Switching Investments

Wouter M. Koolen* and Steven de Rooij†

June 21, 2010

*CWI, Amsterdam, The Netherlands
†University of Cambridge

What We Do

All About A Line

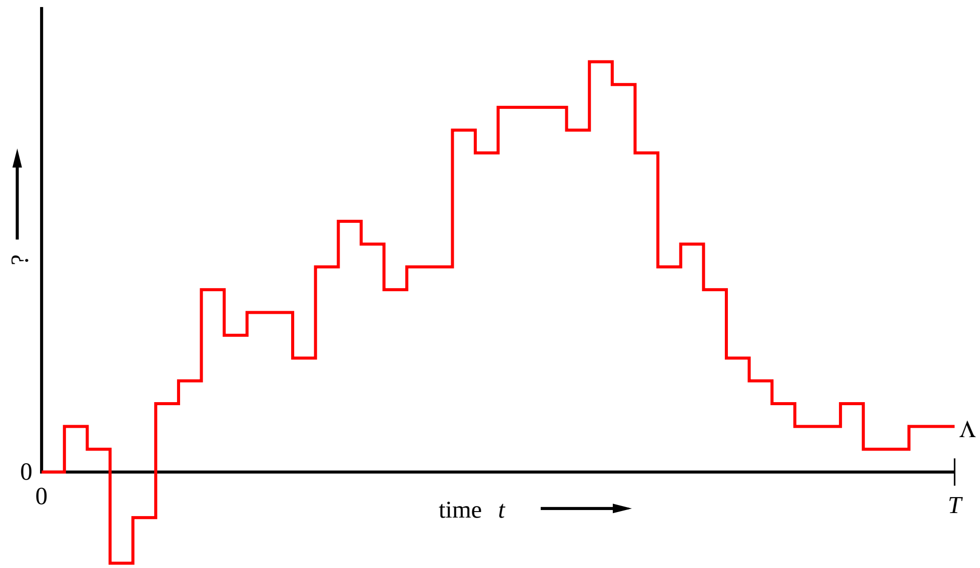

All About A Line

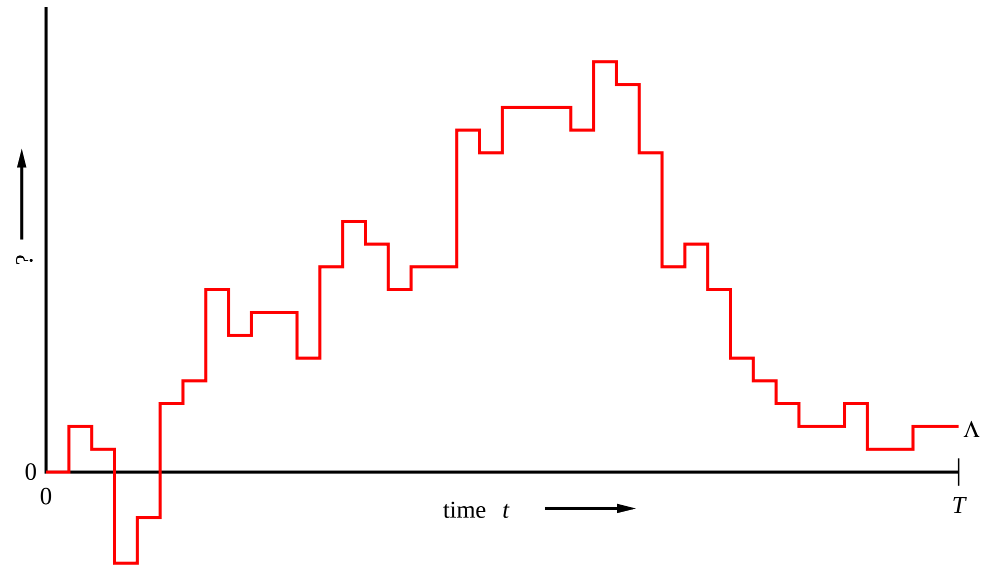

Vertical axis:

- Prediction with expert advice: $L_1(x_{1:t}) - L_2(x_{1:t})$
- Hypothesis testing: $\log(P_1(x_{1:t})/P_0(x_{1:t}))$
- The logarithm of a stock price.

