No-Regret Learning with Unbounded Losses: The Case of Logarithmic Pooling
Supplementary Appendix

May 10, 2023

A. Omitted Proofs

Proof of Lemma 4.4. Fix any t. Note that since the space of possible weights is Δ^m, it is most natural to think of ∇R as a function from Δ^m to $\mathbb{R}^m/T(\mathbf{1}_m)$, i.e. $\mathbb{R}^m$ modulo translation by the all-ones vector (which is orthogonal to Δ^m in $\mathbb{R}^m$). That is, $\nabla R(\mathbf{w}) = -((w_1)^{\alpha-1}, \dots, (w_m)^{\alpha-1})$, where this vector may be thought of as modulo translation by the all-ones vector. Nevertheless, we find it convenient to define $\partial_i R(\mathbf{w}) := -(w_i)^{\alpha-1}$. We define $\partial_i L^t(\mathbf{w})$ similarly (see Section 4.4).

Define $\mathbf{h} \in \mathbb{R}^m$ to have coordinates $h_i := \partial_i R(\mathbf{w}^t) - \eta_t \partial_i L^t(\mathbf{w}^t)$. Per the update rule, we have that $h_i \equiv R(\mathbf{w}^{t+1}) \mod T(\mathbf{1}_m)$. We have

$$-(w_i^t)^{\alpha-1} - \eta_t\zeta = \partial_i R(\mathbf{w}^t) - \eta_t\zeta \le h_i \le \partial_i R(\mathbf{w}^t) + \frac{\eta_t\zeta}{w_i^t} = -(w_i^t)^{\alpha-1} + \frac{\eta_t\zeta}{w_i^t} \tag{1}$$

Applying the first and last claims of Lemma A.1 (below) with $a = \alpha - 1$, $\mathbf{v} = \mathbf{w}^t$, $\kappa = \eta_t\zeta$, and $\mathbf{g} = -\mathbf{h}$, we have that there exists a unique $c \in \mathbb{R}$ such that

$$\sum_{i=1}^m (-h_i + c)^{1/(\alpha-1)} = 1,$$

and in fact that $-\eta_t\zeta \le c \le m\eta_t\zeta$. (Equation 1 is relevant here because it is equivalent to the $v_i^a - \frac{\kappa}{w_i^t} \le g_i \le v_i^a + \kappa$ conditions in Lemma A.1. This is also where we use that $\eta_t\zeta \le (1-\alpha)^2(w_i^t)^\alpha$, which is equivalent to $\kappa \le a^2 v_i^{a+1}$.) The significance of this fact is that $(-h_i + c)^{1/(\alpha-1)}$ is precisely w_i^{t+1}, since (in $\mathbb{R}^m$) we have that $(\partial_i R(\mathbf{w}^{t+1}), \dots, \partial_i R(\mathbf{w}^{t+1})) = \mathbf{h} - c \cdot \mathbf{1}$ for *some* c, and in particular this c must be such that $\sum_i w_i^{t+1} = 1$. In particular, this means that for all i, we have

$$(w_i^{t+1})^{\alpha-1} = -h_i + c \le (w_i^t)^{\alpha-1} + \eta_t\zeta + \frac{1}{\min_k w_k}\eta_t\zeta = (w_i^t)^{\alpha-1} + \left(\frac{1}{\min_k w_k} + 1\right)\eta_t\zeta.$$

Here, the inequality comes from the left inequality of Equation 1 and the fact that $c \le \frac{1}{\min_k w_k}\eta_t\zeta$. If we also have that $\eta_t\zeta \le (1-\alpha)^2(w_i^t)^\alpha$, then the last claim of Lemma A.1 gives us that

$$(w_i^{t+1})^{\alpha-1} = -h_i + c \le (w_i^t)^{\alpha-1} + \eta_t\zeta + m\eta_t\zeta = (w_i^t)^{\alpha-1} + (m+1)\eta_t\zeta.$$

Similarly, we have

$$(w_i^{t+1})^{\alpha-1} = -h_i + c \ge (w_i^t)^{\alpha-1} - \frac{\eta_t\zeta}{w_i^t} - \eta_t\zeta = (w_i^t)^{\alpha-1} - \left(\frac{1}{w_i^t} + 1\right)\eta_t\zeta.$$

□

Lemma A.1. *Let $-1 < a < 0$ and $\mathbf{g} \in \mathbb{R}^m$. There is a unique $c \in \mathbb{R}$ such that $\sum_i (g_i + c)^{1/a} = 1$. Furthermore, let $\mathbf{v} \in \Delta^m$ and $\kappa \ge 0$. Then:*

- *If $g_i \le v_i^a + \kappa$ for all i, then $c \ge -\kappa$.*
- *If $g_i \ge v_i^a - \frac{\kappa}{v_i}$ for all i, then $c \le \frac{\kappa}{\min_i v_i}$.*
  - *And if, furthermore, $\kappa \le a^2 v_i^{a+1}$ for all i, then $c \le m\kappa$.*

Proof. Observe that $\sum_i (g_i + c)^{1/a}$ is a continuous, monotone decreasing function on $c \in (-\min_i g_i, \infty)$; the range of the function on this interval is $(0, \infty)$. Therefore, there is a unique $c \in (-\min_i g_i, \infty)$ such that the sum equals 1.

