Switching Investments

Wouter M. Koolen[1] and Steven de Rooij[2]

[1] Centrum Wiskunde en Informatica (CWI), P.O. Box 94079, NL-1090 GB Amsterdam, wmkoolen@cwi.nl

[2] Statistical Laboratory, DPMMS, Wilberforce Road, Cambridge CB3 0WB, UK, steven@statslab.cam.ac.uk

Abstract. We present a simple online two-way trading algorithm that exploits fluctuations in the unit price of an asset. Rather than analysing worst-case performance under some assumptions, we prove a novel, unconditional performance bound that is parameterised either by the actual dynamics of the price of the asset, or by a simplifying model thereof. The algorithm processes T prices in $O(T^2)$ time and $O(T)$ space, but if the employed prior density is exponential, the time requirement reduces to $O(T)$. The result translates to the prediction with expert advice framework, and has applications in data compression and hypothesis testing.

1 Introduction

We consider a two-player game played between *Investor* and *Nature*. Investor starts out with one unit of cash. At each time, Investor decides which fraction of his current capital to invest in an asset (denoted A), and how much to keep in his boot (denoted B). Nature, on the other hand, chooses the price of the asset.

A play for Nature is a function $\Lambda : [0, T] \to \mathbb{R}$ that specifies the natural logarithm of the unit price of A as a function of time. The end-time T is part of Nature's move and unknown to Investor. An example play is shown in Figure 1.

Investor's *payoff* is defined as the natural logarithm of his capital at the end-time T, where shares owned are valued at the final logprice $\Lambda(T)$. In hindsight, it would have been optimal for Investor to follow the strategy S_Λ that invests all capital in A at local minima of Λ, and liquidates all shares into B at local maxima. Let $\boldsymbol{z} = z_0, \ldots, z_m$ denote the sequence of logprices at local extrema of Λ, with $z_0 = \Lambda(0)$ and $z_m = \Lambda(T)$. The payoff of the strategy S_Λ thus equals

$$S_\Lambda {*} \Lambda \;\coloneqq\; \sum_{1 \le i \le m} \max\{0, z_i - z_{i-1}\}.$$

We construct a foresight-free, computationally efficient strategy π that guarantees payoff $\pi{*}\Lambda$ close to $S_\Lambda{*}\Lambda$. The definition of π relies on the selection of a probability density function on $[0, \infty)$ that for convenience we identify with π itself (see Section 2), and we abbreviate $-\ln \pi(h)$ to $\ell(h)$. We then prove

$$\pi{*}\Lambda \;\ge\; S_\Lambda{*}\Lambda - \sum_{1 \le i \le m} \ell\big(|z_i - z_{i-1}|\big) - (m-1)c_\pi - \ln 2 - 2\epsilon_\pi, \qquad (1)$$

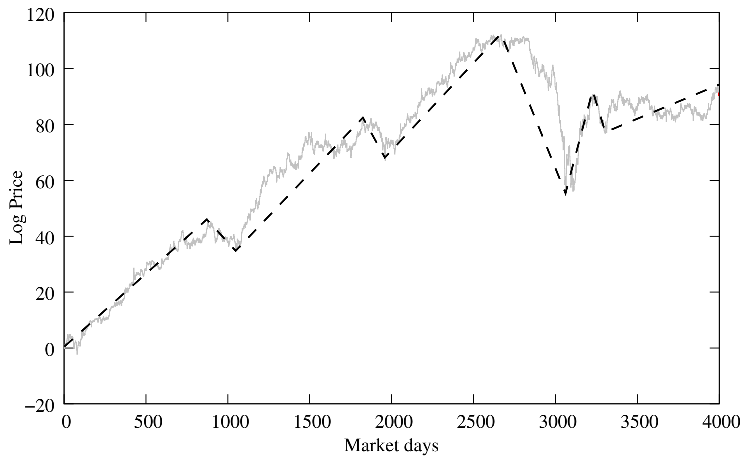

Fig. 1: An example play Λ for Nature, with a regularised trend line Λ'.

where c_π and ϵ_π are two constants that depend on π. Thus the payoff of π on Λ falls short of the optimum by an overhead that depends on the complexity of Λ, measured in terms of both the length of the vector $\boldsymbol{z}$, and the sizes of its entries. The bound is entirely independent of the time scale T.

When Λ is simple, i.e. has few large fluctuations, (1) shows that π exploits almost all achievable payoff. The bound degenerates when Λ sports many small fluctuations, for which the overhead $\ell(x)$ exceeds the benefit x of trading. However, we prove that for any *regularisation* Λ' of Λ, as illustrated by the dashed line in Figure 1 and defined precisely in Section 3.2, π's payoff satisfies

$$\pi{*}\Lambda' \;\leq\; \pi{*}\Lambda. \tag{2}$$

Thus, we may pretend that Nature actually played Λ', and apply the bound (1) with Λ' in place of Λ. In fact the regulariser Λ' may be interpreted as a *model* for Nature's play Λ. The most complex model then yields the bound as presented in (1), but we may now concern other models, that strike a better balance between model complexity and goodness of fit. Such tradeoff models will usually yield better bounds. In conclusion, if in hindsight a simple regulariser can be found with large payoff, then π will collect most of that payoff as well.

Example 1. Let Λ and Λ' be, respectively, the play for Nature and the regulariser shown in Figure 1. The extrema of the regulariser are given by $\boldsymbol{z}' = (0, 42, 36, 82, 68, 112, 57, 90, 77, 90)$. Then $S_{\Lambda'}{*}\Lambda' = (42-0) + (82-36) + \ldots + (90-77) = 178$. Now we select the exponential density listed in Table 1 for the definition of π; the values for c_π and ϵ_π are also listed there. We can now apply bounds (2) and (1) to find

$$\pi{*}\Lambda \;\geq\; \pi{*}\Lambda' \;\geq\; 178 - 64.8 - 8 \cdot 0.034 - \ln 2 - 2 \cdot 3.40 \;\approx\; 105.4.$$

Note that there may be choices of Λ' for which the bound is better, and even for the optimal choice of Λ' the strategy π may perform substantially better than our bound indicates. The actual payoff of π on these data is $\pi{*}\Lambda = 175.4$. $\Diamond$

Applications and Related Work. Our model and its analysis are phrased in financial terms. However, it applies much more widely. We list four examples.

One-Way Trading and Two-Way Trading. This is the most direct example. We let Λ be the logarithm of the exchange rate between any two assets, say dollar and yen. If we forbid selling A, we obtain the setting called *One-way Trading*. Efficient algorithms with minimax payoff for one-way trading under various restrictions on Nature's play Λ are known. E.g. fixed daily price growth range [2], fixed price range [7] and bounded quadratic variation [6]. Two-way trading guarantees are derived in [4] by iterating a unidirectional trading algorithm back and forth. Both the algorithms and the bounds are parametrised by the restrictions placed on Nature's play.

Our results are of a different kind. First, no restrictions are placed on Nature's play. Second, our guarantees are expressed in terms of Nature's actual play (or a regularisation thereof), and hence remain informative when Nature does not play to ruin Investor.

