# No-Regret Algorithms for Unconstrained Online Convex Optimization 

Matthew Streeter<br>Duolingo, Inc.*<br>Pittsburgh, PA 15232<br>matt@duolingo.com

H. Brendan McMahan<br>Google, Inc.<br>Seattle, WA 98103<br>mcmahan@google.com


#### Abstract

Some of the most compelling applications of online convex optimization, including online prediction and classification, are unconstrained: the natural feasible set is $\mathbb{R}^{n}$. Existing algorithms fail to achieve sub-linear regret in this setting unless constraints on the comparator point $\dot{x}$ are known in advance. We present algorithms that, without such prior knowledge, offer near-optimal regret bounds with respect to any choice of $\dot{x}$. In particular, regret with respect to $\dot{x}=0$ is constant. We then prove lower bounds showing that our guarantees are near-optimal in this setting.


## 1 Introduction

Over the past several years, online convex optimization has emerged as a fundamental tool for solving problems in machine learning (see, e.g., [3, 12] for an introduction). The reduction from general online convex optimization to online linear optimization means that simple and efficient (in memory and time) algorithms can be used to tackle large-scale machine learning problems. The key theoretical techniques behind essentially all the algorithms in this field are the use of a fixed or increasing strongly convex regularizer (for gradient descent algorithms, this is equivalent to a fixed or decreasing learning rate sequence). In this paper, we show that a fundamentally different type of algorithm can offer significant advantages over these approaches. Our algorithms adjust their learning rates based not just on the number of rounds, but also based on the sum of gradients seen so far. This allows us to start with small learning rates, but effectively increase the learning rate if the problem instance warrants it.

This approach produces regret bounds of the form $\mathcal{O}(R \sqrt{T} \log ((1+R) T))$, where $R=\|\dot{x}\|_{2}$ is the $L_{2}$ norm of an arbitrary comparator. Critically, our algorithms provide this guarantee simultaneously for all $\dot{x} \in \mathbb{R}^{n}$, without any need to know $R$ in advance. A consequence of this is that we can guarantee at most constant regret with respect to the origin, $\dot{x}=0$. This technique can be applied to any online convex optimization problem where a fixed feasible set is not an essential component of the problem. We discuss two applications of particular interest below:

Online Prediction Perhaps the single most important application of online convex optimization is the following prediction setting: the world presents an attribute vector $a_{t} \in \mathbb{R}^{n}$; the prediction algorithm produces a prediction $\sigma\left(a_{t} \cdot x_{t}\right)$, where $x_{t} \in \mathbb{R}^{n}$ represents the model parameters, and $\sigma: \mathbb{R} \rightarrow Y$ maps the linear prediction into the appropriate label space. Then, the adversary reveals the label $y_{t} \in Y$, and the prediction is penalized according to a loss function $\ell: Y \times Y \rightarrow \mathbb{R}$. For appropriately chosen $\sigma$ and $\ell$, this becomes a problem of online convex optimization against functions $f_{t}(x)=\ell\left(\sigma\left(a_{t} \cdot x\right), y_{t}\right)$. In this formulation, there are no inherent restrictions on the model coefficients $x \in \mathbb{R}^{n}$. The practitioner may have prior knowledge that "small" model vectors are more

[^0]likely than large ones, but this is rarely best encoded as a feasible set $\mathcal{F}$, which says: "all $x_{t} \in \mathcal{F}$ are equally likely, and all other $x_{t}$ are ruled out." A more general strategy is to introduce a fixed convex regularizer: $L_{1}$ and $L_{2}^{2}$ penalties are common, but domain-specific choices are also possible. While algorithms of this form have proved very effective at solving these problems, theoretical guarantees usually require fixing a feasible set of radius $R$, or at least an intelligent guess of the norm of an optimal comparator $\stackrel{\grave{x}}{x}$.

The Unconstrained Experts Problem and Portfolio Management In the classic problem of predicting with expert advice (e.g., [3]), there are $n$ experts, and on each round $t$ the player selects an expert (say $i$ ), and obtains reward $g_{t, i}$ from a bounded interval (say $[-1,1]$ ). Typically, one uses an algorithm that proposes a probability distribution $p_{t}$ on experts, so the expected reward is $p_{t} \cdot g_{t}$.

Our algorithms apply to an unconstrained version of this problem: there are still $n$ experts with payouts in $[-1,1]$, but rather than selecting an individual expert, the player can place a "bet" of $x_{t, i}$ on each expert $i$, and then receives reward $\sum_{i} x_{t, i} g_{t, i}=x_{t} \cdot g_{t}$. The bets are unconstrained (betting a negative value corresponds to betting against the expert). In this setting, a natural goal is the following: place bets so as to achieve as much reward as possible, subject to the constraint that total losses are bounded by a constant (which can be set equal to some starting budget which is to be invested). Our algorithms can satisfy constraints of this form because regret with respect to $\stackrel{\circ}{x}=0$ (which equals total loss) is bounded by a constant.

It is useful to contrast our results in this setting to previous applications of online convex optimization to portfolio management, for example [6] and [2]. By applying algorithms for exp-concave loss functions, they obtain $\log$-wealth within $\mathcal{O}(\log (T))$ of the best constant rebalanced portfolio. However, this approach requires a "no-junk-bond" assumption: on each round, for each investment, you always retain at least an $\alpha>0$ fraction of your initial investment. While this may be realistic (though not guaranteed!) for blue-chip stocks, it certainly is not for bets on derivatives that can lose all their value unless a particular event occurs (e.g., a stock price crosses some threshold). Our model allows us to handle such investments: if we play $x_{i}>0$, an outcome of $g_{i}=-1$ corresponds exactly to losing $100 \%$ of that investment. Our results imply that if even one investment (out of exponentially many choices) has significant returns, we will increase our wealth exponentially.

Notation and Problem Statement For the algorithms considered in this paper, it will be more natural to consider reward-maximization rather than loss-minimization. Therefore, we consider online linear optimization where the goal is to maximize cumulative reward given adversarially selected linear reward functions $f_{t}(x)=g_{t} \cdot x$. On each round $t=1 \ldots T$, the algorithm selects a point $x_{t} \in \mathbb{R}^{n}$, receives reward $f_{t}\left(x_{t}\right)=g_{t} \cdot x_{t}$, and observes $g_{t}$. For simplicity, we assume $g_{t, i} \in$ $[-1,1]$, that is, $\left\|g_{t}\right\|_{\infty} \leq 1$. If the real problem is against convex loss functions $\ell_{t}(x)$, they can be converted to our framework by taking $g_{t}=-\nabla \ell_{t}\left(x_{t}\right)$ (see pseudo-code for REWARD-DoUBLING), using the standard reduction from online convex optimization to online linear optimization [13].
We use the compressed summation notation $g_{1: t}=\sum_{s=1}^{t} g_{s}$ for both vectors and scalars. We study the reward of our algorithms, and their regret against a fixed comparator $\grave{x}$ :

$$
\operatorname{Reward} \equiv \sum_{t=1}^{T} g_{t} \cdot x_{t} \quad \text { and } \quad \operatorname{Regret}(\stackrel{\circ}{x}) \equiv g_{1: T} \cdot \stackrel{\circ}{x}-\sum_{t=1}^{T} g_{t} \cdot x_{t}
$$

Comparison of Regret Bounds The primary contribution of this paper is to establish matching upper and lower bounds for unconstrained online convex optimization problems, using algorithms that require no prior information about the comparator point $\stackrel{\circ}{x}$. Specifically, we present an algorithm that, for any $\dot{x} \in \mathbb{R}^{n}$, guarantees $\operatorname{Regret}(\stackrel{\circ}{x}) \leq \mathcal{O}\left(\|\dot{x}\|_{2} \sqrt{T} \log \left(\left(1+\|\dot{x}\|_{2}\right) \sqrt{T}\right)\right)$. To obtain this guarantee, we show that it is sufficient (and necessary) that reward is $\Omega\left(\exp \left(\left|g_{1: T}\right| / \sqrt{T}\right)\right.$ ) (see Theorem 1). This shift of emphasis from regret-minimization to reward-maximization eliminates the quantification on $\stackrel{x}{x}$, and may be useful in other contexts.

