

Trading in Markovian Price Models

Sham M. Kakade and Michael Kearns

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104

Abstract. We examine a Markovian model for the price evolution of a stock, in which the probability of local upward or downward movement is arbitrarily dependent on the current price itself (and perhaps some auxiliary state information). This model directly and considerably generalizes many of the most well-studied price evolution models in classical finance, including a variety of random walk, drift and diffusion models. Our main result is a “universally profitable” trading strategy — a single fixed strategy whose profitability competes with the optimal strategy (which knows all of the underlying parameters of the infinite and possibly nonstationary Markov process).

1 Introduction

We examine a Markovian model for the price evolution of a stock, in which the probability of local upward or downward movement is arbitrarily dependent on the current price itself (and perhaps some auxiliary state information). Our main result is a “universally profitable” trading strategy — a single fixed strategy whose profitability competes with the optimal strategy (which knows all of the underlying parameters of the infinite and possibly nonstationary Markov process). While we shall make this statement more precise shortly, our strategy is provably profitable whenever the optimal strategy has significant profits.

The strategy itself is efficient and simple, and employs a “best expert” weighting scheme (Cesa-Bianchi et al. [1997]) over two substrategies — one of which attempts to do rudimentary learning from past observations (which may be extremely sparse), and one of which tries to spot significant directional trends in price. Our main technical contribution is a proof that in our model, one of these two strategies must always have a profit that compares favorably with the optimal strategy.

There are several motivations for the model we introduce. The language of Wall Street and finance is riddled with suggestions that the dynamics of price movement may depend strongly on price itself. Professionals and articles discuss “support” and “resistance” levels for a stock — specific prices or ranges of prices below or above which the market will apparently not let the share price fall or rise, respectively. The field of technical analysis is dominated by price patterns whose appearance is thought to signal future behavior. The common notion of

price uptrends or downtrends is predicated on a series of price levels in which the directional bias is nonzero.

There are also many less speculative reasons price dynamics may change dramatically with price. For example, one might expect there to be support for the share price at the level at which market capitalization (share price times number of outstanding shares, which is essentially the cost of buying the entire company) equals the liquid assets of the company. Similarly, many investors become uncomfortable if the ratio of the share price to a company’s earnings (P/E ratio) becomes excessively large compared to its sector average. Note that it in these cases, there may be many factors aside from price influencing trading behavior (market cap, P/E) — but such factors do result in different price dynamics at different prices.

From the perspective of related literature on trading algorithms, we are particularly interested in price models that fall in between the highly adversarial assumptions typical of competitive analysis and universal portfolio work in computer science (Cover and Ordentlich [1996], Blum and Kalai [1999], El-Yaniv et al. [2001], Helmbold et al. [1996]), and the strong statistical assumptions typical of classical finance random walk and diffusion models and their generalizations (reviewed in Section 3). Our model and result can be thought of as exhibiting a “sweet spot” in the pantheon of price models, in the sense that it contains an extremely rich range of statistical behaviors, yet still permits a universally profitable trading strategy.

We emphasize from the outset that while our model is a gross oversimplification of all the complexities that enter into real-world price formation, it directly and considerably generalizes many of the most well-studied price evolution models in classical finance, including a variety of random walk, drift and diffusion models (see Section 3). To our knowledge, our model has not been explicitly considered before in the finance literature, especially in an algorithmic and learning context.

The outline of the paper follows. In Section 2, we provide the formal definition of our model and the optimal trading strategy that knows the process parameters. In Section 3, we briefly review some of the most common price evolution models in the finance and computer science literatures and relate our model to these. In Section 4, we discuss a number of interesting properties of the model and give simulation results for a particular instance that demonstrates these properties. Section 5 contains our main result. In Section 6 we generalize our result to permit simple extensions of the state.

2 Model and Definitions

In the most basic version of our model, the probabilistic dynamics of directional price movement depend only on the current price. More precisely, we assume that for every integer p between $-\infty$ and $+\infty$, there is a *bias* value $\epsilon(p) \in [-\frac{1}{2}, \frac{1}{2}]$. The interpretation of this bias is as follows: if the price at time t is p_t , then with probability $\frac{1}{2} + \epsilon(p_t)$ we have $p_{t+1} = p_t + 1$, and with probability $\frac{1}{2} - \epsilon(p_t)$ we

have $p_{t+1} = p_t - 1$. Note that in this model, the $|p_{t+1} - p_t| = 1$ always; it will be clear that all of our results hold with only slight degradation in a more general setting in which p_{t+1} must only remain in a bounded range around p_t , including the possibility of no movement. In our model, prices movements are additive, prices are always integer values, and negative prices are allowed for convenience, as all that will matter are the profits made from price movements. (In the long version of this paper discuss a generalization of our results when p_t represents the log price.) Without loss of generality, we always assume the initial price p_1 is 0.

The complete probabilistic dynamics of price movement at all possible prices is given by the infinite vector of biases $\epsilon(p)$ for all integers p , which we shall denote simply by ϵ . A model in our class is thus a countably infinite-state Markov process. (Note that $\epsilon = \mathbf{0}$ corresponds to an unbiased random walk.) We emphasize that this Markov process may be nonstationary and non-recurrent — an infinite walk may never return to its origin, and may forever visit new prices.