All About A Line

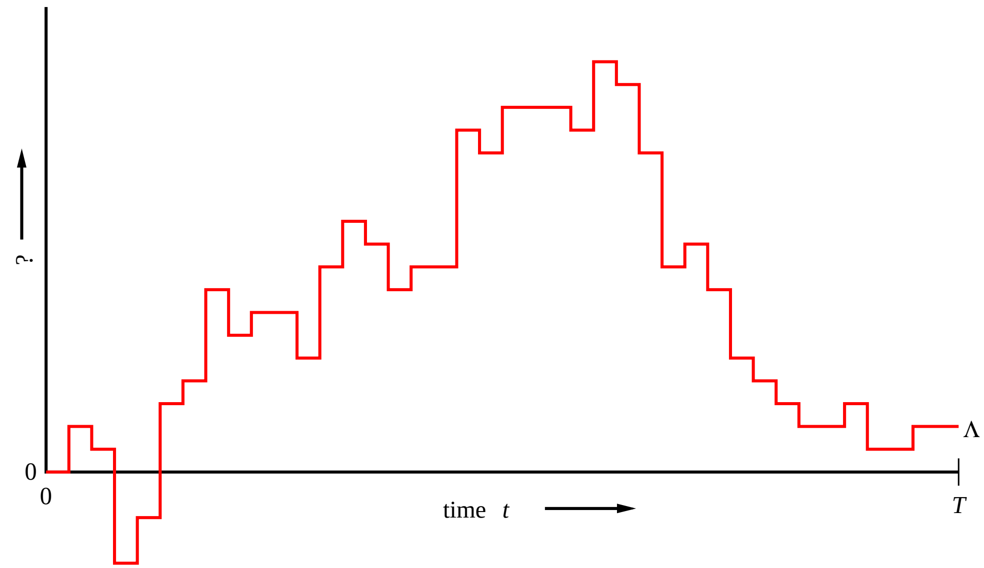

Vertical axis:

- Prediction with expert advice: $L_1(x_{1:t}) - L_2(x_{1:t})$
- Hypothesis testing: $\log(P_1(x_{1:t})/P_0(x_{1:t}))$
- The logarithm of a stock price.

Goal: predict whether the line will go up or down.

Basic Investment Strategies

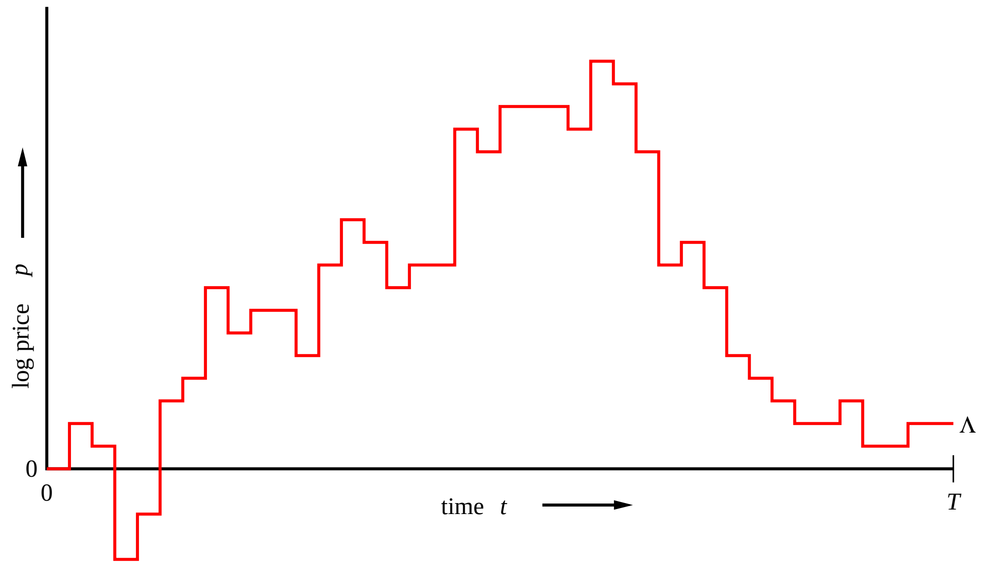

A basic investment strategy σ_t is to sell at a predetermined time t.