Table 1 compares the bounds for REWARD-DoUBLING (this paper) to those of two previous algorithms: online gradient descent [13] and projected exponentiated gradient descent [8, 12]. For each

[^1]Assuming $\left\|g_{t}\right\|_{2} \leq 1$ :

|  | $\stackrel{\circ}{x}=0$ | $\\|\stackrel{\circ}{x}\\|_{2} \leq R$ | Arbitrary $\stackrel{\circ}{x}$ |
| :--- | :---: | :---: | :---: |
| Gradient Descent, $\eta=\frac{R}{\sqrt{T}}$ | $R \sqrt{T}$ | $R \sqrt{T}$ | $\\|\dot{x}\\|_{2} T$ |
| REWARD-DoUbLING | $\epsilon$ | $R \sqrt{T} \log \left(\frac{n(1+R) T}{\epsilon}\right)$ | $\\|\stackrel{\circ}{x}\\|_{2} \sqrt{T} \log \left(\frac{n\left(1+\\|\check{x}\\|_{2}\right) T}{\epsilon}\right)$ |

Assuming $\left\|g_{t}\right\|_{\infty} \leq 1$ :

|  | $\dot{x}=0$ | $\\|\dot{x}\\|_{1} \leq R$ | Arbitrary $\dot{x}$ |
| :--- | :---: | :---: | :---: |
| Exponentiated G.D. | $R \sqrt{T \log n}$ | $R \sqrt{T \log n}$ | $\\|\dot{x}\\|_{1} T$ |
| REWARD-DoUbLING | $\epsilon$ | $R \sqrt{T} \log \left(\frac{n(1+R) T}{\epsilon}\right)$ | $\\|\dot{x}\\|_{1} \sqrt{T} \log \left(\frac{n\left(1+\\|\dot{x}\\|_{1}\right) \sqrt{T}}{\epsilon}\right)$ |

Table 1: Worst-case regret bounds for various algorithms (up to constant factors). Exponentiated G.D. uses feasible set $\left\{x:\|x\|_{1} \leq R\right\}$, and Reward-Doubling uses $\epsilon_{i}=\frac{\epsilon}{n}$ in both cases.
algorithm, we consider a fixed choice of parameter settings and then look at how regret changes as we vary the comparator point $\dot{x}$.

Gradient descent is minimax-optimal [1] when the comparator point is contained in a hypershere whose radius is known in advance $\left(\|\grave{x}\|_{2} \leq R\right)$ and gradients are sparse ( $\left\|g_{t}\right\|_{2} \leq 1$, top table). Exponentiated gradient descent excels when gradients are dense ( $\left\|g_{t}\right\|_{\infty} \leq 1$, bottom table) but the comparator point is sparse ( $\|\dot{x}\|_{1} \leq R$ for $R$ known in advance). In both these cases, the bounds for REWARD-DOUBLING match those of the previous algorithms up to logarithmic factors, even when they are tuned optimally with knowledge of $R$.
The advantage of REWARD-Doubling shows up when the guess of $R$ used to tune the competing algorithms turns out to be wrong. When $\dot{x}=0$, REWARD-DoUbLing offers constant regret compared to $\Omega(\sqrt{T})$ for the other algorithms. When $\stackrel{\circ}{x}$ can be arbitrary, only REWARD-DoUBLING offers sub-linear regret (and in fact its regret bound is optimal, as shown in Theorem 8).
In order to guarantee constant origin-regret, REWARD-DoUBLING frequently "jumps" back to playing the origin, which may be undesirable in some applications. In Section 4 we introduce SmOOTH-REWARD-DOUBLING, which achieves similar guarantees without resetting to the origin.

Related Work Our work is related, at least in spirit, to the use of a momentum term in stochastic gradient descent for back propagation in neural networks [7, 11, 9]. These results are similar in motivation in that they effectively yield a larger learning rate when many recent gradients point in the same direction.
In Follow-The-Regularized-Leader terms, the exponentiated gradient descent algorithm with unnormalized weights of Kivinen and Warmuth [8] plays $x_{t+1}=\arg \min _{x \in \mathbb{R}_{+}^{n}} g_{1: t} \cdot x+\frac{1}{\eta}(x \log x-x)$, which has closed-form solution $x_{t+1}=\exp \left(-\eta g_{1: t}\right)$. Like our algorithm, this algorithm moves away from the origin exponentially fast, but unlike our algorithm it can incur arbitrarily large regret with respect to $\grave{x}=0$. Theorem 9 shows that no algorithm of this form can provide bounds like the ones proved in this paper.
Hazan and Kale [5] give regret bounds in terms of the variance of the $g_{t}$. Letting $G=\left|g_{1: t}\right|$ and $H=\sum_{t=1}^{T} g_{t}^{2}$, they prove regret bounds of the form $\mathcal{O}(\sqrt{V})$ where $V=H-G^{2} / T$. This result has some similarity to our work in that $G / \sqrt{T}=\sqrt{H-V}$, and so if we hold $H$ constant, then when $V$ is low, the critical ratio $G / \sqrt{T}$ that appears in our bounds is large. However, they consider the case of a known feasible set, and their algorithm (gradient descent with a constant learning rate) cannot obtain bounds of the form we prove.

## 2 Reward and Regret

In this section we present a general result that converts lower bounds on reward into upper bounds on regret, for one-dimensional online linear optimization. In the unconstrained setting, this result will be sufficient to provide guarantees for general $n$-dimensional online convex optimization.

Theorem 1. Consider an algorithm for one-dimensional online linear optimization that, when run on a sequence of gradients $g_{1}, g_{2}, \ldots, g_{T}$, with $g_{t} \in[-1,1]$ for all $t$, guarantees

$$
\begin{equation*}
\text { Reward } \geq \kappa \exp \left(\gamma\left|g_{1: T}\right|\right)-\epsilon \tag{1}
\end{equation*}
$$

where $\gamma, \kappa>0$ and $\epsilon \geq 0$ are constants. Then, against any comparator $\dot{x} \in[-R, R]$, we have

$$
\begin{equation*}
\operatorname{Regret}(\dot{x}) \leq \frac{R}{\gamma}\left(\log \left(\frac{R}{\kappa \gamma}\right)-1\right)+\epsilon \tag{2}
\end{equation*}
$$

letting $0 \log 0=0$ when $R=0$. Further, any algorithm with the regret guarantee of Eq. (2) must guarantee the reward of Eq. (1).

We give a proof of this theorem in the appendix. The duality between reward and regret can also be seen as a consequence of the fact that $\exp (x)$ and $y \log y-y$ are convex conjugates. The $\gamma$ term typically contains a dependence on $T$ like $1 / \sqrt{T}$. This bound holds for all $R$, and so for some small $R$ the $\log$ term becomes negative; however, for real algorithms the $\epsilon$ term will ensure the regret bound remains positive. The minus one can of course be dropped to simplify the bound further.

## 3 Gradient Descent with Increasing Learning Rates

In this section we show that allowing the learning rate of gradient descent to sometimes increase leads to novel theoretical guarantees.

To build intuition, consider online linear optimization in one dimension, with gradients $g_{1}, g_{2}, \ldots, g_{T}$, all in $[-1,1]$. In this setting, the reward of unconstrained gradient descent has a simple closed form:
Lemma 2. Consider unconstrained gradient descent in one dimension, with learning rate $\eta$. On round $t$, this algorithm plays the point $x_{t}=\eta g_{1: t-1}$. Letting $G=\left|g_{1: t}\right|$ and $H=\sum_{t=1}^{T} g_{t}^{2}$, the cumulative reward of the algorithm is exactly

$$
\text { Reward }=\frac{\eta}{2}\left(G^{2}-H\right)
$$

We give a simple direct proof in Appendix A. Perhaps surprisingly, this result implies that the reward is totally independent of the order of the linear functions selected by the adversary. Examining the expression in Lemma 2, we see that the optimal choice of learning rate $\eta$ depends fundamentally on two quantities: the absolute value of the sum of gradients $(G)$, and the sum of the squared gradients $(H)$. If $G^{2}>H$, we would like to use as large a learning rate as possible in order to maximize reward. In contrast, if $G^{2}<H$, the algorithm will obtain negative reward, and the best it can do is to cut its losses by setting $\eta$ as small as possible.
One of the motivations for this work is the observation that the state-of-the-art online gradient descent algorithms adjust their learning rates based only on the observed value of $H$ (or its upper bound $T$ ); for example $[4,10]$. We would like to increase reward by also accounting for $G$. But unlike $H$, which is monotonically increasing with time, $G$ can both increase and decrease. This makes simple guess-and-doubling tricks fail when applied to $G$, and necessitates a more careful approach.

