Switching Investments

Wouter M. Koolen and Steven de Rooij

October 6, 2010

Centrum Wiskunde & Informatica

What We Do

All About A Line

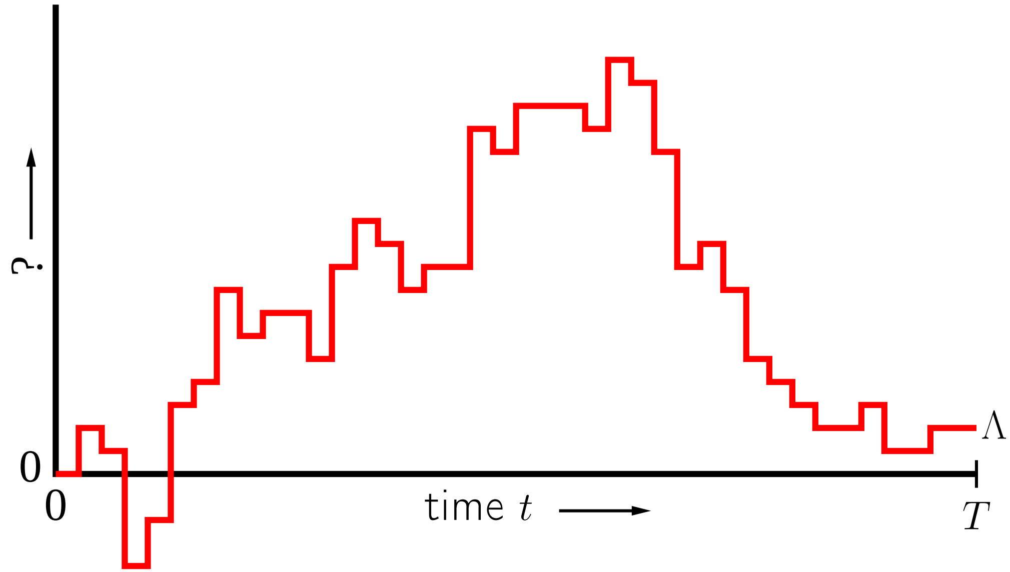

All About A Line

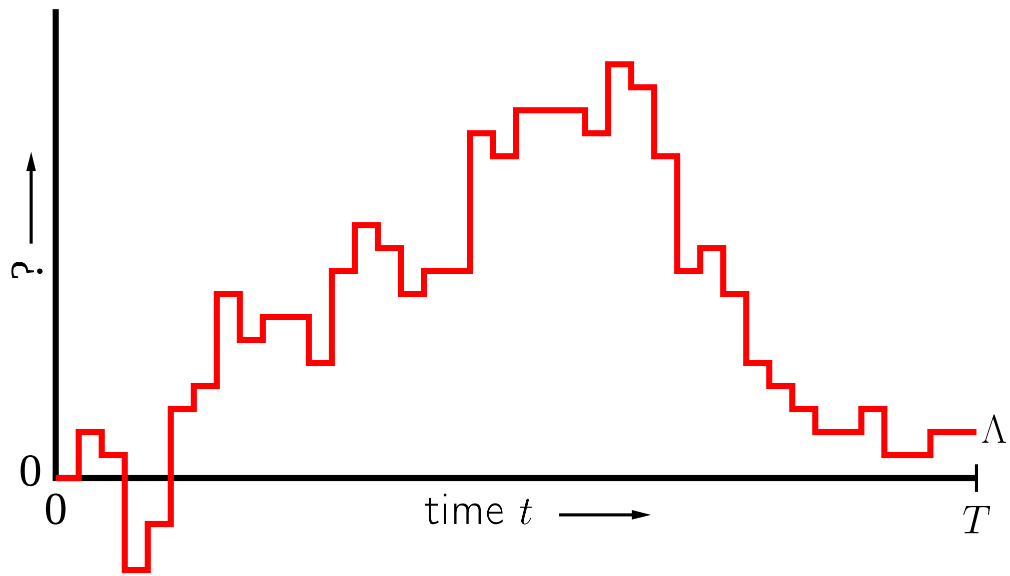

Vertical axis:

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

All About A Line

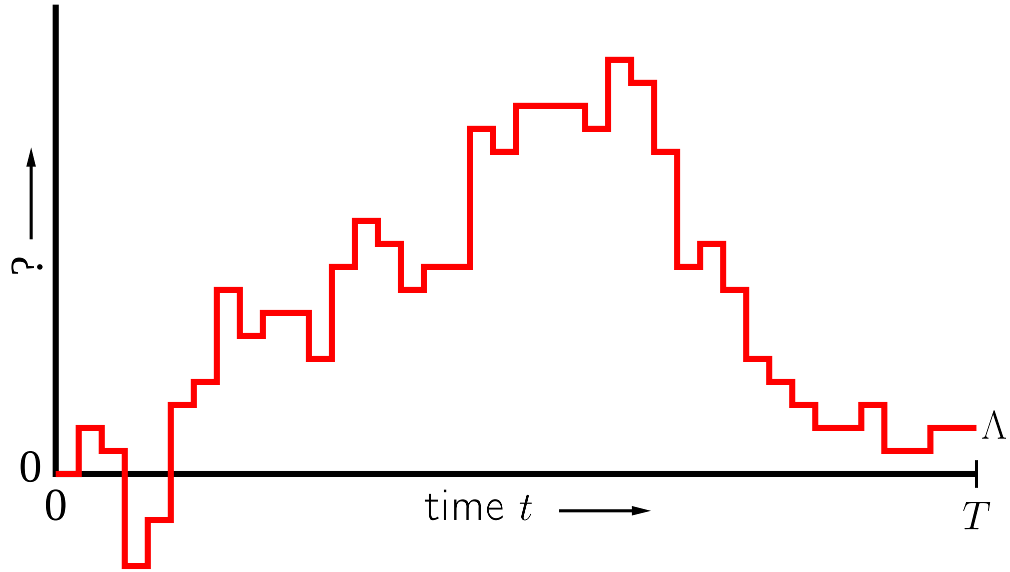

Vertical axis:

