
No-Regret and Incentive-Compatible Online Learning

Rupert Freeman¹ David M. Pennock² Chara Podimata³ Jennifer Wortman Vaughan⁴

Abstract

We study online learning settings in which experts act strategically to maximize their influence on the learning algorithm’s predictions by potentially misreporting their beliefs about a sequence of binary events. Our goal is twofold. First, we want the learning algorithm to be no-regret with respect to the best fixed expert in hindsight. Second, we want incentive compatibility, a guarantee that each expert’s best strategy is to report his true beliefs about the realization of each event. To achieve this goal, we build on the literature on wagering mechanisms, a type of multi-agent scoring rule. We provide algorithms that achieve no regret and incentive compatibility for myopic experts for both the full and partial information settings. In experiments on datasets from FiveThirtyEight, our algorithms have regret comparable to classic no-regret algorithms, which are not incentive-compatible. Finally, we identify an incentive-compatible algorithm for forward-looking strategic agents that exhibits diminishing regret in practice.

1. Introduction

We study an online learning setting in which a learner makes predictions about a sequence of T binary events (Vovk, 1990; Littlestone & Warmuth, 1994; Cesa-Bianchi et al., 1997; Freund & Schapire, 1997; Vovk, 1998; Auer et al., 2002). The learner has access to a pool of K experts, each with beliefs about the likelihood of each event occurring. The standard goal of the learner is to output a sequence of predictions almost as accurate as those of the best fixed expert in hindsight. Such a learner is said to have no regret.

¹Darden School of Business, University of Virginia. ²DIMACS, Rutgers University. ³Harvard University. ⁴Microsoft Research NYC. Correspondence to: Rupert Freeman <freemanr@darden.virginia.edu>, David M. Pennock <dpennock@dimacs.rutgers.edu>, Chara Podimata <podimata@g.harvard.edu>, Jennifer Wortman Vaughan <jenn@microsoft.com>.

But what if the experts that the learner consults are strategic agents, capable of reporting predictions that do not represent their true beliefs? As pointed out by Roughgarden & Shriververs (2017), when the learner is not only making predictions but also (implicitly or explicitly) evaluating the experts, experts might have incentive to misreport. The Good Judgment Project,¹ a competitor in IARPA’s Aggregative Contingent Estimation geopolitical forecasting contest, scored individual forecasters and rewarded the top 2%—dubbed “Superforecasters” (Tetlock & Gardner, 2015)—with perks such as paid conference travel; some are now employed by a spinoff company. Similarly, the website FiveThirtyEight² not only predicts election results by aggregating different pollsters, but also publicly scores the pollsters, in a way that correlates with the amount of influence that the pollsters have over the FiveThirtyEight aggregate. It is natural to expect that forecasters might respond to the competitive incentive structure in these settings by seeking to maximize the influence that they exert on the learner’s prediction.

When an online learning algorithm is designed in such a way that experts are motivated to report their true beliefs, we say it is *incentive-compatible*. Incentive compatibility is desirable for several reasons. First, when experts do not report truthfully, the learner’s prediction may be harmed. Second, learning algorithms that fail incentive compatibility place an additional layer of cognitive burden on the experts, who must now reason about the details of the algorithm and other experts’ reports and beliefs in order to decide how to act optimally. To our knowledge, the standard multiplicative-weights-type algorithms fail incentive compatibility, in the sense that experts can sometimes achieve a greater influence on the algorithm’s prediction by misreporting their beliefs; we illustrate this explicitly through manipulation examples. Our goal in this work is to design incentive-compatible online learning algorithms without compromising on the quality of the algorithm’s predictions. That is, we seek algorithms that are both incentive-compatible and no-regret, for both the full and partial (bandit) information settings.

Towards this goal, we show a novel connection between online learning and *wagering mechanisms* (Lambert et al., 2008; 2015), a type of multi-agent scoring rule that allows

¹<https://goodjudgment.com>

²<https://fivethirtyeight.com/>

a principal to elicit the beliefs of a group of agents without taking on financial risk. Using this connection, we construct online learning algorithms that are incentive-compatible and incur sublinear regret. For the full information setting, we introduce Weighted-Score Update (WSU), which yields regret $O(\sqrt{T \ln K})$, matching the optimal regret achievable for general loss functions, even without incentive guarantees. For the partial information setting, we introduce Weighted-Score Update with Uniform Exploration (WSU-UX), which achieves regret $O(T^{2/3}(K \ln K)^{1/3})$.

We focus primarily on experts that strategize only about their influence at the next timestep. However, we obtain a partial extension for forward-looking experts. Building on a mechanism that was proposed for forecasting competitions (Witkowski et al., 2018), we identify an algorithm, ELF-X, for the full information setting that is incentive-compatible and achieves diminishing regret in simulations.

Our theoretical results are supported by experiments on data gathered from an online prediction contest on FiveThirtyEight. Our algorithms achieve regret almost identical to the classic (and not incentive-compatible) Multiplicative Weights Update (MWU) (Freund & Schapire, 1997) and EXP3 (Auer et al., 2002) algorithms in the full and partial information settings respectively, though WSU falls short of the optimal regret achieved by Hedge for quadratic loss.

Related Work. Other work has drawn connections between online learning and incentive-compatible forecasting, particularly in the context of prediction markets (Abernethy et al., 2013; Abernethy & Frongillo, 2011; Frongillo et al., 2012; Hu & Storkey, 2014). Our work is most closely related to that of Roughgarden & Schrijvers (2017), but differs from theirs in several important ways. Most crucially, Roughgarden and Schrijvers consider algorithms that maintain *unnormalized* weights over the experts, and they assume that an expert’s incentives are only affected by these weights. In our work, incentives are tied to the expert’s *normalized* weight—that is, his probability of being selected by the learning algorithm. We argue that normalized weights better reflect experts’ incentives in reality, since reputation tends to be relative more than absolute; put another way, doubling the unnormalized weight of every expert should not increase an expert’s utility, since his influence over the learner’s prediction remains the same. Under Roughgarden and Schrijvers’ model, the design problem is fairly simple when the loss function is a proper loss (Reid & Williamson, 2009)—that is, one that can be elicited by a proper scoring rule (Savage, 1971; Gneiting & Raftery, 2007), such as the quadratic loss function—and can be solved with a multiplicative weights algorithm. Because of this, they focus primarily on the absolute loss function, which is not a proper loss. In contrast, in our model, the design problem is nontrivial even for these “easier” proper loss functions.

Conceptually, our work builds on work by Witkowski et al. (2018), who use competitive scoring rules—a subclass of wagering mechanisms—to design incentive-compatible forecasting competitions. We discuss their work further in Section 5. Our work also has connections with the work of Orabona & Pál (2016), who introduce a class of *coin-betting* algorithms for online learning. Although Orabona & Pál (2016) do not address incentives and do not make a connection with the wagering mechanisms literature, our WSU algorithm can be interpreted as a coin-betting algorithm.³

2. Model and Preliminaries

We consider a setting in which a learner interacts with a set of K experts, each making probabilistic predictions about a sequence of T binary outcomes.⁴ At each round $t \in [T]$, each expert $i \in [K]$ has a private belief $b_{i,t} \in [0, 1]$, unknown to the learner, about the outcome for that round. Both the experts’ beliefs and the sequence of outcomes may be chosen arbitrarily, and potentially adversarially.

