LOG-OPTIMAL PORTFOLIOS FOR TREE-STRUCTURED WAGERS

NIGEL R. O'BRIAN

ABSTRACT. A well-known result of J.L. Kelly gives the log-optimal portfolio for a horse race in terms of each runner's probability of winning and corresponding betting returns. Here Kelly's result is generalised to combinations of 'exotic' bets, which depend on runners finishing in specified order. The method takes the form of a exact, recursive algorithm and applies more generally to yield a log-optimal portfolio for any gambling situation where the available bets have a natural tree-type structure.

KEYWORDS: Log-optimal, portfolio, gambling, bet, wager, Kelly, convexity, Kuhn-Tucker.

1. INTRODUCTION

Given a set of N risky and interrelated investment opportunities, let Z_i be the random variable representing the value of a unit investment in opportunity i after a fixed time period. Alternatively, the Z_i are the payoffs from various types of wager in a game of chance. If the joint distribution of the Z_i is known, an investor or gambler will typically want to balance return against risk by constructing a diversified portfolio from the available opportunities. If x_i is the fraction of available funds allocated to opportunity i, then

$$\sum_{i=1}^{N} x_i Z_i$$

is the total value of an initial unit investment after the given time period or game. Here we consider *log-optimal portfolios*, as in Kelly [4], where the allocations x_i are chosen to maximise the *expected rate of return*

$$R(x_1, \ldots, x_N) = E\left(\ln \sum_{i=1}^{N} x_i Z_i\right). \tag{1}$$

For repeated reinvestment of funds over a sequence of time periods or games portfolios of this type have a number of asymptotic optimality properties which make them of special interest. See [2] and the references there for further discussion, along with precise statements and proofs.

In [4] Kelly points out the close relation between the expected rate of return and the *information rate* of the underlying random process. This is illustrated by the example of a horse race, where a gambler can make bets

Date: October 16, 2001.

on the various runners to win the race. Kelly also shows how to find the log-optimal portfolio in this case. In the simplest situation the gambler is required to bet all available funds. Then a simple argument based on positivity of cross-entropy shows that the expected growth rate is maximised by betting on each runner in proportion to its probability of winning, regardless of the odds offered.

More realistically the gambler may choose to bet only part of available funds, including the possibility of not betting at all. This leads to a quite different and more complicated problem, and [4] concludes with an argument leading to the log-optimal portfolio in this case. Other examples of non-trivial exact solutions for log-optimal portfolios seem to be rare, but see [1] for a general numerical method.

Here the horse race example is generalised to include heterogeneous combinations of win bets with so-called 'exotic' bet types. Terminology varies from place to place, but in Australia these include the *Exacta* (betting on two horses to finish first and second in specified order) and *Trifecta* (betting on first, second and third in given order). The main result is an exact algorithm determining a log-optimal portfolio based on such bets.

The construction given here in fact applies even more generally, suggesting the possibility of other applications. Let Ω be the set of outcomes in a game of chance, assumed to be disjoint and exhaustive and with given probability distribution. Typically, available bets correspond to certain subsets U_α of Ω. The bet succeeds, returning a specified multiple T_α of the original stake, when the outcome lies in U_α. Otherwise the stake is lost. Here, and elsewhere, the quantity T_α includes the refund of the original stake, so actual profit from a successful unit bet on the event U_α is $T_\alpha - 1$.

For example, in a horse race with N runners offering win and Exacta bets, the relevant set of outcomes is the set of $N(N-1)$ pairs (i,j) giving first and second finishers. The win bet on i is associated with the subset U_i, containing all (i,j) for $j \neq i$.

In this and other cases of interest the available subsets U_α correspond to the vertices α of a rooted tree, with Ω at the root, the singleton subsets of Ω at the leaves and

$$U_\alpha = \bigcup_{\beta \in C(\alpha)} U_\beta\,, \qquad p_\alpha = \sum_{\beta \in C(\alpha)} p_\beta\,,$$

where the union is disjoint, p_α is the probability of event U_α and $C(\alpha)$ is the set of child vertices of α. Thus existence of an edge implies inclusion, and $U_\beta \subseteq U_\alpha$ exactly when vertex β is a descendant of α in the tree.