Prediction with Expert Advice. Two experts, say A and B sequentially issue predictions. We denote their cumulative loss at time t by $L_{\mathrm{A}}(t)$ and $L_{\mathrm{B}}(t)$. We let $\Lambda(t) = L_{\mathrm{B}}(t) - L_{\mathrm{A}}(t)$. In prediction tasks with so-called *mixable loss* [13], guarantees for our financial game directly translate to expert performance bounds and vice versa. Efficient strategies include the seminal *Fixed Share* [9], the *Switching Method* [12], and the *Switch Distribution* and its derivatives [8, 10]. These algorithms guarantee payoff $\rho{*}\Lambda \geq S_{\Lambda'}{*}\Lambda' - O(m' \ln T)$ for each Λ' with m' blocks. The logarithmic dependence of the bound on the time T of these algorithms means that for any arbitrary number h, if by switching just a single time the payoff could be improved by h, there is a sample size T such that these algorithms are not able to exploit this.

Variable Share [9] switches based on the losses L^{A} and L^{B}. Its payoff guarantee depends logarithmically on the loss of the best reference strategy with m' blocks. However, its analysis assumes so-called *bounded loss*, and does not apply to financial games (which involve logarithmic loss, which is unbounded).

Prefix Coding/Compression. Fix two prefix codes A and B for a sequence of outcomes $x_1, \ldots, x_T$. Let $L_{\mathrm{A}}(t)$ and $L_{\mathrm{B}}(t)$ denote the code-length of A and B on the outcomes $x_1, \ldots, x_t$ measured in nats. Now let $\Lambda(t) = L_{\mathrm{B}}(t) - L_{\mathrm{A}}(t)$. It is well-known that we can build a prefix code that attains code length $\ln(2) + \min\{L_{\mathrm{A}}(T), L_{\mathrm{B}}(T)\}$ on the data. When different codes are good for different segments of the data, we observe fluctuation in Λ. Using standard information-theoretic methods, e.g. [3], our financial prediction scheme can be transformed into a prefix code that exploits these fluctuations.

Hypothesis Testing. We are given a *null hypothesis* P_0 and an *alternative hypothesis* P_1. Both candidate hypotheses are probabilistic models for some sequence of observations $x_1, x_2, \ldots, x_T$. Let $\Lambda(t) = \ln\bigl(P_1(x_1, \ldots, x_t)/P_0(x_1, \ldots, x_t)\bigr)$ be the loglikelihood ratio between P_1 and P_0. Thus Λ measures the amount of evidence against the null hypothesis and can be used as a test statistic. Traditionally [1], we choose a threshold $\tau > 0$ and reject the null hypothesis when $\Lambda(T) \geq \tau$, an event that is extremely unlikely under P_0. The case where $\Lambda(T)$ is below the threshold τ, while $\Lambda(t) \geq \tau$ at some earlier time t is considered in [11, 5], and tests are presented that lose as little evidence as possible while remaining unbiased. These tests are based on strategies that switch only once, and resemble strategies for one-way trading. By the same method, our strategy induces a fair test statistic that can be used to reject P_0 whenever Λ fluctuates heavily; an event that is also unlikely under P_0.

Outline. We explicate the setting and describe the strategy π for Investor in Section 2. We analyse the payoff of π and prove our payoff guarantee in Section 3. We then show how to implement the strategy π efficiently in Section 4.

2 Setting

We introduce the details of our financial game. We first review Nature's play Λ. We then construct strategies for Investor, culminating in the definition of the strategy π. We conclude this section with a lemma that simplifies all later proofs by exploiting the symmetry between A and B.

2.1 Nature's Play $\boldsymbol{\Lambda}$

A play for Nature is a logprice function $\Lambda : [0, T] \to \mathbb{R}$. The end-time T is part of Nature's move, and unknown to Investor.

For simplicity, we restrict attention to the setting where Λ is discrete, i.e. piecewise constant with jumps at integer times. This is sufficient for the practical scenario where Λ is monitored intermittently (albeit possibly very often). Later in the analysis it will be convenient for technical reasons to generalise to piecewise continous plays for Nature with finitely many local extrema and finitely many discontinuities; nevertheless ultimately we remain concerned with the discrete setting only.

Our results do extend quite readily to the wide class of càdlàg logprice functions (right-continuous with left limits). These encompass continuous time models that are often considered in the financial literature, such as Brownian motion with drift, etc. Such theoretically interesting generalisations are deferred to future publications.

2.2 Investor's Strategy $\boldsymbol{\pi}$

We now construct the strategy π for Investor in three stages. Two basic strategies exist. Strategy A invests the initial unit capital in the asset, whereas strategy

B keeps all capital in the boot. At the end of the game, all shares are valued at the final logprice $\Lambda(T)$. The payoffs, defined as Investor's final logcapital, of the basic strategies equal

$$\mathrm{A}{*}\Lambda \;\coloneqq\; \Lambda(T) - \Lambda(0) \qquad\qquad \text{and} \qquad\qquad \mathrm{B}{*}\Lambda \;\coloneqq\; 0.$$

Since we use logprice differences extensively, we abbreviate $\Lambda(t) - \Lambda(s)$ to $\Lambda|_s^t$.

Time-switched Strategies. From these basic strategies A and B we construct more interesting strategies. Let $\boldsymbol{t} = t_0, t_1, t_2, t_3, \ldots$ be a sequence of times such that $0 = t_0 \leq t_1 \leq t_2 \leq \ldots$ The strategy $\boldsymbol{t}^{\mathrm{A}}$ switches at times $\boldsymbol{t}$ starting with A. That is, $\boldsymbol{t}^{\mathrm{A}}$ invests all capital in A until time t_1. At that time it sells all shares, and keeps all money in B until time t_2. Then it again invests all capital in A until time t_3 etc. Symmetrically, $\boldsymbol{t}^{\mathrm{B}}$ is the strategy that switches at times $\boldsymbol{t}$ starting with B. Thus the payoffs of $\boldsymbol{t}^{\mathrm{A}}$ and $\boldsymbol{t}^{\mathrm{B}}$ when Nature plays Λ are

$$\boldsymbol{t}^{\mathrm{A}}{*}\Lambda \;\coloneqq\; \sum_{i=0}^{\infty} \Lambda|_{T\wedge t_{2i}}^{T\wedge t_{2i+1}} \qquad\qquad \text{and} \qquad\qquad \boldsymbol{t}^{\mathrm{B}}{*}\Lambda \;\coloneqq\; \sum_{i=0}^{\infty} \Lambda|_{T\wedge t_{2i+1}}^{T\wedge t_{2i+2}}.$$

Of course, a good time switch sequence $\boldsymbol{t}$ for Investor depends on Nature's unknown move Λ. However, Investor may hedge by dividing his initial capital according to some prior distribution ρ on the switch time sequence $\boldsymbol{t}$, and construct time-switched strategies ρ^{A} and ρ^{B} with payoffs

$$\rho^{\mathrm{A}}{*}\Lambda \;\coloneqq\; \ln \int \exp\bigl(\boldsymbol{t}^{\mathrm{A}}{*}\Lambda\bigr)\, \mathrm{d}\rho(\boldsymbol{t}) \quad \text{and} \quad \rho^{\mathrm{B}}{*}\Lambda \;\coloneqq\; \ln \int \exp\bigl(\boldsymbol{t}^{\mathrm{B}}{*}\Lambda\bigr)\, \mathrm{d}\rho(\boldsymbol{t}),$$

and the meta strategy ρ with payoff $\rho{*}\Lambda \coloneqq \ln\left(\frac{1}{2}\exp\bigl(\rho^{\mathrm{A}}{*}\Lambda\bigr) + \frac{1}{2}\exp\bigl(\rho^{\mathrm{B}}{*}\Lambda\bigr)\right)$.