Basic Investment Strategies

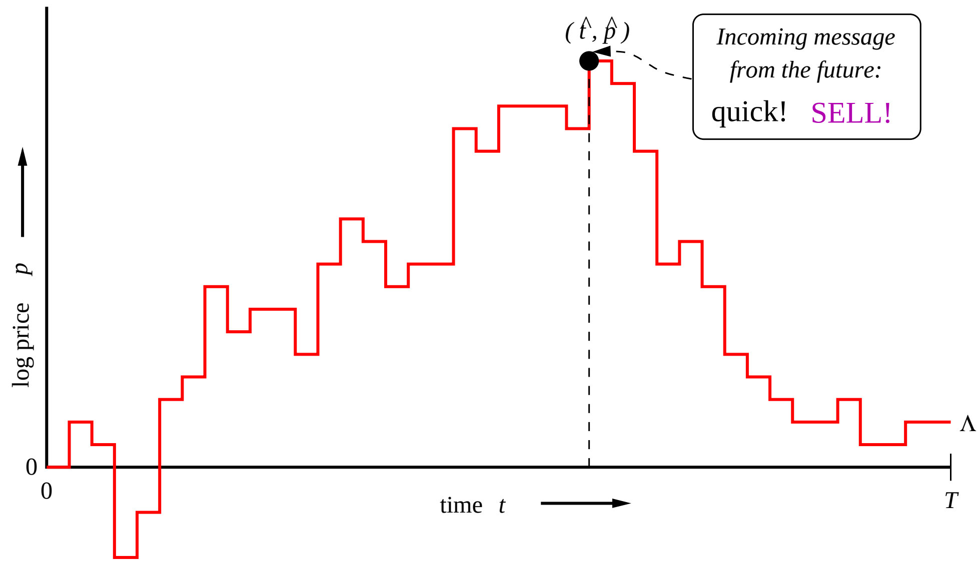

A basic investment strategy σ_t is to sell at a predetermined time t.

Problem: *in hindsight* we know when the oil started leaking!

Hedging

We distribute initial capital over $\sigma_1, \ldots, \sigma_T$.
Let $\pi(t)$ denote the fraction of capital assigned to strategy σ_t.
Let $\Lambda(0) = 0$. We obtain payoff:

$$\log \sum_{t=0}^{T} e^{\Lambda(t)} \pi(t) \geq \log \left(e^{\Lambda(\hat{t})} \pi(\hat{t}) \right) = \Lambda(\hat{t}) - (-\log \pi(\hat{t})).$$

Hedging

We distribute initial capital over $\sigma_1, \ldots, \sigma_T$.
Let $\pi(t)$ denote the fraction of capital assigned to strategy σ_t.
Let $\Lambda(0) = 0$. We obtain payoff:

$$\log \sum_{t=0}^{T} e^{\Lambda(t)} \pi(t) \geq \log \left(e^{\Lambda(\hat{t})} \pi(\hat{t}) \right) = \Lambda(\hat{t}) - (\underbrace{-\log \pi(\hat{t})}_{\text{regret}}).$$

Regret may be relatively large or small, depending on

- How much the stock price drops
- The granularity of measurement

Hedging

We distribute initial capital over $\sigma_1, \ldots, \sigma_T$.
Let $\pi(t)$ denote the fraction of capital assigned to strategy σ_t.
Let $\Lambda(0) = 0$. We obtain payoff:

$$\log \sum_{t=0}^{T} e^{\Lambda(t)} \pi(t) \geq \log \left(e^{\Lambda(\hat{t})} \pi(\hat{t}) \right) = \Lambda(\hat{t}) - (\underbrace{-\log \pi(\hat{t})}_{\text{regret}}).$$

Regret may be relatively large or small, depending on

- How much the stock price drops
- The granularity of measurement $\longleftarrow$ undesirable!