Given p_α and T_α, the method described below yields the allocation x_α for each vertex α in a log-optimal portfolio. The computation of the x_α begins at the root vertex and moves downwards, involving essentially identical calculations at each vertex of the tree. This permits a compact implementation in any computer language which supports recursion.

Notice that the root vertex always corresponds to the certain event. For games of chance the associated fraction of the portfolio therefore represents the proportion of funds not placed at risk. Typically the return is then unity, since this portion of the portfolio is unaffected by the outcome of the game. Here we allow the possibility of other values, perhaps reflecting payment of interest on retained funds. For any vertex α, unavailability of bets of that type may be signified by taking $T_\alpha = 0$.

For a horse race with 3 runners, the situation considered by Kelly is represented in this way by Figure 1. The root 0 of the tree represents retained funds, and the three child vertices correspond to bets on each of the runners to win. Figure 2 similarly shows the situation when Exacta bets are also available. The leaves now represent the six possibilities for first and second place, the root vertex 0 again corresponds to retained funds and the intermediate vertices to win bets on each of the three runners.

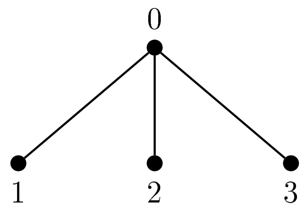

Figure 1. Win

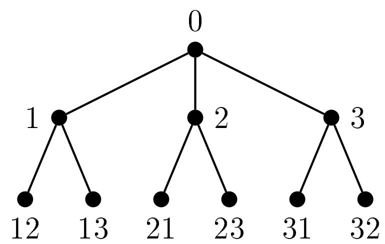

Figure 2. Win/Exacta

Combination of several betting markets in this way clearly increases the chance of finding and exploiting market inefficiencies. Some quantitative evidence for this is given in [3], which demonstrates (at least according to published final odds) the existence of occasional arbitrage opportunities between exotic betting markets in the New South Wales totalisator system.

Not all bet types fit into this tree structure. Exceptions include the *Place* bet on an individual runner (which succeeds if the horse finishes in the first three) and *Quinella* (specifying first and second without regard to order).

After describing the construction of the log-optimal portfolio we consider uniqueness properties of the result, and conclude with examination of some simple situations and illustrative examples. In particular the specialisation to win bets gives a new derivation of Kelly's original formula.

2. Formulation of the Optimisation Problem

We first review the general features of the optimisation problem for the quantity $R(x_1, \ldots, x_N)$ defined by (1). The negative of R is easily seen to be convex in the variables $x_1, \ldots, x_N$ and the set of available portfolios is also convex, consisting of the points of the $(N-1)$-simplex,

$$\sum_{i=1}^{N} x_i = 1 , \quad x_i \geq 0 .$$

General convexity theory [5, 6] gives necessary and sufficient conditions for a maximum of R in the form of a set of equations and inequalities known as the *Kuhn-Tucker conditions.*

In the present case it is easy to derive the Kuhn-Tucker conditions directly. The following is an abbreviated form of the argument given in [2]. Let $\mathbf{x} = (x_1, \ldots, x_N)$ be a log-optimal portfolio, and $\mathbf{y} = (y_1, \ldots, y_N)$ any other portfolio. For $0 \leq t \leq 1$ set

$$\mathbf{z}(t) = (1-t)\mathbf{x} + t\mathbf{y}\,.$$

A straightforward calculation gives

$$\left.\frac{dR(\mathbf{z}(t))}{dt}\right|_{t=0} = E\left(\frac{\sum_i y_i Z_i}{\sum_i x_i Z_i}\right) - 1\,.$$

Since $\mathbf{x}$ is a maximum this derivative must be ≤ 0, so

$$E\left(\frac{\sum_i y_i Z_i}{\sum_i x_i Z_i}\right) \leq 1$$

for all portfolios $\mathbf{y}$. By considering extreme points $y_j = 1$ of the $(N-1)$-simplex we obtain

$$E\left(\frac{Z_j}{\sum_i x_i Z_i}\right) \leq 1 \tag{2}$$

for all $1 \leq j \leq N$. Moreover, if $x_j > 0$ the vector $\mathbf{z}(t)$ can be extended to values $t < 0$ without leaving the simplex. This gives the reverse inequality, and shows that

$$E\left(\frac{Z_j}{\sum_i x_i Z_i}\right) = 1 \tag{3}$$

when $x_j > 0$. Together (2) and (3) make up the Kuhn-Tucker conditions for this problem. General convexity theory implies that the conditions are both necessary and sufficient for a maximum of R. If $-R$ is strictly convex, then the log-optimal portfolio is unique. Otherwise the set of such portfolios forms a convex subset of the $(N-1)$-simplex.