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

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

Basic Investment Strategies

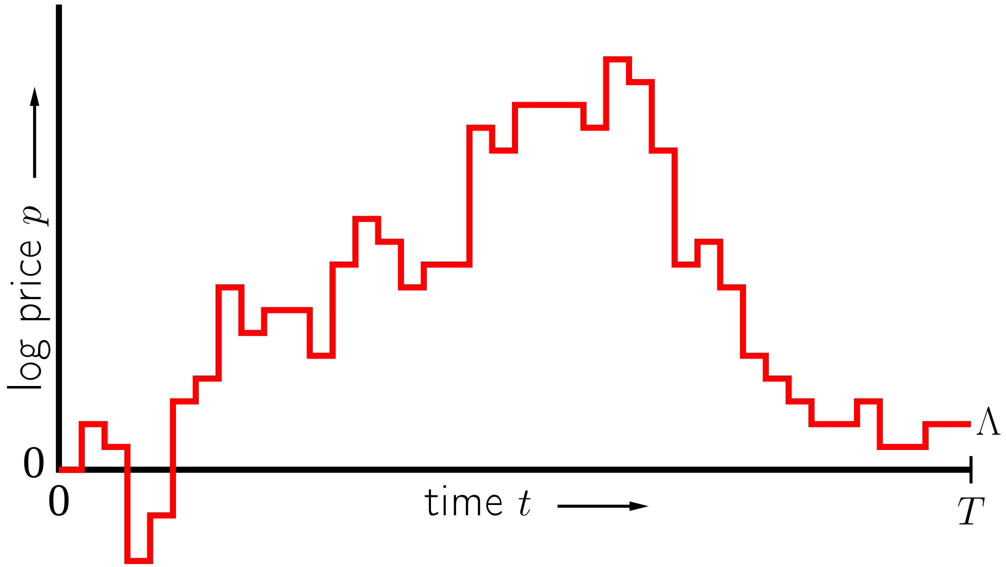

A basic investment strategy σ_t is to sell at a predetermined time t.

Basic Investment Strategies

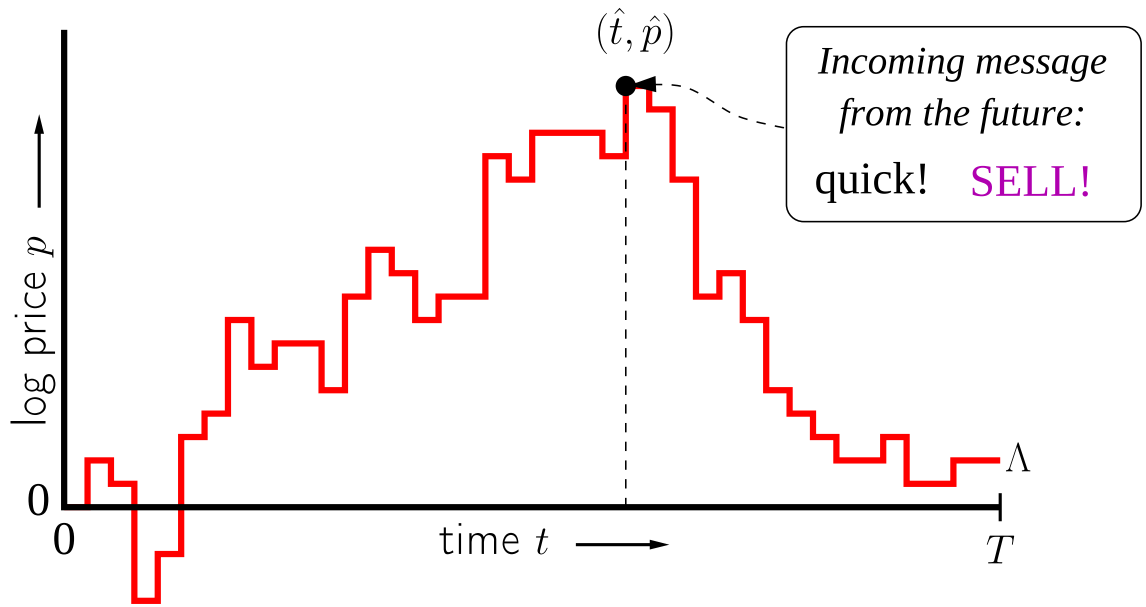

A basic investment strategy σ_t is to sell at a predetermined time t.

