Buy Low, Sell High

Wouter M. Koolen[a,b], Vladimir Vovk[a]

[a] *Computer Learning Research Centre, Department of Computer Science, Royal Holloway, University of London, Egham, Surrey, TW20 0EX, United Kingdom*
[b] *Centrum Wiskunde & Informatica (CWI), P.O. Box 94079, NL-1090 GB Amsterdam, The Netherlands*

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 any upcrossing present. A financial advisor, to express that her secret strategy approximates this ideal, may publish a function $G(a, b)$ and promise that her strategy will

Initial price $\omega_0 = 1$.
Starting capital $K_0 = 1$.
For $t = 1, 2, \dots$

- Investor takes position $S_t \in \mathbb{R}$.
- Market reveals price $\omega_t \geq 0$.
- $K_t := K_{t-1} + S_t(\omega_t - \omega_{t-1})$

Figure 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?

Email addresses: wouter@cs.rhul.ac.uk (Wouter M. Koolen), v.vovk@rhul.ac.uk (Vladimir Vovk)

3. Do we need her? 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.

1.1.1. Sell high

Guarantees for trading once (selling at the maximum) were completely characterised by Dawid et al. (2011). 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 $\omega_0, \ldots, \omega_t$. An adjuster F *strictly dominates* another adjuster F' if $F(y) \geq F'(y)$ for all y, with $F(y) > F'(y)$ for some y. 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

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

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 admissible adjuster that is witnessed by the strategy S_u that takes position 1 until the price first exceeds u and takes position 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*

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

again with equality iff F is admissible.

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

1.1.2. Iterated trading

Intuitive trading strategies for iterated trading were proposed by Koolen and De Rooij (2010, 2013), 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 $S_Q \coloneqq \int S_{\alpha,\beta} Q(\mathrm{d}\alpha, \mathrm{d}\beta)$, 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

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

and so the full strategy S_Q witnesses

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

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

1.2. Climax

Intuitively, the two-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

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

(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. We do not formally 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

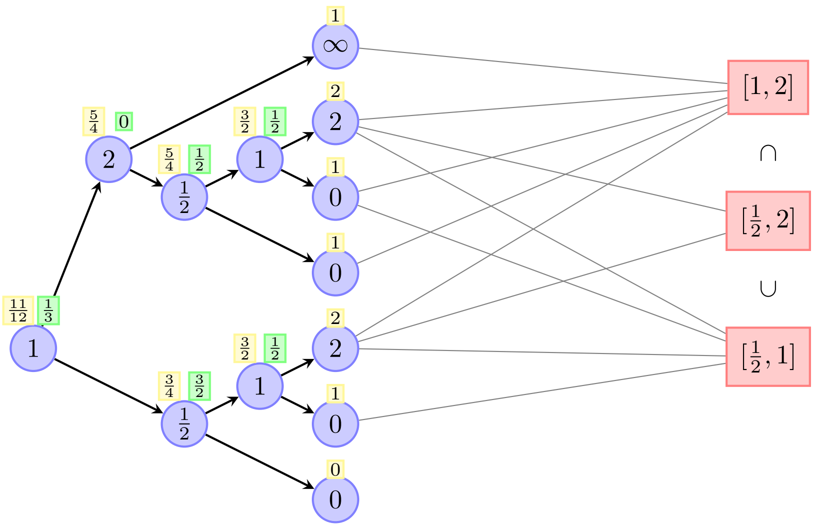

(a) Prices, capitals, and positions (b) Upcrossing relation (c) Intervals

Figure 2: Toy world

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 shown in Figure 2a, we assign to each leaf (identified by the price path ω leading to it) 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 given by

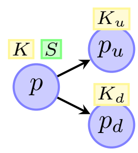

$$K \;=\; \frac{p - p_d}{p_u - p_d} K_u + \frac{p_u - p}{p_u - p_d} K_d \qquad \text{and} \qquad S \;=\; \frac{K_u - K_d}{p_u - p_d}. \tag{3}$$

To derive this, we minimise the objective K w.r.t. variables K and S subject to the constraints $K + S(p_u - p) \geq K_u$ and $K + S(p_d - p) \geq K_d$. These constraints may be combined to $\frac{K_u - K}{p_u - p} \leq S \leq \frac{K - K_d}{p - p_d}$. This form allows us to eliminate the variable S, resulting in the constraint $\frac{K_u - K}{p_u - p} \leq \frac{K - K_d}{p - p_d}$, which we may reorganise to $(p - p_d)K_u + (p_u - p)K_d \leq (p_u - p_d)K$. Minimising K and back-substitution to obtain S now result in (3).