Note that these conditions also imply $\sum_i x_i = 1$. Both (2) and (3) give

$$E\left(\frac{x_j Z_j}{\sum_i x_i Z_i}\right) = x_j$$

when multiplied by x_j, and the sum over $1 \leq j \leq N$ gives

$$\sum_j x_j = E\left(\frac{\sum_j x_j Z_j}{\sum_i x_i Z_i}\right) = E(1) = 1\,,$$

as required. Hence any solution of the Kuhn-Tucker conditions with $x_i \geq 0$ will automatically lie in the $(N-1)$-simplex.

We now specialise to the case of the tree structure described above. We start with an arbitrary finite tree with a distinguished vertex, called the *root*, denoted by the symbol 0. A vertex β is a *descendant* of another vertex α if α lies on the path joining β to the root. We write this relation as $\alpha < \beta$.

Conversely α is an *ancestor* of β in this case. A vertex without descendants is called a *leaf.* In addition we use the following notation and terminology:

$L(\alpha)$**:** the set of leaves β with $\beta \geq \alpha$.
$C(\alpha)$**:** the set of immediate descendants, or *child* vertices of α.
$\pi(\alpha)$**:** the immediate ancestor or *parent* vertex of α.

We identify the set $L(0)$ of all leaves with the points of a probability space, with distribution given by p_β for $\beta \in L(0)$. Then a general vertex α corresponds to the event $L(\alpha)$, with probability of occurrence $p_\alpha = \sum_{\beta \in L(\alpha)} p_\beta$.

Finally we assume that, for each vertex α, a bet on the occurrence of event $L(\alpha)$ is available. We call this a bet of type α, and investigate the properties of a portfolio which allocates fractions $x_\alpha \geq 0$ of available funds to such bets. We regard the value of a unit bet of type α as a random variable Z_α, taking value $T_\alpha \geq 0$ if successful and zero otherwise. We proceed to apply the general arguments above to find a log-optimal portfolio.

At each vertex α of the tree, define

$$X_\alpha = \sum_{0 \leq \beta \leq \alpha} x_\beta T_\beta \,.$$

This quantity is a function of the portfolio, and represents the total value of bets of type β, with $0 \leq \beta \leq \alpha$, when α itself succeeds. If β is a leaf then X_β will be the total value $\sum_\gamma x_\gamma Z_\gamma$ of the portfolio in case of outcome β. But then

$$Z_\alpha = \begin{cases} T_\alpha & \text{if } \beta \in L(\alpha), \\ 0 & \text{otherwise,} \end{cases}$$

so we have

$$E\left(\frac{Z_\alpha}{\sum_\gamma x_\gamma Z_\gamma}\right) = \sum_{\beta \in L(0)} p_\beta \times \frac{Z_\alpha}{X_\beta} = \sum_{\beta \in L(\alpha)} p_\beta \times \frac{T_\alpha}{X_\beta}\,.$$

and the Kuhn-Tucker conditions on the x_α for a log-optimal portfolio take the form

$$\sum_{\beta \in L(\alpha)} \frac{p_\beta T_\alpha}{X_\beta} \quad \begin{cases} = 1 & \text{if } x_\alpha > 0, \\ \leq 1 & \text{if } x_\alpha = 0. \end{cases} \tag{4}$$

Alternatively, when T_α is nonzero,

$$\sum_{\beta \in L(\alpha)} \frac{p_\beta}{X_\beta} \quad \begin{cases} = T_\alpha^{-1} & \text{if } x_\alpha > 0, \\ \leq T_\alpha^{-1} & \text{if } x_\alpha = 0. \end{cases}\,. \tag{5}$$

If $T_\alpha = 0$ then a bet on α is a certain loss and (4) forces $x_\alpha = 0$. Note also that the expected growth rate itself is given by the formula

$$\sum_{\alpha \in L(0)} p_\alpha \ln X_\alpha \,.$$

3. Reformulation of the Kuhn-Tucker Conditions

We seek values $x_\alpha \geq 0$ satisfying the Kuhn-Tucker conditions (5). As already noted the condition $\sum_\alpha x_\alpha = 1$ follows automatically. Since

$$x_\alpha T_\alpha = X_\alpha - X_{\pi(\alpha)}\,,$$

and $x_\alpha = 0$ when $T_\alpha = 0$, it is enough to determine the quantities X_α. Although $\pi(0)$ is undefined, we can extend this formula to $\alpha = 0$ by adopting the convention that $X_{\pi(0)} = 0$.