In the full information setting, each expert reports his prediction $p_{i,t} \in [0, 1]$ to the learner. The learner then chooses her own prediction $\bar{p}_t \in [0, 1]$ and observes the outcome realization $r_t \in \{0, 1\}$. Finally, the learner and the experts incur losses $\ell_t = \ell(\bar{p}_t, r_t)$ and $\ell_{i,t} = \ell(p_{i,t}, r_t)$, $\forall i \in [K]$, where $\ell : [0, 1] \times \{0, 1\} \rightarrow [0, 1]$ is a bounded loss function.⁵ As is common in the literature, we restrict our attention to algorithms in which the learner maintains a timestep-specific probability distribution $\pi_t = (\pi_{1,t}, \dots, \pi_{K,t})$ over the experts, and chooses her prediction \bar{p}_t according to this distribution. Unless specified, this means that the learner predicts $\bar{p}_t = p_{i,t}$ with probability $\pi_{i,t}$; some of our results additionally apply when $\bar{p}_t = \sum_{i \in [K]} \pi_{i,t} p_{i,t}$.

Under partial information, the protocol remains the same except that the learner is explicitly restricted to choosing a single expert I_t on each round t (according to distribution π_t) and does not observe the predictions of other experts.

The goal of the learner is twofold. First, she wishes to incur a total loss that is not too much worse than the loss of the best fixed expert in hindsight. This is captured using the classic notion of *regret*, given by

$$R = \mathbb{E} \left[\sum_{t \in [T]} \ell_t - \min_{i \in [K]} \sum_{t \in [T]} \ell_{i,t} \right],$$

³In the language of coin betting, in WSU experts wager an η fraction of their wealth on the positive realization of the event, and the actual outcome of each coin flip is the expert’s loss minus the weighted average loss of all other experts.

⁴We focus on binary outcomes to simplify the presentation of our results, but our techniques could be applied more broadly.

⁵The loss function taking values in $[0, 1]$ is without loss of generality since any bounded loss function could be scaled.

where the expectation is taken with respect to randomness in the learner's choice of \bar{p}_t .

No-regret algorithms have been proposed in both the full and partial information settings. Many, such as Hedge (Freund & Schapire, 1997) and MWU (Arora et al., 2012), achieve regret of $O(\sqrt{T \ln K})$ for general loss functions by maintaining unnormalized weights $w_{i,t}$ for each expert i that are updated multiplicatively at each timestep. Hedge uses the update rule $w_{i,t+1} = w_{i,t} \exp(-\eta \ell_{i,t})$, while MWU uses $w_{i,t+1} = w_{i,t}(1 - \eta \ell_{i,t})$ for appropriately chosen values of η . These weights are then normalized to arrive at the distribution π_t . For the case of exp-concave loss functions, such as the quadratic loss, Kivinen & Warmuth (1999) showed that by aggregating experts' predictions and tuning η appropriately, Hedge can achieve regret $O(\ln K)$.

For the partial information setting, the EXP3 algorithm of Auer et al. (2002) achieves a regret of $O(\sqrt{TK \ln K})$. EXP3 maintains a set of expert weights similar to those of Hedge. However, since the learner can only observe the prediction of the chosen expert, she uses an unbiased estimator $\hat{\ell}_{i,t}$ of each expert i 's loss in her updates in place of $\ell_{i,t}$. The update rule then becomes $w_{i,t+1} = w_{i,t} \exp(-\eta \hat{\ell}_{i,t})$.

The second goal of the learner is to incentivize experts to truthfully report their private beliefs. In our model, at each timestep t , each expert i chooses his report $p_{i,t}$ strategically to maximize the probability $\pi_{i,t+1}$ that he is chosen at timestep $t+1$. An algorithm is *incentive-compatible* if experts maximize this probability by reporting $p_{i,t} = b_{i,t}$, irrespective of the reports of the other experts.

Definition 2.1 (Incentive Compatibility). *An online learning algorithm is incentive-compatible if for every timestep $t \in [T]$, every expert i with belief $b_{i,t}$, every report $p_{i,t}$, every vector of reports of the other experts $\mathbf{p}_{-i,t}$, and every history of reports $(\mathbf{p}_{t'})_{t' < t}$ and outcomes $(r_{t'})_{t' < t}$,*

$$\begin{aligned} & \mathbb{E}_{r_t \sim \text{Bern}(b_{i,t})} [\pi_{i,t+1} | (b_{i,t}, \mathbf{p}_{-i,t}), r_t, (r_{t'})_{t' < t}, (\mathbf{p}_{t'})_{t' < t}] \\ & \geq \mathbb{E}_{r_t \sim \text{Bern}(b_{i,t})} [\pi_{i,t+1} | (p_{i,t}, \mathbf{p}_{-i,t}), r_t, (r_{t'})_{t' < t}, (\mathbf{p}_{t'})_{t' < t}]. \end{aligned}$$

where by $r \sim \text{Bern}(b)$ we denote a random variable r taking value 1 with probability b and 0 otherwise.

Incentive compatibility guarantees that any regret bounds apply not only with respect to the reports of the experts, but also with respect to their beliefs. This notion of regret is often called *strategic regret*, and in general may be higher or lower than standard regret. For an incentive-compatible algorithm, the two notions coincide.

To achieve incentive compatibility, we restrict attention to proper loss functions (Reid & Williamson, 2009), referred to in the forecasting literature as proper scoring rules (McCarthy, 1956; Savage, 1971; Gneiting & Raftery, 2007).

Definition 2.2. *A loss function ℓ is said to be proper if*

$$\mathbb{E}_{r \sim \text{Bern}(b)} [\ell(p, r)] \geq \mathbb{E}_{r \sim \text{Bern}(b)} [\ell(b, r)], \forall p \neq b.$$

Restricting attention to proper loss functions, we are guaranteed that an expert who cares only about his expected loss would truthfully report his beliefs. However, this does not apply for experts who care about their probability of being chosen by the learner, as in our setting. Indeed, known online learning algorithms fail to be incentive-compatible even for proper loss functions. We illustrate this in the following example for MWU with the (proper) quadratic loss function $\ell(p, r) = (p - r)^2$. Here the normalization of weights by the factor $\sum_{j \in [K]} w_{j,t}$, which depends on both $p_{i,t}$ and r_t , can create incentives for agent i to deviate. We note that a similar counterexample can be proved for Gradient Descent too, and we include it in Appendix A.