Using binomial pricing, we have labelled each internal situation in Figure 2a 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 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 the price reaches 1, we cash in 1 dollar leaving a position of at least $\frac{1}{2}$. If 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 the price reaches $\frac{1}{2}$, we take a position of $\frac{1}{2}$, which leaves at least 1 dollar in cash. If 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 say that a price sequence $\omega = (\omega_0, \ldots, \omega_t)$ *upcrosses* $[a, b]$ if there are times $0 \le i \le j \le t$ such that $\omega_i \le a$ and $\omega_j \ge b$. We call a function $G : D \to [0, \infty)$ with domain

$$D \;\coloneqq\; \big\{(a, b) \;\big|\; 0 < a \le 1, b \ge a\big\}$$

a *candidate guarantee* if it is upper semi-continuous, decreasing in its first argument and increasing in its second argument. A candidate guarantee G is an *adjuster* if there is a strategy that ensures $K_t \ge G(a, b)$ for every upcrossing $[a, b]$ of every price evolution ω. An adjuster that is not strictly dominated is *admissible*. In our characterisation of the admissible adjusters we will need the following technical condition. We say that a candidate guarantee G is *saturated* if

$$G \;\ge\; s_G \;\coloneqq\; \inf\left\{h \ge 0 \;\middle|\; \int_{G(a,b) \ge h} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2} < \infty\right\}, \tag{4}$$

where the integral is over D and $\inf \emptyset = \infty$ (by the σ-additivity of measures, the integral in (4) is infinite when $h = s_G$). It will follow from Theorem 3 below that $G \vee s_G$ (we abbreviate binary minima to $\wedge$ and maxima to $\vee$) is a saturated adjuster for any adjuster G; in this sense any adjuster can be saturated by increasing it to the *point of saturation* s_G.

This allows us to demarcate precisely the achievable and optimal guarantees. The proof is split into Fact 5, Theorem 6 and Corollary 7 below.

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

$$\int_0^\infty 1 - \exp\left(- \int_{G(a,b) \geq h} \frac{\mathrm{d}a \, \mathrm{d}b}{(b-a)^2}\right) \mathrm{d}h \;\leq\; 1. \tag{5}$$

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

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 can rewrite (5) in terms of the *second-argument upper inverse* of G

$$G^{-1}(a,h) \;\coloneqq\; \inf\{b \geq a \mid G(a,b) \geq h\} \tag{6}$$

as follows:

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

We also provide a canonical representation of each adjuster as a convex combination of elementary guarantees. These elementary guarantees are analogous to the threshold strategies of the univariate case in the sense that they have just one non-zero payoff level. However, they do have richer geometric structure. A closed set $I \subseteq D$ is called *north-west* if $(a,b) \in I$ implies $(0,a] \times [b,\infty) \subseteq I$. Some examples of north-west sets are displayed in Figure 3. We associate to each north-west set its *frontier*

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

By the previous theorem, the following guarantee is an adjuster when $I \neq \emptyset$:

$$G_I(a,b) \;\coloneqq\; \frac{\mathbf{1}_{\{(a,b) \in I\}}}{1 - \exp\left(- \int_I \frac{\mathrm{d}a \, \mathrm{d}b}{(b-a)^2}\right)} \;=\; \frac{\mathbf{1}_{\{f_I(a) \leq b\}}}{1 - \exp\left(- \int_0^1 \frac{1}{f_I(a') - a'} \, \mathrm{d}a'\right)}.$$

We call a non-empty set $I \subseteq D$ *admissible* if either $I = D$ or $\int_I \frac{\mathrm{d}a \, \mathrm{d}b}{(b-a)^2} < \infty$. So the adjuster G_I is admissible iff the set I is admissible. A family $(I_h)_{h \geq 0}$ of north-west sets is called *nested* if $x \leq y$ implies $I_x \supseteq I_y$. It is called *closely nested* if it is nested and for each $(a,b) \in D$ the set $\{h \mid (a,b) \in I_h\}$ is closed (remember that each I_h is closed by the definition of nested families). It is called *admissible* if each I_h is either admissible or empty.

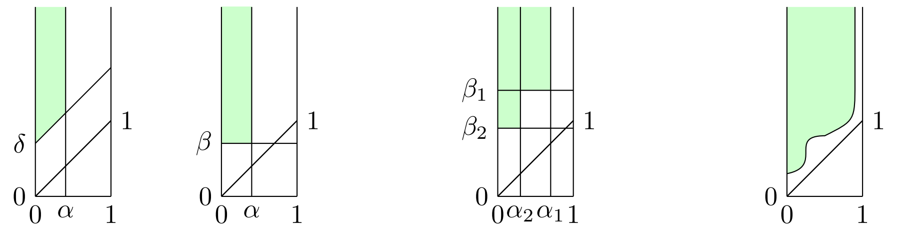

(a) Skewed (b) Square (c) Disjunction of squares (d) Curved set

Figure 3: Examples of north-west sets

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

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

G is admissible iff this holds with equality and the family (I_h) 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 closely nested family of north-west sets $(I_h)_{h \geq 0}$ (with $I_0 = D$ playing no useful role). 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 $S_G \coloneqq \int S_{I_h} Q(\mathrm{d}h)$.

