{64}. \end{aligned}$$

Hence \tilde{p} is not a $(\ell_2, \text{Lin}(\mathcal{C}), \delta)$ -swap agnostic learner for $\delta < 1/64$.

To see that f is not multicalibrated for small α , observe that conditioned on $x_2 \in \{\pm 1\}$, the correlation between x_1 and $\tilde{p} - p^*$ is $1/8$. \blacksquare

Note that item (1) above separates swap omniprediction from omniprediction and agnostic learning. This separation can also be derived from [GKR⁺22, Theorem 7.5] which separated (standard) omniprediction from agnostic learning, since swap omniprediction implies standard omniprediction.

Comparing notions for GLM losses. When we restrict our attention to \mathcal{L}_{GLM} , in fact, the notions of swap loss OI and swap omniprediction are equivalent. The key observation here is that $\partial \mathcal{L}_{\text{GLM}} \circ \mathcal{C} = \mathcal{C}$, as shown in [GHK⁺23]. Paired with Theorem 4.3 and Theorem 5.1, we obtain the following collapse.

Claim 6.5. *The notions of $(\mathcal{L}_{\text{GLM}}, \mathcal{C}, \alpha_1)$ -swap loss OI and $(\mathcal{L}_{\text{GLM}}, \mathcal{C}, \alpha_2)$ -swap omniprediction are equivalent.*

Proof. To see this, note that by Theorem 5.1, $(\mathcal{L}_{\text{GLM}}, \mathcal{C}, \alpha_1)$ -swap loss OI is equivalent to $(\partial \mathcal{L}_{\text{GLM}} \circ \mathcal{C}, \alpha'_1)$ -swap multicalibration. We know from Theorem 4.3 that this is also equivalent to $(\partial \mathcal{L}_{\text{GLM}} \circ \mathcal{C}, \alpha_2)$ -swap omniprediction. So, by the fact that $\partial \mathcal{L}_{\text{GLM}} \circ \mathcal{C} = \mathcal{C}$, we have the claimed equivalence. \blacksquare

Finally, we know that loss OI implies omniprediction for \mathcal{L}_{GLM} , since this holds true for all \mathcal{L} . We do not know if these notions are equivalent for \mathcal{L}_{GLM} , since the construction in Lemma 6.1 used the ℓ_4 loss which does not belong to \mathcal{L}_{GLM} .

Acknowledgements. We thank Sumegha Garg, Christopher Jung, Salil Vadhan, and Udi Wieder for their helpful comments on this work.

References

- [BL20] Avrim Blum and Thodoris Lykouris. Advancing Subgroup Fairness via Sleeping Experts. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151, pages 55:1–55:24, 2020.
- [BLM01] Shai Ben-David, Philip M. Long, and Yishay Mansour. Agnostic boosting. In *14th Annual Conference on Computational Learning Theory, COLT*, 2001.
- [BM07a] Avrim Blum and Yishay Mansour. From external to internal regret. *J. Mach. Learn. Res.*, 8:1307–1324, 2007.
- [BM07b] Avrim Blum and Yishay Mansour. Learning, regret minimization, and equilibria. In *Algorithmic Game Theory, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani editors*. Cambridge University Press, 2007.
- [CBL06] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [DDZ23] Zhun Deng, Cynthia Dwork, and Linjun Zhang. Happymap : A generalized multi-calibration method. In *14th Innovations in Theoretical Computer Science Conference, ITCS 2023*, volume 251 of *LIPICs*, pages 41:1–41:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [DKR⁺21] Cynthia Dwork, Michael P. Kim, Omer Reingold, Guy N. Rothblum, and Gal Yona. Outcome indistinguishability. In *ACM Symposium on Theory of Computing (STOC'21)*, 2021.
- [DKR⁺22] Cynthia Dwork, Michael P. Kim, Omer Reingold, Guy N. Rothblum, and Gal Yona. Beyond bernoulli: Generating random outcomes that cannot be distinguished from nature. In *The 33rd International Conference on Algorithmic Learning Theory*, 2022.
- [DLLT23] Cynthia Dwork, Daniel Lee, Huijia Lin, and Pranay Tankala. New insights into multi-calibration, 2023.
- [FV98] Dean P. Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- [FV99] Dean Foster and Rakesh Vohra. Regret in the on-line decision problem. *Games and Economic Behavior*, 29:7–35, 1999.
- [GHHK⁺23] Ira Globus-Harris, Declan Harrison, Michael Kearns, Aaron Roth, and Jessica Sorrell. Multicalibration as boosting for regression, 2023.
- [GHK⁺23] Parikshit Gopalan, Lunjia Hu, Michael P. Kim, Omer Reingold, and Udi Wieder. Loss minimization through the lens of outcome indistinguishability. In *Innovations in theoretical computer science (ITCS'23)*, 2023.