Example 2.1. *Let $\ell(p, r) = (p - r)^2$. Under standard initialization for MWU, $w_{i,1} = 1$ for all $i \in [K]$. Suppose that $b_{1,1} = 0.5$ and $p_{i,1} = 0$ for all $i \in \{2, \dots, K\}$. Then $\mathbb{E}[\pi_{1,2}]$, the expected probability that expert 1 is chosen at time 2 under MWU with respect to his own beliefs, is*

$$0.5 \left(\frac{1 - \eta(1 - p_{1,1})^2}{K - \eta(1 - p_{1,1})^2 - \eta(K-1)} \right) + 0.5 \left(\frac{1 - \eta p_{1,1}^2}{K - \eta p_{1,1}^2} \right).$$

For $K \geq 3$ and $T \geq 9 \ln(3)$, the denominator in the first term is less than the denominator in the second term, independent of $p_{1,1}$. The derivative of $\mathbb{E}[\pi_{1,2}]$ with respect to $p_{1,1}$ is therefore strictly positive at 0.5, implying that expert 1 maximizes his utility by reporting some $p_{1,1} > 0.5$.

Thus, unlike in the setting of Roughgarden & Schrijvers (2017), using a proper loss function with a standard algorithm is not enough, and new algorithmic ideas are needed. To derive our algorithms, we draw a connection between online learning and *wagering mechanisms*, one-shot elicitation mechanisms that allow experts to bet on the quality of their predictions relative to others. In the one-shot wagering setting introduced by Lambert et al. (2008), each agent $i \in [K]$ holds a belief $b_i \in [0, 1]$ about the likelihood of an event. Agent i reports a probability p_i and a wager $w_i \geq 0$. A wagering mechanism, Γ , maps the reports $\mathbf{p} = (p_1, \dots, p_K)$, wagers $\mathbf{w} = (w_1, \dots, w_K)$, and the realization r of the binary event to payments $\Gamma_i(\mathbf{p}, \mathbf{w}, r)$ for each agent i . The purpose of the wager is to allow each agent to set a maximum allowable loss, which is captured by imposing the constraint that $\Gamma_i(\mathbf{p}, \mathbf{w}, r) \geq 0, \forall i \in [K]$. We restrict our attention to *budget-balanced* wagering mechanisms for which $\sum_{i \in [K]} \Gamma_i(\mathbf{p}, \mathbf{w}, r) = \sum_{i \in [K]} w_i$.

A wagering mechanism Γ is said to be *incentive-compatible* if for every agent $i \in [K]$ with belief $b_i \in [0, 1]$, every report $p_i \in [0, 1]$, every vector of reports of the other agents \mathbf{p}_{-i} , and every vec-

tor of wagers \mathbf{w} , $\mathbb{E}_{r \sim \text{Bern}(b_i)} [\Gamma_i((b_i, \mathbf{p}_{-i}), \mathbf{w}, r)] \geq \mathbb{E}_{r \sim \text{Bern}(b_i)} [\Gamma_i((p_i, \mathbf{p}_{-i}), \mathbf{w}, r)]$.

One class of budget-balanced, incentive-compatible wagering mechanisms is the Weighted Score Wagering Mechanisms (WSWMs) of Lambert et al. (2008; 2015). Fixing any proper loss function ℓ bounded in $[0, 1]$, agent i receives

$$\Gamma_i^{\text{WSWM}}(\mathbf{p}, \mathbf{w}, r) = w_i \left(1 - \ell(p_i, r) + \sum_{j \in [K]} w_j \ell(p_j, r) \right).$$

WSWMs are incentive-compatible because the payment an agent receives is a linear function of his loss, measured by a proper loss function. An agent makes a profit (i.e., receives payment greater than his wager), whenever his loss is smaller than the wager-weighted average agent loss, so accurate agents are more likely to increase their wealth.

3. The Full Information Setting

In this section, we present and analyze an online prediction algorithm, Weighted-Score Update (WSU), for the full information setting. We show that WSU is incentive-compatible and achieves regret $O(\sqrt{T \ln K})$.

Our key observation is that we can define a black-box reduction that transforms any budget-balanced wagering mechanism Γ to an online learning algorithm by setting $\pi_{t+1} = \Gamma(\mathbf{p}_t, \pi_t, r_t)$. Here we can interpret an expert's weight according to distribution π_t as their currency. Each expert "wagers" π_t at time t and receives a payoff π_{t+1} , which depends on the reports of the experts \mathbf{p} and the realization r_t . It is easy to see that any online prediction algorithm that is derived from an incentive-compatible wagering mechanism will in turn be incentive-compatible, because any misreport that increases weight π_{t+1} would also be a successful misreport in the wagering setting.

One might hope that applying this reduction to the WSWM would directly yield a no-regret online learning algorithm. But this is not the case, due to the fact that an expert who makes an inaccurate prediction can lose too much of his wealth (probability) if all other experts have low loss, and it can take a long time to recover from this. To handle this, we allow experts to "wager" only an η fraction of their current probability at each timestep for some $\eta \in (0, 0.5]$. This guarantees that no expert can obtain a probability $\pi_{i,t}$ close to zero without having made a long series of inaccurate predictions. Formally, the update rule of our algorithm, the Weighted-Score Update (WSU), is defined by:

$$\pi_{i,t+1} = \eta \Gamma_i^{\text{WSWM}}(\mathbf{p}_t, \pi_t, r_t) + (1 - \eta) \pi_{i,t}, \quad (1)$$

with weights $\pi_{i,1}$ initialized to $\pi_{i,1} = 1/K$ for all i .

We must show that π_t is a valid probability distribution

over experts at each t . This follows from the WSWM being budget-balanced; the proof is in the appendix (Lemma B.1).

By rewriting the WSU update rule in terms of relative loss $L_{i,t} = \ell_{i,t} - \sum_{j \in [K]} \pi_{j,t} \ell_{j,t}$, we can see that the form of the update is quite familiar. In particular, from Equation 1,

$$\begin{aligned} \pi_{i,t+1} &= \eta \pi_{i,t} \left(1 - \ell_{i,t} + \sum_{j \in [K]} \pi_{j,t} \ell_{j,t} \right) + (1 - \eta) \pi_{i,t} \\ &= \pi_{i,t} (1 - \eta L_{i,t}). \end{aligned} \quad (2)$$

This resembles the update rule for the (unnormalized) weights maintained by MWU, but with the relative loss $L_{i,t}$ in place of $\ell_{i,t}$. The D-PROD algorithm of Even-Dar et al. (2008) involves a similar update, but using loss relative to a single fixed distribution over experts instead of π_t .

We are now ready to prove our guarantees. The proof of Theorem 3.1 proceeds in a similar manner to the standard proof that MWU satisfies no regret. However, our proof is slightly simpler because we do not need to make a distinction between (unnormalized) weights and (normalized) probabilities. We can therefore avoid introducing the standard potential function used in proofs of no regret.

Theorem 3.1. *WSU is incentive-compatible and for step size $\eta = \sqrt{\ln(K)/T}$ yields regret $R \leq 2\sqrt{T \ln K}$.*

Proof. For incentive compatibility, note that from Equation (1), $\pi_{i,t+1}$ is a convex combination of a WSWM payment and $\pi_{i,t}$, which cannot be influenced by i 's report at time t . Since truthful reporting (at least weakly) maximizes each of these components, it also maximizes the sum.