These two theorems parallel those of Dawid et al. (2011) with a twist. Whereas Dawid et al. (2011) decompose 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 price ω_0 is one, and that the investor starts with initial capital $K_0 = c$. 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 $K_T = c + \sum_{t=1}^{T} S_t(\omega_t - \omega_{t-1})$. A *trading strategy* S

assigns to each sequence of past prices $\omega_{<t} = (\omega_0, \dots, \omega_{t-1})$ a *position* $S(\omega_{<t}) \in \mathbb{R}$. Let $S *_c \omega$ denote the payoff (i.e. final capital) of strategy S on price sequence ω. That is

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

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

$$\overline{\mathbb{E}}\,[X] \;\coloneqq\; \inf\big\{c \;\big|\; \exists \text{ strategy } S \;\forall \text{ price sequence } \omega : \; S *_c \omega \geq X(\omega) \;\big\}.$$

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 $\omega' = (\omega'_0, \dots \omega'_t)$. This so-called *conditional upper price* is given by

$$\overline{\mathbb{E}}\,[X|\omega'] \;\coloneqq\; \inf\big\{c \;\big|\; \exists \text{ strategy } S \;\forall \text{ price sequence } \omega : \; S *_c \omega \geq X(\omega'_{<t}\omega) \;\big\},$$

where ω ranges over price sequences 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 sequence ω the minimal payoff necessary to guarantee G on it:

$$X_G(\omega) \;\coloneqq\; \sup_{[a,b]\,:\,\omega \text{ upcrosses } [a,b]} G(a,b) \;=\; \max_{\substack{0 \leq i \leq j \\ 1 \geq \omega_i \leq \omega_j}} G(\omega_i, \omega_j) \tag{7}$$

We now connect adjusters and pricing:

Fact 5. *A candidate guarantee G is an adjuster iff $\overline{\mathbb{E}}[X_G] \leq 1$.*

This observation 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*

$$\overline{\mathbb{E}}[X_G] \;=\; \int_0^\infty 1 - \exp\left(- \int_{G(a,b) \geq h} \frac{\mathrm{d}a \,\mathrm{d}b}{(b-a)^2} \right) \mathrm{d}h. \tag{8}$$

In terms of G^{-1} this can be rewritten as

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

Using this pricing formula, we are also able to characterise admissibility.

Corollary 7. *A candidate G is an admissible adjuster iff $\overline{\mathbb{E}}[X_G] = 1$ and G is saturated.*

Proof. If $\overline{\mathbb{E}}[X_G] > 1$ then G is not an adjuster. If $\overline{\mathbb{E}}[X_G] < 1$ then G is strictly dominated by the renormalised adjuster $G/\overline{\mathbb{E}}[X_G]$. If $\overline{\mathbb{E}}[X_G] = 1$ and G is not saturated it is strictly dominated by the (saturated) adjuster $G \vee s_G$, where s_G is the point of saturation defined in (4).

For the reverse, if $\overline{\mathbb{E}}[X_G] = 1$ and G is saturated, then any strictly dominating candidate guarantee must have upper price > 1 by the previous theorem, and cannot be an adjuster. Indeed, suppose $G' \geq G$ and $G'(a,b) > G(a,b)$. Since G is upper semi-continuous, $G'(a,b) > c+\epsilon > c-\epsilon > G(a',b')$ for some $c, \epsilon > 0$, $a' < a$, and $b' > b$. Our monotonicity assumption about G and G' now implies $G'(\alpha,\beta) > c+\epsilon > c-\epsilon > G(\alpha,\beta)$ for all $\alpha \in (a',a)$ and $\beta \in (b,b')$. To prove that $\overline{\mathbb{E}}[X_{G'}] > \overline{\mathbb{E}}[X_G] = 1$ it suffices to establish that

$$\int_{G'(\alpha,\beta)\geq h} \frac{\mathrm{d}\alpha \,\mathrm{d}\beta}{(\beta-\alpha)^2} > \int_{G(\alpha,\beta)\geq h} \frac{\mathrm{d}\alpha \,\mathrm{d}\beta}{(\beta-\alpha)^2}$$

on a set of h of positive measure. This is obvious since $\{G' \geq h\} \supseteq \{G \geq h\}$ for all h, the difference $\{G' \geq h\} \setminus \{G \geq h\}$ contains $(a',a) \times (b,b')$ for all $h \in (c-\epsilon, c+\epsilon)$, and both integrals are finite for $h \in (c-\epsilon, c+\epsilon)$. □

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 Dawid et al. (2011) (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 saturated guarantee $G(a,b) := F(b \vee 1)$ that ignores its first argument. By Theorem 6