For each vertex α we define a function $F_\alpha(z)$ of a real variable $z \geq 0$, which plays a key part in the calculation. If α is a leaf then

$$F_\alpha(z) = \frac{p_\alpha}{z}$$

if $z > 0$, and $+\infty$ if $z = 0$. The remaining F_α are defined inductively by the relation

$$F_\alpha(z) = \sum_{\beta \in C(\alpha)} \min\bigl(F_\beta(z), T_\beta^{-1}\bigr)\,.$$

If $T_\beta = 0$ we take $T_\beta^{-1} = +\infty$ and choose the minimum accordingly. Similarly the sum is $+\infty$ if any of the constituent terms are infinite. The following proposition gives some insight into the significance of this definition.

Proposition 1. *The $F_\alpha(z)$ have the following properties:*

(a) *Each $F_\alpha(z)$ is a strictly positive, continuous, decreasing function of z.*
(b) *$F_\alpha(z) \to 0$ as $z \to \infty$.*
(c) *There exists a constant $c_\alpha \geq 0$ such that $F_\alpha(z)$ is constant on the interval $[0, c_\alpha]$ and strictly decreasing for $z \geq c_\alpha$.*
(d) *$F_\alpha(z) \leq p_\alpha/z$ for all α and z.*
(e) *$F_\alpha(0)$ is an expression involving only T_β^{-1} for $\beta > \alpha$.*
(f) *The value $\min(F_\alpha(0), T_\alpha^{-1})$ is the minimum outlay on bets of type $\beta \geq \alpha$ needed to guarantee a return of 1 unit in the event that α succeeds. A value of $+\infty$ indicates that a non-zero return cannot be guaranteed.*

Proof. All the properties follow by induction, starting at the leaves. Clearly (a), (b), (c) hold when α is a leaf, provided we interpret continuity at $z = 0$ as the condition $F_\alpha(0) = +\infty$. For the general case, note that these properties are preserved under finite sums and taking the minimum with a strictly positive constant. Hence if F_β satisfies (a), (b), (c) for all $\beta \in C(\alpha)$ then so will F_α.

If α is a leaf then $F_\alpha(z) = p_\alpha/z$, so (d) is satisfied. For the induction step,

$$F_\alpha(z) \leq \sum_{\beta \in C(\alpha)} F_\beta(z) \leq \sum_{\beta \in C(\alpha)} \frac{p_\beta}{z} = \frac{p_\alpha}{z}\,.$$

Property (e) is again immediate from the definition, since at a leaf $F_\alpha(0) = T_\alpha^{-1}$, and the inductive construction involves only formation of sums and comparison with T_β^{-1} for $\beta > \alpha$.

If α is a leaf then the outlay for unit return from a successful bet on α is T_α^{-1} and

$$\min(F_\alpha(0), T_\alpha^{-1}) = \min(+\infty, T_\alpha^{-1}) = T_\alpha^{-1},$$

giving (f) in this case. For the induction step suppose that (f) holds for all $\beta \in C(\alpha)$. If α succeeds then exactly one of the β will succeed. By induction the minimum wager to return one unit and excluding α is

$$\sum_{\beta \in C(\alpha)} \min(F_\beta(0), T_\beta^{-1}) = F_\alpha(0)$$

Alternatively an allocation of T_α^{-1} to outcome α has the same effect. The overall minimum outlay for unit return is the smaller of the two, or

$$\min(F_\alpha(0), T_\alpha^{-1}).$$

Note that $F_0(0) < 1$ implies the possibility of a guaranteed profit. In gambling terminology this situation is called a *lock* or *Dutch Book*. At the opposite extreme the condition $F_0(0) = +\infty$ indicates the existence of some outcome $\alpha \in L(0)$ with $T_\beta = 0$ for all $0 \leq \beta \leq \alpha$. Then $X_\alpha = 0$ for all portfolio choices, and outcome α will leave the gambler bankrupt. In this case the logarithm, and hence the expected growth rate, is undefined.