We now prove the first bullet. Since $x^{1/a}$ decreases in x and $g_i \le v_i^a + \kappa$, we have that

$$1 = \sum_i (g_i + c)^{1/a} \ge \sum_i (v_i^a + \kappa + c)^{1/a}.$$

Suppose for contradiction that $c < -\kappa$. Then $v_i^a + \kappa + c < v_i^a$ for all i, so

$$\sum_i (v_i^a + \kappa + c)^{1/a} > \sum_i (v_i^a)^{1/a} = \sum_i v_i = 1.$$

This is a contradiction, so in fact $c \ge -\kappa$.

The first claim of the second bullet is analogous. Since $x^{1/a}$ decreases in x and $g_i \ge v_i^a - \kappa v_i$, we have that

$$1 = \sum_i (g_i + c)^{1/a} \le \sum_i \left(v_i^a - \frac{\kappa}{v_i} + c\right)^{1/a}. \tag{2}$$

Suppose for contradiction that $c > \frac{\kappa}{v_i}$ for every i. Then $v_i^a - \frac{\kappa}{v_i} + c > v_i^a$ for all i, so

$$\sum_i \left(v_i^a - \frac{\kappa}{v_i} + c\right)^{1/a} < \sum_i (v_i^a)^{1/a} = \sum_i v_i = 1.$$

This is a contradiction, so in fact $c \le \frac{\kappa}{\min_i v_i}$.

We now prove the second claim of the second bullet. To do so, we note the following technical lemma (proof below).

Lemma A.2. *For $-1 < a < 0$ and $\kappa, c \geq 0$, the function $f(x) = \left(x^a - \frac{\kappa}{x} + c\right)^{1/a}$ is defined and concave at any value of $x > 0$ such that $a^2 x^{a+1} \geq \kappa$.*

Since for a general concave function f it holds that $\frac{1}{m}\sum_{i=1}^m f(x_i) \leq f\left(\frac{1}{m}\sum_{i=1}^m x_i\right)$, the following inequality follows from Lemma A.2:

$$\sum_i \left(v_i^a - \frac{\kappa}{v_i} + c\right)^{1/a} \leq m\left(\left(\frac{1}{m}\right)^a - \kappa m + c\right)^{1/a}.$$

(Here we are using the fact that $\sum_i v_i = 1$.) Now, combining this fact with Equation 2, we have that

$$m\left(\left(\frac{1}{m}\right)^a - \kappa m + c\right)^{1/a} \geq 1$$
$$\left(\frac{1}{m}\right)^a - \kappa m + c \leq \left(\frac{1}{m}\right)^a$$

so $c \leq m\kappa$, as desired. $\square$

Proof of Lemma A.2. To show that f is defined for any x such that $a^2x^{a+1} \geq \kappa$, we need to show that $x^a - \frac{\kappa}{x} + c > 0$ for such values of x. This is indeed the case:

$$x^a - \frac{\kappa}{x} + c \geq x^a - a^2x^a + c = (1-a^2)x^a + c > c \geq 0.$$

Now we show concavity. We have

$$f''(x) = \frac{1}{-a}\left(\left(1+\frac{1}{-a}\right)\left(x^a - \frac{\kappa}{x} + c\right)^{1/a-2}\left(ax^{a-1} + \frac{\kappa}{x^2}\right)^2 - \left(x^a - \frac{\kappa}{x} + c\right)^{1/a-1}\left(a(a-1)x^{a-2} - \frac{2\kappa}{x^3}\right)\right)$$

so we wish to show that

$$\left(1+\frac{1}{-a}\right)\left(x^a - \frac{\kappa}{x} + c\right)^{1/a-2}\left(ax^{a-1} + \frac{\kappa}{x^2}\right)^2 \leq \left(x^a - \frac{\kappa}{x} + c\right)^{1/a-1}\left(a(a-1)x^{a-2} - \frac{2\kappa}{x^3}\right)$$

for every x such that $a^2x^{a+1} \geq \kappa$. Fix any such x, and let $d = \frac{\kappa}{x^{a+1}}$ (so $0 \leq d \leq a^2$). We have

$$\begin{aligned}
d &\leq a^2 \\
(1+a)(a^2-d)d &\geq 0 \\
(1-a)(a+d)^2 &\leq -a(1-d)(a(a-1)-2d) && \text{(rearrange terms)} \\
\left(1-\frac{1}{a}\right)(a+d)^2x^a &\leq ((1-d)x^a)(a(a-1)-2d) && \text{(multiply by } \tfrac{x^a}{-a}\text{)} \\
\left(1-\frac{1}{a}\right)(a+d)^2x^a &\leq ((1-d)x^a + c)(a(a-1)-2d) && (c(a(a-1)-2d) \geq 0) \\
\left(1-\frac{1}{a}\right)((a+d)x^{a-1})^2 &\leq ((1-d)x^a + c)(a(a-1)-2d)x^{a-2} && \text{(multiply by } x^{a-2}\text{)} \\
\left(1-\frac{1}{a}\right)\left(ax^{a-1} + \frac{\kappa}{x^2}\right)^2 &\leq \left(x^a - \frac{\kappa}{x} + c\right)\left(a(a-1)x^{a-2} - \frac{2\kappa}{x^3}\right) && \text{(substitute } d = \kappa x^{-a-1}\text{).}
\end{aligned}$$