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

To evaluate the rightmost integral, we use the variable substitution $h = F(y)$ (for $y \geq 1$ and $h \geq F(1)$). We then employ integration by parts, and obtain

$$\begin{aligned}
\overline{\mathbb{E}}[X_G] &= F(1) + \int_1^{\infty} \frac{1}{y} \,\mathrm{d}F(y) && (9)\\
&= F(1) + \left.\frac{F(y)}{y}\right|_1^{\infty} + \int_1^{\infty} \frac{F(y)}{y^2} \,\mathrm{d}y \\
&= \int_1^{\infty} \frac{F(y)}{y^2} \,\mathrm{d}y && (10)
\end{aligned}$$

This derivation assumes that $F(\infty)/\infty = 0$. If $F(\infty)/\infty$ exists and is strictly positive, both (9) and (10) are equal to ∞, and so $\overline{\mathbb{E}}[X_G]$ is still equal to (10). And if $F(\infty)/\infty$ does not exist, both (9) and (10) are again equal to ∞: if one or both of them were finite, $F(\infty)/\infty$ would exist as their difference. $\square$

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 (10), 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 (10) 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 highly desirable guarantees are unfortunately way too good to be true: they cannot 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

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

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

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

for any $0 \leq q < p < 1$ (e.g. $0.0145\ (b-a)^{0.9}/a^{0.8}$ will be an adjuster). The results in Section 5.3 below imply that this guarantee is witnessed by the strategy that in situation ω with minimum price m takes position

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

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

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

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

5. Proof of Theorem 6

In this section we prove the characterisation theorem. We first give the intuition, and then prove the lower bound by constructing a strategy for Market and the upper bound by constructing a strategy for Investor.

5.1. Intuitive picture

Say we restrict Market to price sequences of some fixed length T that either increase or decrease each time step by $\pm 2^{-n}$. Then we can label each possible price path ω by $X_G(\omega)$, the required capital to guarantee G on it, and apply binomial pricing as before in backward fashion to determine the upper price of G in this restricted world. We find that the upper price is the expected value of X_G under a random walk that takes steps $\pm 2^{-n}$ with equal probability, and that is stopped at zero.

By letting $T \to \infty$ and $n \to \infty$ we get sharper lower bounds on the upper price. Even though we do not consider trading in continuous time, we do find that the upper price is the expected value of X_G when price paths are sampled from Brownian motion (again stopped at zero).

This limit still has the built-in restriction that Market does not produce price paths with jumps. Interestingly, as we will show in this section, this limit-based lower bound is met (by the limit strategy for Investor). This proves in particular that Market cannot make adversarial use of price jumps. Note that this fact does not hold for option pricing in general, but it does hold for our particular option X_G.

It will be convenient to prove the following more general statement.

Theorem 8. *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 (7)). The conditional upper price of X_G in situation σ is*

$$\overline{\mathbb{E}}[X_G|\sigma] \;=\; C + \int_C^\infty 1 - \exp\left(- \int_{\substack{G(a \wedge m, b) \geq h \\ 0 < a \leq r \;;\; a \leq b}} \frac{\mathrm{d}a \, \mathrm{d}b}{(b-a)^2} \right) \mathrm{d}h. \tag{11}$$

The above equation (11) may also be written using G^{-1} as

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

The proof consists of two parts. For the lower bound ($\geq$) we construct an adversarial Market based on random walks. For the upper bound ($\leq$) 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 Vovk (2012), who derives option prices assuming continuous price paths. We are not aware of general probability-free option pricing results that allow discontinuous price processes.

5.2. 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 $R2^{-n}$ with $R \coloneqq \lfloor \omega_s 2^n \rfloor$ (rather than the real ω_s, so that prices from time $s+1$ are indeed multiples of 2^{-n}). Define the stopping time τ to be the least time such that $\omega_\tau = 0$. On run ω, we desire to superreplicate X_G, which can be rewritten as

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

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

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

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 $M \coloneqq \lceil m 2^n \rceil$. 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

$$\begin{aligned} I_k \; &\coloneqq \; G\big(k2^{-n}, (k+i_k)2^{-n}\big) \qquad \text{for } 1 \leq k < M, \\ I_M \; &\coloneqq \; G\big(m, (M+i_M)2^{-n}\big) \qquad \text{and} \end{aligned}$$

$$L \;\coloneqq\; \max_{k=1,\dots,M-1} I_k$$

so that

$$\begin{aligned}
\tilde{\mathbb{E}}[X_G|\sigma] \;&\geq\; \mathbb{E}\,(C \vee L \vee I_M) \;=\; C + \mathbb{E}\left((L \vee I_M - C)^+\right) \\
&=\; C + \int_C^\infty \mathbb{P}(L \vee I_M \geq h)\,\mathrm{d}h \\
&=\; C + \int_C^\infty 1 - \mathbb{P}(L < h)\,\mathbb{P}(I_M < h)\,\mathrm{d}h, \qquad (12)
\end{aligned}$$

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 $\frac{1}{i+1}$.[1] We have $\mathbb{P}(i_k \leq j) = 1 - \frac{1}{2+j}$. 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 $M \leq R$) equals $\frac{R-M+1}{R-M+i+1}$. We have $\mathbb{P}(M + i_M \leq R + j) = \frac{j+1}{R-M+j+2}$; this formula is also true for $M = R+1$.