In this paper, we will be concerned with trading algorithms that have no a priori information about ϵ , yet can compete with the optimal algorithm that knows the full vector of biases. In order to make such comparisons, it is necessary to somehow limit the amount of risk the optimal algorithm can assume. For instance, if $\epsilon(p) = 1/2$ for some price p , so upwards movement at price p is a certainty, the “optimal” algorithm should purchase an infinite number of shares. We shall thus limit our attention to trading strategies whose share position (number of shares owned (*long*) or owed (*short*)) at any time is at most 1. Other restrictions are possible, but this one has especially natural properties.

With this restriction, then, the optimal algorithm $A_{\text{opt}} = A_{\text{opt}}(\epsilon)$ is straightforward. If the current price is p_t and $\epsilon(p_t) > 0$, then A_{opt} buys one share; and if $\epsilon(p_t) < 0$ then A_{opt} sells (shorts) 1 share. If $\epsilon(p_t) = 0$, then A_{opt} takes no action. Whichever action A_{opt} takes at time t , at the next time step $t + 1$, A_{opt} reverses its action by selling the share bought or buying the share sold at time t , and then repeating the process on p_{t+1} . Thus after each time step, A_{opt} either earns +1 (if it bought a share and the price rose, or it sold a share and the price fell), or loses -1. Thus, we can view A_{opt} as an algorithm for 1-step binary prediction of price movements on the probabilistic sequence of prices. Note that if the price enters a long period of upwards price movement (for example) that A_{opt} correctly predicts, then A_{opt} will be repeatedly buying a share, selling it at the next step and immediately buying another share, etc. This behavior is formally equivalent to buying a single share and holding it for the same period.

For any given bias vector ϵ and number of steps T , we let $\mathbf{p} = (p_1, p_2, \dots, p_T)$ be a random variable that is a sequence of T prices generated according to ϵ . Without loss of generality, we assume $p_1 = 0$. For any trading algorithm A and price sequence \mathbf{p} , we let $V(A, \mathbf{p})$ denote the total amount earned or lost by A on \mathbf{p} divided by T (so that earnings are normalized to per-step averages) and $V(A, \epsilon, T) = \mathbf{E}_{\epsilon, T}[V(A, \mathbf{p})]$ is thus the expected per-step earnings or losses of A over T -step sequences \mathbf{p} distributed according to ϵ . We limit ourselves to only consider algorithms A which limit their share position to at most 1 share, and so

it is easy to see that $V(A, \epsilon, T)$ is between -1 and 1 . We note that $V(A, \epsilon, T)$ can be highly dependent on the specific value of T , since we are in an infinite-state Markov process: larger values of T may cause us to visit new price levels whose dynamics are entirely unlike those seen on smaller time scales.

With these definitions, it is easy to show that $V(A_{\text{opt}}, \epsilon, T)$ is in fact the optimal expected value among all trading algorithms whose position must be at most 1 share at all times, which we shall thus also denote with the shorthand $V^*(\epsilon, T)$. Note that $V^*(\epsilon, T) \in [0, 1]$ always.

For any sequence \mathbf{p} , we define $\#(\mathbf{p})$ to denote the number of *unique* prices appearing in \mathbf{p} . Thus $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})]$ is the expected number of unique prices, and $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$ is the expected fraction of steps that are first visits to some price. This expectation will play a crucial role in our analysis.

3 Related Models

There is a rich history of mathematical models for the evolution of price time series. Perhaps the most basic and well-studied of these are variants of standard random walk or diffusion processes, often referred to in the literature as *Wiener processes*. Among others, this category includes pure unbiased random walks of price and random walks with overall upward or downward drifts (for instance, to model the overall growth of the securities markets historically). Perhaps the most general in this line of models is the *Ito process*, in which the instantaneous drift and variance may depend arbitrarily on both the current price and the time. A good overview of all of these models can be found in Hull [1993].

Our model can be viewed being considerably more general than a Wiener process with drift, but considerably less general than a general Ito process. In particular, it will be easy to see that our results will not hold in such a model. Broadly speaking, if the price process is allowed to depend arbitrarily on time, it is impossible to compete with the profitability of an omniscient party who knows exactly the nature of this time dependence.

The popularity of the various random walk models stems in part from their consistency with broader economic theory, most notably the Efficient Market Hypothesis (EMH), the thesis that individual trader rationality should drive all (expected) profit and arbitrage opportunities out of the market for a stock (or at least all those opportunities beyond those implied by consistent long-term growth, inflation, or drift). However, a long line of relatively recent works have carefully questioned and refuted random walk models and their variants, primarily on the basis of observed conflicts between historical price data and model predictions (Lo and MacKinlay [1999]). Some of these studies have suggested behavioral explanations for the deviations between historical prices and the EMH, an explanation certainly in the spirit of our model, where the market may react differently to different prices for psychological reasons.

The extensive field of *technical analysis* (Murphy [1999]), which suggests that certain price (and other) patterns may presage market behavior, is also clearly at odds with the EMH, at least in its strongest form. The long-term

statistical profitability of certain technical indicators has been argued based on historical data (Brock et al. [1992]). The implicit assumptions of many technical strategies is that price dynamics are largely determined by the current price and some simple auxiliary state information (such as whether the recent price has shown an uptrend or downtrend, or the high and low prices over some recent time window). While our basic model permits only the current price as the Markovian state, in Section 6 we generalize our main result to hold for simple generalizations that incorporate many common technical indicators.