Note that the fifth line is justified by the fact that $c \geq 0$ and $a(a-1) \geq 2d$ (because $a^2 \geq d$ and $-a > a^2 \geq d$). Now, multiplying both sides by $\left(x^a - \frac{\kappa}{x} + c\right)^{1/a-2}$ completes the proof. $\square$

Proof of Corollary 4.5. Note that $\eta\gamma = \frac{1}{\sqrt{T}m^{(1+\alpha)/2}}$ and also that $\eta_t \leq \eta$ for all t; we will be using these facts.

To prove (#1), we proceed by induction on t. In the case of $t = 1$, all weights are $1/m$, so the claim holds for sufficiently large T. Now assume that the claim holds for a generic $t < T$; we show it for $t + 1$.

By the small gradient assumption, we may use Lemma 4.4 with $\zeta = \gamma$. By the inductive hypothesis (and the fact that $\eta_t \leq \eta$), we may apply the second part of Lemma 4.4:

$$\begin{aligned}(w_i^{t+1})^{\alpha-1} &\leq (w_i^t)^{\alpha-1} + (m+1)\eta\gamma \leq \cdots \leq (1/m)^{\alpha-1} + t(m+1)\eta\gamma. \\ &\leq (1/m)^{\alpha-1} + \frac{(T-1)(m+1)}{m^{(1+\alpha)/2}\sqrt{T}} \leq 3m^{(1-\alpha)/2}\sqrt{T}.\end{aligned}$$

Since $\frac{-1}{2} < \alpha - 1 < 0$, this means that $w_i^t \geq \frac{1}{10\sqrt{m}} T^{1/(2(\alpha-1))}$.

We also have that

$$(w_i^{t+1})^{\alpha} \geq \frac{1}{(10\sqrt{m})^{\alpha}} T^{\alpha/(2(\alpha-1))} \geq \frac{4}{m^{(1+\alpha)/2}} T^{-1/2} = 4\eta\gamma$$

for T sufficiently large, since $\frac{\alpha}{2(\alpha-1)} > \frac{-1}{2}$. This completes the inductive step, and thus the proof of (#1).

To prove (#2), we use the following technical lemma (see below for the proof).

Lemma A.3. *Fix $x > 0$ and $-1 < a < 0$. Let $f(y) = (x^a + y)^{1/a}$. Then for all $y > -x^a$, we have*

$$x - f(y) \leq \frac{-1}{a} x^{1-a} y \tag{3}$$

and for all $-1 < c \leq 0$, for all $cx^a \leq y \leq 0$, we have

$$f(y) - x \leq \frac{1}{a}(1+c)^{1/a-1} x^{1-a} y. \tag{4}$$

We apply Equation 3 to $x = w_i^t$, $y = (m+1)\eta\gamma$, and $a = \alpha - 1$. This tells us that

$$w_i^t - w_i^{t+1} \leq w_i^t - ((w_i^t)^{\alpha-1} + (m+1)\eta\gamma)^{1/(\alpha-1)} \leq 2(w_i^t)^{2-\alpha}(m+1)\eta\gamma.$$

The first step follows by the second part of Lemma 4.4 and the fact that $\eta_t \leq \eta$. The second step follows from Equation 3 and uses the fact that $\frac{1}{1-\alpha} > 2$.

For the other side of (#2), we observe that since by (#1) we have $(w_i^t)^{\alpha} \geq 4\eta\gamma$, it follows that $\frac{1}{2}(w_i^t)^{\alpha} \geq (w_i^t + 1)\eta\gamma$, and so $\left(\frac{1}{w_i^t} + 1\right)\eta\gamma \leq \frac{1}{2}(w_i^t)^{\alpha-1}$. Therefore, we can apply Equation 4 to $x = w_i^t$, $y = -\left(\frac{1}{w_i^t} + 1\right)\eta\gamma$, $a = \alpha - 1$, and $c = -\frac{1}{2}$. This tells us that

$$w_i^{t+1} - w_i^t \leq \left((w_i^t)^{\alpha-1} - \left(\frac{1}{w_i^t} + 1\right)\eta\gamma\right)^{1/(\alpha-1)} - w_i^t \leq 16(w_i^t)^{2-\alpha}\left(\frac{1}{w_i^t} + 1\right)\eta\gamma$$

$$\leq 32(w_i^t)^{1-\alpha}\eta\gamma.$$

This completes the proof. □

Proof of Lemma A.3. For all $y > -x^a$, we have

$$f'(y) = \frac{1}{a}(x^a + y)^{1/a-1}$$

and

$$f''(y) = \frac{1}{a}\left(\frac{1}{a} - 1\right)(x^a + y)^{1/a-2} > 0,$$

so f' is increasing. Thus, for positive values of y we have

$$f'(0) \leq \frac{f(y) - f(0)}{y} = \frac{f(y) - x}{y} \leq f'(y)$$

and for negative values of y we have

$$f'(y) \leq \frac{f(y) - f(0)}{y} = \frac{f(y) - x}{y} \leq f'(0).$$

Regardless of whether y is positive or negative, this means that $x - f(y) \leq -yf'(0) = \frac{-1}{a}x^{1-a}y$.