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

Hedging

We distribute our initial capital $\$1$ over strategies $\sigma_0, \ldots, \sigma_T$.
Let $\tau(t)$ denote the fraction of capital assigned to σ_t.
Let $\Lambda(0) = 0$. We obtain payoff:

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

Hedging

We distribute our initial capital $\$1$ over strategies $\sigma_0, \ldots, \sigma_T$.
Let $\tau(t)$ denote the fraction of capital assigned to σ_t.
Let $\Lambda(0) = 0$. We obtain payoff:

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

Regret may be relatively large or small, depending on

- ✔ The granularity of measurement

Hedging

We distribute our initial capital $\$1$ over strategies $\sigma_0, \ldots, \sigma_T$.
Let $\tau(t)$ denote the fraction of capital assigned to σ_t.
Let $\Lambda(0) = 0$. We obtain payoff:

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

Regret may be relatively large or small, depending on

- ✔ The granularity of measurement $\leftarrow$ undesirable!

Price Switched Strategies

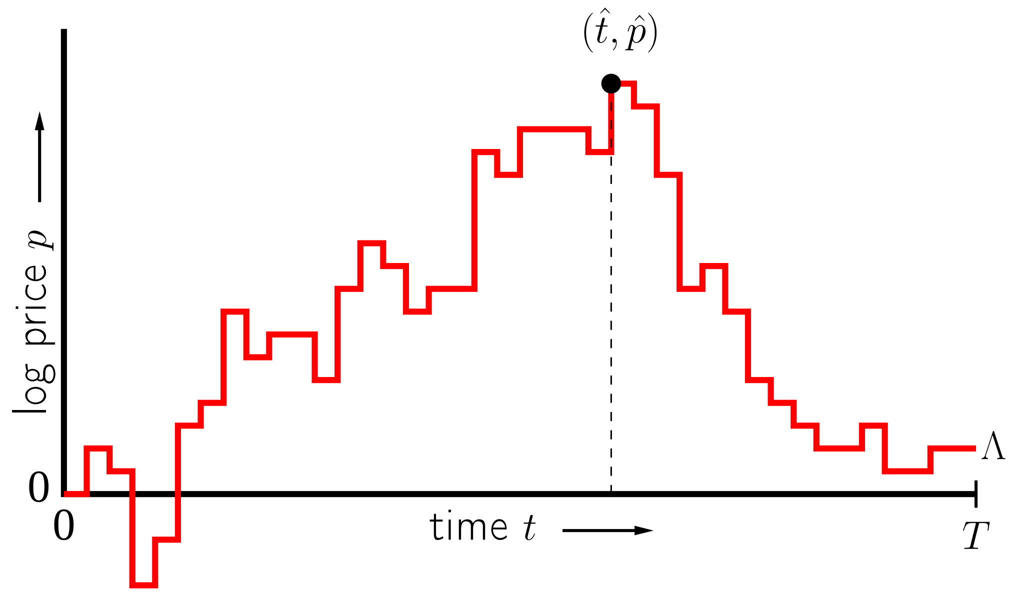

We parameterised the strategy to sell by time t...

Price Switched Strategies

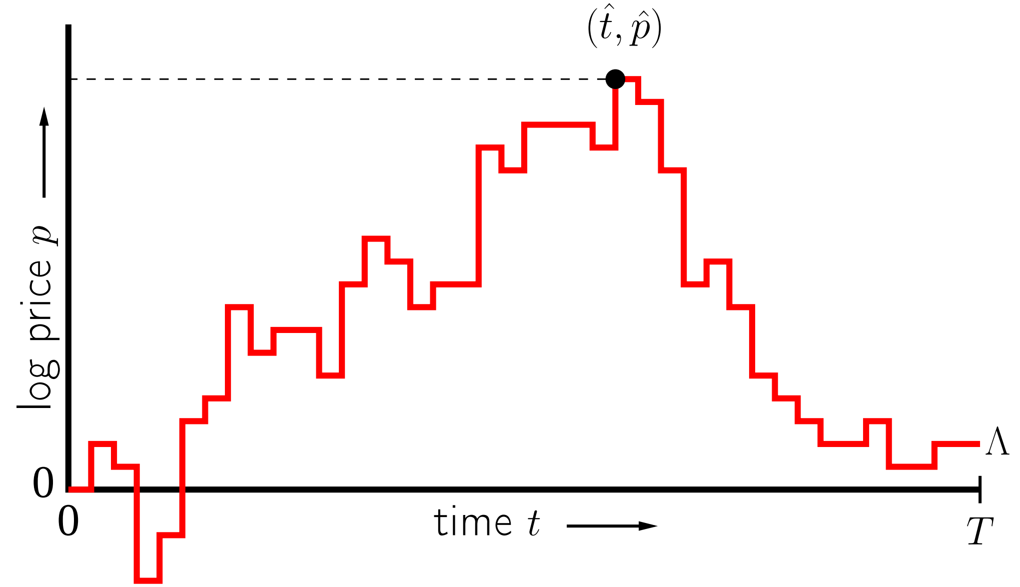

Let us now define σ_p to sell when $\Lambda(t) \geq p$.