As noted in the Introduction, our model is also considerably more specialized (in terms of the allowed price behavior) than the worst-case price models often examined in computer science and related fields (Cover and Ordentlich [1996], Blum and Kalai [1999], El-Yaniv et al. [2001], Helmbold et al. [1996]). Indeed, in such models, one could never prove that *any* fixed collection of strategies always contained one competing with the optimal strategy that knew the price generation process. The precise point of our model and result is the introduction of a more limited but still quite powerful statistical model for price evolution, along with the proof that a fixed and simple strategy that mixes rudimentary learning and trend-spotting must always be competitive.

4 Properties of the Model, and an Example

Let us now enumerate a few properties (or in some cases, non-properties) of our model that are noteworthy and that distinguish it from some of the more classical models discussed in Section 3:

- As already noted, setting all $\epsilon(p) = 0$ yields a standard (additive) random walk, while all $\epsilon(p) = \alpha$ for some nonzero α yields a random walk with drift.
- One can also program rich mixtures of uptrends, downtrends, unbiased random walks, support and resistance levels, and other features in a single instance of our model.
- While our model does not allowed the detailed specification of time-dependent events (and indeed, our main result would not hold in models such as a general Ito process), one can program rich temporal behaviors *in expectation*. We shall see examples shortly.
- None of the standard random variables of interest — such as the price after T steps, the maximum and minimum prices over T steps, or the profitability of fixed trading strategies — are (necessarily) unimodal in distribution or sharply peaked around their means.
- The optimal per-step profitability $V^*(\epsilon, T)$ may be nonmonotonic in T .

We now examine a concrete instance of our model. Since all we are concerned with is the additive movements of the price, without loss of generality we assume that the initial price is zero. Now consider the following instance of our model, which will be shown in simulation in Figures 1 through 3. In these Figures, the upper two horizontal lines are green (at y-values of 25 and 40) and the lower two horizontal lines are red (at y-values of -25 and -100).

- For $p = -25, \dots, 25$, $\epsilon(p) = 0$. Thus, on either side of the initial price, there is a region of unbiased random walk. In the figures, this region is between the upper red line and the lower green line.
- For $p = 26, \dots, 40$, $\epsilon(p) = 0.1$. Thus, above the random walk region near the initial price, there is a small region of uptrend. This region is between the two green lines.
- For $p > 40$, $\epsilon(p) = -0.1$. Thus, above the uptrend there is infinite resistance to further upward movement. This region is above the upper green line.
- For $p = -100, \dots, -26$, $\epsilon(p) = -0.1$. Thus, below the random walk region near the initial price, there is an extensive region of downtrend. This region is between the two red lines.
- For $p < -101$, $\epsilon(p) = 0$. Thus there is an infinite region of unbiased random walk below the lower red line.

Figures 1 through 3 each show 16 randomly sampled time series from the model ϵ described above. Each figure shows samples from one of the four time scales $T = 100, 1000, 10000$. All of the behaviors identified at the beginning of this section are clearly exhibited, and are discussed in the figure captions.

5 Main Result

In this section, we develop our main result: a trading strategy that knows nothing about the underlying model parameters ϵ , but whose per-step profitability can be provably related to $V^*(\epsilon, T)$. While the analysis is rather involved, the strategy itself and the intuition behind it are appealingly simple and are now sketched briefly.

The key to the analysis is the quantity $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$, the expected fraction of first visits to prices. The first insight is that if this expectation is “small”, then we make “enough” repeat visits to prices to obtain a slight advantage in estimating the biases. For the final result to work out, we must show that this intuition holds even when the average number of visits per state is far too small (such as a constant) to apply concentration inequalities such as the Chernoff bound. Essentially, while we may not have large enough samples to assert an advantage in estimating the bias of any *particular* price, we prove that an advantage exists *on average* across the prices visited. In this case a rather rudimentary learning strategy fares well.

The second insight is that if $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$ is “large”, the price must be following a strong trend that is driven by an overall directional bias, and cannot be easily reversed *on a comparable time scale* (even though it may be reversed on much longer time scales). In this case a simple trend-following or momentum strategy is profitable.

The challenge in the analysis is to make these intuitions precise, and to prove that profitability is possible for all values of $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$. In Section 5.1 we provide the analysis for the case of “small” values for $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$, and in Section 5.2 we consider the case of large values. Section 5.3 stitches the pieces together to give our main result, which we now state:

Theorem 1. (Main Result) Let $\gamma > 0$, and let T' satisfy $T'e^{-\gamma^2 T'/100} < \gamma/32$ (which is satisfied for $T' = \Omega((1/\gamma^2)\log(1/\gamma))$). There exists an algorithm A_{master} , taking input γ , such that for all ϵ and $T \geq T'$, as long as $V^*(\epsilon, T) \geq 2\sqrt{\gamma}$, we have

$$V(A_{\text{master}}(\gamma), \epsilon, T) \geq \frac{\gamma}{4}V^*(\epsilon, T) - \sqrt{2\log(2)/T}. \quad (1)$$

Let us interpret this result briefly. If the profitability of the optimal algorithm is too small (quantified as being below $2\sqrt{\gamma}$), we simply “give up” and are not competitive. The parameter γ thus provides a trade-off to the user of A_{master} . Smaller values of γ will cause the lower bound on $V(A_{\text{master}}, \epsilon, T)$ to take effect at smaller values of $V^*(\epsilon, T)$, but the competitive ratio (which is essentially $\gamma/4$) degrades accordingly. Larger values of γ cause us to not compete at all for a wider range of $V^*(\epsilon, T)$, but give a better competitive ratio when $V^*(\epsilon, T)$ is sufficiently large.