$\square$

The next step is to reformulate the Kuhn-Tucker conditions in terms of the $F_\alpha(z)$.

Proposition 2. *The Kuhn-Tucker conditions* (5) *are equivalent to*

$$F_\alpha(X_\alpha) \begin{cases} = T_\alpha^{-1} & \text{if } x_\alpha > 0, \\ \leq T_\alpha^{-1} & \text{if } x_\alpha = 0. \end{cases} \tag{6}$$

For a log-optimal portfolio we also have

$$\sum_{\beta \in L(\alpha)} \frac{p_\beta}{X_\beta} = F_\alpha(X_\alpha) = \min(F_\alpha(X_{\pi(\alpha)}), T_\alpha^{-1}). \tag{7}$$

for all vertices α.

Proof. If α is a leaf vertex then $F_\alpha(X_\alpha) = p_\alpha / X_\alpha$ by definition, so conditions (5) and (6) coincide in this case. Note also that if $F_\alpha(X_\alpha)$ satisfies (6) for a particular vertex α then $x_\alpha = 0$ implies $X_\alpha = X_{\pi(\alpha)}$ and

$$T_\alpha^{-1} \geq F_\alpha(X_\alpha) = F_\alpha(X_{\pi(\alpha)}).$$

Alternatively, if $x_\alpha > 0$ then $X_\alpha \geq X_{\pi(\alpha)}$ and

$$T_\alpha^{-1} = F_\alpha(X_\alpha) \leq F_\alpha(X_{\pi(\alpha)}),$$

since F_α is a decreasing function. In either case

$$F_\alpha(X_\alpha) = \min(F_\alpha(X_{\pi(\alpha)}), T_\alpha^{-1}).$$

The general result now follows by induction. Assume (7) and the equivalence of (5) and (6) for all child vertices $\beta > \alpha$. Then

$$\begin{aligned}\sum_{\beta \in L(\alpha)} \frac{p_\beta}{X_\beta} &= \sum_{\beta \in C(\alpha)} \sum_{\gamma \in L(\beta)} \frac{p_\gamma}{X_\gamma} \\ &= \sum_{\beta \in C(\alpha)} \min(F_\beta(X_\alpha), T_\beta^{-1}) \\ &= F_\alpha(X_\alpha)\,.\end{aligned}$$

This gives the equivalence of (5) and (6) for vertex α. Then (7) for α follows from (6) as above. □

4. Solution of the Kuhn-Tucker Conditions

The solution of the Kuhn-Tucker conditions in the form (6) proceeds inductively. The process starts at the root 0, again using the convention that $X_{\pi(0)} = 0$. Assume then that we seek X_α satisfying (6), and that $X_{\pi(\alpha)}$ has already been determined. If

$$F_\alpha(X_{\pi(\alpha)}) \le T_\alpha^{-1}$$

we can simply take $X_\alpha = X_{\pi(\alpha)}$. Then $x_\alpha = 0$ and (6) is satisfied.

This leaves the case $F_\alpha(X_{\pi(\alpha)}) > T_\alpha^{-1}$. Since $F_\alpha(z)$ is continuous, decreasing and approaches zero as $z \to \infty$, a solution X_α to $F_\alpha(X_\alpha) = T_\alpha^{-1}$ must exist in the range $(X_{\pi(\alpha)}, +\infty)$. Moreover F_α is strictly decreasing when $F_\alpha(z) < F_\alpha(0)$, so the solution must be unique.

More generally, for any value a with $0 < a < F_\alpha(0)$ we can consider the problem of finding the unique z^* for which $F_\alpha(z^*) = a$. One obstacle to finding an exact solution here is the fact that the formula for $F_\alpha(z)$ involves expressions of the form $\min(F_\beta(z), T_\beta^{-1})$, where the choice for the minimum depends on the unknown z. It turns out that we can use F_α itself, and in particular its decreasing property, to avoid the apparent circularity.

The calculation of z^* uses a second induction, starting at the leaves $L(\alpha)$ and associating a pair of real numbers A_β, B_β (depending on a and α) with each vertex $\beta \ge \alpha$. In effect we retrace the definition of F_α, at each stage specifying which of $F_\beta(z)$ and T_β^{-1} will be included in the sum (3). The aim is to transform the equation $F_\alpha(z^*) = a$ into a simple algebraic relation involving z^*.