Price-switched Strategies. Price-switched strategies decide when to trade based on the logprice $\Lambda(t)$ instead of the time t itself. This renders their payoff independent of the time-scale. Fix a sequence of nonnegative reals $\boldsymbol{\delta} = \delta_1, \delta_2, \ldots$ We denote by $\boldsymbol{\delta}^{\mathrm{A}}$ the strategy that initially invests all capital in A, and waits until the first time s_1 where the logprice difference $\Lambda|_0^{s_1}$ is at least δ_1. It then sells all shares and puts the money into B, until the first subsequent time s_2 that the logprice difference $\Lambda|_{s_1}^{s_2}$ is at most $-\delta_2$. Then it invests all capital into A again, until the logprice difference $\Lambda|_{s_2}^{s_3}$ is at least δ_3, etc. The strategy $\boldsymbol{\delta}^{\mathrm{B}}$ is defined symmetrically, with switching times $r_0, r_1, \ldots$ The switching time sequences $\boldsymbol{s}$ and $\boldsymbol{r}$ are obtained as follows. First $s_0 = r_0 = 0$. Then recursively

$$\begin{aligned}
s_i &\coloneqq \min\bigl\{t \geq s_{i-1} \bigm| \Lambda|_{s_{i-1}}^t \geq +\delta_i\bigr\} \quad & r_i &\coloneqq \min\bigl\{t \geq r_{i-1} \bigm| \Lambda|_{r_{i-1}}^t \leq -\delta_i\bigr\} \quad & &i \text{ even},\\
s_i &\coloneqq \min\bigl\{t \geq s_{i-1} \bigm| \Lambda|_{s_{i-1}}^t \leq -\delta_i\bigr\} \quad & r_i &\coloneqq \min\bigl\{t \geq r_{i-1} \bigm| \Lambda|_{r_{i-1}}^t \geq +\delta_i\bigr\} \quad & &i \text{ odd}.
\end{aligned}$$

Both $\boldsymbol{s}$ and $\boldsymbol{r}$ are a function of $\boldsymbol{\delta}$ and Λ and satisfy $\boldsymbol{s}(\boldsymbol{\delta}, \Lambda) = \boldsymbol{r}(\boldsymbol{\delta}, -\Lambda)$. By convention, the minimum is infinite if no suitable successor time exists in the domain of Λ, i.e before time T. The payoffs of $\boldsymbol{\delta}^{\mathrm{A}}$ and $\boldsymbol{\delta}^{\mathrm{B}}$ are given by

$$\boldsymbol{\delta}^{\mathrm{A}}{*}\Lambda \;\coloneqq\; \boldsymbol{s}^{\mathrm{A}}{*}\Lambda \qquad\qquad \text{and} \qquad\qquad \boldsymbol{\delta}^{\mathrm{B}}{*}\Lambda \;\coloneqq\; \boldsymbol{r}^{\mathrm{B}}{*}\Lambda.$$

The strategy $\boldsymbol{\delta}^{\mathrm{A}}$ has the following property. Whenever it sells its shares, say at time s_i for some odd i, the asset price, and hence its capital, has multiplied by *at least* $\exp(\delta_i) \geq 1$ since the acquisition at time s_{i-1}. This holds irrespective of Nature's play. In particular, between time s_i and s_{i+1} for odd i, the logarithm of its capital equals

$$\Lambda|_{s_0}^{s_1} + \Lambda|_{s_2}^{s_3} + \Lambda|_{s_4}^{s_5} + \ldots + \Lambda|_{s_{i-1}}^{s_i} \;\geq\; \delta_1 + \delta_3 + \delta_5 + \ldots + \delta_i.$$

Of course, for each logprice difference sequence $\boldsymbol{\delta}$, the number of switches that is executed, and hence the quality of $\boldsymbol{\delta}^{\mathrm{A}}$ depends on Nature's move Λ. Let $\mathcal{D} = \left\{\boldsymbol{\delta}^{\mathrm{A}}, \boldsymbol{\delta}^{\mathrm{B}} \;\middle|\; \boldsymbol{\delta} \in [0,\infty)^\infty\right\}$ be the set of price-switched strategies for Investor.

The Strategy $\boldsymbol{\pi}$. Again, we may hedge by dividing our initial capital according to some prior π on $\boldsymbol{\delta}$, and obtain strategies π^{A} and π^{B} with payoffs

$$\pi^{\mathrm{A}}{*}\Lambda \;\coloneqq\; \ln \int \exp\bigl(\boldsymbol{\delta}^{\mathrm{A}}{*}\Lambda\bigr)\, \mathrm{d}\pi(\boldsymbol{\delta}) \quad \text{and} \quad \pi^{\mathrm{B}}{*}\Lambda \;\coloneqq\; \ln \int \exp\bigl(\boldsymbol{\delta}^{\mathrm{B}}{*}\Lambda\bigr)\, \mathrm{d}\pi(\boldsymbol{\delta}),$$

and the meta strategy π with payoff $\pi{*}\Lambda \coloneqq \ln\left(\frac{1}{2}\exp\bigl(\pi^{\mathrm{A}}{*}\Lambda\bigr) + \frac{1}{2}\exp\bigl(\pi^{\mathrm{B}}{*}\Lambda\bigr)\right)$. Note that the price-switched strategies in $\mathcal{D}$ are independent of the time scale, and so are these strategies based on them.

Requirements on π. The above construction works for any prior π. In this paper we analyse the behaviour of strategies π that satisfy these requirements:

1. π is the independent infinite product distribution of some probability density function on $[0,\infty)$. Since the distinction is always clear, we also denote the univariate density by π.
2. the function $x \mapsto \mathrm{e}^x \pi(x)$ is increasing.
3. the density π is log-convex.

The first requirement ensures that we can hedge capital according to π. The second requirement ensures that paying $-\ln \pi(x)$ to gain x is a better deal when x is larger. The third requirement ensures that we rather pay $-\ln \pi(x+y)$ than $-\ln \pi(x) - \ln \pi(y)$ to gain $x+y$. We use the following consequences in our bounds.

Lemma 1. *Let π satisfy the requirements 1–3 above. Then*

1. π is strictly positive.
2. π is strictly decreasing.
3. $\int_h^\infty \pi(x)\,\mathrm{d}x \geq \pi(h)$ for each $h \geq 0$.

Proof. Since π is a convex probability density, it is decreasing and thus $0 < \pi(0) = \mathrm{e}^0\pi(0)$. Since $\mathrm{e}^x\pi(x)$ increases, we have $\pi(x) > 0$ for all x. Then, since π is a non-zero convex probability density, it must be *strictly* decreasing. Finally, for $0 \leq h \leq x$ we have $\pi(x) = \pi(x)\mathrm{e}^x\mathrm{e}^{-x} \geq \pi(h)\mathrm{e}^h\mathrm{e}^{-x}$. Therefore $\int_h^\infty \pi(x)\,\mathrm{d}x \;\geq\; \pi(h)\int_h^\infty \mathrm{e}^{h-x}\,\mathrm{d}x \;=\; \pi(h)$. □

The last fact implies that the density $\pi(x) \leq 1$ for all x. Throughout this paper, we abbreviate $-\ln \pi(x)$ to $\ell(x)$. Thus ℓ is nonnegative, concave and increasing.