Note also that Theorem 1 provides an “anytime” result, in that the strategy A_{master} is competitive *simultaneously* on all time scales, and does not require T as an input. This is important in light of the fact that $V^*(\epsilon, T)$ may be nonmonotonic in T .

The remainder of this section is devoted to developing A_{master} and proving Theorem 1.

5.1 A Statistical Strategy

We now define a simple trading algorithm that makes minimal use of past observations. We shall denote this algorithm A_{stat} . If the current price p_t is being visited for the first time (that is, time t is the earliest appearance of price p_t in the sequence \mathbf{p}), A_{stat} makes no trade. In this case, after p_{t+1} is revealed, A_{stat} stores a *first-visit record* consisting of the price p_t along with an indication of whether p_{t+1} went up or down from p_t .

If t is *not* the first time price p_t has been visited, then A_{stat} looks up the first-visit record for p_t and trades according to this record — that is, if after the first visit to p_t the price went up, A_{stat} buys one share, otherwise it sells one share, respectively. To obey the 1-share position limit, at time $t + 1$ A_{stat} sells off or buys back the position it accumulated and repeats the process on p_{t+1} .

Thus, A_{stat} is the algorithm that makes perhaps the least possible use of statistical history, simply predicting that what happened after the very first visit to a price will continue to happen. Obviously, it would make more intuitive sense to collect statistics on *all* the visits to a given price, and trade based on these cumulative statistics. But it turns out that A_{stat} must operate with sample sizes that are far too small to usefully apply large-deviation bounds such as the Chernoff inequality. Hence, one can’t provide a general bound in which our expected value is a linear fraction of the optimal value. Instead, we compete against the square of the optimal value (which we conjecture is the best one could hope for). More formally, we have:

Theorem 2. (*Statistical Strategy*) For any biases ϵ and any T ,

$$V(A_{\text{stat}}, \epsilon, T) \geq V^*(\epsilon, T)^2 - \mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]. \quad (2)$$

Proof. Let us first write down an explicit expression for the T -step optimal value $V^*(\epsilon, T)$. At each time step t , the optimal algorithm examines the bias $\epsilon(p_t)$ to decide how to trade. Abusing notation slightly, let us denote $\epsilon(t) = \epsilon(p_t)$ when t is a time value and not a price value and thus there is no risk of confusion. The expected profit of A_{opt} at time t is then

$$\left(\frac{1}{2} + |\epsilon(t)|\right) (+1) + \left(\frac{1}{2} - |\epsilon(t)|\right) (-1) = 2|\epsilon(t)|. \quad (3)$$

Now recall that $V^*(\epsilon, T) = \mathbf{E}_{\epsilon, T}[V(A_{\text{opt}}, \mathbf{p})]$. Since $V(A_{\text{opt}}, \mathbf{p})$ is a sum of T 1-step returns, by linearity of expectation we may write

$$V^*(\epsilon, T) = \frac{1}{T} \sum_{t=1}^T \sum_{\mathbf{p}: |\mathbf{p}|=t} \Pr_{\epsilon, t}[\mathbf{p}](2|\epsilon(t)|) \quad (4)$$

where each inner sum over sequences \mathbf{p} of length $t \leq T$ is the expected profit of A_{opt} on the step t .

Let us now analyze the 1-step expected profit of algorithm A_{stat} at time t . If t is the first visit to the price p_t in the sequence \mathbf{p} , then the profit of A_{stat} is 0. Otherwise, the profit depends on whether the first visit to p_t revealed the correct or incorrect sign of $\epsilon(p_t)$. More precisely, the expected return of A_{stat} on non-first visits to p_t may be written

$$\left(\frac{1}{2} + |\epsilon(t)|\right) (2|\epsilon(t)|) + \left(\frac{1}{2} - |\epsilon(t)|\right) (-2|\epsilon(t)|) = 4|\epsilon(t)|^2. \quad (5)$$

The logic here is that with probability $\frac{1}{2} + |\epsilon(t)|$, the first visit to p_t reveals the correct sign of the bias, in which case on all subsequent visits, A_{stat} will behave the same as A_{opt} and receive $2|\epsilon(t)|$ in expected profits; and with probability $\frac{1}{2} - |\epsilon(t)|$, the first visit to p_t reveals the incorrect sign, in which case on all subsequent visits, A_{stat} will receive $-2|\epsilon(t)|$. Thus the expectation is taken over *both* the randomization on the current visit to p_t , and the randomization on the first visit.

We would now like to apply this observation on the 1-step profit of A_{stat} to obtain an expression for $V(A_{\text{stat}}, \epsilon, T)$; the main challenge is in dealing with the dependencies introduced by conditioning on the number of visits to each price level. The following inequality can be shown (details omitted):

$$V(A_{\text{stat}}, \epsilon, T) \geq \frac{1}{T} \sum_{t=1}^T \sum_{\mathbf{p}: |\mathbf{p}|=t} \Pr_{\epsilon, t}[\mathbf{p}](4|\epsilon(t)|^2) - \mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T] \quad (6)$$

Combining Equation (4) and Equation (6), we now have:

$$V(A_{\text{stat}}, \epsilon, T) - V^*(\epsilon, T)^2 = \frac{4}{T} \sum_{t=1}^T \sum_{\mathbf{p}: |\mathbf{p}|=t} \Pr_{\epsilon, t}[\mathbf{p}] |\epsilon(t)|^2 - \left(\frac{2}{T} \sum_{t=1}^T \sum_{\mathbf{p}: |\mathbf{p}|=t} \Pr_{\epsilon, t}[\mathbf{p}] |\epsilon(t)| \right)^2 - \mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T] \quad (7)$$

It remains to show that the first two terms are positive.