Starting at a leaf $\beta \in L(\alpha)$, set

$$A_\beta = p_\beta\,, \qquad B_\beta = 0\,.$$

For a general vertex $\beta \ge \alpha$, suppose that A_γ, B_γ have been defined for all child vertices γ of β. The key step is now to identify the subset $C_1(\beta)$ of $C(\beta)$ consisting of those γ for which

$$\frac{A_\gamma}{z^*} + B_\gamma \le T_\gamma^{-1}\,. \tag{8}$$

Since z^* is still unknown we cannot do the comparison directly, at least when $A_\gamma > 0$ and $B_\gamma < T_\gamma^{-1}$. However in this case the condition is equivalent to

$$\frac{A_\gamma T_\gamma}{1 - B_\gamma T_\gamma} \leq z^* .$$

Since $F_\alpha(z)$ is decreasing, and strictly decreasing around z^*, this in turn is the same as

$$F_\alpha\left(\frac{A_\gamma T_\gamma}{1 - B_\gamma T_\gamma}\right) \geq F_\alpha(z^*) = a ,$$

which may be checked directly by evaluation of F_α.

Once $C_1(\beta)$ has been identified in this way, the inductive definition continues with

$$A_\beta = \sum_{\gamma \in C_1(\beta)} A_\gamma ,$$

$$B_\beta = \sum_{\gamma \in C_1(\beta)} B_\gamma + \sum_{\gamma \in C_2(\beta)} T_\gamma^{-1} ,$$

where $C_2(\beta)$ is the complement of $C_1(\beta)$ in $C(\beta)$. Note that $T_\gamma = 0$ implies $\gamma \in C_1(\beta)$, so the A_β and B_β are never required to take infinite values. The significance of this construction is given by the next result.

Proposition 3. *For all $\beta \geq \alpha$ the A_β, B_β satisfy*

$$\frac{A_\beta}{z^*} + B_\beta = F_\beta(z^*) .$$

Proof. If β is a leaf vertex this is immediate. Otherwise assume the result for all children γ of β and apply induction, using

$$\begin{aligned} F_\beta(z^*) &= \sum_{\gamma \in C(\beta)} \min(F_\gamma(z^*), T_\gamma^{-1}) \\ &= \sum_{\gamma \in C(\beta)} \min\left(\frac{A_\gamma}{z^*} + B_\gamma, T_\gamma^{-1}\right) \\ &= \sum_{\gamma \in C_1(\beta)} \left(\frac{A_\gamma}{z^*} + B_\gamma\right) + \sum_{\gamma \in C_2(\beta)} T_\gamma^{-1} \\ &= \frac{A_\beta}{z^*} + B_\beta . \end{aligned}$$

$\square$

In particular

$$\frac{A_\alpha}{z^*} + B_\alpha = F_\alpha(z^*) = a$$

and

$$z^* = \frac{A_\alpha}{a - B_\alpha}$$

is the required solution. Hence we can find X_α satisfying $F_\alpha(X_\alpha) = T_\alpha^{-1}$.

Note that intermediate results in this process can be re-used in the calculation of X_β for $\beta > \alpha$. If $\gamma \geq \alpha$ the preceding result shows that $C_1(\gamma)$ is also the subset consisting of those $\beta \in C(\gamma)$ for which

$$F_\beta(X_\alpha) \leq T_\beta^{-1}\,.$$

Hence we may take $X_\beta = X_\alpha$ for $\beta \in C_1(\alpha)$, with a similar check available on the child vertices of β. Otherwise $\beta \in C_2(\alpha)$ and $F_\beta(X_\alpha) > T_\beta^{-1}$. The same procedure then gives $X_\beta > X_\alpha$ satisfying $F_\beta(X_\beta) = T_\beta^{-1}$.

For a large tree the solution of $F_\alpha(z^*) = a$ involves many evaluations of $F_\alpha(z)$, but most of these are only for the purpose of comparing z and z^*. We can make the process more efficient by storing and updating the values

$$z_1 = \max_{F_\alpha(z) > a} z\,, \qquad z_2 = \min_{F_\alpha(z) < a} z\,.$$

as the calculation proceeds. Then $z_1 < z^* < z_2$ and if the next argument z lies outside the interval (z_1, z_2) the comparison with z^* is available immediately. Otherwise we can use the new z, along with the value of $F_\alpha(z)$, to further shrink the interval.