### 3.1 Analysis in One Dimension

In this section we analyze algorithm REWARD-DoUbLING-1D (Algorithm 1), which consists of a series of epochs. We suppose for the moment that an upper bound $\bar{H}$ on $H=\sum_{t=1}^{T} g_{t}^{2}$ is known in advance. In the first epoch, we run gradient descent with a small initial learning rate $\eta=\eta_{1}$. Whenever the total reward accumulated in the current epoch reaches $\eta \bar{H}$, we double $\eta$ and start a new epoch (returning to the origin and forgetting all previous gradients except the most recent one).
Lemma 3. Applied to a sequence of gradients $g_{1}, g_{2}, \ldots, g_{T}$, all in $[-1,1]$, where $H=\sum_{t=1}^{T} g_{t}^{2} \leq$ $\bar{H}$, REWARD-DoUBLING-1D obtains reward satisfying

$$
\begin{equation*}
\text { Reward }=\sum_{t=1}^{T} x_{t} g_{t} \geq \frac{1}{4} \eta_{1} \bar{H} \exp \left(a \frac{\left|g_{1: T}\right|}{\sqrt{\bar{H}}}\right)-\eta_{1} \bar{H} \tag{3}
\end{equation*}
$$

for $a=\log (2) / \sqrt{3}$.

```
Algorithm 1 REWARD-DOUBLING-1D
    Parameters: initial learning rate \(\eta_{1}\), upper
    bound \(\bar{H} \geq \sum_{t=1}^{T} g_{t}^{2}\).
    Initialize \(x_{1} \leftarrow 0, i \leftarrow 1\), and \(Q_{1} \leftarrow 0\).
    for \(t=1,2, \ldots, T\) do
        Play \(x_{t}\), and receive reward \(x_{t} g_{t}\).
        \(Q_{i} \leftarrow Q_{i}+x_{t} g_{t}\).
        if \(Q_{i}<\eta_{i} \bar{H}\) then
            \(x_{t+1} \leftarrow x_{t}+\eta_{i} g_{t}\).
        else
            \(i \leftarrow i+1\).
            \(\eta_{i} \leftarrow 2 \eta_{i-1} ; Q_{i} \leftarrow 0\).
            \(x_{t+1} \leftarrow 0+\eta_{i} g_{t}\).
```

```
Algorithm 2 REWARD-DOUBLING
    Parameters: maximum origin-regret \(\epsilon_{i}\)
    for \(1 \leq i \leq n\).
    for \(i=1,2, \ldots, n\) do
        Let \(A_{i}\) be a copy of algorithm
        REWARD-DoUBLING-1D-GUESS
        (see Theorem 4), with parameter \(\epsilon_{i}\).
    for \(t=1,2, \ldots, T\) do
        Play \(x_{t}\), with \(x_{t, i}\) selected by \(A_{i}\).
        Receive gradient vector \(g_{t}=-\nabla f_{t}\left(x_{t}\right)\).
        for \(i=1,2, \ldots, n\) do
            Feed back \(g_{t, i}\) to \(A_{i}\).
```

Proof. Suppose round $T$ occurs during the $k$ 'th epoch. Because epoch $i$ can only come to an end if $Q_{i} \geq \eta_{i} \bar{H}$, where $\eta_{i}=2^{i-1} \eta_{1}$, we have

$$
\begin{equation*}
\text { Reward }=\sum_{i=1}^{k} Q_{i} \geq\left(\sum_{i=1}^{k-1} 2^{i-1} \eta_{1} \bar{H}\right)+Q_{k}=\left(2^{k-1}-1\right) \eta_{1} \bar{H}+Q_{k} \tag{4}
\end{equation*}
$$

We now lower bound $Q_{k}$. For $i=1, \ldots, k$ let $t_{i}$ denote the round on which $Q_{i}$ is initialized to 0 , with $t_{1} \equiv 1$, and define $t_{k+1} \equiv T$. By construction, $Q_{i}$ is the total reward of a gradient descent algorithm that is active on rounds $t_{i}$ through $t_{i+1}$ inclusive, and that uses learning rate $\eta_{i}$ (note that on round $t_{i}$, this algorithm gets 0 reward and we initialize $Q_{i}$ to 0 on that round). Thus, by Lemma 2 , we have that for any $i$,

$$
Q_{i}=\frac{\eta_{i}}{2}\left(\left(g_{t_{i}: t_{i+1}}\right)^{2}-\sum_{s=t_{i}}^{t_{i+1}} g_{s}^{2}\right) \geq-\frac{\eta_{i}}{2} \bar{H}
$$

Applying this bound to epoch $k$, we have $Q_{k} \geq-\frac{1}{2} \eta_{k} \bar{H}=-2^{k-2} \eta_{1} \bar{H}$. Substituting into (4) gives

$$
\begin{equation*}
\text { Reward } \geq \eta_{1} \bar{H}\left(2^{k-1}-1-2^{k-2}\right)=\eta_{1} \bar{H}\left(2^{k-2}-1\right) \tag{5}
\end{equation*}
$$

We now show that $k \geq \frac{\left|g_{1: T}\right|}{\sqrt{3 \bar{H}}}$. At the end of round $t_{i+1}-1$, we must have had $Q_{i}<\eta_{i} \bar{H}$ (otherwise epoch $i+1$ would have begun earlier). Thus, again using Lemma 2,

$$
\frac{\eta_{i}}{2}\left(\left(g_{t_{i}: t_{i+1}-1}\right)^{2}-\bar{H}\right) \leq \eta_{i} \bar{H}
$$

so $\left|g_{t_{i}: t_{i+1}-1}\right| \leq \sqrt{3 \bar{H}}$. Thus,

$$
\left|g_{1: T}\right| \leq \sum_{i=1}^{k}\left|g_{t_{i}: t_{i+1}-1}\right| \leq k \sqrt{3 \bar{H}}
$$

Rearranging gives $k \geq \frac{\left|g_{1: T}\right|}{\sqrt{3 \bar{H}}}$, and combining with Eq. (5) proves the lemma.
We can now apply Theorem 1 to the reward (given by Eq. (3)) of REWARD-Doubling-1D to show

$$
\begin{equation*}
\operatorname{Regret}(\stackrel{\circ}{x}) \leq b R \sqrt{\bar{H}}\left(\log \left(\frac{4 R b \sqrt{\bar{H}}}{\eta_{1}}\right)-1\right)+\eta_{1} \bar{H} \tag{6}
\end{equation*}
$$

for any $\stackrel{\circ}{x} \in[-R, R]$, where $b=a^{-1}=\sqrt{3} / \log (2)<2.5$. When the feasible set is also fixed in advance, online gradient descent with a fixed learning obtains a regret bound of $\mathcal{O}(R \sqrt{T})$. Suppose we use the estimate $\bar{H}=T$. By choosing $\eta_{1}=\frac{1}{T}$, we guarantee constant regret against the origin, $\stackrel{\circ}{x}=0$ (equivalently, constant total loss). Further, for any feasible set of radius $R$, we still have
worst-case regret of at most $\mathcal{O}(R \sqrt{T} \log ((1+R) T))$, which is only modestly worse than that of gradient descent with the optimal $R$ known in advance.
The need for an upper bound $\bar{H}$ can be removed using a standard guess-and-doubling approach, at the cost of a constant factor increase in regret (see appendix for proof).
Theorem 4. Consider algorithm REWARD-Doubling-1D-Guess, which behaves as follows. On each era $i$, the algorithm runs REWARD-DOUBLING-1D with an upper bound of $\bar{H}_{i}=2^{i-1}$, and initial learning rate $\eta_{1}^{i}=\epsilon 2^{-2 i}$. An era ends when $\bar{H}_{i}$ is no longer an upper bound on the sum of squared gradients seen during that era. Letting $c=\frac{\sqrt{2}}{\sqrt{2}-1}$, this algorithm has regret at most

$$
\text { Regret } \leq c R \sqrt{H+1}\left(\log \left(\frac{R}{\epsilon}(2 H+2)^{5 / 2}\right)-1\right)+\epsilon
$$

### 3.2 Extension to $n$ dimensions

To extend our results to general online convex optimization, it is sufficient to run a separate copy of Reward-Doubling-1D-Guess for each coordinate, as is done in Reward-Doubling (Algorithm 2). The key to the analysis of this algorithm is that overall regret is simply the sum of regret on $n$ one-dimensional subproblems which can be analyzed independently.
Theorem 5. Given a sequence of convex loss functions $f_{1}, f_{2}, \ldots, f_{T}$ from $\mathbb{R}^{n}$ to $\mathbb{R}$, REWARD-DOUBLING with $\epsilon_{i}=\frac{\epsilon}{n}$ has regret bounded by

$$
\begin{aligned}
\operatorname{Regret}(\stackrel{\circ}{x}) & \leq \epsilon+c \sum_{i=1}^{n}\left|\grave{x}_{i}\right| \sqrt{H_{i}+1}\left(\log \left(\frac{n}{\epsilon}\left|\grave{x}_{i}\right|\left(2 H_{i}+2\right)^{5 / 2}\right)-1\right) \\
& \leq \epsilon+c\|\grave{x}\|_{2} \sqrt{H+n}\left(\log \left(\frac{n}{\epsilon}\|\grave{x}\|_{2}^{2}(2 H+2)^{5 / 2}\right)-1\right)
\end{aligned}
$$

for $c=\frac{\sqrt{2}}{\sqrt{2}-1}$, where $H_{i}=\sum_{t=1}^{T} g_{t, i}^{2}$ and $H=\sum_{t=1}^{T}\left\|g_{t}\right\|_{2}^{2}$.
Proof. Fix a comparator $\stackrel{\dot{x}}{ }$. For any coordinate $i$, define

$$
\operatorname{Regret}_{i}=\sum_{t=1}^{T} \stackrel{\circ}{x}_{i} g_{t, i}-\sum_{t=1}^{T} x_{t, i} g_{t, i}
$$