Since $G(a,b)$ is right-continuous in b for each a, the infimum in (6) is attained for each $h \geq 0$. We then have $G(a,b) < h$ for all $b < G^{-1}(a,h)$ and $G(a,b) \geq h$ for all $b \geq G^{-1}(a,h)$. And we have $G\left(a, G^{-1}(a,h)\right) \geq h$, with $>$ if the level h does not occur at all. Then, for $h \geq C$,

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

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,

$$\ln \mathbb{P}(L < h) \;=\; \ln \prod_{k=1}^{M-1} \mathbb{P}(I_k < h) \;=\; \ln \prod_{k=1}^{M-1} \left(1 - \frac{1}{2 - k + 2^n G^{-1}(k2^{-n},h)}\right)$$

[1] A cute way to obtain this and similar statements is as follows. The price is a martingale, and hence so is the price stopped at either level $(k+i)2^{-n}$ or $(k-1)2^{-n}$. Furthermore, such stopping happens almost surely. Writing p for the probability of hitting the high level first, the fact that the expected stopped price equals the current price translates to the equation $k2^{-n} = p(k+i)2^{-n} + (1-p)(k-1)2^{-n}$, which is results in $p = \frac{1}{i+1}$.

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

Plugging these inequalities into (12) results in the lower bound

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

for $\overline{\mathbb{E}}[X_G|\sigma]$. Letting $n \to \infty$, we obtain the inequality $\geq$ in (11). (Notice that we only need the convergence of the above outer integral to the outer integral in (11) 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.)

5.3. Upper bound from Investor strategy

Now we prove the inequality $\leq$ in (11). To prove this statement of the form $\overline{\mathbb{E}}\left[X_G|\sigma\right] \leq \phi(\sigma)$ we need to (see Section 2) exhibit a trading strategy that, starting in situation $\sigma = (\sigma_0, \ldots, \sigma_t)$ with capital $\phi(\sigma)$ maintains its capital above $X_G(\sigma\omega)$ for every finite sequence ω of future prices. Our approach will be to find a strategy that in fact maintains the capital above $\phi(\sigma\omega) \geq X_G(\sigma\omega)$. To accomplish this, we only need to consider how the capital and ϕ change over a single trading round. That is, in situation σ with capital equal to $\phi(\sigma)$, we need to choose a position S that for every next price r ensures

$$
\phi(\sigma) + S(r - \sigma_t) \;\geq\; \phi(\sigma, r).
$$

To argue that such an S exists, we first observe that a repeated last price does not change ϕ, so that $\phi(\sigma) = \phi(\sigma, \sigma_t)$. The core of the proof (presented below) lies in establishing that $\phi(\sigma, r)$ is concave in r. The desired position S may then be taken equal to any super-derivative of $\phi(\sigma, r)$ at the current price $r = \sigma_t$.

Fix a situation $\sigma = (\sigma_0, \ldots, \sigma_t)$ with minimum $m > 0$ and a next price $r > 0$ (we deal with the situation $r = 0$ separately at the end). We now proceed to the concavity argument. For convenience, we rewrite the right-hand side of (11) evaluated at σ, r as

$$
\phi(\sigma, r) \;=\; \int_0^\infty 1 - \exp\left(-\int_{G^{\sigma,r}(a,b)\geq h} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right) \mathrm{d}h \tag{13}
$$

where the constant term C is incorporated as a diverging inner integral and the complex range of the inner integral is absorbed into the renormalised guarantee

$$G^{\sigma,r}(a,b) \;:=\; \max\bigl\{X_G(\sigma), G(r\wedge m, r), G(ra\wedge m, rb)\bigr\}.$$

Since arbitrary sums and integrals of concave functions are concave, it is sufficient to show that

$$\exp\left(-\int_{G^{\sigma,r}(a,b)\geq h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right) \tag{14}$$

is convex in $r>0$ for all $h\geq 0$. First, if $h\in[0,X_G(\sigma)]$ the integral in (14) diverges, and (14) is identically 0 and hence convex. So now assume that $h>X_G(\sigma)$. We then argue that the expression in (14) equals

$$f(r) \;:=\; \exp\left(-\int_{G(ra\wedge m, rb)\geq h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right) \;=\; \exp\left(-\int_{\substack{G(a\wedge m,b)\geq h\\ 0<a\leq r}}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right).$$