5. Uniqueness Properties of the Portfolio

The argument of the previous section shows that in most cases the solution to the Kuhn-Tucker conditions, and hence the log-optimal portfolio, is unique. Exceptions occur in the solution of $F_\alpha(X_\alpha) = T_\alpha^{-1}$ when $F_\alpha(0) = T_\alpha^{-1}$. Then we must have $F_\alpha(X_{\pi(\alpha)}) = T_\alpha^{-1}$ also. In the above approach, we take the obvious solution $X_\alpha = X_{\pi(\alpha)}$ when this happens, giving $x_\alpha = 0$. But since $F_\alpha(z)$ is constant on an interval $[0, c_\alpha]$ any choice of X_α in the range

$$X_{\pi(\alpha)} \leq X_\alpha \leq c_\alpha$$

is possible.

For example, if the tree contains just the root vertex and its children then

$$F_0(0) = \sum_{\alpha \in C(0)} \min(F_\alpha(0), T_\alpha^{-1}) = \sum_{\alpha \in C(0)} T_\alpha^{-1}$$

and the ambiguity appears when

$$T_0^{-1} = \sum_{\alpha \in C(0)} T_\alpha^{-1}\,.$$

Then a combination of bets allocating T_α^{-1} to each α in $C(0)$ is exactly equivalent to a single allocation of T_0^{-1} to the root vertex, since both ensure a return of one unit, regardless of outcome.

Another natural way to resolve the ambiguity would be to take X_α as large as possible, so $X_\alpha = c_\alpha$. In contrast to the original method, this yields a portfolio in which funds are concentrated as much as possible towards the root of the tree. This involves additional calculation to find the c_α, which can be arranged inductively as follows.

If α is a leaf then $c_\alpha = 0$. For a general vertex α, define

$$G_\alpha(z) = \min(F_\alpha(z), T_\alpha^{-1})$$

and note that G_α has the same property of being constant on an interval $[0, b_\alpha]$ and strictly decreasing for $z \geq b_\alpha$. If $F_\alpha(0) \leq T_\alpha^{-1}$ then $b_\alpha = c_\alpha$. Otherwise b_α is the unique solution to $F_\alpha(b_\alpha) = T_\alpha^{-1}$, which we can find by the method of the previous section. Since

$$F_\alpha(z) = \sum_{\beta \in C(\alpha)} G_\beta(z)$$

we complete the induction step with the observation that

$$c_\alpha = \min_{\beta \in C(\alpha)} b_\beta \,.$$

6. Examples

In this section we see how various basic examples fit into the present framework. The first illustrates the log-optimal strategy in its most elementary form: the exploitation of a single favourable betting opportunity. We then revisit the example used by Kelly in [4], both with and without the restriction on betting all available funds. The section concludes with a numerical example based on a mixture of win and Exacta bets.

6.1. **Single Bet and Fixed Return.** Suppose a gambler has a choice between a bet with probability p of success, returning T per unit invested, and a safe investment with guaranteed return R. This situation can be formulated in terms of the tree shown in Figure 3, where the labels specify the values $(\alpha, p_\alpha, T_\alpha)$ at each vertex α. Thus $T_0 = R$ and and the two child

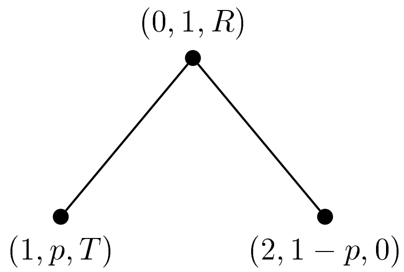

FIGURE 3. Bet vs. fixed return

vertices 1 and 2 represent success and failure of the bet respectively. It is easy to find the log-optimal portfolio by elementary calculus. We know that x_2 must be zero, so there is essentially only one unknown, either x_0 or x_1, and the problem reduces to finding the maximum of a concave function of one variable on the interval $[0, 1]$.