Observe that

$$
\sum_{i=1}^{n} \operatorname{Regret}_{i}=\sum_{t=1}^{T} \stackrel{\circ}{x} \cdot g_{t}-\sum_{t=1}^{T} x_{t} \cdot g_{t}=\operatorname{Regret}(\stackrel{\circ}{x})
$$

Furthermore, Regret $_{i}$ is simply the regret of REWARD-DoUbling-1D-GUESS on the gradient sequence $g_{1, i}, g_{2, i}, \ldots, g_{T, i}$. Applying the bound of Theorem 4 to each Regret ${ }_{i}$ term completes the proof of the first inequality. For the second inequality, let $\vec{H}$ be a vector whose $i^{\text {th }}$ component is $\sqrt{H_{i}+1}$, and let $\vec{x} \in \mathbb{R}^{n}$ where $\vec{x}_{i}=\left|\grave{x}_{i}\right|$. Using the Cauchy-Schwarz inequality, we have

$$
\sum_{i=1}^{n}\left|\grave{x}_{i}\right| \sqrt{H_{i}+1}=\vec{x} \cdot \vec{H} \leq\|\grave{x}\|_{2}\|\vec{H}\|_{2}=\| \overbrace{2} \sqrt{H+n}
$$

This, together with the fact that $\log \left(\left|\dot{x}_{i}\right|\left(2 H_{i}+2\right)^{5 / 2}\right) \leq \log \left(\|i\|_{2}^{2}(2 H+2)^{5 / 2}\right)$, suffices to prove second inequality.

In some applications, $n$ is not known in advance. In this case, we can set $\epsilon_{i}=\frac{\epsilon}{i^{2}}$ for the $i$ th coordinate we encounter, and get the same bound up to constant factors.

## 4 An Epoch-Free Algorithm

In this section we analyze Smóth-Reward-Doubling, a simple algorithm that achieves bounds comparable to those of Theorem 4, without guessing-and-doubling. We consider only the 1-d problem, as the technique of Theorem 5 can be applied to extend to $n$ dimensions. Given a parameter
$\eta>0$, we achieve

$$
\begin{equation*}
\text { Regret } \leq R \sqrt{T}\left(\log \left(\frac{R T^{3 / 2}}{\eta}\right)-1\right)+1.76 \eta \tag{7}
\end{equation*}
$$

for all $T$ and $R$, which is better (by constant factors) than Theorem 4 when $g_{t} \in\{-1,1\}$ (which implies $T=H$ ). The bound can be worse on a problems where $H<T$.
The idea of the algorithm is to maintain the invariant that our cumulative reward, as a function of $g_{1: t}$ and $t$, satisfies Reward $\geq N\left(g_{1: t}, t\right)$, for some fixed function $N$. Because reward changes by $g_{t} x_{t}$ on round $t$, it suffices to guarantee that for any $g \in[-1,1]$,

$$
\begin{equation*}
N\left(g_{1: t}, t\right)+g x_{t+1} \geq N\left(g_{1: t}+g, t+1\right) \tag{8}
\end{equation*}
$$

where $x_{t+1}$ is the point the algorithm plays on round $t+1$, and we assume $N(0,1)=0$.
This inequality is approximately satisfied (for small $g$ ) if we choose

$$
x_{t+1}=\frac{\partial N\left(g_{1: t}+g, t\right)}{\partial g} \approx \frac{N\left(g_{1: t}+g, t\right)-N\left(g_{1: t}, t\right)}{g} \approx \frac{N\left(g_{1: t}+g, t+1\right)-N\left(g_{1: t}, t\right)}{g}
$$

This suggests that if we want to maintain reward at least $N\left(g_{1: t}, t\right)=\frac{1}{t}\left(\exp \left(\left|g_{1: t}\right| / \sqrt{t}\right)-1\right)$, we should set $x_{t+1} \approx \operatorname{sign}\left(g_{1: t}\right) t^{-3 / 2} \exp \left(\frac{\left|g_{1: t}\right|}{\sqrt{t}}\right)$. The following theorem (proved in the appendix) provides an inductive analysis of an algorithm of this form.
Theorem 6. Fix a sequence of reward functions $f_{t}(x)=g_{t} x$ with $g_{t} \in[-1,1]$, and let $G_{t}=\left|g_{1: t}\right|$. We consider Smooth-Reward-Doubling, which plays 0 on round 1 and whenever $G_{t}=0$; otherwise, it plays

$$
\begin{equation*}
x_{t+1}=\eta \operatorname{sign}\left(g_{1: t}\right) B\left(G_{t}, t+5\right) \tag{9}
\end{equation*}
$$

with $\eta>0$ a learning-rate parameter and

$$
\begin{equation*}
B(G, t)=\frac{1}{t^{3 / 2}} \exp \left(\frac{G}{\sqrt{t}}\right) \tag{10}
\end{equation*}
$$

Then, at the end of each round $t$, this algorithm has

$$
\operatorname{Reward}(t) \geq \eta \frac{1}{t+5} \exp \left(\frac{G_{t}}{\sqrt{t+5}}\right)-1.76 \eta
$$

Two main technical challenges arise in the proof: first, we prove a result like Eq. (8) for $N\left(g_{1: t}, t\right)=$ $(1 / t) \exp \left(\left|g_{1: t}\right| / \sqrt{t}\right)$. However, this Lemma only holds for $t \geq 6$ and when the sign of $g_{1: t}$ doesn't change. We account for this by showing that a small modification to $N$ (costing only a constant over all rounds) suffices.

By running this algorithm independently for each coordinate using an appropriate choice of $\eta$, one can obtain a guarantee similar to that of Theorem 5.

## 5 Lower Bounds

As with our previous results, it is sufficient to show a lower bound in one dimension, as it can then be replicated independently in each coordinate to obtain an $n$ dimensional bound. Note that our lower bound contains the factor $\log (|\dot{x}| \sqrt{T})$, which can be negative when $\stackrel{\circ}{x}$ is small relative to $T$, hence it is important to hold $\grave{x}$ fixed and consider the behavior as $T \rightarrow \infty$. Here we give only a proof sketch; see Appendix A for the full proof.
Theorem 7. Consider the problem of unconstrained online linear optimization in one dimension, and an online algorithm that guarantees origin-regret at most $\epsilon$. Then, for any fixed comparator $\dot{x}$, and any integer $T_{0}$, there exists a gradient sequence $\left\{g_{t}\right\} \in[-1,1]^{T}$ of length $T \geq T_{0}$ for which the algorithm's regret satisfies

$$
\operatorname{Regret}(\stackrel{\circ}{x}) \geq 0.336|\dot{x}| \sqrt{T \log \left(\frac{\mid \stackrel{x}{x} \sqrt{T}}{\epsilon}\right)} .
$$

Proof. (Sketch) Assume without loss of generality that $\dot{x}>0$. Let $Q$ be the algorithm's reward when each $g_{t}$ is drawn independently uniformly from $\{-1,1\}$. We have $\mathrm{E}[Q]=0$, and because the algorithm guarantees origin-regret at most $\epsilon$, we have $Q \geq-\epsilon$ with probability 1 . Letting $G=g_{1: T}$, it follows that for any threshold $Z=Z(T)$,

$$
\begin{aligned}
0 & =\mathrm{E}[Q] \\
& =\mathrm{E}[Q \mid G<Z] \cdot \operatorname{Pr}[G<Z]+\mathrm{E}[Q \mid G \geq Z] \cdot \operatorname{Pr}[G \geq Z] \\
& \geq-\epsilon \operatorname{Pr}[G<Z]+\mathrm{E}[Q \mid G \geq Z] \cdot \operatorname{Pr}[G \geq Z] \\
& >-\epsilon+\mathrm{E}[Q \mid G \geq Z] \cdot \operatorname{Pr}[G \geq Z] .
\end{aligned}
$$

Equivalently,

$$
\mathrm{E}[Q \mid G \geq Z]<\frac{\epsilon}{\operatorname{Pr}[G \geq Z]}
$$

We choose $Z(T)=\sqrt{k T}$, where $k=\left\lfloor\log \left(\frac{R \sqrt{T}}{\epsilon}\right) / \log \left(p^{-1}\right)\right\rfloor$. Here $R=|\dot{x}|$ and $p>0$ is a constant chosen using binomial distribution lower bounds so that $\operatorname{Pr}[G \geq Z] \geq p^{k}$. This implies

$$
\mathrm{E}[Q \mid G \geq Z]<\epsilon p^{-k}=\epsilon \exp \left(k \log p^{-1}\right) \leq R \sqrt{T}
$$

This implies there exists a sequence with $G \geq Z$ and $Q<R \sqrt{T}$. On this sequence, regret is at least $G \dot{x}-Q \geq R \sqrt{k T}-R \sqrt{T}=\Omega(R \sqrt{k T})$.