Price Switched Strategies

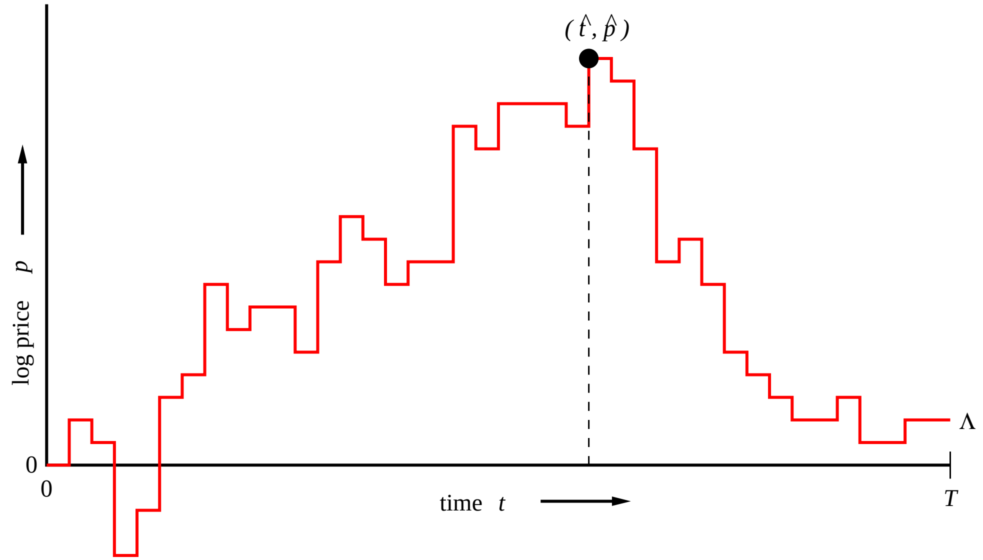

We parameterised the strategy to sell by time t...

Price Switched Strategies

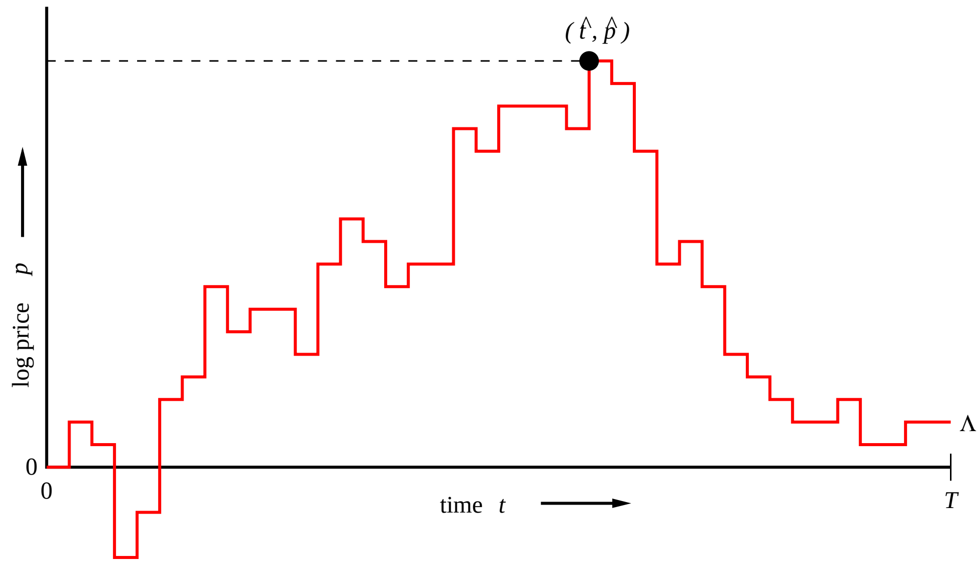

We parameterised the strategy to sell by time t...
Let us now define σ_p to sell when $\Lambda(t) \geq p$.