Example 2. The densities shown in Table 1, ordered from heavy to light tails, satisfy all the requirements. ◊

Table 1: Example priors.

	Fat tail	Pareto	Exponential
$\pi(x)$	$\dfrac{\log(o)}{(x+o)(\log(x+o))^2}$	$(c-1)o^{c-1}(x+o)^{-c}$	$\alpha e^{-\alpha x}$
Condition	$2 \leq (o-1)\log o$ (Sufficient: $o \geq 2.89$)	$1 < c \leq o$	$0 < \alpha \leq 1$
Parameters	$o = 3$	$c = 2,\, o = 3$	$\alpha = 1/3$
ϵ_π	4.10396	3.55884	3.39788
c_π	0.016645	0.0288849	0.034016

2.3 Exploiting Symmetry

Payoff is measured as (the natural logarithm of) Investor's final amount of cash. Of course, cash and asset are intrinsically symmetric. We make this precise as follows. We say that the following pairs of strategies are *dual*

$$\mathrm{A}, \mathrm{B} \qquad \boldsymbol{t}^{\mathrm{A}}, \boldsymbol{t}^{\mathrm{B}} \qquad \rho^{\mathrm{A}}, \rho^{\mathrm{B}} \qquad \rho, \rho \qquad \boldsymbol{\delta}^{\mathrm{A}}, \boldsymbol{\delta}^{\mathrm{B}} \qquad \pi^{\mathrm{A}}, \pi^{\mathrm{B}} \qquad \pi, \pi$$

and vice versa in each case. The meta strategies ρ and π are self-dual.

Lemma 2 (Duality). *Let S and S' be dual strategies. Then for each Λ*

$$S{*}\Lambda \;=\; S'{*}(-\Lambda) + \Lambda|_0^T.$$

Proof. The lemma is trivial for the dual pair A and B. We proceed to prove the lemma for the dual strategies $\boldsymbol{t}^{\mathrm{A}}$ and $\boldsymbol{t}^{\mathrm{B}}$, the other cases follow simply by definition. Recall that $\exp(\Lambda)$ is the asset price in cash per share, so that $\exp(-\Lambda)$ is the price in shares per cash. Thus $\boldsymbol{t}^{\mathrm{B}}{*}(-\Lambda)$ is the log-number of shares resulting from investing one share according to the strategy $\boldsymbol{t}^{\mathrm{A}}$. Finally, $\Lambda|_0^T = \Lambda(T) - \Lambda(0)$ is the result of exchanging cash to asset initially, and asset to cash at the end. □

3 Payoff Bound

In this section we prove the payoff guarantees for the strategy π that were given in the introduction. We build towards the statement and proof of a more precise version of the bounds in the following subsections. First, in Section 3.1 we show that Nature's worst-case logprice functions are continuous. Then, in Section 3.2 we show that Investor's payoff decreases when Nature plays more regular. In Section 3.3 we analyse Investor's payoff under a regularity assumption on Λ called γ-separation. Finally, in Section 3.4 we show how to establish γ-separation if it does not obtain and establish the bound in the form of Theorem 5.

3.1 Nature Plays a Continuous Logprice Function Λ

We now prove that it is sub-optimal for Nature to play a discontinuous Λ. To do so, we show that Investor's payoff is reduced when Nature eliminates a jump by inserting a linear interpolation. Let Λ have a discontinuity at t. We define Λ', *the t-ironing of Λ*, by $\Lambda'(s) \coloneqq \Lambda(s)$ for $s < t$, $\Lambda'(s+1) \coloneqq \Lambda(s)$ for $s > t$, and $\Lambda'(s) \coloneqq (1+t-s)\Lambda(t-)+(s-t)\Lambda(t)$ for $t \leq s \leq t+1$, where $\Lambda(t-) \coloneqq \lim_{s\uparrow t} \Lambda(s)$. This definition is illustrated by Figure 2.

Theorem 1 (Continuous Free Lunch). *Fix any play for Nature Λ with a discontinuity at time t, and let Λ' be the t-ironing of Λ. Then*

$$\pi{*}\Lambda' \;\leq\; \pi{*}\Lambda.$$

Proof. See Figure 2. By duality (Lemma 2), we may assume that the jump is upward. Obviously, any strategy $\boldsymbol{\delta}'$ that does not switch at time t on Λ has identical payoff on Λ and Λ'. Now consider any strategy $\boldsymbol{\delta}' = (\dots, h, h-l, \dots)$, where h prompts a switch at time t on Λ. We now modify the strategy to $\boldsymbol{\delta} = (\dots, h, u-l, \dots)$ and we compare the term corresponding to $\boldsymbol{\delta}'$ in the integral for $\pi{*}\Lambda'$ to the term corresponding to $\boldsymbol{\delta}$ in the integral for $\pi{*}\Lambda$:

$$\frac{\exp(\boldsymbol{\delta}'{*}\Lambda')\pi(\boldsymbol{\delta}')}{\exp(\boldsymbol{\delta}{*}\Lambda)\pi(\boldsymbol{\delta})} \;=\; \frac{\exp(h-l)\pi(h-l)}{\exp(u-l)\pi(u-l)} \;\leq\; 1,$$

where the inequality uses that $\mathrm{e}^h\pi(h)$ is increasing (see Section 2.2). The proof follows by observing that the mapping that takes $\boldsymbol{\delta}'$ to $\boldsymbol{\delta}$ is a translation. $\square$

When Investor follows the strategy π, there is no benefit for Nature to playing a logprice function Λ with jumps. Without loss of generality we henceforth restrict Nature to continuous plays. This simplifies analysis considerably, as it allows us to assume that switches specified by any $\boldsymbol{\delta}$ occur at *exactly* the specified logprices.

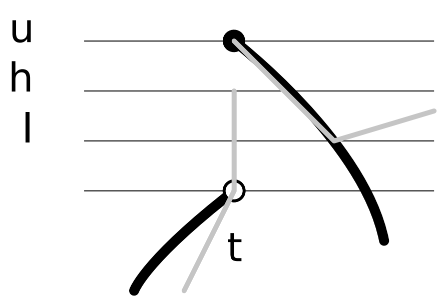

(a) Nature's move Λ with jump at time t, and strategy $\boldsymbol{\delta} = (\dots, h, u-l, \dots)$.

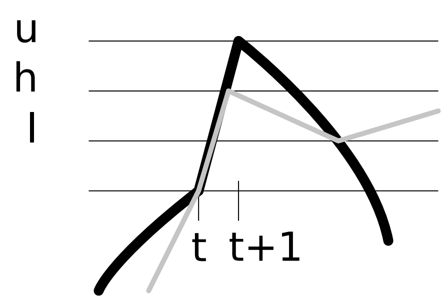

(b) The t-ironing Λ' of Λ, and strategy $\boldsymbol{\delta}' = (\dots, h, h-l, \dots)$

Fig. 2: Worst-Case Plays for Nature are Continuous.