Theorem 8. Consider the problem of unconstrained online linear optimization in $\mathbb{R}^{n}$, and consider an online algorithm that guarantees origin-regret at most $\epsilon$. For any radius $R$, and any $T_{0}$, there exists a gradient sequence gradient sequence $\left\{g_{t}\right\} \in\left([-1,1]^{n}\right)^{T}$ of length $T \geq T_{0}$, and a comparator $\stackrel{\circ}{x}$ with $\|\stackrel{\circ}{x}\|_{1}=R$, for which the algorithm's regret satisfies

$$
\operatorname{Regret}(\stackrel{\circ}{x}) \geq 0.336 \sum_{i=1}^{n}\left|\dot{x}_{i}\right| \sqrt{T \log \left(\frac{\left|\dot{\circ}_{i}\right| \sqrt{T}}{\epsilon}\right)} .
$$

Proof. For each coordinate $i$, Theorem 7 implies that there exists a $T \geq T_{0}$ and a sequence of gradients $g_{t, i}$ such that

$$
\sum_{t=1}^{T} \stackrel{\circ}{x}_{i} g_{t, i}-\sum_{t=1}^{T} x_{t, i} g_{t, i} \geq 0.336\left|\stackrel{\circ}{x}_{i}\right| \sqrt{T \log \left(\frac{\left|\grave{x}_{i}\right| \sqrt{T}}{\epsilon}\right)}
$$

(The proof of Theorem 7 makes it clear that we can use the same $T$ for all $i$.) Summing this inequality across all $n$ coordinates then gives the regret bound stated in the theorem.

The following theorem presents a stronger negative result for Follow-the-Regularized-Leader algorithms with a fixed regularizer: for any such algorithm that guarantees origin-regret at most $\epsilon_{T}$ after $T$ rounds, worst-case regret with respect to any point outside $\left[-\epsilon_{T}, \epsilon_{T}\right]$ grows linearly with $T$.
Theorem 9. Consider a Follow-The-Regularized-Leader algorithm that sets

$$
x_{t}=\underset{x}{\arg \min }\left(g_{1: t-1} x+\psi_{T}(x)\right)
$$

where $\psi_{T}$ is a convex, non-negative function with $\psi_{T}(0)=0$. Let $\epsilon_{T}$ be the maximum origin-regret incurred by the algorithm on a sequence of $T$ gradients. Then, for any $\dot{x}$ with $|\dot{x}|>\epsilon_{T}$, there exists a sequence of $T$ gradients such that the algorithm's regret with respect to $\stackrel{\circ}{x}$ is at least $\frac{T-1}{2}\left(|\dot{x}|-\epsilon_{T}\right)$.

In fact, it is clear from the proof that the above result holds for any algorithm that selects $x_{t+1}$ purely as a function of $g_{1: t}$ (in particular, with no dependence on $t$ ).

## 6 Future Work

This work leaves open many interesting questions. It should be possible to apply our techniques to problems that do have constrained feasible sets; for example, it is natural to consider the unconstrained experts problem on the positive orthant. While we believe this extension is straightforward, handling arbitrary non-axis-aligned constraints will be more difficult. Another possibility is to develop an algorithm with bounds in terms of $H$ rather than $T$ that doesn't use a guess and double approach.

## References

[1] Jacob Abernethy, Peter L. Bartlett, Alexander Rakhlin, and Ambuj Tewari. Optimal strategies and minimax lower bounds for online convex games. In COLT, 2008.
[2] Amit Agarwal, Elad Hazan, Satyen Kale, and Robert E. Schapire. Algorithms for portfolio management based on the Newton method. In ICML, 2006.
[3] Nicolò Cesa-Bianchi and Gabor Lugosi. Prediction, Learning, and Games. Cambridge University Press, New York, NY, USA, 2006. ISBN 0521841089.
[4] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. In COLT, 2010.
[5] Elad Hazan and Satyen Kale. Extracting certainty from uncertainty: Regret bounded by variation in costs. In COLT, 2008.
[6] Elad Hazan and Satyen Kale. On stochastic and worst-case models for investing. In Advances in Neural Information Processing Systems 22. 2009.
[7] Robert A. Jacobs. Increased rates of convergence through learning rate adaptation. Neural Networks, 1987.
[8] Jyrki Kivinen and Manfred Warmuth. Exponentiated Gradient Versus Gradient Descent for Linear Predictors. Journal of Information and Computation, 132, 1997.
[9] Todd K. Leen and Genevieve B. Orr. Optimal stochastic search and adaptive momentum. In NIPS, 1993.
[10] H. Brendan McMahan and Matthew Streeter. Adaptive bound optimization for online convex optimization. In COLT, 2010.
[11] Barak Pearlmutter. Gradient descent: Second order momentum and saturating error. In NIPS, 1991.
[12] Shai Shalev-Shwartz. Online learning and online convex optimization. Foundations and Trends in Machine Learning, 4(2):107-194, 2012.
[13] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In ICML, 2003.

## A Proofs

This appendix gives the proofs omitted in the body of the paper, with the corresponding lemmas and theorems restated for convenience.
Theorem 1. Consider an algorithm for one-dimensional online linear optimization that, when run on a sequence of gradients $g_{1}, g_{2}, \ldots, g_{T}$, with $g_{t} \in[-1,1]$ for all $t$, guarantees

$$
\begin{equation*}
\operatorname{Reward} \geq \kappa \exp \left(\gamma\left|g_{1: T}\right|\right)-\epsilon \tag{1}
\end{equation*}
$$

where $\gamma, \kappa>0$ and $\epsilon \geq 0$ are constants. Then, against any comparator $\stackrel{\circ}{x} \in[-R, R]$, we have

$$
\begin{equation*}
\operatorname{Regret}(\dot{x}) \leq \frac{R}{\gamma}\left(\log \left(\frac{R}{\kappa \gamma}\right)-1\right)+\epsilon \tag{2}
\end{equation*}
$$

letting $0 \log 0=0$ when $R=0$. Further, any algorithm with the regret guarantee of Eq. (2) must guarantee the reward of Eq. (1).

Proof. Let $G_{T}=\left|g_{1: T}\right|$. By definition, given the reward guarantee of Eq. (1) we have

$$
\begin{equation*}
\text { Regret } \leq R G_{T}-\kappa \exp \left(\gamma G_{T}\right)+\epsilon \tag{11}
\end{equation*}
$$

If $R=0$, then Eq. (2) follows immediately. Otherwise, note this is a concave function in $G_{T}$, and setting the first derivative equal to zero shows

$$
G^{*}=\frac{1}{\gamma} \log \left(\frac{R}{\gamma \kappa}\right)
$$

maximizes regret (for large enough $R$ we could have $G^{*}>T$, and so this $G^{*}$ is not actually achievable by the adversary, but this is fine for lower bounding regret). Plugging $G^{*}$ into Eq. (11) and simplifying yields the bound of Eq. (2). For the second claim, suppose Eq. (2) holds. Then, again by definition, we must have

$$
\begin{equation*}
\text { Reward } \geq R G-\frac{R}{\gamma} \log \left(\frac{R}{\gamma \kappa}\right)+\frac{R}{\gamma}-\epsilon \tag{12}
\end{equation*}
$$

This bound is a concave function of $R$, and since it holds for any $R \geq 0$ by assumption, we can choose the $R$ that maximizes the bound, namely $R^{*}=\gamma \kappa \exp (\gamma G)$. Note

$$
\frac{R^{*}}{\gamma} \log \left(\frac{R^{*}}{\gamma \kappa}\right)=\frac{R^{*}}{\gamma} \log (\exp (\gamma G))=R^{*} G
$$

and so plugging $R^{*}$ into Eq. (12) yields

$$
\text { Reward } \geq \frac{1}{\gamma} R^{*}-\epsilon=\kappa \exp (\gamma G)-\epsilon
$$