Recall that each $\epsilon(t) = \epsilon(p_t)$ is actually the bias at some price level p_t . Let us define for each price q

$$w(q) = \frac{1}{T} \sum_{t=1}^T \sum_{\mathbf{p}: |\mathbf{p}|=t, p(t)=q} \Pr_{\epsilon, t}[\mathbf{p}]. \quad (8)$$

It is easy to see that since the price must remain in the range $[-T, T]$ on sequences of length at most T , the values $w(-T), \dots, w(T)$ sum to 1 and are all positive, and thus can be interpreted as a distribution. The first two terms in Equation (7) may be rewritten as

$$4 \left(\sum_{q=-T}^T w(q) |\epsilon(q)|^2 - \left(\sum_{q=-T}^T w(q) |\epsilon(q)| \right)^2 \right) \quad (9)$$

This difference is non-negative as desired, by the convexity of the function $f(x) = x^2$. (Interestingly, note that this difference has the form of the variance of $\epsilon(q)$ with respect to the distribution $w(q)$.) \square

5.2 A Momentum Strategy

We now turn attention to a strategy that will succeed for large values of the quantity $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$. For any given values of γ and T , the momentum strategy $A_{\text{mom}}(\gamma, T)$ can be described as follows:

1. For all $p \in (-\gamma T/4, \gamma T/4)$, take no action.
2. For all $p \geq \gamma T/4$, purchase one share and sell it back at the next time step.
3. For all $p \leq -\gamma T/4$, sell one share and purchase it back at the next time step.

Note this strategy uses knowledge of the time T ; however, this dependency can be removed (details omitted) to yield an algorithm that is competitive on all time scales simultaneously.

The following definitions will be necessary in our analysis of A_{mom} . For the remainder of this subsection, \mathbf{p} will denote a price sequence of length T for some fixed T . Let $\max(\mathbf{p})$ ($\min(\mathbf{p})$, respectively) be the maximum (minimum, respectively) price reached on \mathbf{p} . Let $\text{drop}(\mathbf{p})$ be the absolute value of the difference between $\max(\mathbf{p})$ and the *smallest* price reached on \mathbf{p} *after* the first visit to $\max(\mathbf{p})$. Thus, $\text{drop}(\mathbf{p})$ measures the “fall” from the high price. Similarly, we

define $\text{rise}(\mathbf{p})$ to be the absolute value of the difference between $\min(\mathbf{p})$ and the largest price reached on \mathbf{p} after the first visit to $\min(\mathbf{p})$.

A_{mom} enjoys the following performance guarantee.

Theorem 3. (*Momentum Strategy*) Let $\gamma > 0$, and let T' satisfy $T'e^{-\gamma^2 T'/48} < \gamma/16$. If $T > T'$ and if either $\mathbf{E}_{\epsilon, T}[\max(\mathbf{p})/T] \geq \gamma$ or $\mathbf{E}_{\epsilon, T}[|\min(\mathbf{p})/T|] \geq \gamma$ then

$$V(A_{\text{mom}}(\gamma, T), \epsilon, T) \geq \frac{\gamma}{2} \geq \frac{\gamma}{2} V^*(\epsilon, T) \quad (10)$$

Note that unlike the guarantee for A_{stat} , A_{mom} must be run for a time larger than some threshold time. Essentially, this time is the time long enough to discover the trend. Also, note that we always must have either $\mathbf{E}_{\epsilon, T}[\max(\mathbf{p})] \geq \mathbf{E}_{\epsilon, T}[\#\mathbf{p}]/2$ or $\mathbf{E}_{\epsilon, T}[|\min(\mathbf{p})|] \geq \mathbf{E}_{\epsilon, T}[\#\mathbf{p}]/2$.

At the heart of the proof of Theorem 3 is the following simple probabilistic lemma. The lemma essentially states that if the price makes large moves on some time scale, then with high probability it cannot return to its starting value on a comparable time scale.

Lemma 1. For any constant $a > 0$, we have

1. For all ϵ, T and $z \geq aT$, $\Pr_{\epsilon, T}[\max(\mathbf{p}) = z \text{ and } \text{drop}(\mathbf{p}) \geq aT/2] \leq e^{-a^2 T/12}$.
2. For all ϵ, T and $z \leq -aT$, $\Pr_{\epsilon, T}[\min(\mathbf{p}) = z \text{ and } \text{rise}(\mathbf{p}) \geq aT/2] \leq e^{-a^2 T/12}$.

Proof. (Sketch) We sketch only Part 1, as Part 2 is entirely symmetric. First let us suppose that among the biases $\epsilon(0), \dots, \epsilon(z)$ there are more than $aT/4$ which are negative. In this case we show that the probability of $\max(\mathbf{p})$ even reaching the price z is small. In order for the price to reach z , it clearly must “get through” these negative biases — in other words, the price must have a net upwards movement of at least $aT/4$ even when restricted only to those visits to prices with negative bias. If we modify all of these negative biases to be equal to 0 (unbiased), we can clearly only *increase* the probability that $\max(\mathbf{p})$ reaches the price z .