- Time-switched strategy: decision to sell depends on t
- Price-switched strategy: decision to sell depends on $\Lambda(t)$

Price Switched Strategies

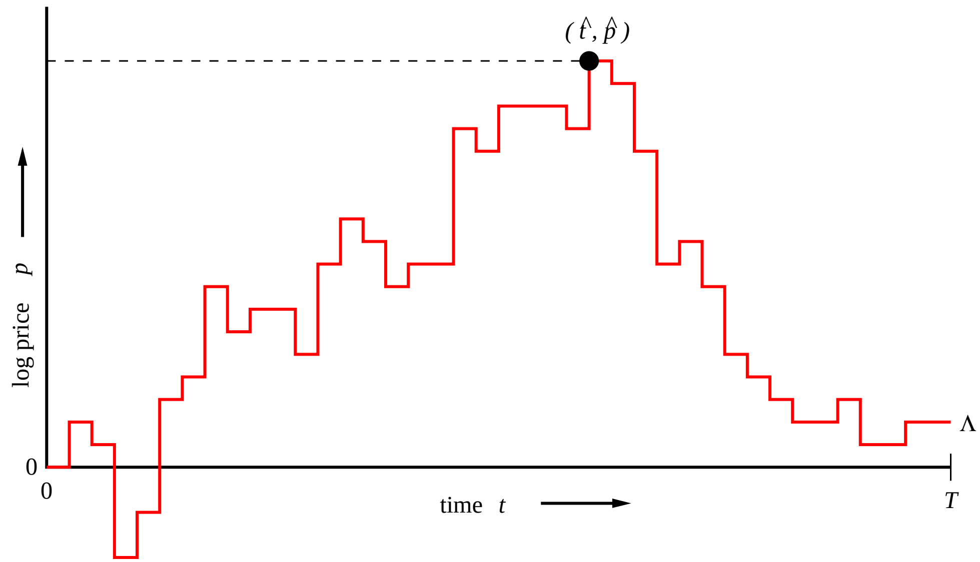

We parameterised the strategy to sell by time t...
Let us now define σ_p to sell when $\Lambda(t) \geq p$.

- Time-switched strategy: decision to sell depends on t
- Price-switched strategy: decision to sell depends on $\Lambda(t)$

We can no longer sell at every moment. But that's OK.