Now, let $-1 < c \leq 0$ and suppose that $cx^a \leq y \leq 0$. Since f' is increasing, we have that

$$f'(y) \geq f'(cx^a) = \frac{1}{a}((1+c)x^a)^{1/a-1} = \frac{1}{a}(1+c)^{1/a-1}x^{1-a},$$

so

$$f(y) - x \leq yf'(y) \leq \frac{1}{a}(1+c)^{1/a-1}x^{1-a}y.$$

□

Proof of Lemma 4.8. We first derive an expression for $\partial_i L(\mathbf{w})$ given expert reports $\mathbf{p}^1, \ldots, \mathbf{p}^m$, where $L(\mathbf{w})$ is the log loss of the logarithmic pool $\mathbf{p}^*(\mathbf{w})$ of $\mathbf{p}^1, \ldots, \mathbf{p}^m$ with weights $\mathbf{w}$, and j is the realized outcome. We have[1]

$$\begin{aligned}
\partial_i L(\mathbf{w}) &= -\partial_i \ln \frac{\prod_{k=1}^m (p_j^k)^{w_k}}{\sum_{\ell=1}^n \prod_{k=1}^m (p_\ell^k)^{w_k}} = \partial_i \ln\left(\sum_{\ell=1}^n \prod_{k=1}^m (p_\ell^k)^{w_k}\right) - \partial_i \ln\left(\prod_{k=1}^m (p_j^k)^{w_k}\right) \\
&= \frac{\sum_{\ell=1}^n \ln p_\ell^i \cdot \prod_{k=1}^m (p_\ell^k)^{w_k}}{\sum_{\ell=1}^n \prod_{k=1}^m (p_\ell^k)^{w_k}} - \ln p_j^i = \sum_{\ell=1}^n p_\ell^*(\mathbf{w}) \ln p_\ell^i - \ln p_j^i. \qquad (5)
\end{aligned}$$

[1]It should be noted that $\nabla L(\mathbf{w})$ is most naturally thought of as living in $\mathbb{R}^m / T(\mathbf{1}_m)$, i.e. m-dimensional space modulo translation by the all-ones vector, since $\mathbf{w}$ lives in a place that is orthogonal to the all-ones vector. As an arbitrary but convenient convention, we define $\partial_i L(\mathbf{w})$ to be the specific value derived below, and define the small gradient assumption accordingly.

Equation (2) now follows fairly straightforwardly. Equation 5 tells us that $\partial_i L(\mathbf{w}) \leq -\ln p_J^i$, where J is the random variable corresponding to the realized outcome. Therefore, we have

$$\mathbb{P}\left[\partial_i L(\mathbf{w}) \geq \zeta\right] \leq \mathbb{P}\left[-\ln p_J^i \geq \zeta\right] = \mathbb{P}\left[p_J^i \leq e^{-\zeta}\right] = \sum_{j=1}^{n} \mathbb{P}\left[J = j \ \& \ p_j^i \leq e^{-\zeta}\right]$$

$$= \sum_{j=1}^{n} \mathbb{P}\left[p_j^i \leq e^{-\zeta}\right] \mathbb{P}\left[J = j \mid p_j^i \leq e^{-\zeta}\right] \leq \sum_{j=1}^{n} \mathbb{P}\left[J = j \mid p_j^i \leq e^{-\zeta}\right] \leq ne^{-\zeta},$$

where the last step follows by the calibration property. This proves Equation (2).

We now prove Equation (3). The proof has a similar idea, but is somewhat more technical. We begin by proving the following lemma; we again use the calibration property in the proof.

Lemma A.4. *For all q, we have*

$$\mathbb{P}\left[\forall j \exists i : p_j^i \leq q\right] \leq mnq.$$

Proof. Let J be the random variable corresponding to the index of the outcome that ends up happening. We have

$$\begin{aligned}\mathbb{P}\left[\forall j \exists i : p_j^i \leq q\right] &\leq \mathbb{P}\left[\exists i : p_J^i \leq q\right] = \sum_{j \in [n]} \mathbb{P}\left[J = j \ \& \ \exists i : p_j^i \leq q\right] \\ &\leq \sum_{j \in [n]} \sum_{i \in [m]} \mathbb{P}\left[J = j \ \& \ p_j^i \leq q\right] \\ &= \sum_{j \in [n]} \sum_{i \in [m]} \mathbb{P}\left[p_j^i \leq q\right] \mathbb{P}\left[J = j \mid p_j^i \leq q\right] \leq \sum_{j \in [n]} \sum_{i \in [m]} 1 \cdot q = mnq,\end{aligned}$$

where the fact that $\mathbb{P}\left[J = j \mid p_j^i \leq q\right] \leq q$ follows by the calibration property. □