Lemma 2. Consider unconstrained gradient descent in one dimension, with learning rate $\eta$. On round $t$, this algorithm plays the point $x_{t}=\eta g_{1: t-1}$. Letting $G=\left|g_{1: t}\right|$ and $H=\sum_{t=1}^{T} g_{t}^{2}$, the cumulative reward of the algorithm is exactly

$$
\text { Reward }=\frac{\eta}{2}\left(G^{2}-H\right)
$$

Proof. The algorithm's cumulative reward after $T$ rounds is

$$
\begin{equation*}
\sum_{t=1}^{T} x_{t} g_{t}=\sum_{t=1}^{T} g_{t} \eta g_{1: t-1}=\frac{\eta}{2}\left(\left(g_{1: T}\right)^{2}-\sum_{t=1}^{T} g_{t}^{2}\right) \tag{13}
\end{equation*}
$$

To verify the second equality, note that $\left(g_{1: T}\right)^{2}-\left(g_{1: T-1}\right)^{2}=g_{T}^{2}+2 g_{T}\left(g_{1: T-1}\right)$, so on round $T$ the right hand side increases by $\eta g_{T}\left(g_{1: T-1}\right)$, as does the left hand side. The equality then follows by induction on $T$.

It is worth noting that the standard $R \sqrt{T}$ bound can be derived from the above result fairly easily. We have

$$
\begin{aligned}
\text { Regret } & \leq R G-\frac{\eta}{2}\left(G^{2}-H\right) \\
& \leq \frac{\eta}{2} H+\max _{G}\left(R G-\frac{\eta G^{2}}{2}\right) \\
& \leq \frac{\eta}{2} H+\frac{R^{2}}{2 \eta}
\end{aligned}
$$

where the max is achieved by taking $G=R / \eta$. Taking $\eta=R / \sqrt{T}$ then gives the standard bound. However, this bound significantly underestimates the performance of constant-learning-rate gradient descent when $G$ is large. This is in contrast to our regret bounds, which are always tight with respect to their matching reward bounds.
Theorem 4. Consider algorithm REWARD-Doubling-1D-Guess, which behaves as follows. On each era $i$, the algorithm runs REWARD-DoUBLING-1D with an upper bound of $\bar{H}_{i}=2^{i-1}$, and initial learning rate $\eta_{1}^{i}=\epsilon 2^{-2 i}$. An era ends when $\bar{H}_{i}$ is no longer an upper bound on the sum of squared gradients seen during that era. Letting $c=\frac{\sqrt{2}}{\sqrt{2}-1}$, this algorithm has regret at most

$$
\text { Regret } \leq c R \sqrt{H+1}\left(\log \left(\frac{R}{\epsilon}(2 H+2)^{5 / 2}\right)-1\right)+\epsilon
$$

Proof. Suppose round $T$ occurs in era $k$, and let $t_{i}$ be the round on which era $i$ starts, with $t_{k+1} \equiv$ $T+1$. Define $H_{i}=\sum_{s=t_{i}}^{t_{i+1}-1} g_{s}^{2}$. To prove the theorem we will need several inequalities. First, note that $H=\sum_{i=1}^{k} H_{i} \geq \sum_{i=1}^{k-1} \bar{H}_{i}=2^{k-1}-1$, or $2^{k-1} \leq H+1$. Thus,

$$
\sum_{i=1}^{k} \sqrt{\bar{H}_{i}}=\sum_{i=0}^{k-1} \sqrt{2^{i}}=\frac{\sqrt{2}^{k}-1}{\sqrt{2}-1} \leq \frac{\sqrt{2^{k}}}{\sqrt{2}-1} \leq \frac{\sqrt{2(H+1)}}{\sqrt{2}-1}=c \sqrt{H+1}
$$

Next, note that for any $i$ we have

$$
\frac{\sqrt{\bar{H}_{i}}}{\eta_{1}^{i}}=\frac{1}{\epsilon} 2^{\frac{i-1}{2}+2 i} \leq \frac{1}{\epsilon} 2^{2.5 k} \leq \frac{1}{\epsilon}(2(H+1))^{(5 / 2)}
$$

Note that the bound of Lemma 3 applies for all $T$ where $H \leq \bar{H}$, and thus so does Eq. (6). Thus, we can apply this bound to the regret in era $k$ on rounds $t_{k}$ through $T$, as well as on the regret in each earlier era. Then, total regret with respect to the best point in $[-R, R]$ is at most the sum of the regret in each era, so

$$
\begin{aligned}
\text { Regret } & \leq \sum_{i=1}^{k} R \sqrt{\bar{H}_{i}}\left(\log \left(\frac{R \sqrt{\bar{H}_{i}}}{\eta_{1}^{i}}\right)-1\right)+\eta_{1}^{i} H_{i} \\
& \leq \sum_{i=1}^{k} R \sqrt{\bar{H}_{i}}\left(\log \left(\frac{R}{\epsilon}(2 H+2)^{5 / 2}\right)-1\right)+\eta_{1}^{i} H_{i} \\
& \leq c R \sqrt{H+1}\left(\log \left(\frac{R}{\epsilon}(2 H+2)^{5 / 2}\right)-1\right)+\sum_{i=1}^{k} \eta_{1}^{i} H_{i}
\end{aligned}
$$

Finally, because $H_{i} \leq \bar{H}_{i}+1 \leq 2 \bar{H}_{i}=2^{i}$, we have $\sum_{i=1}^{k} \eta_{1}^{i} H_{i} \leq \sum_{i=1}^{k} \epsilon 2^{-i} \leq \epsilon$, which completes the proof.
Theorem 6. Fix a sequence of reward functions $f_{t}(x)=g_{t} x$ with $g_{t} \in[-1,1]$, and let $G_{t}=\left|g_{1: t}\right|$. We consider Smooth-Reward-Doubling, which plays 0 on round 1 and whenever $G_{t}=0$; otherwise, it plays

$$
\begin{equation*}
x_{t+1}=\eta \operatorname{sign}\left(g_{1: t}\right) B\left(G_{t}, t+5\right) \tag{9}
\end{equation*}
$$

with $\eta>0$ a learning-rate parameter and

$$
\begin{equation*}
B(G, t)=\frac{1}{t^{3 / 2}} \exp \left(\frac{G}{\sqrt{t}}\right) . \tag{10}
\end{equation*}
$$

Then, at the end of each round $t$, this algorithm has

$$
\operatorname{Reward}(t) \geq \eta \frac{1}{t+5} \exp \left(\frac{G_{t}}{\sqrt{t+5}}\right)-1.76 \eta
$$

Proof. We present a proof for the case where $\eta=1$; since $\eta$ simply scales all of the $x_{t}$ played by the algorithm (and hence, reward), the result for general $\eta$ follows immediately. We use the minimum reward function

$$
\begin{equation*}
N(G, t)=\frac{1}{t} \exp \left(\frac{G}{\sqrt{t}}\right) \tag{14}
\end{equation*}
$$

The proof will be by induction on $t$, with the induction hypothesis that the cumulative reward of the algorithm at the end of round $t$ satisfies

$$
\begin{equation*}
\operatorname{Reward}(t) \geq N\left(G_{t}, t+5\right)-\epsilon_{1: t}, \tag{15}
\end{equation*}
$$

where $\epsilon_{1}=N(1,6)$ and for $t>1, \epsilon_{t+1}=\tilde{\epsilon}(t+5)$ with

$$
\tilde{\epsilon}(\tau)=\frac{1}{\tau+1} \exp \left(\frac{1}{\sqrt{\tau+1}}\right)-\frac{1}{\tau}+\frac{1}{\tau^{3 / 2}}
$$

We will then show that the sum of $\epsilon_{t}$ 's is always bounded by a constant.
For the base case, $t=1$, we play $x=0$ so end the round with zero reward, while the RHS of Eq. (15) is $N\left(\left|g_{1}\right|, 6\right)-N(1,6) \leq 0$.
Now, suppose the induction hypothesis holds at the end of some round $t \geq 1$. Without loss of generality, suppose $g_{1: t} \geq 0$ so $G_{t}=g_{1: t}$. We consider two cases. First, suppose $G_{t}>0$ and $G_{t}+g_{t+1}>0\left(\right.$ so $\left.g_{t+1}>-G_{t}\right)$. In this case, $g_{1: t}$ does not change sign when we add $g_{t+1}$; thus, an invariant like that of Eq. (8) is sufficient; we prove such a result in Lemma 10 (given below). More precisely, we play $x_{t+1}$ according to Eq. (9), and

$$
\begin{aligned}
\text { Reward }(t+1) & \geq N\left(G_{t}, t+5\right)-\epsilon_{1: t}+g_{t+1} x_{t+1} & & \text { IH and update rule } \\
& \geq N\left(G_{t}+g_{t+1}, t+5+1\right)-\epsilon_{1: t} & & \text { Lemma } 10 \text { with } \tau=t+5 . \\
& \geq N\left(G_{t+1}, t+5+1\right)-\epsilon_{1: t+1}, & & \text { since } \epsilon_{t+1}>0 .
\end{aligned}
$$