More Price Switching

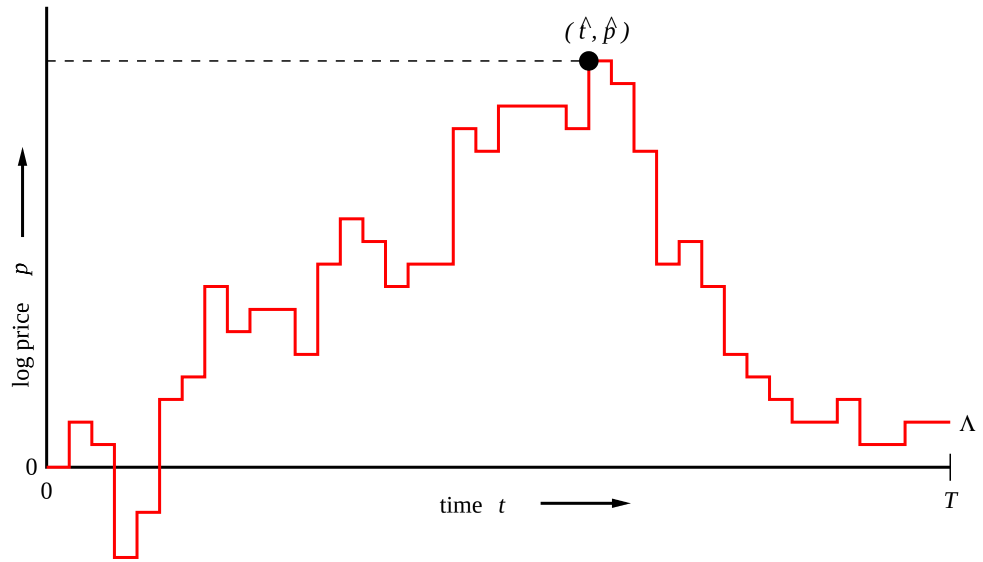

We can hedge as before to obtain at least

$$\log \sum_{p=0}^{\hat{p}} e^p \pi(p) \geq \log \left(e^{\hat{p}} \pi(\hat{p}) \right) = \hat{p} - (-\log \pi(\hat{p})).$$

For sufficiently large $\hat{p}$, the regret is relatively small!

What We Actually Do

Continuous Price

Actually, logprices are not integers and we do not pretend they are.

We can get very close to the previous bound:
if π is a decreasing density on the positive reals, then

$$\log \int_0^{\hat{p}} e^p \pi(p)\,\mathrm{d}p \geq \log \left(\pi(\hat{p}) \int_0^{\hat{p}} e^p \,\mathrm{d}p \right) = \log(e^{\hat{p}} - 1) - (-\log \pi(\hat{p})).$$

We cannot sell at $\hat{p}$ exactly anymore $\rightarrow$ small additional overhead

Multiple Switches

Actually, we are interested in exploiting multiple switches.

Let $\boldsymbol{\delta} = (\delta_1, \delta_2, \ldots)$. A strategy $\sigma_{\boldsymbol{\delta}}$:

- initially invests all capital
- sells all stock when the logprice goes up δ_1 or more, then
- invests all capital again as it goes down δ_2 or more,
- etcetera.

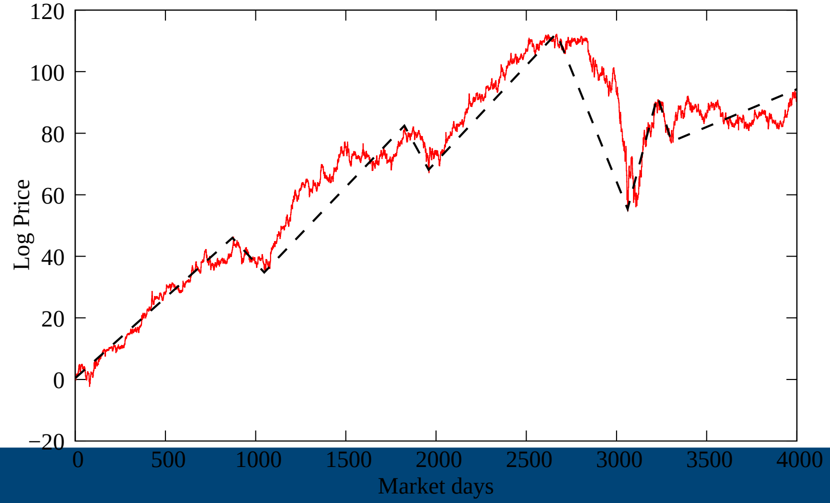

Continuous Time

Intuition: Discontinuities in Λ are helpful.

Continuous Time

Intuition: Discontinuities in Λ are helpful.

Let the logprice function be $\Lambda : [0, T] \to \mathbb{R}$.
A discrete time scenario can be modelled by a step function.

Theorem 1 *The payoff of our strategy decreases when a discontinuity in the logprice function is replaced by a continuous monotonic transition.*

Continuous Time

Intuition: Discontinuities in Λ are helpful.

Let the logprice function be $\Lambda : [0, T] \to \mathbb{R}$.
A discrete time scenario can be modelled by a step function.