3.2 Ordering by Regularity

Given a move for Nature $\Lambda : [0, T] \to \mathbb{R}$, we say that another move $\Lambda' : [0, T'] \to \mathbb{R}$ is *more regular* than Λ, denoted $\Lambda' \preccurlyeq \Lambda$, if there is a monotonic function $f : [0, T'] \to [0, T]$ such that $f(0) = 0$, $f(T') = T$ and $\Lambda' = \Lambda \circ f$. That is, the price levels of the *regularisation* Λ' are a subsequence of the price levels of Nature's move Λ, with the same initial and final price, but potentially less fluctuation. We now show that by following a fixed price-switched strategy, Investor gets richer whenever Nature's move is less regular.

Theorem 2 (Monotonicity). *For each price-switched strategy $S \in \mathcal{D}$ and continuous logprice functions Λ and Λ'*

$$\Lambda' \preccurlyeq \Lambda \qquad \textit{implies} \qquad S{*}\Lambda' \leq S{*}\Lambda.$$

Proof. First note that $\Lambda' \preccurlyeq \Lambda$ iff $-\Lambda' \preccurlyeq -\Lambda$. So by symmetry (Lemma 2) it suffices to prove the theorem for the strategies in $\mathcal{D}$ that start with A. We proceed by induction on the number of switches executed by the strategy $\boldsymbol{\delta}^{\mathrm{A}}$ on the regulariser Λ'. For the base case, suppose this number is zero, i.e $\Lambda'|_0^t < \delta_1$ for each $0 \leq t \leq T'$. Let $m \geq 1$ denote the number of blocks of $\boldsymbol{\delta}^{\mathrm{A}}$ on Λ. There are two cases. If m is even then $\boldsymbol{\delta}^{\mathrm{A}}$ follows B on the last block. Since $\delta_1 > \Lambda'|_0^{T'}$

$$\boldsymbol{\delta}^{\mathrm{A}}{*}\Lambda \;=\; \sum_{1 \leq i < m \text{ odd}} \delta_i \;\geq\; \delta_1 \;>\; \Lambda'|_0^{T'} \;=\; \boldsymbol{\delta}^{\mathrm{A}}{*}\Lambda'.$$

If m is odd, then $\boldsymbol{\delta}^{\mathrm{A}}$ follows A on the last block. Again using Lemma 2, we get

$$\boldsymbol{\delta}^{\mathrm{A}}{*}\Lambda \;=\; \sum_{1 \leq i < m \text{ even}} \delta_i + \Lambda|_0^T \;\geq\; \Lambda|_0^T \;=\; \Lambda'|_0^{T'} \;=\; \boldsymbol{\delta}^{\mathrm{A}}{*}\Lambda'.$$

To prove the induction step, suppose a switch is executed, i.e. the first difference δ_1 is present in the regulariser Λ', and hence also in Nature's play Λ, then the strategy $\boldsymbol{\delta}^{\mathrm{A}}$ switches at price level $\Lambda(0) + \delta_1$ on either play, resulting in the same capital. The switches may occur at different times on Λ and Λ'. Nevertheless, the induction hypothesis applies to the tails of the plays since the remainder of the regulariser Λ' is more regular than the remainder of Nature's move Λ. $\square$

Since the theorem holds pointwise in $\mathcal{D}$, it also holds for the mixture strategy π.

3.3 With γ-Separation

Fix a logprice function Λ. Throughout this section, we use the following notation:

Definition 1. *We denote by $\boldsymbol{z} = z_0, z_1, \ldots, z_m$ the sequence of logprices at the local extrema of Λ (attained or not), with $z_0 = \Lambda(0)$ and $z_m = \Lambda(T)$, and we say that Λ has m* blocks. *Let $\boldsymbol{\Delta} = \Delta_1, \ldots, \Delta_m$ denote the sequence of absolute logprice differences, i.e. $\Delta_i := |z_i - z_{i-1}|$.*

Definition 2. *We say that $\mathit{\Lambda}$ has γ-separation if $\mathit{\Delta}_1, \mathit{\Delta}_m \geq \gamma$ and $\mathit{\Delta}_i \geq 2\gamma$ for each $1 < i < m$. That is, the border optima have logprice difference at least γ with the border and each subsequent pair of local extrema has at least logprice difference 2γ.*

We now analyse the payoff of the strategy π, assuming that $\mathit{\Lambda}$ has γ-separation.

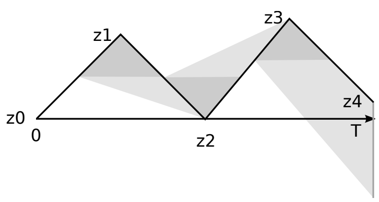

Fig. 3: Domain of Integration Example. Some $\mathit{\Lambda}$, with $m = 4$, is shown in black. The height of the dark gray triangles equals γ. This $\mathit{\Lambda}$ has γ-separation. In particular $\mathit{\Delta}_2 = z_1 - z_2 = 2\gamma$. Theorem 3 integrates over the strategies that are optimal for logprice functions in the light gray region.

Theorem 3 (γ-Separation Payoff). *For each $\mathit{\Lambda}$ with γ-separation*

$$\pi * \mathit{\Lambda} \;\geq\; \underbrace{\sum_{1 \leq i \leq m} (z_i - z_{i-1})_+}_{\textit{gain}} \;-\; \underbrace{\sum_{1 \leq i \leq m} \ell(\mathit{\Delta}_i)}_{\textit{complexity penalty}} \;+\; (m-1)\, \underbrace{\ln(1 - \mathrm{e}^{-\gamma})}_{\textit{overhead per switch}} \;-\; \underbrace{\ln 2}_{\textit{parity}} \, .$$

Proof. We saw in Section 2.3 that π is self-dual, so by symmetry (Lemma 2) we may assume that $z_0 \leq z_1$. As our first bound, we use $\pi * \mathit{\Lambda} \geq \pi^{\mathrm{A}} * \mathit{\Lambda} - \ln 2$. Recall that the payoff $\pi^{\mathrm{A}} * \mathit{\Lambda}$ is defined as $\ln \int \exp\big(\boldsymbol{\delta}^{\mathrm{A}} * \mathit{\Lambda}\big) \, \mathrm{d}\pi(\boldsymbol{\delta})$. As the next step, we re-parameterise the integral by introducing variables $\boldsymbol{h}$, with $h_i \coloneqq z_0 - \sum_{1 \leq j \leq i} (-1)^j \delta_j$. That is, h_i is the logprice at the ith switch of $\boldsymbol{\delta}^{\mathrm{A}}$. Then we obtain a lower bound by restricting the domain of integration. For $1 \leq i < m$ we restrict $h_i \in [z_i - \gamma, z_i]$ for odd i and $h_i \in [z_i, z_i + \gamma]$ for even i. Thus, we keep all prior mass on strategies that switch at logprices h_i that are at most γ nats short of the optimal switching logprice level z_i. We restrict the last logprice to $h_m \in [z_m, \infty)$ for even m and $h_m \in (-\infty, z_m]$ for odd m. This ensures that we do not switch between h_{m-1} and z_m. Thus, we only integrate over those strategies that closely follow $\mathit{\Lambda}$, as illustrated by Figure 3. We first consider even m. Then

$$\pi^{\mathrm{A}} * \mathit{\Lambda} \;\geq\; \ln \int_{z_1 - \gamma}^{z_1} \mathrm{e}^{h_1 - h_0} \pi(h_1 - h_0) \int_{z_2}^{z_2 + \gamma} \pi(h_1 - h_2) \int_{z_3 - \gamma}^{z_3} \mathrm{e}^{h_3 - h_2} \pi(h_3 - h_2) \cdots$$
$$\cdots \int_{z_{m-1}}^{z_{m-1} + \gamma} \pi(h_{m-2} - h_{m-1}) \int_{z_m}^{\infty} \mathrm{e}^{z_m - h_{m-1}} \pi(h_m - h_{m-1}) \, \mathrm{d}\boldsymbol{h}$$

