Buy Low, Sell High

Wouter M. Koolen and Vladimir Vovk

Computer Learning Research Centre, Department of Computer Science, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, United Kingdom

Abstract. We consider online trading in a single security with the objective of getting rich when its price ever exhibits a large upcrossing without risking bankruptcy. We investigate payoff guarantees that are expressed in terms of the extremity of the upcrossings. We obtain an exact and elegant characterisation of the guarantees that can be achieved. Moreover, we derive a simple canonical strategy for each attainable guarantee.

Keywords: Online investment, worst-case analysis, probability-free option pricing

1 Introduction

We consider the simplest trading setup, where an investor trades in a single security as specified in Figure 1. An intuitive rule of thumb is to buy when the price is low, say a, and sell later when the price is high, say b. Trading successfully in such a manner exploits the so-called *upcrossing* $[a, b]$ and secures payoff b/a. In practice we do not know in advance when a stiff upcrossing will occur. Still, we can ask for a strategy whose payoff scales nicely with the extremity of the upcrossing present. A financial advisor, to express that her secret strategy approximates this ideal, may publish a function $G : \mathbb{R}^2 \to \mathbb{R}$ and promise that her strategy will

Initial price $\omega_0 > 0$
Starting capital $K_0 = 1$.
For $t = 1, 2, \dots$
- Investor takes position $S_t \in \mathbb{R}$.
- Market reveals price $\omega_t \geq 0$.
- $K_t := K_{t-1} + S_t(\omega_t - \omega_{t-1})$

Fig. 1: Simple trading protocol

keep our capital above $G(a, b)$ for each upcrossing $[a, b]$

Before trusting her to manage our capital, we would like to answer the following questions:

1. Should we believe her? Is it actually *possible* to guarantee G?
2. Is she ambitious *enough*? Or can one guarantee strictly more than G?
3. Can we *reverse-engineer* a strategy to guarantee G ourselves?

The contribution of this paper is a complete resolution of these questions. We characterise the achievable guarantees, and the admissible (or Pareto optimal, i.e. not strictly dominated) guarantees. We construct, for each achievable G, a relatively simple strategy that achieves it.

1.1 Related Work

This work is a joint sequel to two lines of work. We think of the first line as a complete treatment of the goal of selling high (without buying low first), and of the second line as intuitive strategies for iterated trading. Let us summarise the material that we will use from each.

Sell high Guarantees for trading once (selling at the maximum) were completely characterised in [1]. The results are as follows. We call an increasing right-continuous function $F : [1, \infty) \to [0, \infty)$ a *candidate guarantee*. A candidate guarantee F is an *adjuster* if there is a strategy that ensures $K_t \geq F(\max_s \omega_s)$ for every price evolution $\omega_0, \dots, \omega_t$. An adjuster that is not strictly dominated is called *admissible*. The goal is to find adjusters that are close to the unachievable $F_{\text{ideal}}(y) \coloneqq y$. What can be achieved is characterised as follows:

Theorem 1 (Characterisation). *A candidate guarantee F is an adjuster iff*

$$\int_1^\infty \frac{F(y)}{y^2} \, \mathrm{d}y \;\leq\; 1 \tag{1}$$

Moreover, it is admissible iff (1) *holds with equality.*

This elegant characterisation gives a simple test for adjusterhood. We can get reasonably close to F_{ideal}, for example using the adjusters

$$F(y) \;\coloneqq\; \alpha y^{1-\alpha} \quad \text{for some } 0 < \alpha < 1 \qquad \text{or} \qquad F(y) \;\coloneqq\; \frac{y^2 \ln(2)}{(1+y)\ln(1+y)^2}.$$

The following decomposition allows us to reverse engineer a canonical strategy for each adjuster F. For each price level $u \geq 1$, consider the *threshold guarantee* $F_u(y) \coloneqq u \mathbf{1}_{\{y \geq u\}}$, which is an adjuster witnessed by the strategy S_u that takes position 1 until the price first exceeds u and 0 afterwards. With this definition we have

Theorem 2 (Representation). *A candidate guarantee F is an adjuster iff there is a probability measure P on $[1, \infty)$ such that*

$$F(y) \;\leq\; \int F_u(y) \, \mathrm{d}P(u),$$

again with equality iff F is admissible.

In other words, we can witness any admissible adjuster F by the strategy $S_P \coloneqq \int S_u \, \mathrm{d}P(u)$, that is by splitting the initial capital according to the associated measure $P(u)$ over threshold strategies S_u and never rebalancing.

Iterated trading Intuitive trading strategies for iterated trading were proposed in [2,3], and their worst-case performance guarantees were analysed. We briefly review the construction and guarantees specialised to the case of trading twice. The proposed strategies are of the form $S_Q := \int S_{\alpha,\beta} \, \mathrm{d}Q(\alpha, \beta)$, where Q is some bivariate probability measure and $S_{\alpha,\beta}$ is the threshold strategy that does not invest initially, subsequently invests all capital when the price first drops below α, and finally liquidates the position when the price first exceeds β. Clearly $S_{\alpha,\beta}$ witnesses the guarantee

$$G_{\alpha,\beta}(a,b) \;:=\; \frac{\beta}{\alpha} \mathbf{1}_{\{a \,\leq\, \alpha \text{ and } b \,\geq\, \beta\}},$$

and so the full strategy S_Q witnesses

$$G_Q(a,b) \;:=\; \int G_{\alpha,\beta}(a,b) \, \mathrm{d}Q(\alpha,\beta). \tag{2}$$

(We omit the iterated trading bounds and run-time analysis, they are outside the scope of this paper.)