For the regret, denoting by i^* the best expert in hindsight,

$$\begin{aligned} 1 &\geq \pi_{i^*,T+1} = \pi_{i^*,T} (1 - \eta L_{i^*,T}) \\ &= \pi_{i^*,1} \prod_{t \in [T]} (1 - \eta L_{i^*,t}) = \frac{1}{K} \prod_{t \in [T]} (1 - \eta L_{i^*,t}). \end{aligned}$$

Taking the logarithm for both sides of this inequality, we get

$$\begin{aligned} 0 &\geq -\ln K + \sum_{t \in [T]} \ln(1 - \eta L_{i^*,t}) \\ &\geq -\ln K + \sum_{t \in [T]} (-\eta L_{i^*,t} - \eta^2 L_{i^*,t}^2), \end{aligned}$$

where the last inequality comes from the fact that for $x \leq 1/2$, $\ln(1-x) \geq -x - x^2$ (see Lemma B.2). Rearranging and dividing both sides by η yields

$$-\sum_{t \in [T]} L_{i^*,t} \leq \frac{\ln K}{\eta} + \eta \sum_{t \in [T]} L_{i^*,t}^2.$$

Since we have $\sum_{t \in [T]} L_{i^*, t} = \sum_{t \in [T]} \ell_{i^*, t} - \sum_{t \in [T]} \sum_{j \in [K]} \pi_{j, t} \ell_{j, t} = -R$, this becomes

$$R \leq \frac{\ln K}{\eta} + \eta \sum_{t \in [T]} L_{i^*, t}^2 \leq \frac{\ln K}{\eta} + \eta T.$$

Finally, tuning $\eta = \sqrt{\ln(K)/T}$ gives us the result. \square

If T is not known in advance, a standard doubling trick (Auer et al., 2002) can be applied with only a constant factor increase in regret; see Appendix B.3 for details.

The regret and incentive-compatibility guarantees of WSU presented in Theorem 3.1 hold for all $[0, 1]$ -bounded proper loss functions ℓ . If ℓ is additionally convex, then these guarantees carry over to a (possibly more practical) variant of WSU, termed WSU-Aggr, that uses the same update rule but sets $\bar{p}_t = \sum_{i \in [K]} \pi_{i, t} p_{i, t}$ rather than choosing a single expert. Incentive compatibility is immediate. The regret bound follows from the fact that, by Jensen's inequality,

$$\sum_{t \in [T]} \ell \left(\sum_{i \in [K]} \pi_{i, t} p_{i, t}, r_t \right) \leq \sum_{t \in [T]} \sum_{i \in [K]} \pi_{i, t} \ell(p_{i, t}, r_t).$$

4. The Partial Information Setting

The encouraging results of the previous section apply only when the learner has access to the reports of all experts. But what if the learner has only partial information regarding these reports and still wants to incentivize all experts to report their predictions truthfully? In this section, we provide and analyze a novel algorithm, Weighted-Score Update with Uniform Exploration (WSU-UX), that is simultaneously no-regret and incentive-compatible in the bandit setting in which the learner chooses a single expert I_t at each round and observes only that expert's prediction. We show this algorithm has regret $O(T^{2/3}(K \ln K)^{1/3})$. This guarantee is weaker than that of EXP3, but we see in Section 6 that WSU-UX can perform similarly to EXP3 in practice with the additional advantage of incentive compatibility.

One might think that the standard trick of replacing the loss $\ell_{i, t}$ with an unbiased estimator $\hat{\ell}_{i, t}$ in the WSU update rule would suffice in order to guarantee both incentive compatibility and a regret rate of $O(\sqrt{T \ln K})$. Specifically, following Auer et al. (2002), we might consider setting $\hat{\ell}_{i, t} = 0$ for all experts $i \neq I_t$ whose predictions we do not observe, and $\hat{\ell}_{I_t, t} = \ell_{I_t, t} / \pi_{I_t, t}$ for the chosen expert. However, since these estimated losses are unbounded, this could lead to weights $\pi_{i, t}$ moving outside of $[0, 1]$, and we would no longer have a valid algorithm.

To solve this, we mix a distribution generated via WSU-style updates with a small amount of the uniform distribution.

This does not affect incentives, since the experts have no way of altering the uniform distribution, and has the convenient property that the estimated loss function is now bounded. By carefully tuning parameters, we are able to guarantee a valid probability distribution over experts. The resulting updates are given in Algorithm 1.

Algorithm 1 WSU-UX with parameters η and γ such that $0 < \eta, \gamma < 1/2$ and $\eta K / \gamma \leq 1/2$.

- 1: Set $\pi_{i, 1} = \frac{1}{K}, \forall i \in [K]$
- 2: **for** $t \in [T]$ **do**
- 3: Choose expert $I_t \sim \tilde{\pi}_{i, t} = (1 - \gamma)\pi_{i, t} + \frac{\gamma}{K}$
- 4: Compute: $\hat{\ell}_{I_t, t} = \frac{\ell_{I_t, t}}{\pi_{I_t, t}}$ and $\hat{\ell}_{i, t} = 0, \forall i \neq I_t$
- 5: Update $\pi_{i, t+1} = \pi_{i, t} \left(1 - \eta \left(\hat{\ell}_{i, t} - \sum_{j \in [K]} \pi_{j, t} \hat{\ell}_{j, t} \right) \right)$

We first prove that this is a valid algorithm, that is, that the distributions $\tilde{\pi}_t$ from which an expert is selected are valid, under appropriate settings of η and γ .

Lemma 4.1. *If $\eta K / \gamma \leq 1/2$, the WSU-UX weights π_t and $\tilde{\pi}_t$ are valid probability distributions for all $t \in [T + 1]$.*

Proof. We prove this inductively for π_t and $\tilde{\pi}_t$ simultaneously. The base case is trivial since at time $t = 1$, $\forall i \in [K], \pi_{i, 1} = \tilde{\pi}_{i, 1} = 1/K$. Now assume that for some t both π_t and $\tilde{\pi}_t$ are valid probability distributions. We distinguish two cases. First, suppose $i \neq I_t$. Then, since $\hat{\ell}_{i, t} = 0$, the WSU-UX update rule becomes

$$\pi_{i, t+1} = \pi_{i, t} \left(1 - \eta \left(0 - \pi_{I_t, t} \frac{\ell_{I_t, t}}{\pi_{I_t, t}} \right) \right) \geq 0.$$