Corollary A.5. *For any reports $\mathbf{p}^1, \dots, \mathbf{p}^m$, weight vector $\mathbf{w}$, $i \in [m]$, and $j \in [n]$, we have*

$$\mathbb{P}\left[p_j^*(\mathbf{w}) \geq \frac{(p_j^i)^{w_i}}{q}\right] \leq mnq.$$

Proof. We have

$$p_j^*(\mathbf{w}) = \frac{\prod_{k=1}^{m} (p_j^k)^{w_k}}{\sum_{\ell=1}^{n} \prod_{k=1}^{m} (p_\ell^k)^{w_k}} \leq \frac{(p_j^i)^{w_i}}{\sum_{\ell=1}^{n} \prod_{k=1}^{m} (p_\ell^k)^{w_k}}.$$

Now, assuming that there is an ℓ such that for every k we have $p_\ell^k > q$, the denominator is greater than q, in which case we have $p_j^*(\mathbf{w}) < \frac{(p_j^i)^{w_i}}{q}$. Therefore, if $p_j^*(\mathbf{w}) \geq \frac{(p_j^i)^{w_i}}{q}$, it follows that for every ℓ there is a k such that $p_\ell^k \leq q$. By Lemma A.4, this happens with probability at most mnq. □

We now use Corollary A.5 to prove Equation (3). Note that the equation is trivial for $\zeta < n$, so we assume that $\zeta \geq n$. By setting $q := e^{\frac{-\zeta}{n}}$, we may restate Equation (3) as follows: for any $q \leq \frac{1}{e}$, any $i \in [m]$, and any weight vector $\mathbf{w}$,

$$\mathbb{P}\left[\partial_i L(\mathbf{w}) \leq -\frac{n \ln 1/q}{w_i}\right] \leq mn^2 q.$$

(Note that the condition $q \le \frac{1}{e}$ is equivalent to $\zeta \ge n$.) We prove this result.

From Equation 5, we have

$$\partial_i L(\mathbf{w}) = \sum_{j=1}^{n} p_j^*(\mathbf{w}) \ln p_j^i - \ln p_j^i \ge \sum_{j=1}^{n} p_j(\mathbf{w}) \ln p_j^i.$$

Now, it suffices to show that for each $j \in [n]$, the probability that $p_j(\mathbf{w}) \ln p_j^i \le -\frac{\ln 1/q}{w_i} = \frac{\ln q}{w_i}$ is at most mnq; the desired result will then follow by the union bound. By Corollary A.5, for each j we have that

$$\mathbb{P}\left[p_j(\mathbf{w}) \ln p_j^i \le \frac{(p_j^i)^{w_i}}{q} \ln p_j^i\right] \le mnq.$$

Additionally, we know for a fact that $p_j(\mathbf{w}) \ln p_j^i \ge \ln p_j^i$ (since $p_j(\mathbf{w}) \le 1$), so in fact

$$\mathbb{P}\left[p_j(\mathbf{w}) \ln p_j^i \le \max\left(\frac{(p_j^i)^{w_i}}{q} \ln p_j^i, \ln p_j^i\right)\right] \le mnq.$$

It remains only to show that $\max\left(\frac{(p_j^i)^{w_i}}{q} \ln p_j^i, \ln p_j^i\right) \ge \frac{\ln q}{w_i}$. If $p_j^i \ge q^{1/w_i}$ then this is clearly true, since in that case $\ln p_j^i \ge \frac{\ln q}{w_i}$. Now suppose that $p_j^i < q^{1/w_i}$. Observe that $\frac{x^{w_i}}{q} \ln x$ decreases on $(0, e^{-1/w_i})$, and that (since $q \le \frac{1}{e}$) we have $q^{1/w_i} \le e^{-1/w_i}$. Therefore,

$$\frac{(p_j^i)^{w_i}}{q} \ln p_j^i \le \frac{(q^{1/w_i})^{w_i}}{q} \ln q^{1/w_i} = \frac{\ln q}{w_i}.$$

This completes the proof of Equation (3), and thus of Lemma 4.8. $\square$

The following lemma lower bounds the regret of Algorithm 1 as a function of ζ.

Lemma A.6. *Consider a run of Algorithm 1. Let ζ be such that $-\frac{\zeta}{w_i^t} \le \partial_i L^t(\mathbf{w}^t) \le \zeta$ for all i, t. The total regret is at most*

$$O\left(\zeta^{2(2-\alpha)/(1-\alpha)} T^{(5-\alpha)/(1-\alpha)}\right).$$

Proof of Lemma A.6. We first bound w_i^t for all i, t. From Lemma 4.4, we have that

$$(w_i^{t+1})^{\alpha-1} \le (w_i^t)^{\alpha-1} + \left(\frac{1}{\min_i w_i^t} + 1\right) \eta_t \zeta \le (w_i^t)^{\alpha-1} + 2\zeta.$$

Here we use that $\frac{1}{\min_i w_i} + 1 \le \frac{2}{\min_i w_i}$ and that $\eta_t \le \min_i w_i$. Therefore, we have that

$$(w_i^t)^{\alpha-1} \le (w_i^{t-1})^{\alpha-1} + 2\zeta \le \cdots \le m^{1-\alpha} + 2\zeta(t-1) \le m^{1-\alpha} + 2\zeta T.$$

Thus, $w_i^t \ge (m^{1-\alpha} + 2\zeta T)^{1/(\alpha-1)} \ge \Omega((\zeta T)^{1/(\alpha-1)})$ for all i, t.