1.2 Climax

Intuitively, the dual threshold strategies $S_{\alpha,\beta}$ are the natural generalisation of the single threshold strategies S_u. Since any univariate admissible adjuster is a convex combination of threshold guarantees, it is natural to conjecture that a bivariate candidate guarantee G is an admissible adjuster iff $G = G_Q$ for some Q.

Interestingly however, it turns out that mixture guarantees of the form (2) are typically *strictly dominated!* Let us illustrate what goes awry with a simple example. Consider the mixture-of-thresholds guarantee G defined by

$$G(a,b) \;:=\; \frac{1}{2} G_{1,2}(a,b) + \frac{1}{2} G_{\frac{1}{2},1}(a,b) \;=\; \mathbf{1}_{\{a \,\leq\, 1 \text{ and } b \,\geq\, 2\}} + \mathbf{1}_{\left\{a \,\leq\, \frac{1}{2} \text{ and } b \,\geq\, 1\right\}}.$$

(These weights and thresholds are chosen for simplicity and are by no means essential.) We now argue that G is strictly dominated, by showing that G can be guaranteed from initial capital $\frac{11}{12} < 1$, and hence that G is strictly dominated by the adjuster $\frac{12}{11} G$.

The smallest initial capital required to satisfy the guarantee G can be found from the tree of situations shown in Figure 2a. We restrict Market to the seven price paths that can be obtained by moving starting from the root (the left-most node labelled by 1) to the right along a branch of the tree to a leaf and reading off the price labels inside the circles. Formally, we do not allow price ∞, but it can be replaced by a sufficiently large number. The three intervals mentioned in the guarantee G and the inclusion relation between them are displayed in Figure 2c. Figure 2b indicates which price paths upcross which intervals. We can now compute the capital needed to guarantee G in each situation. First, as

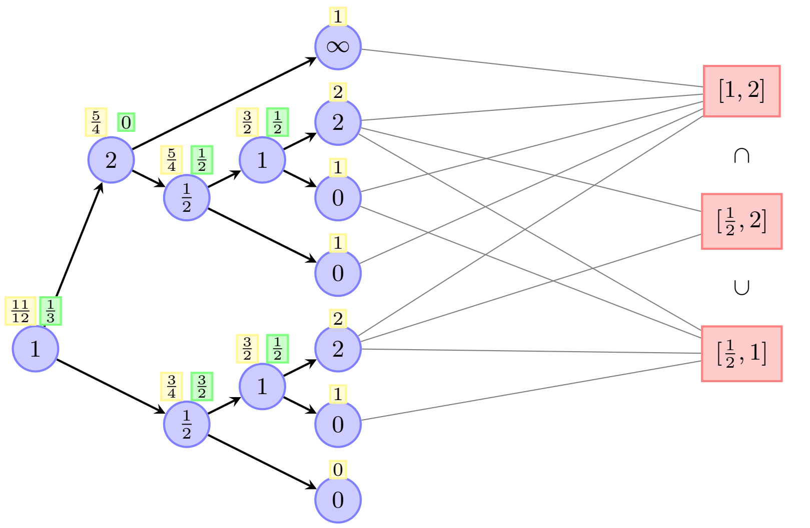

(a) Prices, capitals, and positions (b) Upcrossing relation (c) Intervals

Fig. 2: Toy world

shown in Figure 2a, we assign to each leaf ω the capital necessary to guarantee G on it, which is by definition

$$X_G(\omega) \;\coloneqq\; \max\big\{G(a,b) \;\big|\; a,b \text{ s.t. } \omega \text{ upcrosses } [a,b]\big\}$$

To label the intermediate situations we use backward induction. Let us first explain a single induction step, which is known as binomial pricing.

Binomial pricing tutorial Consider a toy world in which the current price is p, and the future price is either $p_u > p$ in which case we want to guarantee payoff K_u, or $p_d < p$ in which case we want to guarantee K_d. The minimal initial capital K from which this is possible and the position S that achieves this are

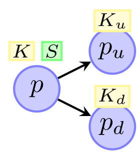

$$K \;=\; \frac{p-p_d}{p_u-p_d}K_u + \frac{p_u-p}{p_u-p_d}K_d \qquad \text{and} \qquad S \;=\; \frac{K_u-K_d}{p_u-p_d}. \tag{3}$$

Using binomial pricing, we have labelled each internal situation in 2b by the price (left) and position (right) obtained in backward fashion. Formally, this argument only shows that an initial capital of $\frac{11}{12}$ is necessary (although intuitively it is clear that the tree exhausts all possibilities, and so $\frac{11}{12}$ is also sufficient). Indeed, it is now easy to check that an initial capital of $\frac{11}{12}$ is sufficient: the strategy that witnesses G from initial capital $\frac{11}{12}$ can be read off Figure 2a. Namely, we take position $\frac{1}{3}$ at time 0 leaving $\frac{7}{12}$ in cash. There are two cases:

- If and when the price reaches $\frac{1}{2}$ before reaching 2, we invest all our cash. This will make our position at least $\frac{7}{6}+\frac{1}{3}=\frac{3}{2}$. If and when the price reaches 1, we cash in 1 dollar leaving a position of at least $\frac{1}{2}$. If and when the price reaches 2, we cash in another dollar. In all cases, we are left with at least $X_G(\omega)$ at the end, where ω is the realized price path.
- Now suppose the price reaches 2 before reaching $\frac{1}{2}$. Cashing in our position, we get at least $\frac{7}{12}+\frac{2}{3}=\frac{5}{4}$ dollars. If and when the price reaches $\frac{1}{2}$, we take a position of $\frac{1}{2}$, which leaves at least 1 dollar in cash. If and when the price reaches 2, we cash in another dollar. In all cases, we again are left with at least $X_G(\omega)$ at the end.