The difference between (14) and $f(r)$ is the absence of the term $G(r\wedge m,r)$ from the maximum. The equality is hence obvious for r such that $G(r\wedge m,r)<h$. For r such that $G(r\wedge m,r)\geq h$ the integral in (14) diverges, but so does that in $f(r)$, and both are hence equal to zero. It remains to establish the convexity of $f(r)$ on $r>0$. Our approach will be to find a sub-derivative f' (that is show $f(p)-f(r)\geq f'(r)(p-r)$ for all $p,r>0$). First if $f(r)=0$ we may take $f'(r)=0$ since $f(p)\geq 0$ for all p. For $f(r)>0$ we formally calculate the first derivative in r, which equals

$$f'(r) \;=\; -\left(\int_{G(r\wedge m,b)\geq h}\frac{\mathrm{d}b}{(b-r)^2}\right)f(r) \;=\; \frac{-f(r)}{G^{-1}(r\wedge m,h)-r},$$

where the last equality uses $f(r)>0$, so that $G(r\wedge m,r)<h$ and hence $G^{-1}(r\wedge m,h)>r$. (We might have made some other assumptions in this calculation, such as the continuity of G in its first argument, but this does not matter as we will never use the fact that $f'(r)$ is f's derivative.) To show that $f'(r)$ is a sub-derivative we may split the integral inside $f(p)$ (since $f(r)>0$ the integral to r is finite) and rewrite

$$f(p)=f(r)\exp\left(-\int_r^p\int_{b:G(a\wedge m,b)\geq h}\frac{\mathrm{d}b}{(b-a)^2}\,\mathrm{d}a\right)$$

with the convention that $\int_r^p=-\int_p^r$ for reversed limits $r>p$. We hence need to show

$$\left(\exp\left(-\int_r^p\int_{b:G(a\wedge m,b)\geq h}\frac{\mathrm{d}b}{(b-a)^2}\,\mathrm{d}a\right)-1\right)f(r) \;\geq\; (p-r)\frac{-1}{G^{-1}(r\wedge m,h)-r}f(r)$$

which may be reorganised as

$$\exp\left(-\int_r^p\int_{b:G(a\wedge m,b)\geq h}\frac{\mathrm{d}b}{(b-a)^2}\,\mathrm{d}a\right) \;\geq\; \frac{G^{-1}(r\wedge m,h)-p}{G^{-1}(r\wedge m,h)-r}.$$

Since the left-hand side is nonnegative and $G^{-1}(r\wedge m,h)>r$, this is trivial for $G^{-1}(r\wedge m,h)\leq p$. If on the other hand $G^{-1}(r\wedge m,h)>p$ we may rewrite the requirement further to

$$\int_r^p\int_{b:G(a\wedge m,b)\geq h}\frac{\mathrm{d}b}{(b-a)^2}\,\mathrm{d}a \;\leq\; \log\left(\frac{G^{-1}(r\wedge m,h)-r}{G^{-1}(r\wedge m,h)-p}\right)$$

But this follows from a simple property of G. For if $r \leq p$ then we may make the left-hand side larger by replacing $G(a \wedge m, b)$ with the larger $G(r \wedge m, b)$. And if $r > p$ then the sign is negative, and we may make the left-hand side larger by replacing $G(a \wedge m, b)$ with the smaller $G(r \wedge m, b)$. In both cases, computing the integral gives

$$\int_r^p \int_{b:G(r\wedge m,b)\geq h} \frac{\mathrm{d}b}{(b-a)^2}\,\mathrm{d}a \;=\; \int_r^p \frac{\mathrm{d}a}{G^{-1}(r\wedge m,h)-a} \;=\; \log\left(\frac{G^{-1}(r\wedge m,h)-r}{G^{-1}(r\wedge m,h)-p}\right)$$

as desired. (The last calculation uses the fact that $G^{-1}(r \wedge m, h)$ is outside the interval of integration in a.)

Finally, we need to argue that this strategy works when the next price drops to $r = 0$. In that case $X_G(\sigma, 0) = X_G(\sigma)$. Since the strategy ensured capital $\geq X_G(\sigma, r) \geq X_G(\sigma)$ for all $r > 0$, the resulting capital at $r = 0$ must still exceed $X_G(\sigma)$ as required.

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 closely nested family of north-west sets and Q is a probability measure on $(0, \infty)$ concentrated on $\{h \mid I_h \neq \emptyset\}$. We now argue that

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

is an adjuster (when we write $\int_x^y$ we mean $\int_{[x,y]}$ unless $x = 0$ or $y = \infty$: 0 and ∞ are not included into the interval of integration). 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}. If, in addition, all non-empty I_h are admissible, Lemmas 9 and 10 below establish that G has unit upper price and is saturated. Admissibility then follows from Theorem 3.