We can thus bound $\Pr_{\epsilon, T}[\max(\mathbf{p}) = z]$ by the probability that in T independent flips of a fair coin, we would see an excess of heads over tails of at least $aT/4$. By the standard Chernoff bound, the probability of seeing such an excess in T flips is at most $e^{-(a^2/2)^2 T/3} = e^{-a^2 T/12}$. Since the probability of $\max(\mathbf{p})$ even reaching z has been thus bounded, the lemma holds in this case.

Otherwise, we must have that at most $aT/4$ of the biases $\epsilon(0), \dots, \epsilon(z)$ are negative. In this case we show that $\Pr_{\epsilon, T}[\text{drop}(\mathbf{p}) \geq aT/2]$ is small. Since the price can drop by a net amount of at most $aT/4$ when restricted only to visits to prices with negative biases, in order to drop by a total of at least $aT/2$, it must drop a further net amount of at least $aT/4$ when restricted only to visits to prices with positive biases. Using a similar argument, it is straightforward to see that this probability being bounded by $e^{-a^2 T/12}$. \square

Lemma 1 is used to prove the following result (and a similar result holds in terms of $\mathbf{E}_{\epsilon, T}[\min(\mathbf{p})/T]$).

Lemma 2. *Let $\gamma > 0$, and let T' be such that $T'e^{-\gamma^2 T'/48} < \gamma/16$. If $T > T'$, then for any biases ϵ*

$$V(A_{\text{mom}}(\gamma, T), \epsilon, T) \geq \mathbf{E}_{\epsilon, T}[\max(\mathbf{p})/T] - \gamma/2 \quad (11)$$

Proof. (Sketch) First, using a Markov inequality argument, one can show:

$$\sum_{x > \gamma T/4} \Pr_{\epsilon, T}[\max(\mathbf{p}) = x](x - \gamma T/4) \geq \mathbf{E}_{\epsilon, T}[\max(\mathbf{p})] - \gamma T/4. \quad (12)$$

Informally, this summation is the expected profit from the cases in which the maximum price exceeds $\gamma T/4$, conditioned on the subsequent drop being at most $\gamma T/4$.

Now one can use Lemma 1 to show that the value of A_{mom} is close to the above. \square

Theorem 3 follows from Lemma 2 under the assumption that $\mathbf{E}_{\epsilon, T}[\max(\mathbf{p})] > \gamma T$, and noting that $V^*(\epsilon, T) \leq 1$.

5.3 Putting the Pieces Together

Theorems 2 and 3 establish that for any biases ϵ and any T , *at least one* of the two strategies A_{stat} and A_{mom} must have a expected profit that compares “favorably” with that of the optimal algorithm that knows ϵ . We now wish to define a *single* strategy accomplishing this same criterion. Of course, one way of doing this is to have a strategy that simply flips a fair coin at the outset of trading to decide with A_{stat} and A_{mom} to use, at a cost of a factor of 2 in our expected return in comparison to $V^*(\epsilon, T)$. While this cost is insignificant in light of the other constant factors we are already absorbing, we prefer to apply the so-called “experts” methods of worst-case on-line analysis. When we generalize our results to permit the biases ϵ to depend on an underlying state variable more complex than just the current price, the experts methodology will be necessary.

In order to apply the experts framework, it is important to recall the observation made in Section 2 that our trading model can really be viewed as an instance of on-line binary prediction. We view trading strategies (and A_{stat} , A_{mom} and the optimal trading algorithm in particular) as making a series of trades or predictions, each of which wins or loses immediately. We can thus immediately apply the on-line weighting scheme of Cesa-Bianchi et al. [1997] to the strategies A_{stat} and A_{mom} (in this case, an especially small set of experts); let us call the resulting strategy A_{master} , since it can be viewed as a “master” strategy allocating capital between the two subordinate strategies. Combining Theorem 16 of Cesa-Bianchi et al. [1997] with Theorems 2 and 3 allows one to show that

$$V(A_{\text{master}}, \epsilon, T) \geq \min \left(V^*(\epsilon, T)^2 - \gamma, \frac{\gamma}{4} V^*(\epsilon, T) \right) - \sqrt{2 \log(2)/T} \quad (13)$$

always holds. Notice that this lower bound may actually be near zero for small values of $V^*(\epsilon, T)$. From this equation, Theorem 1 follows.

6 Extending the State

So far we have focused exclusively on a model in which the directional bias of the price movement may depend arbitrarily on the current price itself. In this section we generalize our results to a considerably richer class of models, in which the directional bias $\epsilon = \epsilon(p, s)$ may depend on both price p and some auxiliary information s . For example, one might posit that a more realistic model for price dynamics is that the directional bias depends not only on the current price, but also on whether the current price was arrived at from below or above. We can model this letting the probability that $p_{t+1} = p_t + 1$ be $\frac{1}{2} + \epsilon(p_t, s_t)$, where $\epsilon(p_t, s_t) \in [-\frac{1}{2}, \frac{1}{2}]$ and $s_t \in \{0, 1\}$ equals 1 if $p_t = p_{t-1} + 1$ (uptrend) and 0 if $p_t = p_{t-1} - 1$ (downtrend). We again have an infinite Markov process, but with the Markovian state now being the pairs (p, s) rather than just p alone. We will continue to use the notation ϵ to denote the infinite set of biases $\epsilon(p, s)$ for all integer prices p and binary trend indicators s .