Second, suppose $i = I_t$. Then

$$\begin{aligned} \pi_{i, t+1} &= \pi_{i, t} \left(1 - \eta \left(\frac{\ell_{i, t}}{\pi_{i, t}} - \pi_{i, t} \frac{\ell_{i, t}}{\pi_{i, t}} \right) \right) \\ &= \pi_{i, t} \left(1 - \eta \frac{\ell_{i, t}}{\pi_{i, t}} (1 - \pi_{i, t}) \right) \\ &\geq \pi_{i, t} \left(1 - \frac{\eta}{\pi_{i, t}} \right) \geq \pi_{i, t} \left(1 - \eta \frac{K}{\gamma} \right) \geq 0, \end{aligned}$$

where the penultimate inequality comes from the fact that $\tilde{\pi}_{i, t} \geq \gamma/K$, since by the inductive assumption $\pi_{i, t} \geq 0$. The last follows from the assumption that $\eta K / \gamma \leq 1/2$. Moreover, for the sum of probabilities we get:

$$\begin{aligned} \sum_{i \in [K]} \pi_{i, t+1} &= \sum_{i \in [K]} \pi_{i, t} \left(1 - \eta \left(\hat{\ell}_{i, t} - \sum_{j \in [K]} \pi_{j, t} \hat{\ell}_{j, t} \right) \right) \\ &= \sum_{i \in [K]} \pi_{i, t} - \eta \left(\sum_{i \in [K]} \pi_{i, t} \hat{\ell}_{i, t} - \sum_{i \in [K]} \pi_{i, t} \sum_{j \in [K]} \pi_{j, t} \hat{\ell}_{j, t} \right) \\ &= 1 - \eta \left(\sum_{i \in [K]} \pi_{i, t} \hat{\ell}_{i, t} - \sum_{j \in [K]} \pi_{j, t} \hat{\ell}_{j, t} \right) = 1. \end{aligned}$$

Thus π_{t+1} is valid. Since $\tilde{\pi}_{t+1}$ is a convex combination of two probability distributions, it is also a probability distribution, completing the inductive argument. \square

We are now ready to state the main theorem. The requirement that $T \geq K \ln K$ ensures that the precondition of Lemma 4.1 is satisfied for the settings of η and γ used.

Theorem 4.1. *For $T \geq K \ln K$ and parameters $\eta = \left(\frac{\ln K}{4K^{1/2}T}\right)^{2/3}$ and $\gamma = \left(\frac{K \ln K}{4T}\right)^{1/3}$, WSU-UX is incentive compatible and yields regret $R \leq 2(4T)^{2/3}(K \ln K)^{1/3}$.*

The proof of the theorem will follow from a series of claims and lemmas. We first examine the moments of $\hat{\ell}_{i,t}$ and verify that it is an unbiased estimator of $\ell_{i,t}$; the proof is direct and in Appendix C.

Lemma 4.2 (Moments). *Taking expectation with respect to the choice of expert at round t and keeping all else fixed, $\forall i \in [K], t \in [T], \mathbb{E}_{I_t \sim \tilde{\pi}_t} [\hat{\ell}_{i,t}] = \ell_{i,t}$. Furthermore,*

$$\mathbb{E}_{I_t \sim \tilde{\pi}_t} [\hat{\ell}_{i,t}^2] = \frac{\ell_{i,t}^2}{\tilde{\pi}_{i,t}} \leq \frac{1}{\tilde{\pi}_{i,t}}. \quad (3)$$

We next provide a second-order regret bound. It differs from the standard second-order regret bounds presented for bandit algorithms (see e.g., Bubeck et al. (2012, Chapter 3)) because it relates the ‘‘estimated regret’’ of the learner to the second moment of the estimated loss of the best-fixed expert in hindsight. The proof can be found in Appendix C.

Lemma 4.3 (Second-Order Bound). *For WSU-UX, the probability vectors π_1, \dots, π_T and the estimated losses $\hat{\ell}_{i,t}$ for $i \in [K], t \in [T]$ induce the following second-order bound:*

$$\sum_{t \in [T]} \sum_{i \in [K]} \pi_{i,t} \hat{\ell}_{i,t} - \sum_{t \in [T]} \hat{\ell}_{i^*,t} \leq \frac{\ln K}{\eta} + \eta \sum_{t \in [T]} \hat{\ell}_{i^*,t}^2 + \eta \sum_{t \in [T]} \sum_{i \in [K]} \pi_{i,t} \hat{\ell}_{i,t}^2$$

where $i^* = \arg \min_{i \in [K]} \sum_{t \in [T]} \ell_{i,t}$.

Proof. Since π_{T+1} is a valid probability distribution (Lemma 4.1), we have

$$\begin{aligned} 1 &\geq \pi_{i^*,T+1} = \pi_{i^*,T} \left(1 - \eta \left(\hat{\ell}_{i^*,T} - \sum_{j \in [K]} \pi_{j,T} \hat{\ell}_{j,T} \right) \right) \\ &= \pi_{i^*,1} \prod_{t \in [T]} \left(1 - \eta \left(\hat{\ell}_{i^*,t} - \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right) \right) \end{aligned}$$

Taking the logarithm for both sides, and using the fact that

$\pi_{i,1} = 1/K, \forall i \in [K]$, we get

$$0 \geq -\ln K + \sum_{t \in [T]} \ln \left(1 - \eta \left(\hat{\ell}_{i^*,t} - \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right) \right). \quad (4)$$

We next show that for all $t \in [T]$ and any $i \in [K]$ $\eta \left(\hat{\ell}_{i,t} - \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right) \leq 1/2$. We distinguish two cases. First, if $i \neq I_t$, then the inequality holds since $\hat{\ell}_{i,t} = 0$ and as a result the expression becomes $-\eta \cdot \pi_{I_t,t} \hat{\ell}_{I_t,t} \leq 0$. Second, if $i = I_t$, then the expression becomes

$$\begin{aligned} \eta \frac{\ell_{I_t,t}}{\tilde{\pi}_{I_t,t}} - \eta \pi_{I_t,t} \frac{\ell_{I_t,t}}{\tilde{\pi}_{I_t,t}} &= \eta \frac{\ell_{I_t,t}}{\tilde{\pi}_{I_t,t}} (1 - \pi_{I_t,t}) \\ &\leq \eta \frac{1}{\tilde{\pi}_{I_t,t}} \quad (\pi_{i,t} \geq 0, \ell_{i,t} \leq 1) \\ &\leq \eta \frac{K}{\gamma} \quad (\tilde{\pi}_{i,t} \geq \gamma/K, \text{ since } \pi_{i,t} \geq 0) \\ &\leq \frac{1}{2} \quad (\text{by definition}) \end{aligned}$$

We can now lower bound Equation (4) using the fact that for $z \leq 1/2$ it holds that: $\ln(1-z) \geq -z - z^2$ (Lemma B.2).

$$\begin{aligned} 0 &\geq -\ln K + \sum_{t \in [T]} \left[-\eta \left(\hat{\ell}_{i^*,t} - \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right) \right] - \\ &\quad - \sum_{t \in [T]} \left[\eta^2 \left(\hat{\ell}_{i^*,t} - \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right)^2 \right] \\ &\geq -\ln K - \eta \left[\sum_{t \in [T]} \hat{\ell}_{i^*,t} - \sum_{t \in [T]} \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right] \\ &\quad - \eta^2 \left[\sum_{t \in [T]} \left(\hat{\ell}_{i^*,t} - \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right)^2 \right] \\ &\geq -\ln K - \eta \left[\sum_{t \in [T]} \hat{\ell}_{i^*,t} - \sum_{t \in [T]} \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right] \\ &\quad - \eta^2 \sum_{t \in [T]} \hat{\ell}_{i^*,t}^2 - \eta^2 \sum_{t \in [T]} \left(\sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right)^2 \\ &\geq -\ln K - \eta \left[\sum_{t \in [T]} \hat{\ell}_{i^*,t} - \sum_{t \in [T]} \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t} \right] \\ &\quad - \eta^2 \sum_{t \in [T]} \hat{\ell}_{i^*,t}^2 - \eta^2 \sum_{t \in [T]} \sum_{j \in [K]} \pi_{j,t} \hat{\ell}_{j,t}^2 \end{aligned}$$

where the second inequality uses the fact that for a, b non-negative, $(a-b)^2 \leq a^2 + b^2$ and the last inequality uses Jensen’s inequality for function $f(x) = x^2$. Rearranging the latter and dividing both sides by η gives the result. \square

With that we can complete the proof of Theorem 4.1.