Theorem 1 *The payoff of our strategy decreases when a discontinuity in the logprice function is replaced by a continuous monotonic transition.*

- We can simplify the analysis by assuming continuity.

Monotonicity

Intuition: The more fluctuations in Λ, the better.

Monotonicity

Intuition: The more fluctuations in Λ, the better.

Theorem 2 *Fix the value of Λ at arbitrary times $0 = t_1 < \ldots < t_m = T$. Our strategy achieves minimal payoff if the successive intervals of Λ are monotone.*

Monotonicity

Intuition: The more fluctuations in Λ, the better.

Theorem 2 *Fix the value of Λ at arbitrary times $0 = t_1 < \ldots < t_m = T$. Our strategy achieves minimal payoff if the successive intervals of Λ are monotone.*

In summary, the regret compared to a specific δ is maximised if

- Λ is continuous (Thm 1)
- Λ is monotonic in-between switches (Thm 2)

The worst case regret is realised in the ideal case for analysis!

Regret Bound

Theorem 3 *Given any Λ and a basic strategy $\boldsymbol{\delta}$ that performs its mth switch on Λ at time T. The payoff of our strategy is at least*

$$\sum_{1 \leq \text{odd } i \leq m} \delta_i - \left(\sum_{i=1}^{m} -\log \pi(\delta_i) + (m-1)c_\pi + \ln 2 + 2\epsilon_\pi \right).$$

Here, c_π and ϵ_π are small constants that depend on π; values are typically around 0.02 and 3.5, respectively.

Regret Bound

Theorem 3 *Given any Λ and a basic strategy $\boldsymbol{\delta}$ that performs its mth switch on Λ at time T. The payoff of our strategy is at least*

$$\sum_{1 \leq \text{odd } i \leq m} \delta_i - \left(\sum_{i=1}^{m} - \log \pi(\delta_i) + (m-1)c_\pi + \ln 2 + 2\epsilon_\pi \right).$$

Here, c_π and ϵ_π are small constants that depend on π; values are typically around 0.02 and 3.5, respectively.

Thus,

- Small fluctuations are hard to exploit
- The bound is best applied to parsimonious strategies (with small m)

Example

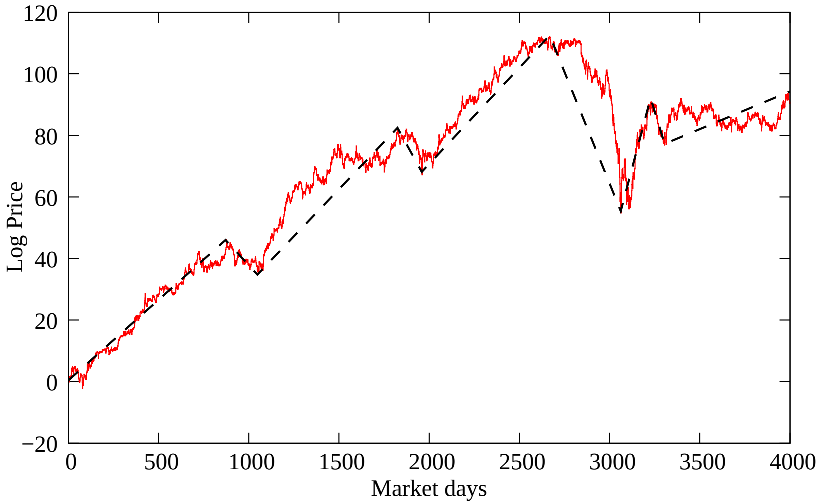

Strategy	Payoff
Invest everything	90
Ideal	1021
Model	178
Bound	105
Actual performance	175