This argument shows that mixture guarantees can be strictly dominated. To get additional insight into why, let us consider the mixture strategy corresponding to G, which evenly divides its capital between $S_{1,2}$ and $S_{\frac{1}{2},1}$. The problem with this strategy is that it secures payoff 2 on price path $\omega=(1,2,1/2,1,0)$, but one only needs $X_G(\omega)=1$ to guarantee G there. The reason is that both small intervals $[\frac{1}{2},1]$ and $[1,2]$ are upcrossed, but their union $[\frac{1}{2},2]$ is not. In other words, the mixture strategy gives an additional payoff in certain circumstances that *does not contribute to the guarantee.* Since the binomial pricing formulas are linear in the payoffs, reducing the payoff at any leaf reduces the required initial capital.

1.3 Overview of results

The previous section shows that the world is not simple, i.e. the intuitive characterization of guarantees is incorrect. We now present our more subtle results. We call a function $G:(0,1]\times(0,\infty)\to[0,\infty)$ a *candidate guarantee* if it is upper semi-continuous, decreasing in its first argument and increasing in its second argument. We define the *second-argument upper inverse* of G by

$$G^{-1}(a,h) \;\coloneqq\; \inf\{b\geq a \mid G(a,b)\geq h\}. \tag{4}$$

Theorem 3 (Characterisation). *A candidate guarantee G is an adjuster iff*

$$\int_0^\infty 1-\exp\left(\int_0^1 \frac{1}{a-G^{-1}(a,h)}\,\mathrm{d}a\right)\mathrm{d}h \;\leq\; 1. \tag{5}$$

Moreover, G is admissible iff (5) *holds with equality.*

We saw in the previous section that a subtle temporal analysis is needed when reasoning about guarantees. Although this is still true for the proof of this theorem, the result itself is elegantly timing-free.

We also have a canonical representation in terms of convex combination of elementary guarantees. These elementary guarantees are analogous to the threshold

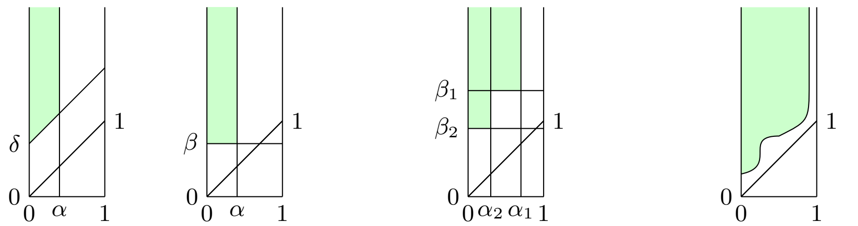

Fig. 3: Example north-west sets

strategies of the univariate case in the sense that they have just two payoff levels. However, they do have richer geometric structure. A closed set $I \subseteq (0,1] \times (0,\infty)$ is called *north-west* if $(a,b) \in I$ implies $(0,a] \times [b,\infty) \subseteq I$. Some example north-west sets are displayed in Figure 3. We associate to each north-west set its *frontier*

$$f_I(a) \;\coloneqq\; \inf\big\{b \geq a \;\big|\; (a,b) \in I\big\}.$$

By the previous theorem, the following guarantee is an admissible adjuster:

$$G_I(a,b) \;\coloneqq\; \frac{\mathbf{1}_{\{f_I(a) \leq b\}}}{1 - \exp\left(\int_0^1 \frac{1}{a' - f_I(a')} \,\mathrm{d}a'\right)}.$$

A family $(I_h)_{h \geq 0}$ of north-west sets is called *nested* if $x \leq y$ implies $I_x \supseteq I_y$.

Theorem 4 (Representation). *A candidate guarantee G is an adjuster iff there are a probability measure Q on $[0,\infty)$ and a nested family $(I_h)_{h\geq 0}$ of north-west sets such that*

$$G(a,b) \;\leq\; \int G_{I_h}(a,b) \,\mathrm{d}Q(h),$$

with equality iff G is admissible.

This theorem gives us a means to construct a canonical strategy for each adjuster G. We first decompose G into a probability measure Q and a nested family of north-west sets $(I_h)_{h\geq 0}$. We then find a strategy S_{I_h} witnessing G_{I_h} for each h. Finally, we recompose these strategies to obtain the full strategy $S_G \coloneqq \int S_{I_h} \,\mathrm{d}Q(h)$.

These two theorems parallel those of [1] with a twist. Whereas [1] decomposes single-argument adjusters in terms of threshold guarantees (which have a single degree of freedom), our elementary guarantees are parametrised by the geometrically much richer north-west sets.