Proof of Theorem 4.1. It follows from incentive compatibility of WSU that an expert maximizes the expected value of $\pi_{i,t+1}$ by minimizing the expected value of $\hat{\ell}_{i,t}$. From the definition of $\hat{\ell}_{i,t}$, it is easy to see that minimizing the expected value of $\hat{\ell}_{i,t}$ is equivalent to minimizing the expected value of $\ell_{i,t}$. By properness of ℓ , this is achieved by truthfully reporting $p_{i,t} = b_{i,t}$.

We now show the regret bound. Taking expectations with respect to the choice of expert at round t for both sides of the equation in Lemma 4.3, we get

$$\begin{aligned} & \sum_{t \in [T]} \sum_{i \in [K]} \pi_{i,t} \mathbb{E}_{I_t \sim \tilde{\pi}_t} [\hat{\ell}_{i,t}] - \sum_{t \in [T]} \mathbb{E}_{I_t \sim \tilde{\pi}_t} [\hat{\ell}_{i^*,t}] \\ & \leq \eta \sum_{t \in [T]} \mathbb{E}_{I_t \sim \tilde{\pi}_t} [\hat{\ell}_{i^*,t}^2] + \frac{\ln K}{\eta} + \eta \sum_{t \in [T]} \sum_{i \in [K]} \pi_{i,t} \mathbb{E}_{I_t \sim \tilde{\pi}_t} [\hat{\ell}_{i,t}^2]. \end{aligned}$$

Using Lemma 4.2, this gives us

$$\begin{aligned} & \sum_{t \in [T]} \sum_{i \in [K]} \pi_{i,t} \ell_{i,t} - \sum_{t \in [T]} \ell_{i^*,t} \leq \\ & \leq \eta \sum_{t \in [T]} \frac{1}{\tilde{\pi}_{i^*,t}} + \frac{\ln K}{\eta} + \eta \sum_{t \in [T]} \sum_{i \in [K]} \pi_{i,t} \frac{1}{\tilde{\pi}_{i,t}} \\ & \leq \eta \sum_{t \in [T]} \frac{K}{\gamma} + \frac{\ln K}{\eta} + 2\eta KT \\ & \leq \frac{\eta KT}{\gamma} + \frac{\ln K}{\eta} + 2\eta KT, \end{aligned}$$

where the second inequality uses the fact that $\pi_{i,t} \leq 2\tilde{\pi}_{i,t}, \forall i \in [K], t \in [T]$ since $\gamma/K \geq 0$ and $\gamma \leq 1/2$. Next, we re-write $\pi_{i,t} = \frac{\tilde{\pi}_{i,t} - \gamma/K}{1-\gamma}$, yielding

$$\begin{aligned} & \sum_{t \in [T]} \sum_{i \in [K]} \frac{\tilde{\pi}_{i,t} - \frac{\gamma}{K}}{1-\gamma} \ell_{i,t} - \sum_{t \in [T]} \ell_{i^*,t} \leq \\ & \leq \frac{\eta KT}{\gamma} + \frac{\ln K}{\eta} + 2\eta KT. \end{aligned}$$

Since $1 - \gamma < 1$ and $\ell(p_{i,t}, r_t) \leq 1$, this can be relaxed to:

$$\begin{aligned} & \sum_{t \in [T]} \sum_{i \in [K]} \tilde{\pi}_{i,t} \ell(p_{i,t}, r_t) - \sum_{t \in [T]} \ell(p_{i^*,t}, r_t) \leq \\ & \leq \gamma T + \frac{\eta KT}{\gamma} + \frac{\ln K}{\eta} + 2\eta KT. \end{aligned}$$

Making $\gamma T = \eta KT/\gamma$ by setting $\gamma = \sqrt{\eta K}$, and $\eta = \left(\frac{\ln K}{4K^{1/2}T}\right)^{2/3}$ we get the regret result.⁶ \square

⁶The last derivation requires that $\eta \leq 1/K$, which is true for large enough horizons $T \geq K \ln K$.

As in the full information setting, a doubling trick can be applied if T is unknown (Appendix C.2).

We note that, unlike the full information setting in which WSU achieves the optimal regret bound for general loss functions, our regret bound in Theorem 4.1 is not as good as what can be achieved without incentive compatibility. Examining our analysis, one can see that if the loss of the best-fixed expert in hindsight is zero at each round, then the regret guarantee achieved by WSU-UX would be the same as EXP3, i.e., $O(\sqrt{T \ln K})$. Closing this gap via a tighter analysis of WSU-UX or via a new incentive-compatible algorithm is a compelling question for future work.

5. Forward-Looking Experts

So far we have assumed that the experts are myopic, aiming at time t to optimize their influence on the algorithm only at time $t + 1$ with no regard for future rounds. It is natural to ask whether it is possible to design learning algorithms that satisfy no regret while incentivizing truthful reports from forward-looking experts who care about their influence $\pi_{i,t'}$ at all $t' > t$. Neither WSU nor WSU-UX achieve this goal; see the appendix for examples that illustrate why.⁷

In order to derive an online learning algorithm that is incentive-compatible for forward-looking experts, we build on work by Witkowski et al. (2018), who studied a forecasting competition setting in which agents make predictions about a series of independent events, competing for a single prize. Unlike in our setting, their goal was to derive an incentive-compatible mechanism for choosing the winning agent; they are agnostic to how the elicited forecasts are aggregated. They defined a mechanism, Event-Lotteries Forecaster Selection Mechanism (ELF), in which, for every predicted event τ , every agent i is assigned a probability of being the event winner based on the quality of their prediction. The winner of the competition is the agent who wins the most events.

We build on this idea to define an online learning algorithm, ELF-X, for the full information setting. Like WSU, ELF-X incorporates WSWM payments, but in a different way. The distribution π_t at time t is defined as the distribution over experts output by the following randomized process:

1. At each round $\tau \in [t]$, pick agent i as the ‘‘winner’’ x_τ with probability $\frac{1}{K} \left(1 - \ell_{i,\tau} + \frac{1}{K} \sum_{j \in [K]} \ell_{j,\tau}\right)$.
2. Select $\arg \max_{i \in [K]} \sum_{\tau \in [t]} \mathbb{1}(x_\tau = i)$, the expert who won the most events, breaking ties uniformly.

It can be shown by a similar argument to that of Witkowski et al. (2018) that ELF-X is incentive-compatible. The proof,

⁷It is worth noting that in these examples, an expert can gain only a negligible amount from misreporting; it is an open question whether WSU satisfies some notion of ϵ -incentive compatibility.

along with a formal definition of incentive compatibility for forward-looking experts, is in the appendix.