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

Price Switched Strategies

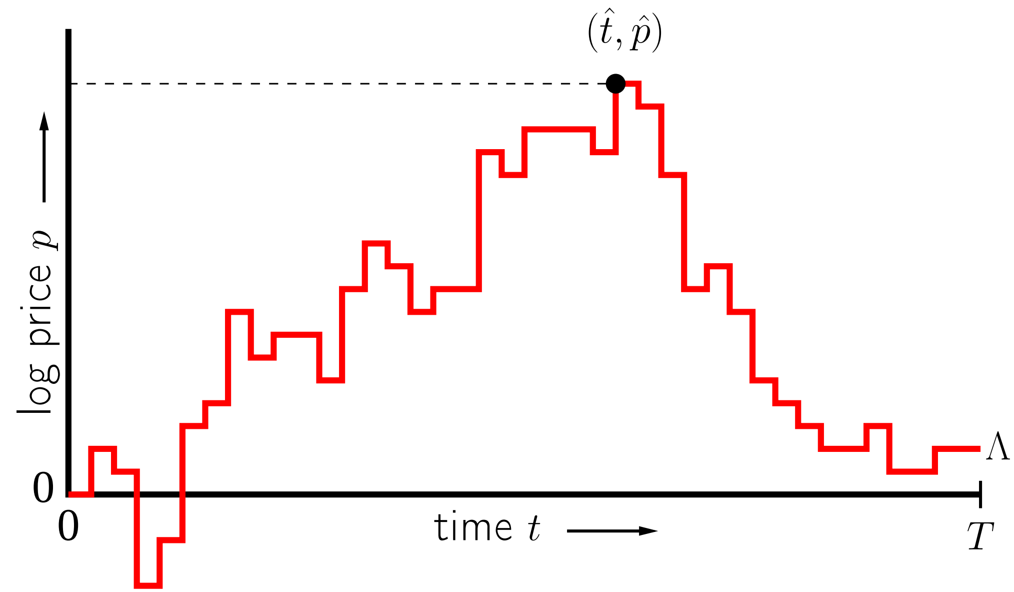

Let us now define σ_p to sell when $\Lambda(t) \geq p$.

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

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

More Price Switching

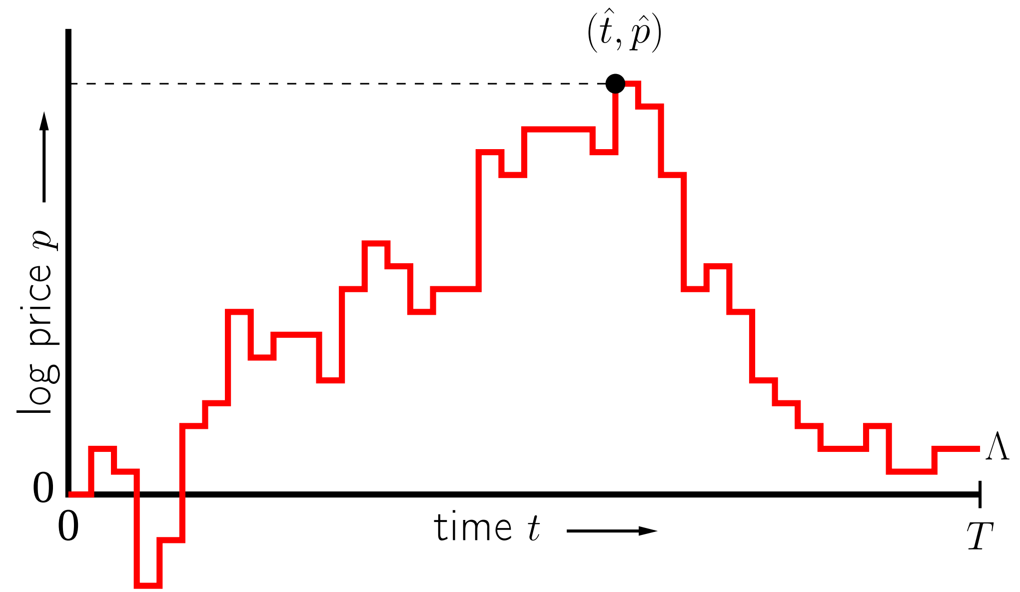

We can hedge, now with π on price levels, to obtain at least

$$\log \sum_{p=0}^{\hat{p}} e^{p} \pi(p) \;\geq\; \log\left(e^{\hat{p}} \pi(\hat{p})\right) \;=\; \underbrace{\hat{p}}_{\text{ideal}} - \underbrace{\left(-\log \pi(\hat{p})\right)}_{\text{regret}}.$$

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

$$\log \int_0^{\hat{p}} e^p \pi(p)\,\mathrm{d}p \geq \log\left(\pi(\hat{p}) \int_0^{\hat{p}} e^p\,\mathrm{d}p\right) = \underbrace{\log(e^{\hat{p}} - 1)}_{\approx \text{ ideal } \hat{p}} - \underbrace{\big(-\log \pi(\hat{p})\big)}_{\text{regret}}.$$

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.

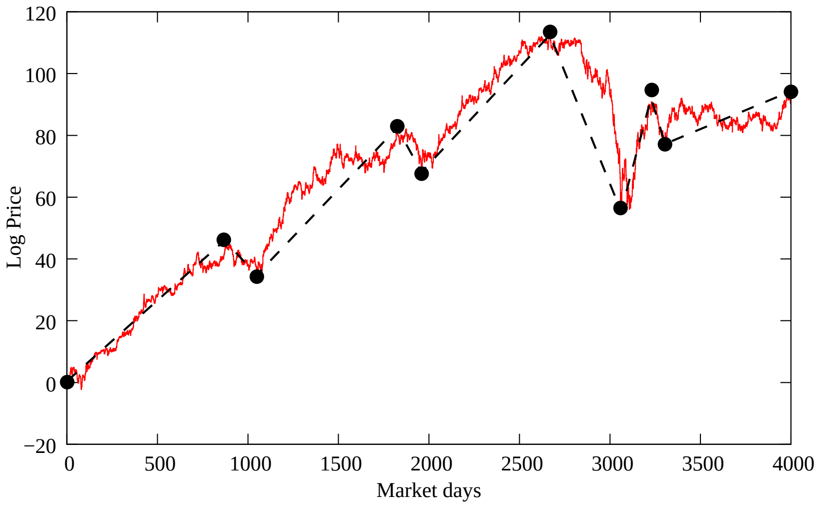

To hedge, take the infinite product distribution of π.