1.4 Outline

The paper is structured as follows. In Section 2 we reduce finding guarantees to a particular instance of probability-free option pricing. The actual option pricing is done in Section 3. Section 4 then discusses simple example guarantees, and in particular proposes an efficiently implementable strategy with an approximately ideal guarantee. The main proofs are delayed to Sections 5 and 6. We discuss the scope and applications of our results in Section 7, where we sketch the implications for online probability prediction and hypothesis testing.

2 Reduction to Option Pricing

We will make use of the definitions of probability-free *option pricing*, which we briefly review here. We assume that the initial asset price ω_0 is one, and that the investor starts with one unit of cash. Trading proceeds in rounds. In trading period t, the investor first chooses his *position* S_t, and then the new price ω_t is revealed. After T iterations, the investor has capital $K_T = 1+\sum_{t=1}^T S_t(\omega_t-\omega_{t-1})$. A *trading strategy* S assigns to each sequence of past prices $\omega_{<t} = (\omega_0, \ldots, \omega_{t-1})$ a *position* $S(\omega_{<t}) \in \mathbb{R}$. Let $S*\omega$ denote the payoff of strategy S on price function ω. That is

$$S * \omega \;:=\; 1 + \sum_{t=1}^{T} S(\omega_{<t})(\omega_t - \omega_{t-1}).$$

We denote by $S *_c \omega$ the payoff obtained by executing strategy S from initial capital c instead of one.

In general, an *option* X assigns to each price function ω a real value $X(\omega)$. (We have already seen one option, namely the payoff functional $\omega \mapsto S * \omega$.) The *upper price* of X, denoted $\overline{\mathbb{E}}[X]$, is the minimal initial capital necessary to super-replicate X, i.e.

$$\overline{\mathbb{E}}[X] \;:=\; \inf\left\{c \;\middle|\; \exists \text{ strategy } S \;\forall \text{ price function } \omega : S *_c \omega \geq X(\omega) \;\right\}.$$

This definition allows us to price options at the start of the game. We may also wonder about the capital necessary to super-replicate X half-way through the game, say after some past $\omega' = (\omega'_0, \ldots \omega'_t)$. This so-called *conditional upper price* is given by

$$\overline{\mathbb{E}}[X|\omega'] \;:=\; \inf\left\{c \;\middle|\; \exists \text{ strategy } S \;\forall \text{ price function } \omega : S *_c \omega \geq X(\omega'_{<t}\omega) \;\right\}.$$

where ω ranges over price functions starting from $\omega_0 = \omega'_t$ the current price. Note how the strategy only trades on the future ω, whereas the option value depends on the past ω'.

3 Characterisation of candidate guarantees

Suppose we conjure up some desirable candidate guarantee G, and wonder whether it is an adjuster, and if so, whether it is admissible. To decide this, we consider the option X_G that assigns to each price function ω the minimal payoff necessary to guarantee G on it:

$$X_G(\omega) \;:=\; \sup_{[a,b]\,:\,\omega \text{ upcrosses } [a,b]} G(a,b) \;=\; \max_{\substack{0\leq i\leq j\\ 1\geq\omega_i\leq\omega_j}} G(\omega_i,\omega_j) \tag{6}$$

We now connect adjusters and pricing

Proposition 5. *A candidate guarantee G is an adjuster iff $\overline{\mathbb{E}}[X_G] \leq 1$. Moreover, G is admissible iff $\overline{\mathbb{E}}[X_G] = 1$.*

Proof. The first equivalence holds by definition, and $\overline{\mathbb{E}}[X_G] < 1$ clearly implies inadmissibility. It follows from the pricing Theorem 6 below that a strictly dominated adjuster must have upper price < 1. □

This result reduces testing for adjusterhood to option pricing. Next we compute the upper price of X_G. Section 5 is dedicated to the proof.

Theorem 6. *The upper price of any candidate guarantee G is*

$$\overline{\mathbb{E}}[X_G] \;=\; \int_0^\infty 1 - \exp\left(\int_0^1 \frac{1}{a - G^{-1}(a,h)} \,\mathrm{d}a\right) \mathrm{d}h.$$

4 Example Adjusters

Before we go into proofs, we have a look at the consequences. We first recover the single-argument adjuster characterisation from the double-argument version. We then consider guarantees expressed in a single-parameter summary of $[a, b]$. Finally we really exploit both arguments, and design admissible adjusters that closely approach the ideal payoff b/a with computationally efficient strategies.

4.1 Selling high: adjusters expressed in the maximum price

Theorem 6 implies the results of [1] (in particular Theorem 1) as a special case.

Proof (Alternative proof of Theorem 1). Let $F : [1, \infty) \to [0, \infty)$ be an increasing right-continuous function. Construct the guarantee $G(a, b) \coloneqq F(b)\mathbf{1}_{\{b \geq 1\}}$ that ignores its first argument. By Theorem 6

$$\overline{\mathbb{E}}[X_G] = \int_0^\infty 1 - \exp\left(\int_0^1 \frac{\mathrm{d}a}{a - \inf\{b \mid F(b) \geq h\}}\right) \mathrm{d}h = \int_0^\infty \frac{\mathrm{d}h}{\inf\{b \mid F(b) \geq h\}}.$$