Theorem 5.1. *ELF-X is incentive-compatible for forward-looking experts.*

While proving that ELF-X is no-regret remains an open problem, in the following section, we present experimental results suggesting that its regret is sublinear in T in practice.

6. Experiments

In this section, we empirically evaluate the performance of our proposed incentive-compatible algorithms, WSU and WSU-UX, compared with standard no-regret algorithms. We also evaluate the performance of ELF-X, which is incentive-compatible for non-myopic experts. Our code and the datasets we use are publicly available online.⁸

We ran each algorithm on publicly available datasets from a forecasting competition run by FiveThirtyEight⁹ in which users (henceforth called “forecasters”) make predictions about the outcomes of National Football League (NFL) games. Before each game, FiveThirtyEight releases information on the past performance of the two opposing teams, and forecasters provide probabilistic predictions about which team will win the game. FiveThirtyEight maintains a public leaderboard with the most accurate forecasters, updated after each game. The datasets for the 2018–2019 and 2019–2020 seasons each include all forecasters’ predictions, labeled with the forecaster’s unique id, information about the corresponding game, and the game’s outcome. Each NFL season has a total of 267 games, so in our setting, $T = 267$. For 2018–2019 (respectively, 2019–2020), while 15,702 (15,140) participated, only 302 (375) made predictions for every game. In order to reduce variance, for each value of K , we sampled 10 groups of K forecasters from the 302 (respectively, 375), and for each such group, ran each algorithm 50 times.

We evaluate performance using quadratic loss. We compare the cumulative loss of each algorithm against the cumulative loss of the best fixed forecaster in hindsight. For the full information setting, we compare WSU and ELF-X against Hedge, which achieves optimal regret guarantees since the quadratic loss is exp-concave, and MWU, which is more similar in form to WSU, in order to evaluate whether anything is lost in terms of regret when incentive compatibility is achieved. For the partial information setting, we compare WSU-UX against EXP3. For each full information algorithm, we run both the variant in which a single expert

is selected at each timestep and the variant in which the learner outputs a weighted combination of expert reports (labeled \star -Aggr). For ELF-X-Aggr, since π_t cannot be computed in closed form, we approximate it via sampling.

We present the results of our experiments on the 2018–2019 dataset in Figure 1; the results on the 2019–2020 dataset are in Appendix E.1, and exhibit similar trends. We note that lines correspond to average regret (across all samples of experts and all repetitions), while the error bands correspond to the 20th and 80th percentiles; this leads to much smaller error bands for larger values of K since the specific sampling of experts has less influence on regret for large K .

Validating our theoretical results, WSU performs almost identically (in terms of the dependence on both K and T) to MWU when fed the same set of reports—this, of course, does not take into account that MWU is not incentive compatible and may lead to misreports in practice, potentially degrading predictions. Interestingly, we also see that WSU-Aggr performs almost identically to MWU-Aggr. This suggests that the performance of WSU-Aggr is considerably better than the bound in Section 3 implies. It is an interesting open question to see whether better regret guarantees can be proved for WSU-Aggr, perhaps with respect to the best fixed distribution of experts. As expected, both WSU and MWU are outperformed by Hedge, which achieves optimal regret bounds for squared loss but no incentive guarantees.

ELF-X appears to exhibit diminishing regret on this dataset, particularly for $K = 5$. However, ELF-X and ELF-X-Aggr perform worse than WSU and WSU-Aggr respectively when fed the same input, particularly when the number of experts is large. Although ELF-X obtains a stronger incentive guarantee, the violations of incentive compatibility for forward-looking experts exhibited by WSU are very small in our examples. In practice, we expect that WSU is a superior choice to ELF-X when balancing regret and incentive properties, even for forward-looking experts.

For the bandit setting, quite encouragingly, we see that the performance of WSU-UX is only slightly worse than that of EXP3, and appears significantly better than the $O(T^{2/3})$ regret bound in Section 4 would suggest. This could be a byproduct of our analysis not being tight, and it remains an open question whether this bound can be improved.

The experiments presented in this section focus on settings with relatively small horizons T since an NFL season has only 267 matches. In Appendix E.2, we present our results (which also validate our theoretical analysis) for Monte Carlo simulations for larger horizons.

⁸Code: <https://github.com/charapod/noregr-and-ic>. Datasets: <https://github.com/fivethirtyeight/nfl-elo-game>