Continuous Time (Theorem 1)

Intuition: Discontinuities in Λ are helpful.

Continuous Time (Theorem 1)

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.)

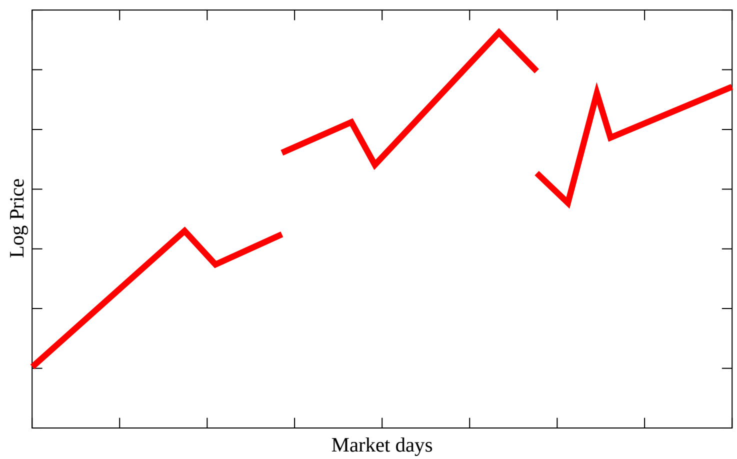

$$\geq$$

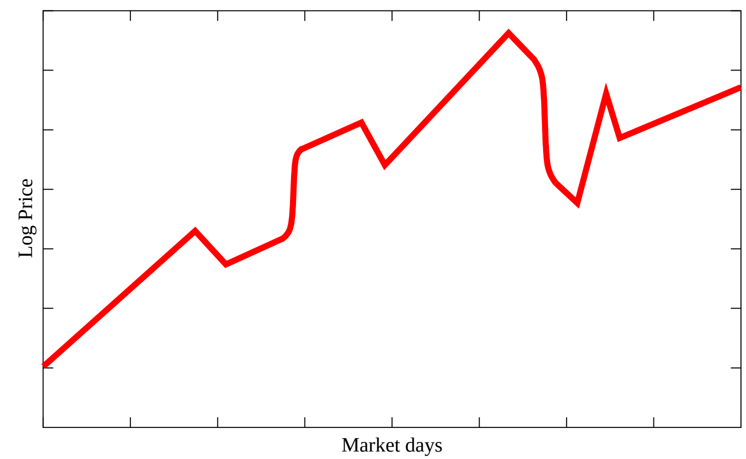

Continuous Time (Theorem 1)

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.)

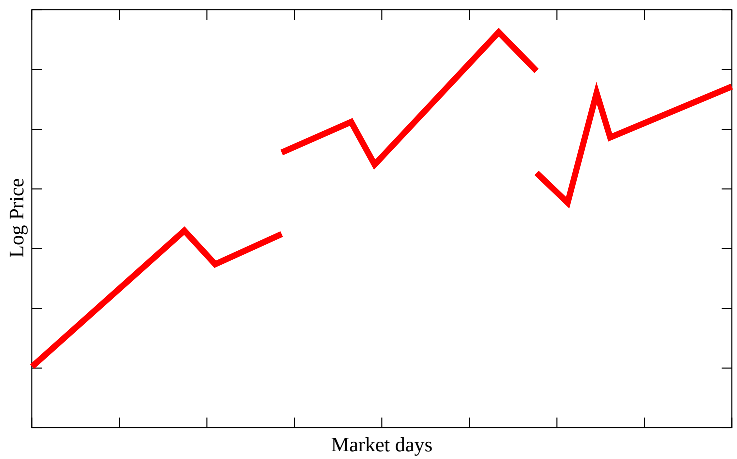

$$\geq$$

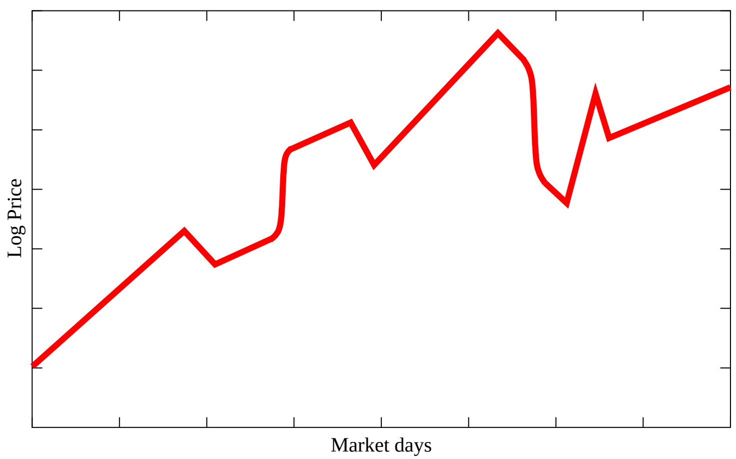

✔ We can simplify the analysis by assuming continuity.

Monotonicity (Theorem 2)

Intuition: The more fluctuations in Λ, the better.