Apply the tail probability bound (Lemma 1(3)) to the innermost integral to get

$$\int_{z_m}^{\infty} \mathrm{e}^{z_m - h_{m-1}} \pi(h_m - h_{m-1}) \,\mathrm{d}h_m \;\geq\; \mathrm{e}^{z_m - h_{m-1}} \pi(z_m - h_{m-1}).$$

Since $|h_i - h_{i-1}| \leq \Delta_i$ and π decreases (Lemma 1(2)) we get

$$\begin{aligned} \pi^{\mathrm{A}} {*} \Lambda \;\geq\; & \ln \prod_{1 \leq i \leq m} \pi(\Delta_i) \;+ \\ & \ln \left(\mathrm{e}^{z_m - z_0} \int\limits_{z_1 - \gamma}^{z_1} \mathrm{e}^{h_1} \int\limits_{z_2}^{z_2+\gamma} \mathrm{e}^{-h_2} \int\limits_{z_3-\gamma}^{z_3} \mathrm{e}^{h_3} \cdots \int\limits_{z_{m-2}-\gamma}^{z_{m-2}} \mathrm{e}^{h_{m-2}} \int\limits_{z_{m-1}}^{z_{m-1}+\gamma} \mathrm{e}^{-h_{m-1}} \,\mathrm{d}\boldsymbol{h} \right) \end{aligned}$$

Now all integrals have become independent. Rewrite odd/even instances like

$$\int_{z_1 - \gamma}^{z_1} \mathrm{e}^{h_1} \,\mathrm{d}h_1 \;=\; \mathrm{e}^{z_1}(1 - \mathrm{e}^{-\gamma}) \quad \text{and} \quad \int_{z_2}^{z_2+\gamma} \mathrm{e}^{-h_2} \,\mathrm{d}h_2 \;=\; \mathrm{e}^{-z_2}(1 - \mathrm{e}^{-\gamma}).$$

By rearranging terms we obtain

$$\pi^{\mathrm{A}} {*} \Lambda \;\geq\; \sum_{1 \leq i \leq m} (z_i - z_{i-1})_+ \;-\; \sum_{1 \leq i \leq m} \ell(\Delta_i) \;+\; (m-1)\ln(1 - \mathrm{e}^{-\gamma}).$$

The case for odd m is analogous. □

3.4 Establishing γ-Separation

Say we have a Λ with γ-separation, and hence a performance guarantee by Theorem 3. If γ is small, then a better bound can be obtained by first regularising Λ to a price function Λ^ϵ with ϵ-separation for some $\epsilon > \gamma$, and only then applying the theorem. In this section we quantify the gain of going from $\gamma = 0$ to ϵ, and then derive our main payoff bound by tuning ϵ.

The regulariser Λ^ϵ is constructed by the algorithm shown in Figure 4a. The key idea of the algorithm, implemented by lines 4–6, is to iteratively remove the smallest fluctuation from $\boldsymbol{z}$. This process is illustrated by Figure 4b. The solid line shows a segment of the logprice function before regularisation. The logprice difference between the two open circles is too small, i.e. $< 2\epsilon$. The dashed line is the logprice function resulting from fluctuation removal. The other lines of the algorithm establish ϵ-separation at the boundaries of Λ.

For any sequence $\boldsymbol{z} = z_0, \ldots, z_m$ we abbreviate the terms in the bound of Theorem 3 that depend on $\boldsymbol{z}$ by defining $\boldsymbol{g} = g_1, \ldots, g_m$ and G by

$$g_i \;\coloneqq\; (z_i - z_{i-1})_+ - \ell(\Delta_i) \qquad \text{and} \qquad G \;\coloneqq\; \sum_{1 \leq i \leq m} g_i.$$

We first study the effect of a single execution of lines 4–6.

1: $u \leftarrow (2\epsilon - \Delta_1)_+ \cdot \operatorname{sign}(z_1 - z_0)$.
2: $v \leftarrow (2\epsilon - \Delta_m)_+ \cdot \operatorname{sign}(z_m - z_{m-1})$.
3: $\boldsymbol{z} \leftarrow (z_0, z_1 + u, z_2 + u, \ldots, z_m + u + v)$ ▷ Ensure $\Delta_1, \Delta_m \geq 2\epsilon$
4: **while** the minimal Δ_i is (strictly) less than 2ϵ **do**
5: $\quad \boldsymbol{z} \leftarrow (z_0, z_1, \ldots, z_{i-2}, z_{i+1}, \ldots, z_m)$ ▷ See (b)
6: **end while**
7: $\boldsymbol{z} \leftarrow (z_0, z_1 - u, z_2 - u, \ldots, z_m - u - v)$ ▷ Reverse line 3
8: **if** $\Delta_1 < \epsilon$ **then** $\boldsymbol{z} \leftarrow (z_0, z_2, z_3, \ldots, z_m)$ ▷ Ensure $\Delta_1 \geq \epsilon$
9: **if** $\Delta_m < \epsilon$ **then** $\boldsymbol{z} \leftarrow (z_0, z_1, \ldots, z_{m-2}, z_m)$ ▷ Ensure $\Delta_m \geq \epsilon$

(a) ϵ-Pruning Algorithm

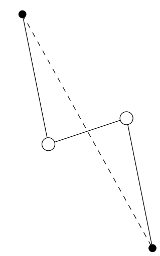

(b) Regularise

Fig. 4: ϵ-Pruning Algorithm and its main regularisation

Lemma 3. *Let $\boldsymbol{z}^\circ$ and $\boldsymbol{z}^\dagger$ be the sequences before and after line 5. Then*

$$G^\dagger - G^\circ \;\geq\; (m^\circ - m^\dagger) \min\big\{0, \ell(2\epsilon) - \epsilon\big\}$$

Proof. Let i be the index of the minimal Δ_i°. Let $l = \Delta_{i-1}^\circ$, $c = \Delta_i^\circ$ and $r = \Delta_{i+1}^\circ$, so that $\Delta_{i-1}^\dagger = l + r - c$ and $2\epsilon > c \leq l, r$. By definition $G^\dagger - G^\circ$ equals

$$\begin{aligned} \big(l + r - c - \ell(l + r - c)\big) - \big(l + r - \ell(l) - \ell(c) - \ell(r)\big) \qquad & \text{if } z_{i-1} \leq z_i, \text{ or} \\ \big(-\ell(l + r - c)\big) - \big(c - \ell(l) - \ell(c) - \ell(r)\big) \qquad & \text{if } z_{i-1} \geq z_i. \end{aligned}$$