For the remaining case, we have $G_{t}+g_{t+1} \leq 0$, implying $g_{t+1} \leq-G_{t} \leq 0$. In this case, we suffer some loss and arrive at $G_{t+1}=\left|G_{t}+g_{t+1}\right|=-g_{t+1}-G_{t}$. Lemma 11 (below) provides the key bound on the additional loss when the sign of $g_{1: t}$ changes. If $G_{t}>0$, we have

$$
\begin{aligned}
\operatorname{Reward}(t+1) & \geq N\left(G_{t}, t+5\right)-\epsilon_{1: t}+g_{t+1} x_{t+1} & & \text { IH and update rule } \\
& \geq N\left(-g_{t+1}-G_{t}, t+5+1\right)-\epsilon_{1: t+1} & & \text { Lemma 11 with } \tau=t+5 \\
& =N\left(G_{t+1}, t+5+1\right)-\epsilon_{1: t+1} . & &
\end{aligned}
$$

If $G_{t}=0$, we can take $g_{t+1}$ non-positive without loss of generality, and playing $x_{t+1}=0$ is no worse than playing $B(0, t+5)$, and so we conclude Eq. (15) holds for all $t$. Finally,

$$
\sum_{t=2}^{\infty} \epsilon_{t} \leq \int_{\tau=6}^{\infty} \tilde{\epsilon}(\tau)=\sqrt{\frac{2}{3}}-2 \gamma+2 \operatorname{Ei}\left(\frac{1}{\sqrt{7}}\right)+\log (6) \leq 1.50
$$

where $\gamma$ is the Euler gamma constant and Ei is the exponential integral. The upper bound can be found easily using numerical methods. Adding $\epsilon_{1}=\exp (1 / \sqrt{6}) / 6 \leq 0.26$ gives $\epsilon_{1: T} \leq 1.76$ for any $T$.
Lemma 10. Let $G>0$ and $\tau \geq 6$. Then, for any $g \in[-1,1]$ such that $G+g \geq 0$,

$$
N(G, \tau)+g B(G, \tau)-N(G+g, \tau+1) \geq 0
$$

where $N$ is defined by Eq. (14) and $B$ is defined by Eq. (10).
Proof. We need to show

$$
\frac{1}{\tau} \exp \left(\frac{G}{\sqrt{\tau}}\right)+\frac{g}{\tau^{3 / 2}} \exp \left(\frac{G}{\sqrt{\tau}}\right)-\frac{1}{\tau+1} \exp \left(\frac{G+g}{\sqrt{\tau+1}}\right) \geq 0
$$

or equivalently, multiplying by $\tau^{3 / 2}(1+\tau) / \exp (G / \sqrt{\tau}) \geq 0$,

$$
\Delta=\sqrt{\tau}(1+\tau)+g(1+\tau)-\tau^{3 / 2} \exp \left(\frac{G+g}{\sqrt{\tau+1}}-\frac{G}{\sqrt{\tau}}\right) \geq 0
$$

Since $\tau+1 \geq \tau$, the exp term is maximized when $G=0$, so

$$
\begin{equation*}
\Delta \geq(g+\sqrt{\tau})(1+\tau)-\tau^{3 / 2} \exp \left(\frac{g}{\sqrt{\tau+1}}\right) \tag{16}
\end{equation*}
$$

Now, we consider the cases where $g \geq 0$ and $g<0$ separately. First, suppose $g>0$, so $g / \sqrt{\tau+1} \in$ $[0,1]$, and we can use the inequality $\exp (x) \leq 1+x+x^{2}$ for $x \in[0,1]$, which gives

$$
\begin{aligned}
\Delta & \geq g+g \tau+\sqrt{\tau}+\tau^{3 / 2}-\tau^{3 / 2}\left(1+\frac{g}{\sqrt{\tau+1}}+\frac{g^{2}}{\tau+1}\right) \\
& \geq g+g \tau+\sqrt{\tau}+\tau^{3 / 2}-\tau^{3 / 2}\left(1+\frac{g}{\sqrt{\tau}}+\frac{1}{\tau}\right) \\
& =g+g \tau+\sqrt{\tau}+\tau^{3 / 2}-\tau^{3 / 2}-g \tau-\sqrt{\tau} \\
& =g>0
\end{aligned}
$$

Now, we consider the case where $g<0$. In order to show $\Delta \geq 0$ in this case, we need a tight upper bound on $\exp (y)$ for $y \in[-1,0]$. To derive one, we note that for $x \geq 0, \exp (x) \geq 1+x+\frac{1}{2} x^{2}$ from the series representation of $e^{x}$, and so $\exp (-x) \leq\left(1+x+\frac{1}{2} x^{2}\right)^{-1}$. Thus, for $y \in[-1,0]$ we have $\exp (y) \leq\left(1-y+\frac{1}{2} y^{2}\right)^{-1}=Q(y)$. Then, starting from Eq. (16),

$$
\Delta \geq(g+\sqrt{\tau})(1+\tau)-\tau^{3 / 2} Q\left(\frac{g}{\sqrt{\tau+1}}\right)
$$

Let $\Delta_{2}=\Delta Q\left(\frac{g}{\sqrt{\tau+1}}\right)^{-1}$. Because $\Delta_{2}$ and $\Delta$ have the same sign, it suffices to show $\Delta_{2} \geq 0$. We have

$$
\begin{aligned}
\Delta_{2} & =\left(1-\frac{g}{\sqrt{\tau+1}}+\frac{g^{2}}{2(t+1)}\right)(g+\sqrt{\tau})(1+\tau)-\tau^{3 / 2} \\
& =\left(1+\tau-g \sqrt{\tau+1}+\frac{1}{2} g^{2}\right)(g+\sqrt{\tau})-\tau^{3 / 2}
\end{aligned}
$$

First, note

$$
\frac{d}{d g} \Delta_{2}=1+\frac{3 g^{2}}{2}+g \sqrt{\tau}+\tau-2 g \sqrt{1+\tau}-\sqrt{\tau} \sqrt{1+\tau}
$$

Since $g \leq 0$, we have $-2 g \sqrt{\tau+1}+g \sqrt{\tau} \geq 0$, and $(t+1)-\sqrt{\tau} \sqrt{\tau+1} \geq 0$, and so we conclude that $\Delta_{2}$ is increasing in $g$, and so taking $g=-1$ we have

$$
\Delta_{2} \geq\left(\frac{3}{2}+\tau+\sqrt{\tau+1}\right)(-1+\sqrt{\tau})-\tau^{3 / 2}
$$

Taking the derivative with respect to $\tau$ reveals this expression is increasing in $\tau$, and taking $\tau=6$ produces a positive value, proving this case.
Lemma 11. For any $g \in[-1,0]$ and $G \geq 0$ such that $G+g \leq 0$, and any $\tau \geq 1$,

$$
N(G, \tau)+g B(G, \tau) \geq N(-g-G, \tau+1)-\tilde{\epsilon}(\tau)
$$

where $N$ is defined by Eq. (14) and $B$ is defined by Eq. (10), and

$$
\tilde{\epsilon}(\tau) \equiv \frac{1}{\tau+1} \exp \left(\frac{1}{\sqrt{\tau+1}}\right)-\frac{1}{\tau}+\frac{1}{\tau^{3 / 2}} .
$$