Monotonicity (Theorem 2)

Intuition: The more fluctuations in Λ, the better.

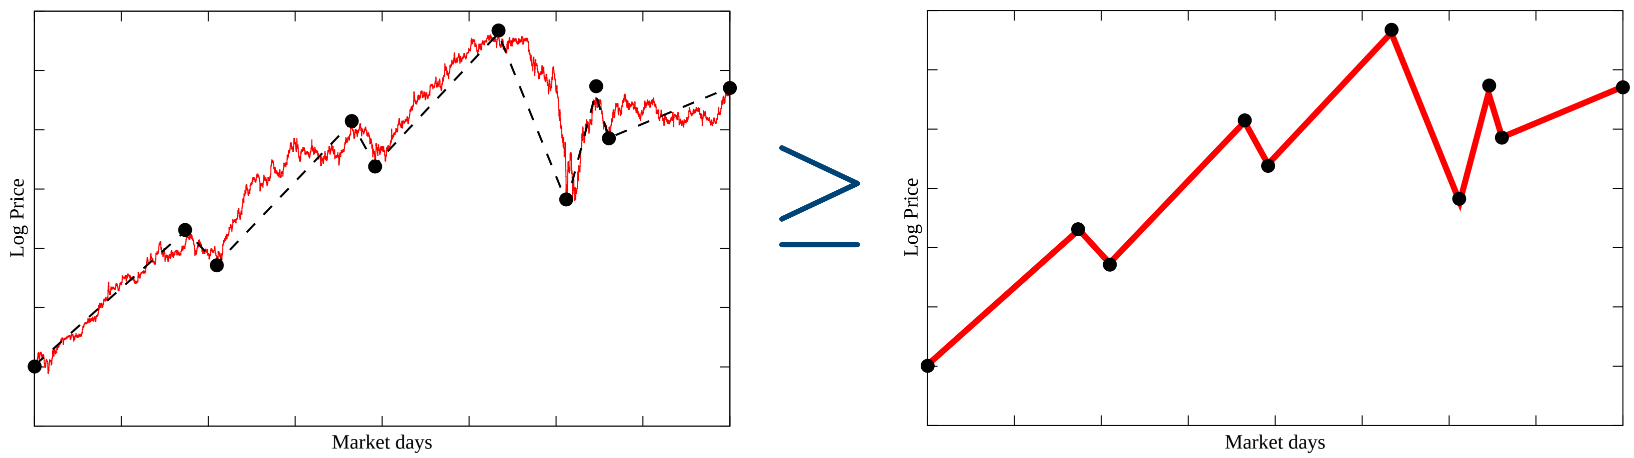

Monotonicity (Theorem 2)

Intuition: The more fluctuations in Λ, the better.

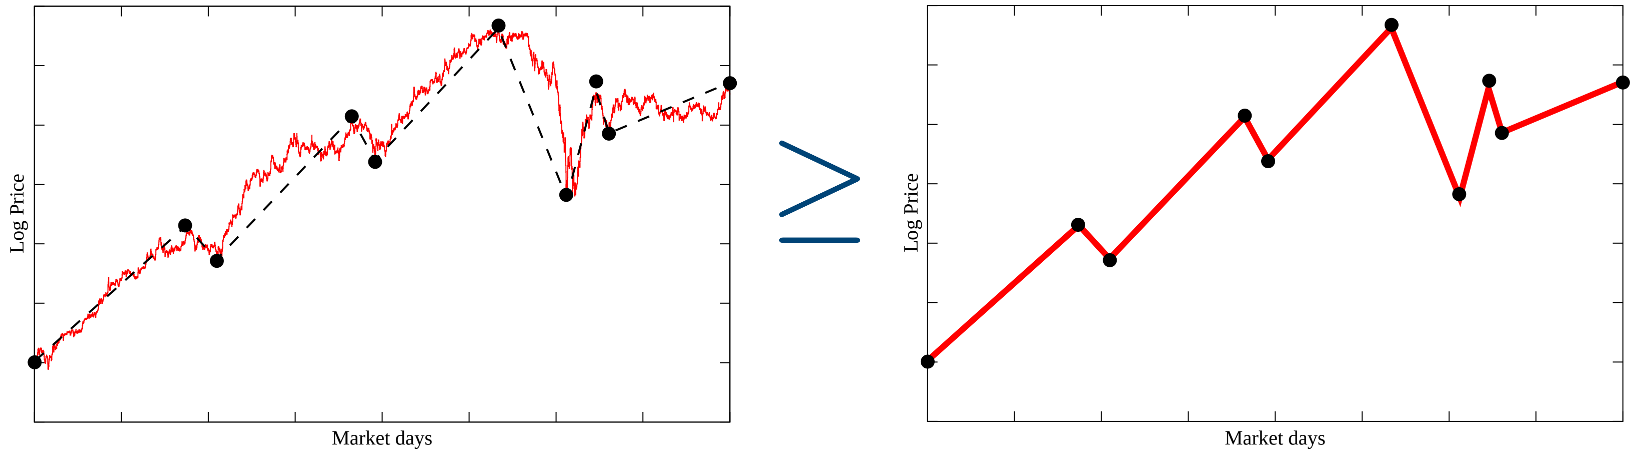

In summary, the regret compared to a specific σ_δ is maximised if

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

The worst case for regret coincides with the ideal case for analysis!