In either case $G^\dagger - G^\circ$ simplifies to $-c - \ell(l + r - c) + \ell(l) + \ell(c) + \ell(r)$. Since ℓ is concave, the worst-case values for l and r are c. For the same reason, the worst-case value for c is either 0 or 2ϵ. Then since ℓ is nonnegative

$$G^\dagger - G^\circ \;\geq\; 2\ell(c) - c \;\geq 2 \min\{\ell(0), \ell(2\epsilon) - \epsilon\} \;\geq 2 \min\{0, \ell(2\epsilon) - \epsilon\}. \qquad \square$$

Now fix $\epsilon \geq 0$. Let $\boldsymbol{z}^\epsilon = z_0^\epsilon, z_1^\epsilon, \ldots, z_{m^\epsilon}^\epsilon$ be the result of applying Algorithm 4a with parameter ϵ to the sequence $\boldsymbol{z}$ of local extrema of Λ, and let Λ^ϵ be any continuous function with local extrema $\boldsymbol{z}^\epsilon$. By construction Λ^ϵ has ϵ-separation and regularises Λ. Theorem 3 gives us a bound on the payoff in terms of Λ^ϵ. We now show how to get a bound in terms of the original Λ.

Theorem 4 (Enforcing ϵ-Separation). *For all $\epsilon \geq 0$ such that $\ell(2\epsilon) < \epsilon$*

$$G^\epsilon - G \;\geq\; \big(m - m^\epsilon\big)\big(\ell(2\epsilon) - \epsilon\big) - 2\ell(2\epsilon).$$

Proof. Let $\boldsymbol{z}^+, \boldsymbol{z}^\star, \boldsymbol{z}^-$ be the sequences after lines 3, 6 and 7 of Algorithm 4a. Thus the algorithm produces (denoted $\rightarrow$) in order

$$\boldsymbol{z} \;\rightarrow\; \boldsymbol{z}^+ \;\rightarrow\; \boldsymbol{z}^\star \;\rightarrow\; \boldsymbol{z}^- \;\rightarrow\; \boldsymbol{z}^\epsilon.$$

with numbers of blocks $m = m^+ \geq m^\star = m^- \geq m^\epsilon$. By Lemma 3 $G^\star - G^+ \geq \big(m^+ - m^\star\big)\big(\ell(2\epsilon) - \epsilon\big)$. It thus remains to show that

$$(G^+ - G) + (G^- - G^\star) + (G^\epsilon - G^-) \;\geq\; (m^\star - m^\epsilon)\big(\ell(2\epsilon) - \epsilon\big) - 2\ell(2\epsilon).$$

We have $G^+ - G = g_1^+ - g_1 + g_{m^+}^+ - g_m$, $G^- - G^\star = g_1^- - g_1^\star + g_{m^-}^- - g_{m^-}^\star$ and

$$G^\epsilon - G^- \;=\; \begin{cases} g_1^\epsilon - g_1^- - g_2^- & \text{if } \varDelta_1^- < \epsilon, \\ 0 & \text{otherwise,} \end{cases} + \begin{cases} g_{m^\epsilon}^\epsilon - g_{m^-}^- - g_{m^- -1}^- & \text{if } \varDelta_{m^-}^- < \epsilon, \\ 0 & \text{otherwise.} \end{cases}$$

These three expressions are symmetric in the first and last element of the sequences concerned. The contributions of the first elements are

$$g_1^+ - g_1 \;=\; u_+ - \ell\big(\varDelta_1 + |u|\big) + \ell(\varDelta_1), \tag{3}$$

$$g_1^- - g_1^\star \;=\; -\,u_+ - \ell(\varDelta_1^-) + \ell\big(\varDelta_1^- + |u|\big), \tag{4}$$

$$g_1^\epsilon - g_1^- - g_2^- \;=\; -\,\varDelta_1^- + \ell(\varDelta_1^-) + \ell(\varDelta_2^-) - \ell(\varDelta_2^- - \varDelta_1^-). \tag{5}$$

If $\varDelta_1^- \geq \epsilon$ then no element is dropped in line 8. The sum of (3) and (4) equals

$$-\ell\big(\varDelta_1 + |u|\big) + \ell(\varDelta_1) - \ell(\varDelta_1^-) + \ell\big(\varDelta_1^- + |u|\big) \;\geq\; -\ell(2\epsilon).$$

Since ℓ increases the last two terms are positive and can be dropped from the bound; the remaining expression is increasing in $\varDelta_1$ by concavity of ℓ and is decreasing in $|u|$. Substitute the worst-case values $\varDelta_1 = 0$ and $|u| = 2\epsilon$.

If on the other hand $\varDelta_1^- < \epsilon$ then one element was dropped in line 8. In this case the sum of (3)–(5) equals

$$-\ell\big(\varDelta_1 + |u|\big) + \ell\big(\varDelta_1^- + |u|\big) + \ell(\varDelta_1) - \varDelta_1^- + \ell(\varDelta_2^-) - \ell(\varDelta_2^- - \varDelta_1^-) \;\geq\; -\,\varDelta_1^-.$$

The lower bound is obtained by cancelling the first two terms and the last two terms since ℓ is increasing and $0 \leq \varDelta_1 \leq \varDelta_1^-$. Since ℓ is nonnegative, we omit the third term as well. We then use $-\varDelta_1^- \;\geq\; -\,\epsilon \;=\; \big(\ell(2\epsilon) - \epsilon\big) - \ell(2\epsilon)$.

The bound for the contribution of the final elements is analogous. In each case, a dropped intermediate elements contributes at most $\ell(2\epsilon) - \epsilon$, while the borders lose at most $\ell(2\epsilon)$ each. □

We now put everything together, and in particular we optimise the value of ϵ.

Theorem 5 (Payoff Bound). *Fix logprice functions $\varLambda$ and $\varLambda'$, the latter with associated $\boldsymbol{z}'$, m' and $\boldsymbol{\varDelta}'$ as in Definition 1. If $\varLambda' \preccurlyeq \varLambda$ then*

$$\pi{*}\varLambda \;\geq\; \sum_{1 \leq i \leq m'} (z_i' - z_{i-1}')_+ - \sum_{1 \leq i \leq m'} \ell(\varDelta_i') - (m'-1)c_\pi - \ln 2 - 2\epsilon_\pi,$$

where ϵ_π is the unique solution to $\pi(2\epsilon) = \frac{1}{\mathrm{e}^\epsilon - 1}$, and $c_\pi = -\ln(1 - \mathrm{e}^{-\epsilon_\pi})$.

Proof. For each $\epsilon \geq 0$ with $\ell(2\epsilon) < \epsilon$

$$\begin{aligned} \pi{*}\varLambda \;\geq\; \pi{*}\varLambda' \;\geq\; \pi{*}\varLambda^\epsilon \;\geq\; & G^\epsilon + (m^\epsilon - 1)\ln(1 - \mathrm{e}^{-\epsilon}) - \ln 2 \\ \geq\; & G' + (m^\epsilon - 1)\ln(1 - \mathrm{e}^{-\epsilon}) + (m' - m^\epsilon)\big(\ell(2\epsilon) - \epsilon\big) - 2\ell(2\epsilon) - \ln 2 \\ \geq\; & G' + (m'-1)\min\big\{\ln(1 - \mathrm{e}^{-\epsilon}), \ell(2\epsilon) - \epsilon\big\} - 2\ell(2\epsilon) - \ln 2. \end{aligned}$$

The inequalities are twice Theorem 2, then Theorem 3, then Theorem 4. To complete the proof we set ϵ to equalise the arguments of the minimum. □

Typical values for ϵ_π and c_π are shown in Table 1.