We now outline why the results of Section 5 continue to hold with some additional machinery, after which we will provide a more general and formal statement. Let \mathbf{p} be the sequence of prices p_t , and let \mathbf{s} be the corresponding sequence of auxiliary values s_t . Let us define $\#(\mathbf{p}, \mathbf{s})$ to be the number of *unique* states (p_t, s_t) visited on (\mathbf{p}, \mathbf{s}) . Then it is easily verified that Theorem 2 holds with $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p})/T]$ replaced by $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p}, \mathbf{s})/T]$. In this case, the obvious modification of strategy A_{stat} — namely, to always trade according to the observed behavior of the price on the *first* visit to state (p, s) — permits an identical analysis.

The extension of Theorem 3 is slightly more involved. In particular, in our new model Lemma 1 simply no longer holds — we can now easily “program” behavior that (for example) causes the price to deterministically rise to some price and then deterministically fall back to its starting value. In the price-only model, such behavior was excluded by Lemma 1, which states that the probability of the conjunction of a steep rise in price and a subsequent drop is exponentially small.

However, in the new model it remains true that if $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p}, \mathbf{s})/T]$ is larger than γ , then either $\mathbf{E}_{\epsilon, T}[\max(\mathbf{p})]$ or $\mathbf{E}_{\epsilon, T}[\min(\mathbf{p})]$ must be large — namely, one of them must be at least $\gamma T/4$ (as opposed to $\gamma T/2$ in the price-only model). This is because for every n unique *states* we visit, we must visit at least $n/2$ unique *prices* as well, since for each price p there are only two associated states $(p, 0)$ and $(p, 1)$. To exploit this, despite the fact that Lemma 1 no longer holds, we make richer use of the Cesa-Bianchi et al. [1997] results. For each $1 \leq i \leq T$, we introduce two simple trading strategies, A_{+i} and A_{-i} . Strategy A_{+i} buys a single share at the outset of trading, and sells it back if and only if the price reaches the value i above its starting point. Strategy A_{-i} sells a single share at the outset of trading, and buys it back if and only if the price reaches the value i below its starting point. If either $\mathbf{E}_{\epsilon, T}[\max(\mathbf{p})]$ or $\mathbf{E}_{\epsilon, T}[\min(\mathbf{p})]$ is least $\gamma T/4$, then clearly the expected maximum per-step profit among the strategies $\{A_{+i}, A_{-i}\}_{1 \leq i \leq T}$ is at least $\gamma/4$.

The new overall algorithm is thus to apply the weighting scheme of Cesa-Bianchi et al. [1997] to the strategies $\{A_{+i}, A_{-i}\}_{1 \leq i \leq T}$ along with the strategy A_{stat} . Regardless of the value of $\mathbf{E}_{\epsilon, T}[\#(\mathbf{p}, \mathbf{s})/T]$, one of these $2T + 1$ strategies will be profitable.

To generalize the analysis above, note that the only property we required of the state space (p, s) is that each possible price p have only a “small” number of possible extensions s (2 in the analysis above). This motivates the following definition: for any price p , let us define $\kappa(p)$ to be $|\{(p, s) \in S\}|$, where S is the set of possible states. For instance, in the example above, for any given p , only the states $(p, 0)$ and $(p, 1)$ are possible, so $\kappa(p) = 2$ always. We then define $\kappa_{\max} = \max_p \{\kappa(p)\}$. Note that κ_{\max} can be finite and small even though an infinite number of values of s are possible as we range over all values of p . For example, if s is defined to be the maximum price in the last ℓ time steps, then for any p , there are at most 2ℓ possible values for s ; but the domain of s is all the integers.

Let $A_{\text{general}}(T)$ refer to this more general algorithm which takes T as an input and which weights the strategies A_{stat} and $\{A_{+i}, A_{-i}\}_{1 \leq i \leq T}$ as discussed above. Then we have the following theorem.

Theorem 4. (*Main Result, Extended State*) *Let κ_{\max} be as defined above. Let $\gamma > 0$, and let T' be such that $T'e^{-\gamma^2 T'/100} < \gamma/32$ (which is satisfied for $T' = \Omega((1/\gamma^2) \log(1/\gamma))$). If $T \geq T'$, then for any ϵ and as long as*

$$V^*(\epsilon, T) \geq 2\sqrt{(\gamma^2/4\kappa_{\max})} + 4\gamma \quad (14)$$

we have

$$V(A_{\text{general}}(T), \epsilon, T) \geq \frac{\gamma}{2\kappa_{\max}} V^*(\epsilon, T) - \sqrt{2 \log(T)/T}. \quad (15)$$

Note that this differs from the price-only result of Theorem 1 in that our competitive ratio is now proportional to γ/κ_{\max} rather than γ , and the regret term $\sqrt{2 \log(T)/T}$ of the weighting scheme now has $\log(T)$ replacing $\log(2)$. Also, this result is not anytime since A_{general} takes as input the time T .

Thus for constant κ_{\max} , our bound essentially suffers only a constant factor degradation.