Lemma 9. *If $(I_h)_{h\geq 0}$ is a closely nested family of north-west sets and Q is a probability measure on $(0, \infty)$ concentrated on $\{h > 0 \mid I_h \neq \emptyset\}$, then $G(a, b) = \int G_{I_h}(a, b)Q(\mathrm{d}h)$ has* (5) *with equality.*

Proof. ≤ 1 in (5) follows from G being an adjuster. We argue ≥ 1.

For each h the set I_h is contained in the super-level set $\{(a, b) \mid G(a, b) \geq h'\}$ where

$$h' \;:=\; \int_0^h \frac{1}{1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(a-b)^2}\right)}Q(\mathrm{d}h) \tag{15}$$

(with $\frac{1}{0}$ understood to be 0). Indeed, if $(a, b) \in I_h$, we have

$$G(a,b) = \int G_{I_h}(a,b)Q(\mathrm{d}h) = \int \frac{\mathbf{1}_{\{(a,b)\in I_h\}}Q(\mathrm{d}h)}{1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)} \geq \int_0^h \frac{Q(\mathrm{d}h)}{1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)} = h' \tag{16}$$

(the same symbol h plays two different roles in this chain, the upper limit of integration and the dummy variable, but this should not lead to confusion). First we give an informal argument that best conveys the intuition. It follows from (15) that

$$\mathrm{d}h' \;=\; \frac{1}{1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)}Q(\mathrm{d}h).$$

With this substitution

$$\int_0^\infty 1-\exp\left(-\int_{G(a,b)\geq h'}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)\mathrm{d}h' \;\geq$$
$$\int_0^\infty \left(1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)\right)\frac{1}{1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)}Q(\mathrm{d}h) \;=\; 1 \qquad (17)$$

as desired.

Since the family (I_h) is closely nested, the decreasing function $\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}$ is left-continuous in h. Indeed, if it is not, say at a point $h=h_0>0$, we have $I_{h_0}\neq I_{h_0-}$, where $I_{h_0-}:=\cap_{h<h_0}I_h$; for any point $(a,b)\in I_{h_0-}\setminus I_{h_0}$, the set $\{h\mid(a,b)\in I_h\}$ contains all $h<h_0$ but does not contain h_0, and so is not closed.

To formalize the argument (17), we rewrite the second integral $\int_0^\infty$ in (17) as $\int_0^\infty f(h)\mu(\mathrm{d}h)$, where

$$f(h):=1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right), \qquad \mu(\mathrm{d}h):=\frac{1}{1-\exp\left(-\int_{I_h}\frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)}Q(\mathrm{d}h).$$

The function f is left-continuous and decreasing; therefore, it is upper semi-continuous. The measure μ is σ-finite. Analogously, the first integral $\int_0^\infty$ in (17) can be written as $\int_0^\infty g(h')\,\mathrm{d}h'$, where g is a decreasing function. Our goal is to prove the inequality in

$$\int_0^\infty g(h')\,\mathrm{d}h' \geq \int_0^\infty f(h)\mu(\mathrm{d}h)=1.$$

We replace $\int_0^\infty f(h)\mu(\mathrm{d}h)$ by a corresponding Lebesgue sum, which can be made arbitrarily close to 1. Namely, divide the range of f, which we take to be $[0,\infty)$ (rather than, say, $[0,1]$), into the intervals $[0,\epsilon),[\epsilon,2\epsilon),\ldots$ of size ϵ, and let $(x_i,y_i]:=f^{-1}([(i-1)\epsilon,i\epsilon))$ for $i=1,2,\ldots$; $(x_i,y_i]$ will form a partition of $[0,\infty)$ with the rightmost element of the partition, $(x_1,y_1)=(x_1,\infty)$, being the only exception to our notation $(x_i,y_i]$. (Some of the $(x_i,y_i]$ may be empty but otherwise they will be closed on the right and open on the left, by the upper semi-continuity of f, apart from the exception mentioned earlier.) By monotone convergence $\sum_{i=1}^\infty(i-1)\epsilon\mu((x_i,y_i])\to\int f\,\mathrm{d}\mu$ as $\epsilon\to0$ (say geometrically). By the definition (15) of h' we have

$$\sum_{i=1}^\infty(i-1)\epsilon\mu((x_i,y_i])=\sum_{i=1}^\infty(i-1)\epsilon(y_i'-x_i')\leq\sum_{i=1}^\infty g(y_i')(y_i'-x_i')\leq\int_0^\infty g(h')\,\mathrm{d}h'$$

(where all the sums are in fact finite), and it remains to let $\epsilon\to0$. $\square$

Lemma 10. *If* $(I_h)_{h\geq 0}$ *is an admissible closely nested family of north-west sets and* Q *is a probability measure on* $(0,\infty)$ *concentrated on* $\{h>0 \mid I_h \neq \emptyset\}$ *then* $G(a,b) = \int G_{I_h}(a,b) Q(\mathrm{d}h)$ *is saturated.*