Using the variable substitution $h = F(y)$ (for $y \geq 1$ and $h \geq F(1)$) and integration by parts, we obtain

$$\begin{aligned}
\overline{\mathbb{E}}[X_G] \;&=\; \int_0^{F(1)} \frac{1}{\inf\{b \mid F(b) \geq h\}} \,\mathrm{d}h + \int_{F(1)}^\infty \frac{1}{\inf\{b \mid F(b) \geq h\}} \,\mathrm{d}h \\
&=\; F(1) + \int_1^\infty \frac{1}{y} \,\mathrm{d}F(y) \qquad (7) \\
&=\; F(1) + \frac{F(y)}{y}\bigg|_1^\infty + \int_1^\infty \frac{F(y)}{y^2} \,\mathrm{d}y \\
&=\; \int_1^\infty \frac{F(y)}{y^2} \,\mathrm{d}y \qquad (8)
\end{aligned}$$

This derivation assumes that $F(\infty)/\infty = 0$. If $F(\infty)/\infty$ exists and is strictly positive, both (7) and (8) are equal to ∞, and so $\overline{\mathbb{E}}[X_G]$ is still equal to (8). And if $F(\infty)/\infty$ does not exist, both (7) and (8) are again equal to ∞: if one or both of them were finite, $F(\infty)/\infty$ would exist as their difference. □

4.2 Adjusters expressed in the size of the upcrossing

The two natural measures of the size of an upcrossing $[a, b]$ are the length $b - a$ and the ratio b/a. Let us consider guarantees expressed in each statistic.

Length Using the tricks from the previous section we see that candidate guarantees of the form $G(a, b) = F(b - a)$ have upper price

$$\overline{\mathbb{E}}\left[X_G\right] \;=\; \int_0^\infty F(y) \frac{e^{-1/y}}{y^2} \,\mathrm{d}y.$$

This is analogous to (8), but with a twist. In financial terms, the distribution with density $\frac{e^{-1/y}}{y^2} \,\mathrm{d}y$ is the *risk-neutral measure* of the largest upcrossed length. Similarly, $y^{-2} \,\mathrm{d}y$ from (8) is the risk-neutral measure of the maximum price.

Ratio We now show that guarantees of the form $G(a, b) = F(b/a)$ for some increasing and unbounded F have infinite upper price. Such guarantees are way too good to be true: they can not be made adjusters even by re-normalisation. For simplicity assume that F is invertible. Then

$$G^{-1}(a, h) \;=\; aF^{-1}(h),$$

so that $\overline{\mathbb{E}}[X_G] = \infty$, because

$$\int_0^1 \frac{1}{a - G^{-1}(a, h)} \,\mathrm{d}a \;=\; \int_0^1 \frac{1}{a(1 - F^{-1}(h))} \,\mathrm{d}a \;=\; -\infty.$$

Other impossibility results follow from the same argument. For example, the intuitively modest candidate $G(a, b) = b^p/a^q$ has infinite price for any $p, q > 0$.

4.3 Approximately ideal adjusters

Our goal is to secure payoff close to the ideal b/a. The previous section shows that we cannot simply dampen the ratio b/a itself, but must make essential use of both arguments. A simple admissible adjuster that approaches the ideal is

$$G(a, b) \;=\; \frac{(b - a)^p}{a^q} \frac{\left(\frac{p-q}{p}\right)^p}{\Gamma(1 - p)}$$

for any $0 \leq q < p < 1$. The results in Section 5.2 below imply that this guarantee is witnessed by the strategy that in situation ω with minimum price m takes position

$$S(\omega) \;=\; \frac{(p - q)}{m^{1-p+q}} \,\Phi\left(\frac{m^{\frac{p-q}{p}}}{\left(X_G(\omega)\Gamma(1 - p)\right)^{1/p}}\right)$$

where $\Phi(x) = \frac{\int_0^x t^{-p} e^{-t} \, \mathrm{d}t}{\Gamma(1-p)}$ is the cumulative distribution function of the Gamma distribution (with shape $1-p$ and scale 1). This function can be evaluated to arbitrary precision by many computer mathematics support systems. Note that $X_G(\omega)$ and m can be maintained incrementally; when the next price r is revealed

$$\begin{aligned} X_G(\omega, r) &= \max\{X_G(\omega), G(m(\omega), r)\} \\ m(\omega, r) &= \min\{m(\omega), r\}. \end{aligned}$$

This admissible adjuster is hence extremely attractive. It approximates the ideal guarantee, and its strategy can implemented efficiently.

5 Proof of Theorem 6

In this section we prove the characterisation theorem. It will be convenient to prove the following more general statement.

Theorem 7. *Fix any candidate guarantee G and situation $\sigma = (\omega_0, \ldots, \omega_s)$. Let us abbreviate the current price to $r := \omega_s$, the lowest observed price to $m := \min_{i=0,\ldots,s} \omega_i$, and the minimal capital needed to satisfy G at time s to $C := X_G(\sigma)$ (see (6)). The conditional upper price of X_G in situation σ is*

$$\overline{\mathbb{E}}[X_G|\sigma] \;=\; C + \int_C^\infty 1 - \frac{G^{-1}(m,h) - r}{G^{-1}(m,h) - m} \exp\left(\int_0^m \frac{\mathrm{d}a}{a - G^{-1}(a,h)}\right) \mathrm{d}h. \quad (9)$$

The proof consists of two parts. For the lower bound we construct an adversarial Market based on random walks. For the upper bound we construct a strategy for Investor. It is quite surprising that these bounds meet, since these markets are generally highly incomplete. Our method is similar to that of [4], which derives option prices assuming continuous price paths. We are not aware of general probability-free option pricing results that allow discontinuous price processes.