Example

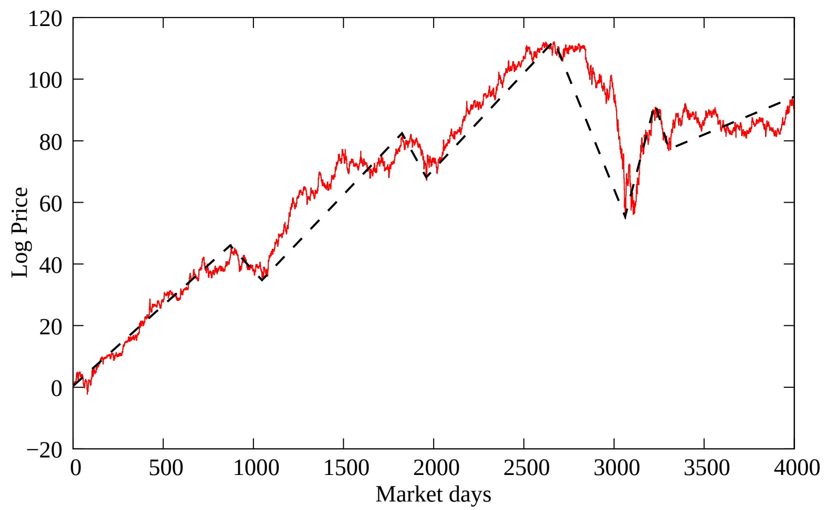

Strategy	Payoff
Invest everything	90
Ideal	1021
Model	178
Bound	105
Actual performance	175

- Performance on real stock: probably not brilliant
- Strategy still useful as a safeguard against excessive loss

Example

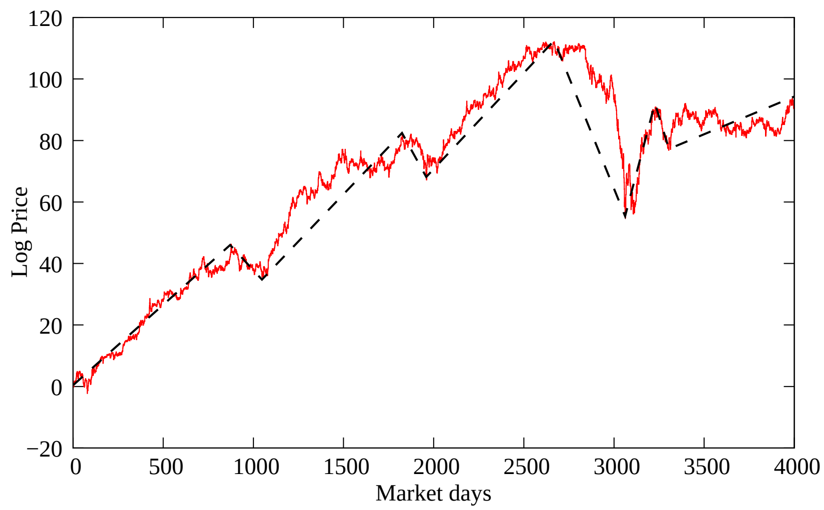

Strategy	Payoff
Invest everything	90
Ideal	1021
Model	178
Bound	105
Actual performance	175

- Performance on real stock: probably not brilliant
- Strategy still useful as a safeguard against excessive loss
- In other applications Λ is usually less adversarial
- Performance is competitive with Fixed Share and typically better than Variable Share for log loss.

Algorithm

A simple algorithm is described in the paper:

- Statisticians: "It's just Bayes"
- Learning Theorists: "It's just the Aggregating Algorithm"

Algorithm

A simple algorithm is described in the paper:

- Statisticians: "It's just Bayes"
- Learning Theorists: "It's just the Aggregating Algorithm"
- Runs in $O(n^2)$ time and $O(n)$ memory.
- If π is memoryless (exponential) running time can be reduced to $O(n)$.

Algorithm

A simple algorithm is described in the paper:

- Statisticians: "It's just Bayes"
- Learning Theorists: "It's just the Aggregating Algorithm"
- Runs in $O(n^2)$ time and $O(n)$ memory.
- If π is memoryless (exponential) running time can be reduced to $O(n)$.
- It buys when you're losing, and sells when you're winning?!

Thanks