- [GJKR18] Stephen Gillen, Christopher Jung, Michael J. Kearns, and Aaron Roth. Online learning with an unknown fairness metric. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 2605–2614, 2018.
- [GJN⁺22] Varun Gupta, Christopher Jung, Georgy Noarov, Mallesh M. Pai, and Aaron Roth. Online multivalid learning: Means, moments, and prediction intervals. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022*, volume 215 of *LIPICs*, pages 82:1–82:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [GKR⁺22] Parikshit Gopalan, Adam Tauman Kalai, Omer Reingold, Vatsal Sharan, and Udi Wieder. Omnipredictors. In *Innovations in Theoretical Computer Science (ITCS’2022)*, 2022.
- [GKR23] Parikshit Gopalan, Michael P. Kim, and Omer Reingold. Swap agnostic learning, or characterizing omniprediction via multicalibration. In *NeurIPS*, 2023.
- [GKSZ22] Parikshit Gopalan, Michael P. Kim, Mihir Singhal, and Shengjia Zhao. Low-degree multicalibration. In *Conference on Learning Theory, 2-5 July 2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pages 3193–3234. PMLR, 2022.
- [GRSW22] Parikshit Gopalan, Omer Reingold, Vatsal Sharan, and Udi Wieder. Multicalibrated partitions for importance weights. In *International Conference on Algorithmic Learning Theory, 29-1 April 2022, Paris, France*, volume 167 of *Proceedings of Machine Learning Research*, pages 408–435. PMLR, 2022.
- [Hau92] David Haussler. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and Computation*, 100(1):78–150, 1992.
- [HKRR18] Úrsula Hébert-Johnson, Michael P. Kim, Omer Reingold, and Guy N. Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning, ICML, 2018*.
- [JLP⁺21] Christopher Jung, Changhwa Lee, Mallesh M. Pai, Aaron Roth, and Rakesh Vohra. Moment multicalibration for uncertainty estimation. In *Conference on Learning Theory, COLT 2021, 15-19 August 2021, Boulder, Colorado, USA*, volume 134 of *Proceedings of Machine Learning Research*, pages 2634–2678. PMLR, 2021.
- [KGZ19] Michael P. Kim, Amirata Ghorbani, and James Zou. Multiaccuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 247–254, 2019.
- [KKG⁺22] Michael P Kim, Christoph Kern, Shafi Goldwasser, Frauke Kreuter, and Omer Reingold. Universal adaptability: Target-independent inference that competes with propensity scoring. *Proceedings of the National Academy of Sciences*, 119(4), 2022.
- [KMR17] Jon M. Kleinberg, Sendhil Mullainathan, and Manish Raghavan. Inherent trade-offs in the fair determination of risk scores. In *8th Innovations in Theoretical Computer Science Conference, ITCS, 2017*.

- [KNRW17] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. *arXiv preprint arXiv:1711.05144*, 2017.
- [KV94] Michael J Kearns and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994.
- [Pie12] Krzysztof Pietrzak. Cryptography from learning parity with noise. In *SOFSEM*, volume 12, pages 99–114. Springer, 2012.
- [Val84] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.