5.1 Lower bound from Market strategy

We will find a lower bound on the conditional upper price $\overline{\mathbb{E}}[X_G|\sigma]$ of the option X_G using a finite up/down scheme. For a natural number n, we discretise the vertical price axis in bins of size 2^{-n}. Consider the following restricted Market starting from time $s+1$. At each discrete time step $t > s$ we have $\omega_t = \omega_{t-1} \pm 2^{-n}$, where ω_s is understood to be $R2^{-n}$, where $R := \lfloor \omega_s 2^n \rfloor$ (rather than the real ω_s). Define the stopping time τ to be least such that $\omega_\tau = 0$. On run ω, we desire to superreplicate X_G, which can be rewritten as

$$X_G(\omega) \;=\; \max_{\substack{0 \le i \le j \le \tau(\omega) \\ 1 \ge \omega_i \le \omega_j}} G(\omega_i, \omega_j)$$

We desire to lower bound the conditional upper price of X_G for the restricted Market. By binomial pricing, this price will be the expected value under a coin flip price process (formally, we explained binomial pricing only for finite games, but the extension to an infinite horizon is easy: consider a game lasting T rounds after which the price ω is frozen and then let $T \to \infty$). That is, the option's price will be at least

$$\mathbb{E}\, X_G\left(\omega_1,\ldots,\omega_s, 2^{-n}(R+\xi_1), 2^{-n}(R+\xi_1+\xi_2),\ldots,2^{-n}(R+\xi_1+\cdots+\xi_\tau)\right),$$

where the regular expectation $\mathbb{E}$ refers to ξ_s being independent random variables taking values ± 1 with equal probabilities and the term ξ_τ should be ignored when $\tau = \infty$. (We say "at least" since ω_s can exceed $R2^{-n}$.) As a first step, observe that what is important are the incremental global minima of ω, and their subsequent maxima. Set $M := \lceil m2^n \rceil$. We have that incremental minima are reached at the levels $k2^{-n}$, $k = 1,\ldots,M-1$, in decreasing order.

Define $i_k = i_k(\omega)$, $k = 1,\ldots,M-1$, to be the largest i such that, after hitting level $k2^{-n}$ at time $t > s$, ω rises to level $(k+i)2^{-n}$ before hitting level $(k-1)2^{-n}$. Define $i_M = i_M(\omega)$ to be the largest i such that, after time s, ω rises to level $(M+i)2^{-n}$ before hitting level $(M-1)2^{-n}$. Now let

$$\begin{aligned} I_k &:= G\big(k2^{-n}, (k+i_k)2^{-n}\big) && \text{for } 1 \le k < M,\\ I_M &:= G\big(m, (M+i_M)2^{-n}\big) && \text{and}\\ L &:= \max_{k=1,\ldots,M-1} I_k \end{aligned}$$

so that

$$\begin{aligned} \tilde{\mathbb{E}}[X_G|\sigma] &\ge \mathbb{E}\,(C \vee L \vee I_M) = C + \mathbb{E}\left((L \vee I_M - C)^+\right)\\ &= C + \int_C^\infty \mathbb{P}(L \vee I_M \ge h)\,\mathrm{d}h\\ &= C + \int_C^\infty 1 - \mathbb{P}(L < h)\,\mathbb{P}(I_M < h)\,\mathrm{d}h, \end{aligned} \tag{10}$$

where $\tilde{\mathbb{E}}$ stands for upper probability under the assumed restrictions on Market. Upon hitting level $k2^{-n}$, where $k < M$, the probability that we rise to level $(k+i)2^{-n}$ (or higher) before we hit level $(k-1)2^{-n}$ equals $\frac{1}{i+1}$. We have $\mathbb{P}(i_k \le j) = 1 - \frac{1}{2+j}$. Starting from the level $R2^{-n}$, the probability that we rise to level $(R+i)2^{-n}$ (or higher) before we hit level $(M-1)2^{-n}$ (where $M \le R$) equals $\frac{R-M+1}{R-M+i+1}$. We have $\mathbb{P}(M + i_M \le R + j) = \frac{j+1}{R-M+j+2}$; this formula is also true for $M = R+1$.

Since $G(a,b)$ is right-continuous in b for each a, the infimum in (4) is attained for each $h \ge 0$. We then have $G(a,b) < h$ for all $b < G^{-1}(a,h)$ and $G(a,b) \ge h$ for all $b \ge G^{-1}(a,h)$. And we have $G\left(a, G^{-1}(a,h)\right) \ge h$, with $>$ if the level h does not occur at all. Then, for $h \ge C$,

$$\begin{aligned} \mathbb{P}(I_M < h) &= \mathbb{P}\left(G\big(m,(M+i_M)2^{-n}\big) < h\right)\\ &= \mathbb{P}\left((M+i_M)2^{-n} < G^{-1}(m,h)\right)\\ &= \mathbb{P}\left(M + i_M < 2^n G^{-1}(m,h)\right) \end{aligned}$$

$$= \frac{1 - R + 2^n G^{-1}(m,h)}{2 - M + 2^n G^{-1}(m,h)}$$

and, for $k = 1, \dots, M-1$,

$$\begin{aligned}
\mathbb{P}(I_k < h) &= \mathbb{P}\left(G\big(k2^{-n}, (k+i_k)2^{-n}\big) < h\right) \\
&= \mathbb{P}\left((k+i_k)2^{-n} < G^{-1}(k2^{-n}, h)\right) \\
&= \mathbb{P}\left(i_k < -k + 2^n G^{-1}(k2^{-n}, h)\right) \\
&= 1 - \frac{1}{2 - k + 2^n G^{-1}(k2^{-n}, h)}
\end{aligned}$$