References

- A. Blum and A. Kalai. Universal portfolios with and without transaction costs. *Machine Learning*, 35(3):193–205, 1999.
- A. Brock, J. Lakonishok, and B. Lebaron. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, (47):1731–1764, 1992.
- Nicolo Cesa-Bianchi, Yoav Freund, David Haussler, David P. Helmbold, Robert E. Schapire, and Manfred K. Warmuth. How to use expert advice. *J. ACM*, 44(3): 427–485, 1997. ISSN 0004-5411.
- T. Cover and E. Ordentlich. Universal portfolios with side information. *IEEE Transactions on Information Theory*, 42(2), 1996.

- R. El-Yaniv, A. Fiat, R. M. Karp, and G. Turpin. Optimal search and one-way trading online algorithms. *Algorithmica*, 30:101–139, 2001.
- David P. Helmbold, Robert E. Schapire, Yoram Singer, and Manfred K. Warmuth. On-line portfolio selection using multiplicative updates. In *International Conference on Machine Learning*, pages 243–251, 1996. URL cite-seer.ist.psu.edu/article/helmbold98line.html.
- J. Hull. *Options, Futures, and Other Derivative Securities*. Prentice Hall, 1993.
- A. Lo and A.C. MacKinlay. *A Non-Random Walk Down Wall Street*. Princeton University Press, 1999.
- J. Murphy. *Technical Analysis of the Financial Markets*. New York Institute of Finance, 1999.




Fig. 1. 16 sampled walks of length $T = 100$. On this short time scale, with high probability, the walk remains in the unbiased region between $p = -25$ and $p = 25$.

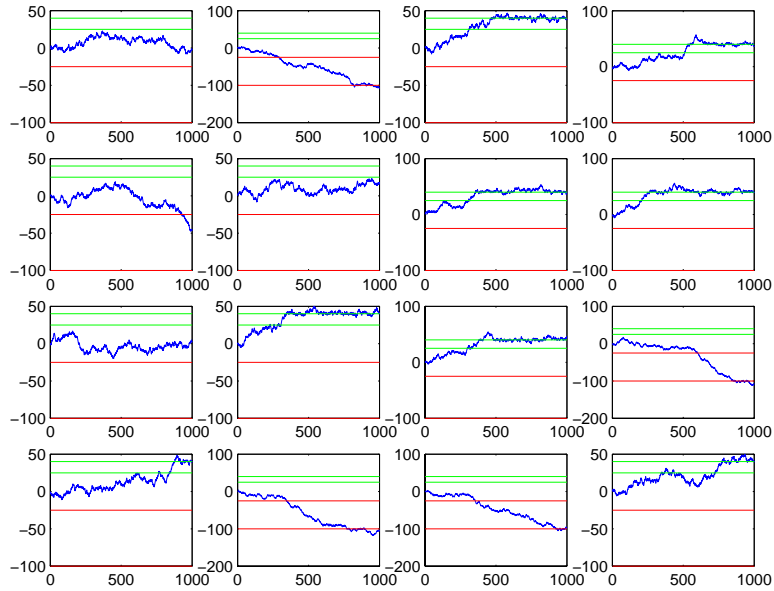


Fig. 2. 16 sampled walks of length $T = 1000$. Most walks either enter the uptrend region above $p = 25$ and are lifted to the resistance level at $p = 40$ (plot (1, 3)), or enter the downtrend region below $p = -25$ (plot (1, 2)). Some walks enter the uptrend or downtrend only very late (plot (2, 1)), or do not even leave the unbiased region (plot (3, 1)).

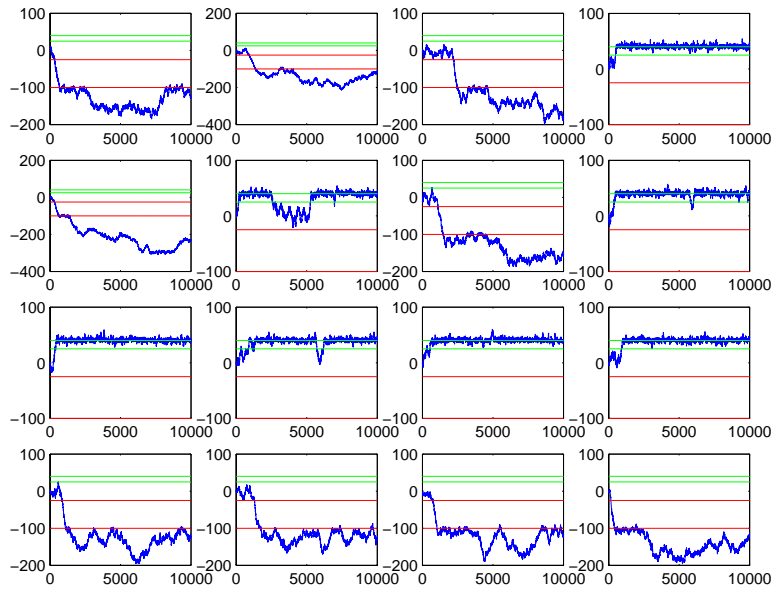


Fig. 3. 16 sampled walks of length $T = 10000$. Now all walks either traverse the uptrend and remain near the $p = 40$ resistance level, or traverse the downtrend and follow an unbiased random walk below $p = -100$, the bottom of the downtrend region. On this time scale, some of the uptrend walks are starting to show “cracks” in the form of dips back into the original unbiased region (plots (2, 2), (2, 4) and (3, 2)). Eventually these cracks will pull the price back through the original unbiased region and into the downtrend; asymptotically, all walks eventually spend most of their lives in the lower unbiased region below $p = -100$.