Theorem 3 *Fix Λ. For any basic strategy σ_δ that performs its m^{th} switch on Λ at time T, the payoff of our strategy is at least*

$$\underbrace{\sum_{1 \le \text{odd } i \le m} \delta_i}_{\textit{ideal}} - \underbrace{\sum_{i=1}^{m} -\log \pi(\delta_i) \; - \; m \cdot \textit{small}}_{\textit{regret}} .$$

Regret Bound

Theorem 3 *Fix Λ. For any basic strategy σ_δ that performs its m^{th} switch on Λ at time T, the payoff of our strategy is at least*

$$\underbrace{\sum_{1 \leq \text{odd } i \leq m} \delta_i}_{\textit{ideal}} - \underbrace{\sum_{i=1}^{m} -\log \pi(\delta_i) - m \cdot \textit{small}}_{\textit{regret}}.$$

Thus,

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

Example

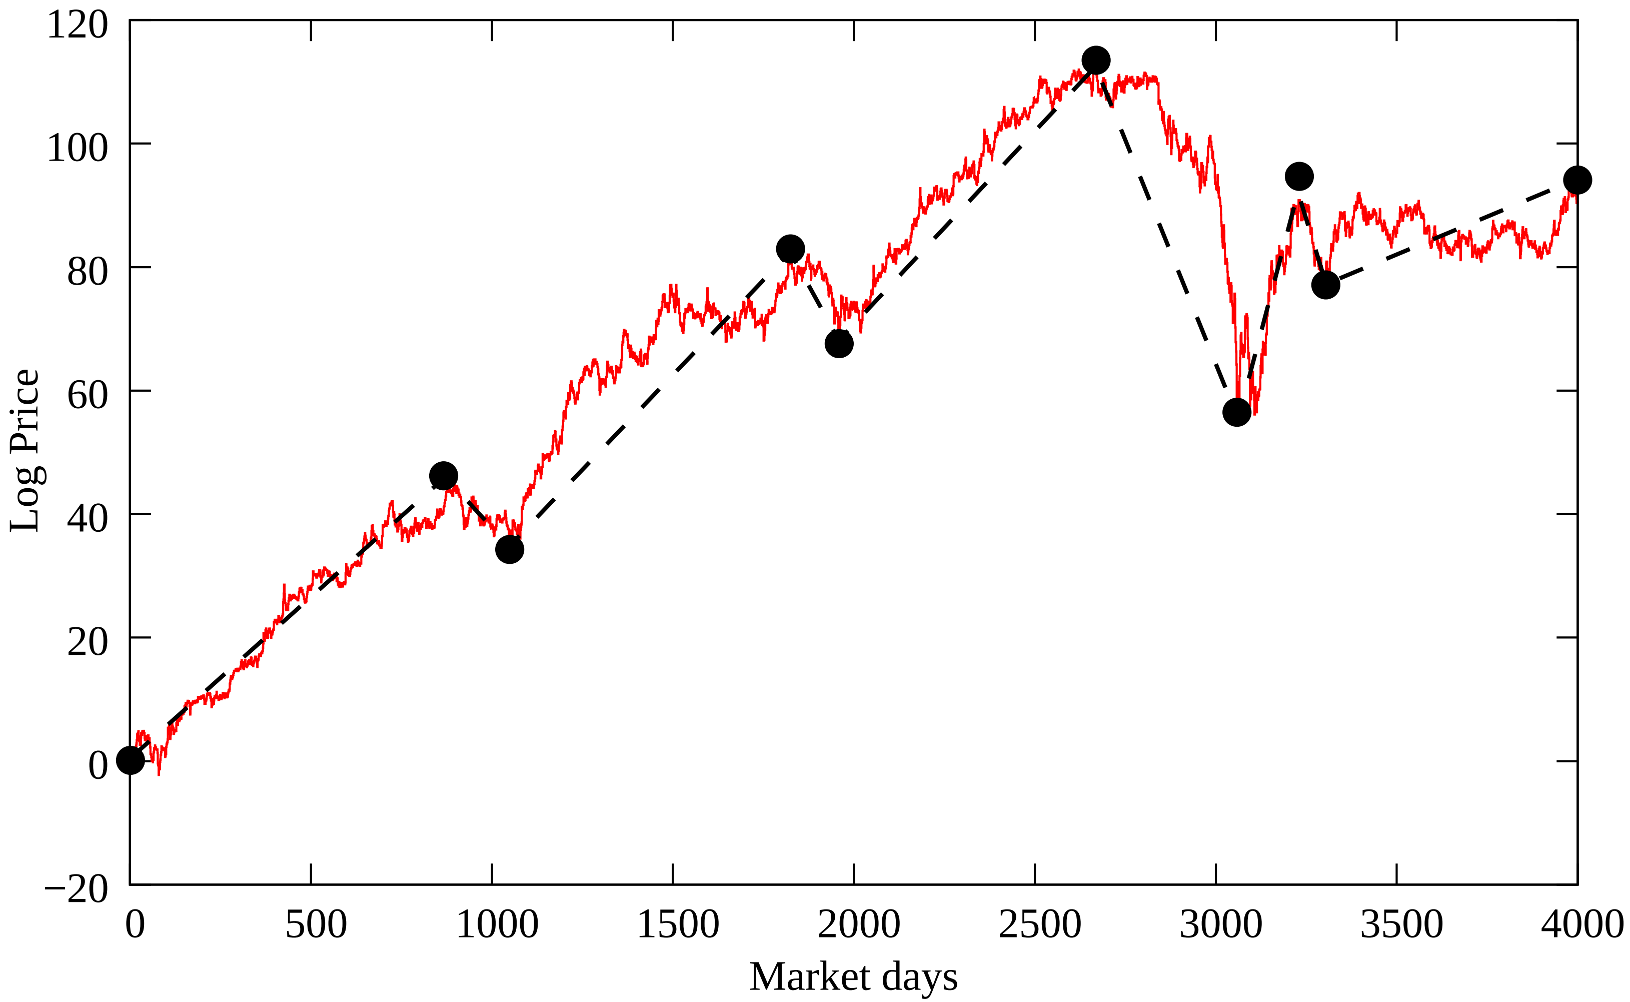

120
100
80
60
40
20
0
−20
0
500
1000
1500
2000
2500
3000
3500
4000
Log Price
Market days