Therefore,

$$\begin{aligned}
\ln \mathbb{P}(L < h) &= \ln \prod_{k=1}^{M-1} \mathbb{P}(I_k < h) = \ln \prod_{k=1}^{M-1} \left(1 - \frac{1}{2 - k + 2^n G^{-1}(k2^{-n}, h)}\right) \\
&= \sum_{k=1}^{M-1} \ln \left(1 - \frac{1}{2 - k + 2^n G^{-1}(k2^{-n}, h)}\right) \\
&\le -\sum_{k=1}^{M-1} \frac{1}{2 - k + 2^n G^{-1}(k2^{-n}, h)} \\
&= -2^{-n} \sum_{k=1}^{M-1} \frac{1}{G^{-1}(k2^{-n}, h) - k2^{-n} + 2 \times 2^{-n}} \\
&\le -\sum_{k=1}^{M-1} \int_{k2^{-n}}^{(k+1)2^{-n}} \frac{\mathrm{d}a}{G^{-1}(a,h) - a + 3 \times 2^{-n}} \\
&\le -\int_{2^{-n}}^{M2^{-n}} \frac{\mathrm{d}a}{G^{-1}(a,h) - a + 3 \times 2^{-n}} \\
&\le -\int_{2^{-n}}^{m} \frac{\mathrm{d}a}{G^{-1}(a,h) - a + 3 \times 2^{-n}}.
\end{aligned}$$

Plugging these inequalities into (10) results in the lower bound

$$C + \int_C^\infty 1 - \frac{G^{-1}(m,h) - 2^{-n}(R-1)}{G^{-1}(m,h) - 2^{-n}(M-2)} \exp\left(\int_{2^{-n}}^{m} \frac{\mathrm{d}a}{a - G^{-1}(a,h) - 3 \times 2^{-n}}\right) \mathrm{d}h$$

for $\overline{\mathbb{E}}[X_G|\sigma]$. Letting $n \to \infty$, we obtain the inequality $\ge$ in (9). (Notice that we only need the convergence of the above outer integral to the outer integral in (9) when the limits of integration C and ∞ are replaced by $C \vee \epsilon$ and $D \in (C \vee \epsilon, \infty)$, respectively, where ϵ is a positive constant.)

5.2 Upper bound from Investor strategy

To prove the inequality $\le$ in (9), we consider the strategy that starts with initial capital equal to the expression in Theorem 6, and then in situation σ takes position (with m and C as defined in Theorem 7.)

$$S(\sigma) \;\coloneqq\; \int_C^\infty \frac{1}{G^{-1}(m,h)-m} \exp\left(\int_0^m \frac{\mathrm{d}a}{a-G^{-1}(a,h)}\right) \mathrm{d}h \tag{11}$$

(this is the derivative of the right-hand side of (9) w.r.t. the current price r). We are required to show that this strategy's capital is always equal to or exceeds the right-hand side of (9). Suppose this condition is satisfied at time t. Since the right-hand side of (9) is linear in r, this condition will still be satisfied at time $t+1$ if neither C nor m change. More generally, if the price becomes p at time $t+1$, the strategy's capital at time $t+1$ is required to be at least

$$f(p) \;\coloneqq\; C \vee G(m,p) + \int_{C\vee G(m,p)}^{\infty} 1 - \frac{G^{-1}(m\wedge p,h)-p}{G^{-1}(m\wedge p,h)-(m\wedge p)} \exp\left(\int_0^{m\wedge p} \frac{\mathrm{d}a}{a-G^{-1}(a,h)}\right) \mathrm{d}h.$$

Since the current capital is at least $f(r)$, it suffices to prove that $f(p)$ lies below our tangent $f(r)+S\cdot(p-r)$ to $f(p)$ at the point $p=r$. Therefore, it suffices to prove that f is concave. There are three regimes:

$$\frac{\partial^2 f(p)}{\partial^2 p} \;=\; \begin{cases} -\displaystyle\int_C^\infty \frac{\exp\left(\int_0^p \frac{\mathrm{d}a}{a-G^{-1}(a,h)}\right)\frac{\partial G^{-1}(p,h)}{\partial p}}{(p-G^{-1}(p,h))^2}\,\mathrm{d}h & \text{if } p<m\\ 0 & \text{if } m<p<G^{-1}(m,C)\\ -\dfrac{\exp\left(\int_0^m \frac{\mathrm{d}a}{a-G^{-1}(a,G(m,p))}\right)\frac{\partial G(m,p)}{\partial p}}{p-m} & \text{if } G^{-1}(m,C)<p \end{cases}$$

The first case is negative as $G^{-1}(p,h)$ increases in p. The last case is negative too, as $p-m$ is positive, and $G(m,p)$ increases in p. In the borderline cases $p=m$ and $p=G^{-1}(m,C)$, the required conditions for concavity on the one-sided first derivatives of f are easy to check.

6 Proof of Theorem 4

In this section we prove the representation theorem.