We now use the standard regret bound for online mirror descent, see e.g. (Orabona, 2021, Theorem 6.8):

$$\text{Regret} \le \max_t \frac{B_R(\mathbf{u}; \mathbf{w}^t)}{\eta_T} + \frac{1}{2\lambda} \sum_{t=1}^{T} \eta_t \left\|\nabla L^t(\mathbf{w}^t)\right\|_*^2 \tag{6}$$

where $B_R(\cdot;\cdot)$ is the Bregman divergence of with respect to R, $\mathbf{u}$ is the optimal (overall loss-minimizing) point, λ is a constant such that R is λ-strongly convex with respect to a norm of our choice over Δ^m, and $\|\cdot\|_*$ is the dual norm of the aforementioned norm.

Note that for any $\mathbf{x} \in \Delta^m$, we have

$$\max_{\mathbf{v}\in\Delta^m} B_R(\mathbf{v}; \mathbf{x}) = \max_{\mathbf{v}\in\Delta^m} R(\mathbf{v}) - R(\mathbf{x}) - \nabla R(\mathbf{x}) \cdot (\mathbf{v} - \mathbf{x}) \le \frac{m^{1-\alpha}}{\alpha} + (\min_i x_i)^{\alpha-1}.$$

In the last step, we use the fact that $-\nabla R(\mathbf{x}) = (x_1^{\alpha-1}, \dots, x_m^{\alpha-1})$ (all of these coordinates are positive), so $-\nabla R(\mathbf{x}) \cdot (\mathbf{v} - \mathbf{x}) \le (x_1^{\alpha-1}, \dots, x_m^{\alpha-1}) \cdot \mathbf{v}$, and that all coordinates of $\mathbf{v}$ are non-negative and add to 1.

Therefore, given our bound on w_i^t, we have that this first component of our regret bound (6) is at most

$$\frac{1}{\eta_T}\left(\frac{m^{1-\alpha}}{\alpha} + m^{1-\alpha} + 2\zeta T\right) \le O\left(\frac{\zeta T}{\eta_T}\right) \le O\left(\frac{\zeta T}{(\zeta T)^{1/(\alpha-1)}}\right) = O\left((\zeta T)^{(2-\alpha)/(1-\alpha)}\right).$$

To bound the second term, we choose to work with the ℓ_1 norm. To show that R is λ-convex it suffices to show that for all $\mathbf{x}, \mathbf{y} \in \Delta^m$ we have $(\nabla^2 R(\mathbf{x})\mathbf{y}) \cdot \mathbf{y} \ge \lambda \|\mathbf{y}\|^2$, where $\nabla^2 R$ is the Hessian of R (Shalev-Shwartz, 2007, Lemma 14) (see also (Orabona, 2021, Theorem 4.3)). Equivalently, we wish to find a λ such that

$$(1-\alpha)\sum_i x_i^{\alpha-2} y_i^2 \ge \lambda.$$

Since $x_i^{\alpha-2} \ge 1$ for all i, the left-hand side is at least $(1-\alpha)\sum_i y_i^2 \ge \frac{1-\alpha}{m}$, so $\lambda = \frac{1-\alpha}{m}$ suffices.

Now, given $\theta \in \mathbb{R}^m$, we have $\|\theta\|_* = \max_{\mathbf{x}:\|\mathbf{x}\|\le 1} \theta \cdot \mathbf{x}$. In the case of the ℓ_1 primal norm, the dual norm is the largest absolute component of θ. Thus, we have

$$\left\|\nabla L^t(\mathbf{x}^t)\right\|_* \le \frac{\zeta}{w_i^t} \le O\left(\zeta(\zeta T)^{1/(1-\alpha)}\right) = O\left(\zeta^{(2-\alpha)/(1-\alpha)} T^{1/(1-\alpha)}\right).$$

Since $\eta_t \le O(T^{-1/2})$, we have that the second component of our regret bound (6) is at most

$$O\left(T \cdot T^{-1/2} \cdot \zeta^{2(2-\alpha)/(1-\alpha)} T^{2/(1-\alpha)}\right) \le O\left(\zeta^{2(2-\alpha)/(1-\alpha)} T^{(5-\alpha)/(1-\alpha)}\right).$$

This component dominates our bound on the regret of the first component, in both ζ and T. This concludes the proof. $\square$

Corollary A.7. *The expected total regret of our algorithm conditional on the small gradient assumption* not *holding, times the probability of this event, is at most* $\tilde{O}(T^{(5-\alpha)/(1-\alpha)-10})$.