4 Implementation

The following algorithm implements the strategy π. For arbitrary prior densities it runs in $O(T^2)$ time. For exponential priors, we reduce the running time to $O(T)$. The key to efficiency is the independent product form of π, which renders the *last* switching price a sufficient statistic.

For concreteness, we measure discrete time in days. As its data structure, the algorithm maintains a set of bank accounts. Each bank account has a *balance*, a *type* that is either A or B, and a *birthday*. The balance of type A accounts is measured in shares, whereas that of type B accounts is measured in cash.

On day zero the initial unit cash is divided evenly into two bank accounts: one account of type B with half a unit of cash, and one account of type A with $\frac{1}{2}\exp(-\Lambda(0))$ shares, i.e. half a unit of cash worth of shares at the initial logprice.

The algorithm then proceeds as follows. Each day $t = 1, 2, \ldots$ the new price $\Lambda(t)$ is announced. The algorithm creates a single new bank account with birthday t. If $\Lambda(t-1) \leq \Lambda(t)$, then the new account is of type B, and a portion of the shares in existing accounts of type A is sold to fill it with cash. On the other hand if $\Lambda(t-1) \geq \Lambda(t)$, then a new account of type A is endowed with shares by investing a fraction of the capital of existing accounts of type B. In either case, the amount traded reestablishes the following invariant. At the end of day t:

- Each account of type A that was created with c shares on birthday i has balance $c\int_\lambda^\infty \pi(h)\,\mathrm{d}h$, where $\lambda = \max_{i \leq j \leq t} \Lambda|_i^j$.
- Each account of type B that was created with capital c on birthday i has balance $c\int_\lambda^\infty \pi(h)\,\mathrm{d}h$, where $\lambda = \max_{i \leq j \leq t} -\Lambda|_i^j$.

To see how this works, consider an A-type account with birthday i and initial balance c, and assume that the invariant was maintained at the end of day $t-1$. First, note that it can only become violated if the maximum changes, that is, if $\Lambda|_i^t$ exceeds the previous maximum $\lambda = \max_{i \leq j < t} \Lambda|_i^j$. Then the balance still is $c\int_\lambda^\infty \pi(h)\,\mathrm{d}h$ but should become $c\int_{\Lambda|_i^t}^\infty \pi(h)\,\mathrm{d}h$. The fraction

$$1 - \frac{\int_{\Lambda|_i^t}^\infty \pi(h)\,\mathrm{d}h}{\int_\lambda^\infty \pi(h)\,\mathrm{d}h} = \frac{\int_\lambda^{\Lambda|_i^t} \pi(h)\,\mathrm{d}h}{\int_\lambda^\infty \pi(h)\,\mathrm{d}h} = \pi\big(H \leq \Lambda|_i^t \big| H \geq \lambda\big) \tag{6}$$

of the balance must be sold to reestablish the invariant, and the resulting cash is transferred to the new account. Note that we only query π via its cumulative distribution function.

Complexity Analysis. After t days, there are $t+2$ bank accounts to maintain, and each bank account potentially requires work each round. Thus, trading for T days takes $O(T^2)$ time and $O(T)$ space.

For exponential priors we can do better by *merging* several bank accounts into a single account with the sum of their balances. This is because for memoryless priors, the fraction (6) to be traded away does not depend on the birthday i, but only on the maximum λ, allowing us to merge bank accounts with the same maximum. Now observe that all bank accounts that are tapped to reestablish the invariant share the same maximum afterwards, and can hence all be merged.

This means that a bank account requires work at most *once*, namely when it is merged away. By maintaining two stacks of bank accounts, one for each type, each ordered by the maximum λ, the running time is brought down to $O(T)$. Since we do not know *when* merges happen, the space requirement is still $O(T)$, and the running time is *amortised* $O(1)$ per day.

5 Conclusion

We presented a simple online algorithm that can be applied to two-way trading, but also to prediction with expert advice, data compression and hypothesis testing (see Section 1). Compared to the many hedging algorithms described in the literature, our approach has two novel properties. First, the overhead of our algorithm is independent of the times at which prices are processed, and second, our bound is free of any conditions on the evolution of the price of the asset, and is parameterised either by the asset price function itself or by a regularised model of it.

The surprisingly simple implementation (Section 4) processes a sequence of T asset prices in $O(T^2)$ time and $O(T)$ space. The algorithm models the scale of the fluctuations of the price using a density function on $[0, \infty)$; if an exponential density is employed, the running time is reduced to $O(T)$.

References

1. Berger, J.O.: Could Fisher, Jeffreys and Neyman have agreed on testing? Statistical Science 18(1), 1–32 (2003)
2. Chen, G.H., Kao, M.Y., Lyuu, Y.D., Wong, H.K.: Optimal buy-and-hold strategies for financial markets with bounded daily returns. In: Proc. of the 31st annual ACM symposium on Theory of computing. pp. 119–128. ACM (1999)
3. Cover, T.M., Thomas, J.A.: Elements of Information Theory. John Wiley (1991)
4. Dannoura, E., Sakurai, K.: An improvement on El-Yaniv-Fiat-Karp-Turpin's money-making bi-directional trading strategy. IPL 66(1), 27–33 (1998)
5. Dawid, A.P., De Rooij, S., Shafer, G., Shen, A., Vereshchagin, N., Vovk, V.: Insuring against loss of evidence in game-theoretic probability, arXiv:1005.1811
6. DeMarzo, P., Kremer, I., Mansour, Y.: Online trading algorithms and robust option pricing. In: Proc. of the 38 annual ACM symposium on Theory of computing. pp. 477–486. ACM (2006)
7. El-Yaniv, R., Fiat, A., Karp, R.M., Turpin, G.: Optimal search and one-way trading online algorithms. Algorithmica 30(1), 101–139 (2001)
8. Van Erven, T., Grünwald, P.D., De Rooij, S.: Catching up faster by switching sooner: a prequential solution to the AIC-BIC dilemma (2008), submitted. Preprint available as arXiv:0807.1005
9. Herbster, M., Warmuth, M.K.: Tracking the best expert. Machine Learning 32, 151–178 (1998)
10. Koolen, W.M., De Rooij, S.: Combining expert advice efficiently. In: Proc. of the 21st Annual Conference on Learning Theory. pp. 275–286 (2008)
11. Shafer, G., Shen, A., Vereshchagin, N., Vovk, V.: Test martingales, Bayes factors, and p-values, arXiv:0912.4269
12. Volf, P., Willems, F.: Switching between two universal source coding algorithms. In: Proc. of the Data Compression Conference, Snowbird, Utah. pp. 491–500 (1998)
13. Vovk, V.: A game of prediction with expert advice. Journal of Computer and System Sciences 56, 153–173 (1998)