6.1 From north-west-sets to adjusters

Say $(I_h)_{h\geq 0}$ is a nested family of north-west sets, and Q is a probability measure on $[0,\infty)$. We now argue that

$$G(a,b) \;\coloneqq\; \int_0^\infty G_{I_h}(a,b)\,\mathrm{d}Q(h)$$

is an adjuster. It is a candidate guarantee; it is upper semi-continuous since all its super-level sets are closed and it is decreasing-increasing since each super-level set is north-west. It is an adjuster, witnessed by the strategy that splits the capital according to Q over strategies S_{I_h}.

6.2 From adjusters to north-west-sets

Say we have an arbitrary adjuster G. We now write it as a convex combination of nested north-west adjusters. Consider the family of super-level sets

$$I_h \;:=\; \left\{(a,b) \;\middle|\; G(a,b) \geq h\right\}$$

Since G is a candidate guarantee, each I_h is closed and north-west. By Theorem 6

$$G_{I_h}(a,b) \;=\; \frac{\mathbf{1}_{\{(a,b)\in I_h\}}}{1-\exp\left(\int_0^1 \frac{1}{a-G^{-1}(a,h)}\,\mathrm{d}a\right)}$$

is an admissible adjuster. Now construct the measure Q on $[0,\infty)$ with

$$Q(\,\mathrm{d}h) \;:=\; \left(1-\exp\left(\int_0^1 \frac{1}{a-G^{-1}(a,h)}\,\mathrm{d}a\right)\right)\,\mathrm{d}h.$$

Obviously Q is non-negative. In addition, since G is an admissible adjuster, Q integrates to 1 and hence is a probability measure. Finally, for each (a,b)

$$\int_0^\infty G_{I_h}(a,b)\,\mathrm{d}Q(u) \;=\; \int_0^\infty \mathbf{1}_{\{(a,b)\in I_h\}}\,\mathrm{d}h \;=\; G(a,b).$$

7 Discussion/Conclusion

We presented strategies for online trading that guarantee a large payoff when the price ever exhibits a large upcrossing, without taking any risk. We obtained an exact and elegant characterisation of the guarantees that can be achieved. We designed a guarantee that is close to ideal, and obtained an efficient strategy.

7.1 Applications

Our results are phrased in terms of finance. However, as we show in Theorem 4, a guarantee can always be achieved by a strategy that neither *sells short*, i.e. takes a negative position $S_t < 0$, or *uses leverage*, i.e. takes a position $S_t \geq K_{t-1}/\omega_{t-1}$ that is more expensive than the capital. So the fraction of capital invested $S_t\omega_{t-1}/K_{t-1} \in [0,1]$ is a proper probability. We can therefore think of our strategies as maintaining weights on two experts. If we substitute, in place of the price, the likelihood ratio between these two experts we obtain online methods for probability prediction with the log loss function.

One application lies is hierarchical modelling, where we want to aggregate at each level of detail the predictions of a model of that complexity, and the recursive combination of more refined models. This construction drives for example the successful data compression method Context Tree Weighting [5].

Another application is hypothesis testing, where a so-called null hypothesis is compared with an alternative hypothesis. Again, substituting the likelihood ratio for the price, securing a high payoff translates to amassing evidence against the null. The presence of a large upcrossing translates back to the existence of a sub-interval of data on which the null looks particularly fishy. Our strategies would report a fair and sharp measure of evidence in the presence of any such anomalous blocks. The advantage of this method is that the loss of evidence (the adjustment) is expressed in terms of the evidential power of the anomaly and not in its timing.

7.2 Downcrossings

A natural question is whether we can exploit the fact that a downcrossing $[a, b]$ occurs, i.e. that the price exceeds b before it drops below a. However, worst-case price paths for the univariate adjuster case always eventually collapse to 0, thus downcrossing any $[a, b]$ for $0 \leq a \leq b \leq \max_t \omega_t$. Hence, the presence of a downcrossing $[a, b]$ only conveys to us the information that the maximum is at least b, and we find ourselves back in the univariate adjuster case.

7.3 Future work

In this paper we focus on two-argument guarantees for buying once, then selling once. We are currently working on a full analysis of multi-argument guarantees for iterated trading: both for a fixed number of times and for arbitrary references.

Acknowledgments The first author is supported by NWO Rubicon grant 680-50-1010.

References

1. Dawid, A.P., de Rooij, S., Grünwald, P., Koolen, W.M., Shafer, G., Shen, A., Vereshchagin, N., Vovk, V.: Probability-free pricing of adjusted American lookbacks. ArXiv e-prints (Aug 2011)
2. Koolen, W.M., de Rooij, S.: Switching investments. In: Hutter, M., Stephan, F., Vovk, V., Zeugman, T. (eds.) Proceedings of the 21st International Conference on Algorithmic Learning Theory (ALT 2010). pp. 239–254. LNAI 6331, Springer, Heidelberg (Oct 2010), `http://www.cwi.nl/~wmkoolen/switching_investments.pdf`
3. Koolen, W.M., de Rooij, S.: Switching investments. Theoretical Computer Science (2012), the Special Issue on Algorithmic Learning Theory for ALT 2010
4. Vovk, V.: Continuous-time trading and the emergence of probability. Tech. rep., Royal Holloway, University of London (2010), `http://arxiv.org/abs/0904.4364`
5. Willems, F., Shtarkov, Y., Tjalkens, T.: The context tree weighting method: basic properties. IEEE Transactions on Information Theory 41(3), 653–664 (1995)