Example

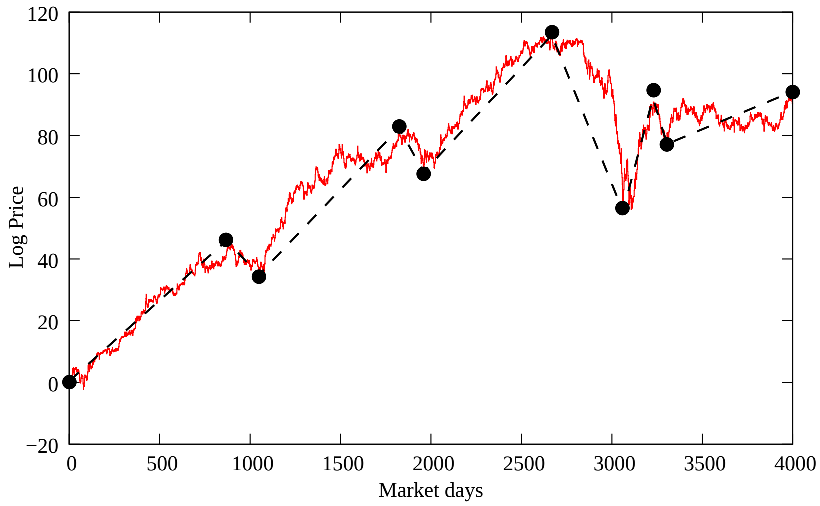

Strategy	Payoff
Invest everything	90
Ideal	1021
Model	178
Bound	105
Actual performance	175

Example

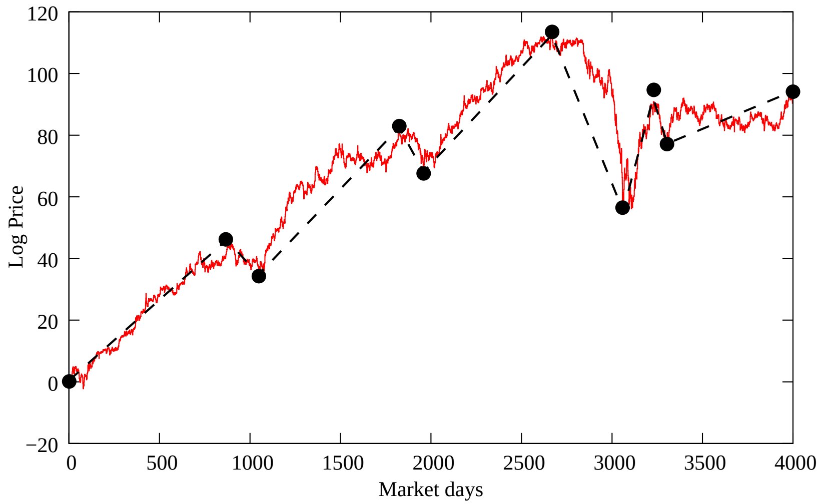

Strategy	Payoff
Invest everything	90
Ideal	1021
Model	178
Bound	105
Actual performance	175

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

Example

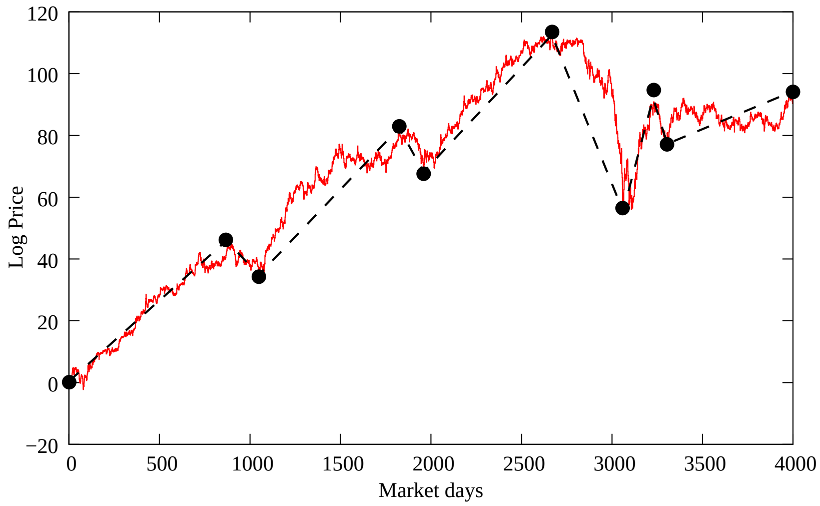

Strategy	Payoff
Invest everything	90
Ideal	1021
Model	178
Bound	105
Actual performance	175

- ✔ 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