Proof. Let Z be the minimum value of ζ such that $-\frac{\zeta}{w_i^t} \leq \partial_i L^t(\mathbf{w}^t) \leq \zeta$ for all i, t. Note that by Lemma 4.8, we have that

$$\mathbb{P}\left[Z \geq x\right] \leq \sum_{i=1}^{m} \sum_{t=1}^{T} (mn^2 e^{-\frac{x}{n}} + ne^{-x}) \leq 2m^2n^2Te^{-\frac{x}{n}}.$$

Let μ be the constant hidden in the big-O of Lemma A.6, i.e. a constant (dependent on m, n, and α) such that

$$\text{Regret} \leq \mu Z^{2(2-\alpha)/(1-\alpha)} T^{(5-\alpha)/(1-\alpha)}.$$

Let $r(Z, T)$ be the expression on the right-hand side. The small gradient assumption not holding is equivalent to $Z > 12n \ln T$, or equivalently, $r(Z, T) > r(12n \ln T, T)$. The expected regret of our algorithm conditional on the small gradient assumption *not* holding, times the probability of this event, is therefore at most the expected value of $r(Z, T)$ conditional on the value being greater than $r(12n \ln T, T)$, times this probability. This is equal to

$$\begin{aligned}
&r(12n \ln T, T) \cdot \mathbb{P}\left[Z > 12n \ln T\right] + \int_{x=r(12n \ln T, T)}^{\infty} \mathbb{P}\left[r(Z, T) \geq x\right] dx \\
&\leq \sum_{k=11}^{\infty} r((k+1)n \ln T, T) \cdot \mathbb{P}\left[Z \geq kn \ln T\right] \\
&\leq \sum_{k=11}^{\infty} \mu \cdot ((k+1)n \ln T)^{2(2-\alpha)/(1-\alpha)} T^{(5-\alpha)/(1-\alpha)} \cdot 2m^2n^2T \cdot T^{-k} \\
&\leq \sum_{k=11}^{\infty} \tilde{O}(T^{1+(5-\alpha)/(1-\alpha)-k}) = \tilde{O}(T^{(5-\alpha)/(1-\alpha)-10}),
\end{aligned}$$

as desired. (The first inequality follows by matching the first term with the $k = 11$ summand and upper-bounding the integral with subsequent summands, noting that $r((k+1)n \ln T, T) \geq 1$.) $\square$

Note that $\frac{5-\alpha}{1-\alpha} - 10 \leq \frac{5-1/2}{1-1/2} - 10 = -1$. Therefore, the contribution to expected regret from the case that the small gradient assumption does not hold is $\tilde{O}(T^{-1})$, which is negligible. Together with Corollary 4.7 (which bounds regret under the small gradient assumption), this proves Theorem 3.2.

We now extend Theorem 3.2 by showing that the theorem holds even if experts are only *approximately* calibrated.

Definition A.8. For $\tau \geq 1$, we say that expert i is τ-*calibrated* if for all $\mathbf{p} \in \Delta^n$ and $j \in [n]$, we have that $\mathbb{P}\left[J = j \mid \mathbf{p}^j = \mathbf{p}\right] \leq \tau p_j$. We say that $\mathbb{P}$ satisfies the τ-*approximate calibration property* if every expert is τ-calibrated.

Corollary A.9. *For any τ, Theorem 3.2 holds even if the calibration property is replaced with the τ-approximate calibration property.*

(Note that the τ is subsumed by the big-O notation in Theorem 3.2; Corollary A.9 does not allow experts to be arbitrarily miscalibrated.)

Technically, Corollary A.9 is a corollary of the *proof* of Theorem 3.2, rather than a corollary of the theorem itself.[2]

Proof. We only used the calibration property in the proofs of Equations (2) and (3). In the proof of Equation (2), we used the fact that $\mathbb{P}\left[J = j \mid p_j^i \leq e^{-\zeta}\right] \leq e^{-\zeta}$; the right-hand side now becomes $\tau e^{-\zeta}$, and so the right-hand side of Equation (2) changes to $\tau n e^{-\zeta}$. Similarly, in the proof of Equation (3), we use the calibration property in the proof of Lemma A.4; the right-hand side of the lemma changes to τmnq, and correspondingly Equation (3) changes to $\tau mn^2 e^{-\zeta/n}$.

Lemma 4.8 is only used in the proof of Corollary A.7, where $2m^2n^2T$ is replaced by $2\tau m^2n^2T$. Since τ is a constant, Corollary A.7 holds verbatim. □

B. $\Omega(\sqrt{T})$ Lower bound

We show that no OMD algorithm with a constant step size[3] substantially outperforms Algorithm 1.

Theorem B.1. *For every strictly convex function $R : \Delta^m \to \mathbb{R}$ that is continuously twice differentiable at its minimum, and $\eta \geq 0$, online mirror descent with regularizer R and constant step size η incurs $\Omega(\sqrt{T})$ expected regret.*

Proof. Our examples will have $m = n = 2$. The space of weights is one-dimensional; let us call w the weight of the first expert. We may treat R as a (convex) function of w, and similarly for the losses at each time step. We assume that $R'(0.5) = 0$; this allows us to assume that $w_1 = 0.5$ and does not affect the proof idea.