⁹<https://projects.fivethirtyeight.com/2019-nfl-forecasting-game/>

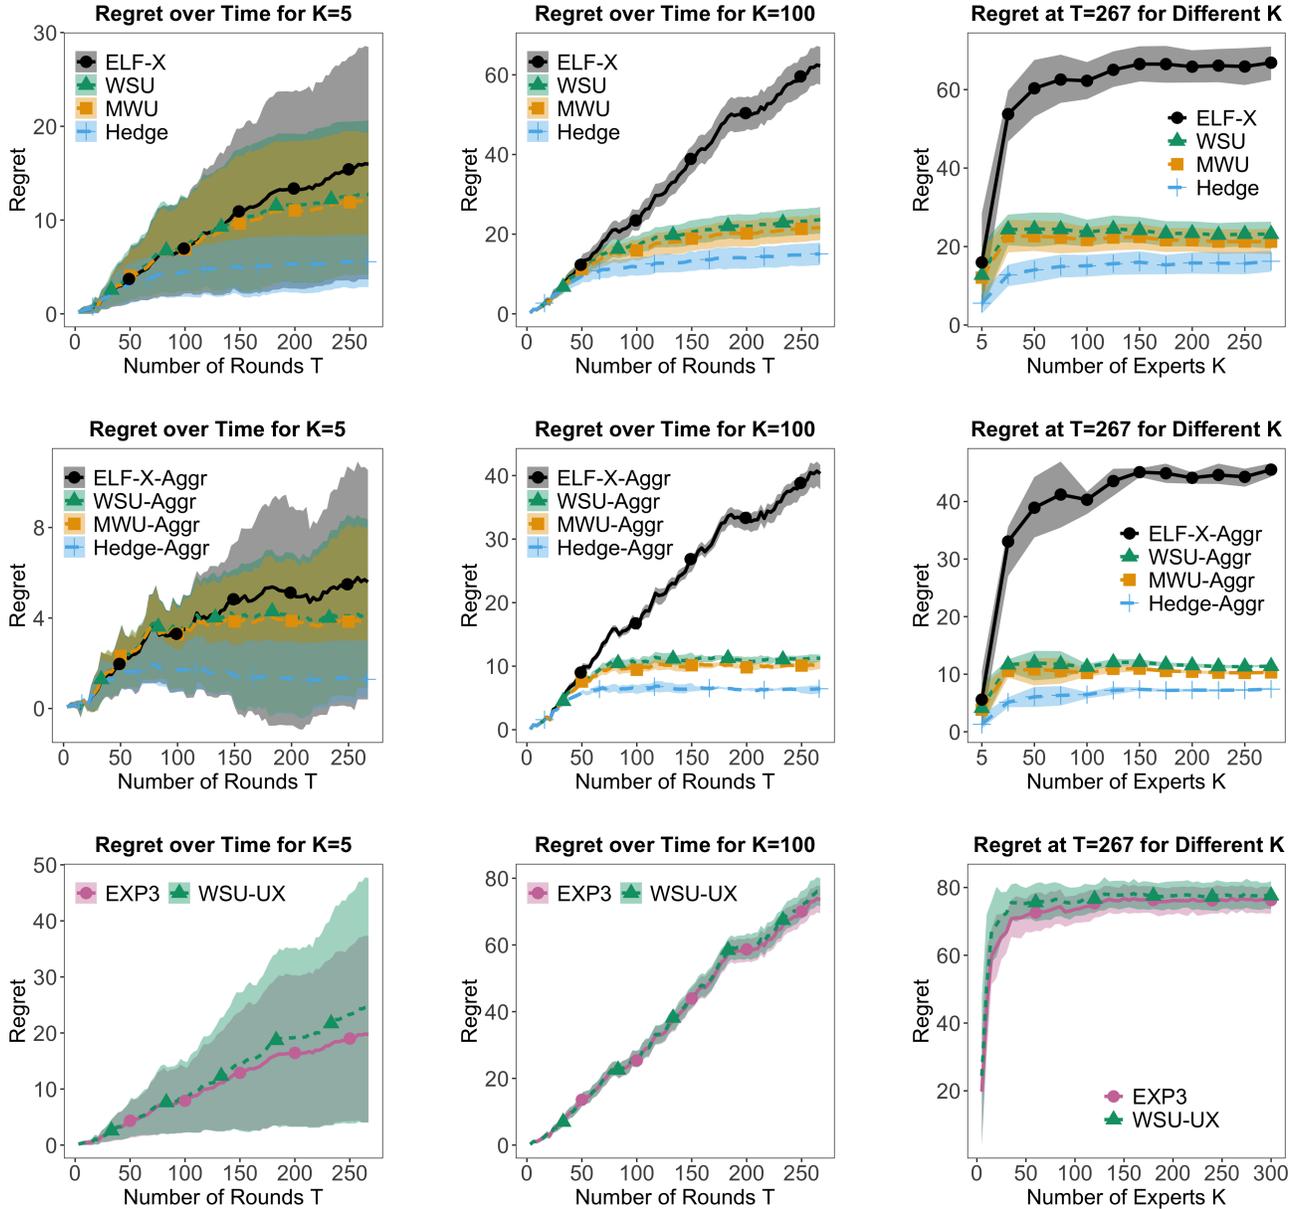


Figure 1. Comparisons on the 2018–2019 FiveThirtyEight NFL dataset. Top: Full-information setting with \hat{p}_t the prediction of a single expert chosen according to π_t . Middle: Full-information setting with $\hat{p}_t = \sum_{i \in [K]} \pi_{i,t} p_{i,t}$. Bottom: Partial information setting.

7. Conclusion and Open Questions

We studied the problem of online learning with strategic experts. We introduced algorithms that are simultaneously no-regret and incentive-compatible, and assessed their performance experimentally on data from FiveThirtyEight. Several open questions arise. In the full-information setting, there is the question of whether an incentive-compatible algorithm exists with better regret bounds for the special case of exp-concave bounded proper loss functions. For the bandit setting, there is the question of whether there

exist incentive-compatible algorithms that bridge the gap between the regret of WSU-UX and that of EXP3, and whether a better regret guarantee could be proved for WSU-UX via a tighter analysis. There is also the question of whether ELF-X is indeed no-regret, as our experimental results might suggest. More broadly, the most important research question that we believe needs to be addressed in online learning from strategic agents is the quantification of the tradeoff between incentive-incompatibility and standard learning guarantees and how to balance these in practice.

Acknowledgments

The authors are grateful to Yiling Chen, Ariel Procaccia, Rob Schapire, Vasilis Syrgkanis, and Jens Witkowski for helpful discussions at different stages of this work, and to the anonymous reviewers for their comments and suggestions. Part of the research was conducted while Rupert Freeman, David M. Pennock, and Chara Podimata were at Microsoft Research NYC as a postdoc, researcher, and intern respectively. Chara Podimata is also supported in part under grant No. CCF-1718549 of the National Science Foundation and the Harvard Data Science Initiative.

References

- Abernethy, J. and Frongillo, R. M. A collaborative mechanism for crowdsourcing prediction problems. In *Advances in Neural Information Processing Systems*, 2011.
- Abernethy, J., Chen, Y., and Vaughan, J. W. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12:1–12:38, 2013.
- Arora, S., Hazan, E., and Kale, S. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- Bubeck, S., Cesa-Bianchi, N., et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- Cesa-Bianchi, N., Freund, Y., Haussler, D., Helmbold, D. P., Schapire, R. E., and Warmuth, M. K. How to use expert advice. *Journal of the ACM (JACM)*, 44(3):427–485, 1997.
- Even-Dar, E., Kearns, M., Mansour, Y., and Wortman, J. Regret to the best vs. regret to the average. *Machine Learning Journal*, 72:21–37, 2008.
- Freund, Y. and Schapire, R. E. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Frongillo, R., Della Penna, N., and Reid, M. D. Interpreting prediction markets: A stochastic approach. In *Advances in Neural Information Processing Systems*, 2012.
- Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Hu, J. and Storkey, A. Multi-period trading prediction markets with connections to machine learning. In *International Conference on Machine Learning*, pp. 1773–1781, 2014.
- Kivinen, J. and Warmuth, M. K. Averaging expert predictions. In *European Conference on Computational Learning Theory*, pp. 153–167. Springer, 1999.
- Lambert, N. S., Langford, J., Wortman, J., Chen, Y., Reeves, D., Shoham, Y., et al. Self-financed wagering mechanisms for forecasting. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 170–179, 2008.
- Lambert, N. S., Langford, J., Vaughan, J. W., Chen, Y., Reeves, D. M., Shoham, Y., and Pennock, D. M. An axiomatic characterization of wagering mechanisms. *Journal of Economic Theory*, 156:389–416, 2015.
- Littlestone, N. and Warmuth, M. K. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.
- McCarthy, J. Measures of the value of information. *Proceedings of the National Academy of Sciences of the United States of America*, 42(9):654, 1956.
- Orabona, F. and Pál, D. Coin betting and parameter-free online learning. In *Advances in Neural Information Processing Systems*, pp. 577–585, 2016.
- Reid, M. D. and Williamson, R. C. Surrogate regret bounds for proper losses. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pp. 897–904. ACM, 2009.
- Roughgarden, T. and Schrijvers, O. Online prediction with selfish experts. In *Advances in Neural Information Processing Systems*, pp. 1300–1310, 2017.
- Savage, L. J. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- Tetlock, P. E. and Gardner, D. *Superforecasting: The Art and Science of Prediction*. Crown, 2015.
- Vovk, V. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2):153–173, 1998.
- Vovk, V. G. Aggregating strategies. *Proc. of Computational Learning Theory, 1990*, 1990.
- Witkowski, J., Freeman, R., Vaughan, J. W., Pennock, D. M., and Krause, A. Incentive-compatible forecasting competitions. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.