Proof. We have

$$
\begin{aligned}
& N(-g-G, \tau+1)-N(G, \tau)-g B(G, \tau) \\
& \quad=\frac{1}{\tau+1} \exp \left(\frac{-g-G}{\sqrt{\tau+1}}\right)-\frac{1}{\tau} \exp \left(\frac{G}{\sqrt{\tau}}\right)-\frac{g}{\tau^{3 / 2}} \exp \left(\frac{G}{\sqrt{\tau}}\right),
\end{aligned}
$$

and since this expression is increasing as $g$ decreases, and $g \geq-1$ in any case,

$$
\leq \frac{1}{\tau+1} \exp \left(\frac{1-G}{\sqrt{\tau+1}}\right)-\frac{1}{\tau} \exp \left(\frac{G}{\sqrt{\tau}}\right)+\frac{1}{\tau^{3 / 2}} \exp \left(\frac{G}{\sqrt{\tau}}\right)
$$

and since $\tau^{3 / 2}>\tau$, taken together the second two terms increase as $G$ decreases, as does the first term, so since $G \geq 0$,

$$
\leq \frac{1}{\tau+1} \exp \left(\frac{1}{\sqrt{\tau+1}}\right)-\frac{1}{\tau}+\frac{1}{\tau^{3 / 2}}=\tilde{\epsilon}(\tau)
$$

and re-arranging proves the lemma.
Theorem 9. Consider a Follow-The-Regularized-Leader algorithm that sets

$$
x_{t}=\underset{x}{\arg \min }\left(g_{1: t-1} x+\psi_{T}(x)\right)
$$

where $\psi_{T}$ is a convex, non-negative function with $\psi_{T}(0)=0$. Let $\epsilon_{T}$ be the maximum origin-regret incurred by the algorithm on a sequence of $T$ gradients. Then, for any $\dot{x}$ with $|\stackrel{\circ}{x}|>\epsilon_{T}$, there exists a sequence of $T$ gradients such that the algorithm's regret with respect to $\stackrel{\circ}{x}$ is at least $\frac{T-1}{2}\left(|\dot{x}|-\epsilon_{T}\right)$.

Proof. For simplicity, we will prove that regret is at least $\frac{T}{2}\left(|\dot{x}|-\epsilon_{T}\right)$ when $T$ is even; if $T$ is odd, we simply take $g_{T}=0$ and consider the first $T-1$ rounds.

Let $T=2 M$. We will consider two gradient sequences. First, suppose $g_{t}=1$ for $t \leq M$, and $g_{t}=-1$ otherwise. Observe that for any $r$, we have $g_{1: M-r}=g_{1: M+r}$, which implies $x_{M-r+1}=$ $x_{M+r+1}$. Thus, the algorithm's total reward is

$$
\begin{aligned}
\sum_{t=1}^{T} x_{t} g_{t} & =\sum_{t=1}^{M} x_{t}-\sum_{t=M+1}^{T} x_{t} \\
& =x_{1}-x_{M+1}+\sum_{r=1}^{M-1} x_{M-r+1}-x_{M+r+1} \\
& =x_{1}-x_{M+1}
\end{aligned}
$$

Because $x_{1}=0$, we get that on this sequence the algorithm has origin-regret $\hat{x} \equiv x_{M+1}$, and so by assumption $\hat{x} \leq \epsilon_{T}$.

Next, suppose $g_{1}=1$ for $t \leq M$, and $g_{t}=0$ otherwise. For this sequence, we will have $x_{t} \leq \hat{x} \leq$ $\epsilon_{T}$ for all $t$, so total reward is at most $M \epsilon_{T}$. For any positive $\dot{x}$ with $\stackrel{\circ}{x}>\epsilon_{T}$, this means that regret with respect to $\dot{x}$ is at least

$$
\stackrel{\circ}{x} M-M \epsilon_{T}=M\left(|\dot{x}|-\epsilon_{T}\right) .
$$

For $\stackrel{\circ}{x}<-\epsilon_{T}$, we can use a similar argument with the sign of the gradients reversed (for both gradient sequences) to get the same bound.

In proving Theorem 7, we will use the following lemma.
Lemma 12. Let $G_{T}=\sum_{i=1}^{T} g_{i}$ be the sum of $T$ random variables, each drawn uniformly from $\{-1,1\}$. Then, for any integer $k$ that is a factor of $T$, we have

$$
\operatorname{Pr}\left[G_{T} \geq \sqrt{k T}\right] \geq p^{k}
$$

where $p=\frac{7}{2^{6}}=0.109375$.
Proof. First, for any $T$ define $p_{T}=\operatorname{Pr}\left[G_{T} \geq \sqrt{T}\right]$, and define

$$
p=\inf _{T \in \mathbb{Z}^{+}} p_{T}
$$

For any $T$, we have $p_{T} \geq 2^{-T}$ trivially, and by the Central Limit Theorem, $\lim _{T \rightarrow \infty} p_{T}=1-$ $\mathcal{N}_{0,1}(1)>0$, where $\mathcal{N}_{0,1}$ is the standard normal cumulative distribution function. It follows that $p>0$, and using numerical methods we find $p=p_{6}=\frac{7}{2^{6}}=0.109375$.

Now, divide the length $T$ sequence into $k$ sequences of length $\frac{T}{k}$. Let $Z_{i}$ be the sum of gradients for the $i$ th of these sequences. Observe that if $Z_{i} \geq \sqrt{\frac{T}{k}}$ for all $i$, then $G_{T}=\sum_{i=1}^{k} Z_{i} \geq k \sqrt{\frac{T}{k}}=$ $\sqrt{k T}$. Furthermore, for any $i$, we have

$$
\operatorname{Pr}\left[Z_{i} \geq \sqrt{\frac{T}{k}}\right]=\operatorname{Pr}\left[G_{\frac{T}{k}} \geq \sqrt{\frac{T}{k}}\right] \geq p
$$

Thus,

$$
\operatorname{Pr}[G \geq \sqrt{k T}] \geq \prod_{i=1}^{k} \operatorname{Pr}\left[Z_{i} \geq \sqrt{\frac{T}{k}}\right] \geq p^{k}
$$

Theorem 7. Consider the problem of unconstrained online linear optimization in one dimension, and an online algorithm that guarantees origin-regret at most $\epsilon$. Then, for any fixed comparator $\stackrel{\dot{x}}{ }$, and any integer $T_{0}$, there exists a gradient sequence $\left\{g_{t}\right\} \in[-1,1]^{T}$ of length $T \geq T_{0}$ for which the algorithm's regret satisfies

$$
\operatorname{Regret}(\stackrel{\circ}{x}) \geq 0.336|\grave{x}| \sqrt{T \log \left(\frac{|\grave{x}| \sqrt{T}}{\epsilon}\right)}
$$

Proof. Let $k=k(T)=\left\lfloor\log \left(\frac{R \sqrt{T}}{\epsilon}\right) / \log \left(p^{-1}\right)\right\rfloor$, and choose $T \geq T_{0}$ large enough so that $4 \leq$ $k \leq T$ and also so that $T$ is a multiple of $k$ (the latter is possible since $k(T)$ grows much more slowly than $T$ ). Let $Q$ be the algorithm's reward when each $g_{t}$ is drawn uniformly from $\{-1,1\}$. Let $G=g_{1: T}$. As shown in the proof sketch, we have

$$
\mathrm{E}[Q \mid G \geq \sqrt{k T}]<\frac{\epsilon}{\operatorname{Pr}[G \geq \sqrt{k T}]}
$$

By Lemma $12, \operatorname{Pr}[G \geq \sqrt{k T}] \geq p^{k}$. Thus,

$$
\mathrm{E}[Q \mid G \geq \sqrt{k T}]<\epsilon p^{-k}=\epsilon \exp \left(k \log p^{-1}\right) \leq R \sqrt{T}
$$

If the algorithm guaranteed $Q \geq R \sqrt{T}$ whenever $G \geq \sqrt{k T}$, then we would have $\mathrm{E}[Q \mid G \geq$ $\sqrt{k T}] \geq R \sqrt{T}$, a contradiction. Thus, there exists a sequence where $G \geq \sqrt{k T}$ and $Q<R \sqrt{T}$, so on this sequence we have

$$
\text { Regret } \geq R \sqrt{k T}-R \sqrt{T}=R \sqrt{T}(\sqrt{k}-1)
$$

Because $k \geq 4$, we have $\frac{1}{2} \sqrt{k} \geq 1$ or $\sqrt{k}-1 \geq \frac{1}{2} \sqrt{k}$, so regret is at least $\frac{1}{2} R \sqrt{k T}=$ $b R \sqrt{T \log \left(\frac{R \sqrt{T}}{\epsilon}\right)}$, where $b=\frac{1}{2} \sqrt{\frac{1}{\log p^{-1}}}>0.336$ (and $p$ is the constant from Lemma 12).


[^0]:    *This work was performed while the author was at Google.

[^1]:    Our bounds are not directly comparable to the bounds cited above: a $\mathcal{O}(\log (T))$ regret bound on $\log$ wealth implies wealth at least $\mathcal{O}(\mathrm{OPT} / T)$, whereas we guarantee wealth like $\mathcal{O}\left(\mathrm{OPT}^{\prime}-\sqrt{T}\right)$. But more importantly, the comparison classes are different.