It is straightforward to check that if Experts 1 and 2 assign probabilities p and $\frac{1}{2}$, respectively, to the correct outcome, then

$$L'(w) = \frac{(1-p)^w}{p^w + (1-p)^w} \ln \frac{1-p}{p}.$$

If roles are reversed (they say $\frac{1}{2}$ and p respectively) then

$$L'(w) = -\frac{(1-p)^{1-w}}{p^{1-w} + (1-p)^{1-w}} \ln \frac{1-p}{p}.$$

We first prove the regret bound if η is small ($\eta \leq T^{-1/2}$). Consider the following setting: Expert 1 always reports $(50\%, 50\%)$; Expert 2 always reports $(90\%, 10\%)$; and Outcome 1 happens with probability 90% at each time step. It is a matter of simple computation that:

- $L'(w) \leq 2$ no matter the outcome or the value of w.
- If $w \geq 0.4$, then $p_1^*(w) \leq 0.8$.

[2]Fun fact: the technical term for a corollary to a proof is a *porism*.

[3]While Algorithm 1 does not always have a constant step size, it does so with high probability. The examples that prove Theorem B.1 cause $\Omega(\sqrt{T})$ regret in the typical case, rather than causing unusually large regret in an atypical case. This makes our comparison of Algorithm 1 to this class fair.

The first point implies that $R'(w_t) \geq -2\eta t$ for all t. It follows from the second point that the algorithm will output weights that will result in an aggregate probability of less than 80% for values of t such that $-2\eta t \geq R'(0.4)$, i.e. for $t \leq \frac{-R'(0.4)}{2\eta}$. Each of these time steps accumulates constant regret compared to the optimal weight vector in hindsight (which with high probability will be near 1). Therefore, the expected total regret accumulated during these time steps is $\Omega(1/\eta) = \Omega(\sqrt{T})$.

Now we consider the case in which η is large ($\eta \geq \sqrt{T}$). In this case our example is the same as before, except we change which expert is "ignorant" (reports $(50\%, 50\%)$ and which is "informed" (reports $(90\%, 10\%)$). Specifically the informed expert will be the one with a lower weight (breaking ties arbitrarily).

We will show that our algorithm incurs $\Omega(\eta)$ regret compared to always choosing weight 0.5. Suppose without loss of generality that at a given time step t, Expert 1 is informed (so $w^t \leq 0.5$). Observe that

$$\begin{aligned} L(w^t) - L(0.5) &= -(0.5 - w^t)L'(0.5) + O((0.5 - w)^2) \\ &= -(0.5 - w^t)\frac{\sqrt{1-p}}{\sqrt{p} + \sqrt{1-p}} \ln \frac{1-p}{p} + O((0.5 - w)^2), \end{aligned}$$

where p is the probability that Expert 1 assigns to the event that happens (so $p = 0.9$ with probability 0.9 and $p = 0.1$ with probability 0.1). This expression is (up to lower order terms) equal to $c(0.5 - w^t)$ if $p = 0.9$ and $-3c(0.5 - w^t)$ if $p = 0.1$, where $c \approx 0.55$. This means that an expected regret (relative to $w = 0.5$) of $0.6c(0.5 - w^t)$ (up to lower order terms) is incurred.

Let D be such that $R''(w) \leq D$ for all w such that $|w - 0.5| \leq \frac{\sqrt{T}}{4D}$. (Such a D exists because R is continuously twice differentiable at 0.5.) If $|w^t - 0.5| \geq \frac{\sqrt{T}}{4D}$, we just showed that an expected regret (relative to $w = 0.5$) of $\Omega\left(\frac{\sqrt{T}}{4D}\right)$ is incurred. On the other hand, suppose that $|w^t - 0.5| \leq \frac{\sqrt{T}}{4D}$. We show that $|w^{t+1} - 0.5| \geq \frac{\sqrt{T}}{4D}$.

To see this, note that $|L'(w^t)| \geq 0.5$, we have that $|R'(w^{t+1}) - R'(w^t)| \geq 0.5\eta$. We also have that $D\,|w^{t+1} - w^t| \geq |R'(w^{t+1}) - R'(w^t)|$, so $D\,|w^{t+1} - w^t| \geq 0.5\eta$. Therefore, $|w^{t+1} - w^t| \geq \frac{\eta}{2D} \geq \frac{\sqrt{T}}{2D}$, which means that $|w^{t+1} - 0.5| \geq \frac{\sqrt{T}}{4D}$.

This means that an expected regret (relative to $w = 0.5$) of $\Omega\left(\frac{\sqrt{T}}{4D}\right)$ is incurred on at least half of time steps. Since D is a constant, it follows that a total regret of at least $\Omega(\sqrt{T})$ is incurred, as desired. □

References

[Orabona 2021] ORABONA, Francesco: A Modern Introduction to Online Learning. In: *CoRR* abs/1912.13213 (2021). – URL `http://arxiv.org/abs/1912.13213`

[Shalev-Shwartz 2007] SHALEV-SHWARTZ, Shai: *Online Learning: Theory, Algorithms, and Applications*, Dissertation, 2007