PROOF. First observe that $\{(a,b) \mid G(a,b) > h'\} \subseteq I_h \subseteq \{(a,b) \mid G(a,b) \geq h'\}$ under the correspondence (15) (the first inclusion follows from (16) with $\leq$ in place of $\geq$, which holds for $(a,b) \notin I_h$). Set $t_G \coloneqq \sup\{h \mid I_h = D\}$ and notice that $I_{t_G} = D$ (we saw in the previous proof that $\int_{I_h} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}$ is a left-continuous function of h). The right inclusion of the sandwich implies that $s_G \geq t'_G$. The left inclusion in combination with the right-continuity of h' as function (15) of h now implies that $s_G = t'_G$: indeed, if $s_G > t'_G$, we have

$$\int_{G(a,b)\geq s_G} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2} \leq \int_{G(a,b)>t'} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2} \leq \int_{I_t} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2} < \infty$$

for any $t > t_G$ such that $t'_G \leq t' < s_G$, which contradicts the fact that $\int_{G(a,b)\geq s_G} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2} = \infty$ (following from the definition of s_G and the σ-additivity of measures). Again applying the right inclusion, we can see that $\{G \geq s_G\} \supseteq I_{t_G} = D$, which means that G is saturated. $\square$

6.2. From adjusters to north-west sets

Say we have an adjuster G. We immediately assume that it is admissible, because the inadmissible case follows by taking a dominating admissible adjuster. We now write it as a convex combination of closely nested north-west adjusters. Consider the family of super-level sets

$$I_h \;\coloneqq\; \big\{(a,b) \bigm| G(a,b) \geq h\big\}.$$

Since G is a candidate guarantee, each I_h is north-west (in particular, closed). The family of I_h is closely nested. As G is saturated, each non-empty I_h is admissible. By Theorem 6

$$G_{I_h}(a,b) \;=\; \frac{\mathbf{1}_{\{(a,b)\in I_h\}}}{1-\exp\left(-\int_{G(a,b)\geq h} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)}$$

is an admissible adjuster when $I_h \neq \emptyset$. Now construct the measure Q on $(0,\infty)$ with

$$Q(\mathrm{d}h) \;\coloneqq\; \left(1-\exp\left(-\int_{G(a,b)\geq h} \frac{\mathrm{d}a\,\mathrm{d}b}{(b-a)^2}\right)\right)\mathrm{d}h.$$

Obviously Q is non-negative and concentrated on $\{h>0 \mid I_h \neq \emptyset\}$. In addition, since G is an admissible adjuster, Q integrates to 1 and hence is a probability measure. Finally, for each (a,b)

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

(with G_{I_h} understood to be 0 when $I_h = \emptyset$).

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 by Willems et al. (1995).

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 (see, e.g. Shafer et al. 2011). 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. A treatment of achievable guarantees parametrised by the block timing can be found in the work of Adamskiy et al. (2012).

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. Another interesting direction is computational efficiency. What resources are required to execute the strategy witnessing an arbitrary admissible adjuster G?

Acknowledgments

We are very grateful to the reviewers of the conference and journal versions of this paper, whose thoughtful comments helped us to eliminate several mistakes and improve the presentation. The first author was supported by NWO Rubicon grant 680-50-1010.

Adamskiy, D., Koolen, W. M., Chernov, A., Vovk, V., Oct. 2012. A closer look at adaptive regret. In: Bshouty, N., Stoltz, G., Vayatis, N., Zeugmann, T. (Eds.), Proceedings of the 23rd International Conference on Algorithmic Learning Theory (ALT). LNAI 7568. Springer, Heidelberg, pp. 290–304.

Dawid, A. P., De Rooij, S., Grünwald, P., Koolen, W. M., Shafer, G., Shen, A., Vereshchagin, N., Vovk, V., Aug. 2011. Probability-free pricing of adjusted American lookbacks. Tech. Rep. 1108.4113 [q-fin.PR], arXiv e-prints.

Koolen, W. M., De Rooij, S., Oct. 2010. Switching investments. In: Hutter, M., Stephan, F., Vovk, V., Zeugman, T. (Eds.), Proceedings of the 21st International Conference on Algorithmic Learning Theory (ALT 2010). LNAI 6331. Springer, Heidelberg, pp. 239–254.

Koolen, W. M., De Rooij, S., 2013. Switching investments. Theoretical Comput. Sci. 473, 61–76, Special Issue on Algorithmic Learning Theory.

Shafer, G., Shen, A., Vereshchagin, N., Vovk, V., 2011. Test martingales, Bayes factors, and p-values. Statistical Science 26, 84–101.

Vovk, V., 2012. Continuous-time trading and the emergence of probability. Finance and Stochastics 16, 561–609.

Willems, F., Shtarkov, Y., Tjalkens, T., 1995. The context tree weighting method: basic properties. IEEE Transactions on Information Theory 41, 653–664.