Alternatively, we can argue as above, starting with

$$F_1(z) = \frac{p}{z}\,, \qquad F_2(z) = \frac{1-p}{z}$$

and

$$F_0(z) = \min\left(\frac{p}{z}, T^{-1}\right) + \min\left(\frac{1-p}{z}, +\infty\right) = \min\left(\frac{p}{z}, T^{-1}\right) + \frac{1-p}{z}.$$

Here $F_0(0) = +\infty$ and there is a unique $X_0 > 0$ with $F_0(X_0) = R^{-1}$. Note that the condition $p/X_0 < T^{-1}$ is the same as $pT < X_0$, which in turn is equivalent to $F_0(pT) > F_0(X_0)$. But $F_0(X_0) = R^{-1}$ and

$$F_0(pT) = \min\left(\frac{p}{pT}, T^{-1}\right) + \frac{1-p}{X_0} = \frac{1}{pT}.$$

Hence $p/X_0 < T^{-1}$ if and only if $1/(pT) > R^{-1}$, or $pT < R$. In this case

$$F_0(X_0) = \frac{p}{X_0} + \frac{1-p}{X_0} = \frac{1}{X_0}$$

and the condition $F_0(X_0) = R^{-1}$ gives $X_0 = R$. Since $X_0 = x_0 R$ we then have $x_0 = 1$. This is not surprising: the condition $pT < R$ implies that the risky investment has lower expectation the risk-free alternative.

Alternatively, when $pT \geq R$,

$$F_0(X_0) = T^{-1} + \frac{1-p}{X_0}$$

and the condition $F_0(X_0) = R^{-1}$ gives

$$x_0 = \frac{T - pT}{T - R}, \qquad x_1 = 1 - x_0 = \frac{pT - R}{T - R}.$$

6.2. **Win bets on a horse race.** Suppose that a horse race has N runners, and a gambler may make bets to win on any combination of runners. Figure 4 shows the situation, where p_i is the probability of winning for runner i and T_i the value of a successful unit bet on i. We look at both the case $T_0 = 0$,

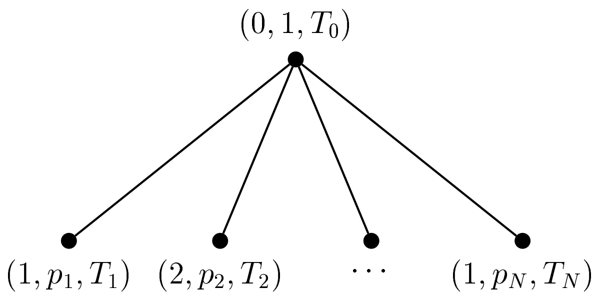

FIGURE 4. Win bets for N runners

where the gambler must bet all available funds, and also $T_0 = 1$, where the gambler may choose to bet only some fraction. In the first case

$$F_i(z) = \frac{p_i}{z}, \qquad F_0(z) = \sum_{i=1}^{N} \min\left(\frac{p_i}{z}, T_i^{-1}\right)$$

and

$$F_0(0) = \sum_{i=1}^{N} T_i^{-1} < T_0^{-1} = +\infty\,.$$

This forces $X_0 = 0$. Also $F_i(X_0) = F_i(0) = +\infty$ so all X_i are non-zero and satisfy $F_i(X_i) = T_i^{-1}$. But this gives

$$\frac{p_i}{x_i T_i} = T_i^{-1}$$

or $x_i = p_i$ for all i. Note that the portfolio does not depend on the returns T_i.

The same result also follows from basic properties of the cross-entropy

$$D\left(\mathbf{p} \;||\; \mathbf{q}\right) = \sum_{i=1}^{N} p_i \ln \frac{p_i}{q_i}$$

of two probability distributions $\mathbf{p} = (p_1, \ldots p_N)$ and $\mathbf{q} = (q_1, \ldots q_N)$. We briefly review the argument, since some points are relevant to the present discussion. The main property of $D\left(\mathbf{p} \;||\; \mathbf{q}\right)$ that we need is that it is always positive, and zero only in case $\mathbf{p} = \mathbf{q}$. See [2] for a proof, as well as further relations between entropy and log-optimal portfolios.

First define

$$Q = \sum_{i=1}^{N} T_i^{-1} \quad \text{and} \quad q_i = T_i^{-1}/Q\,.$$

Then with fraction x_i of total funds allocated to runner i the expected growth rate is

$$\begin{aligned}
\sum_{i=1}^{N} p_i \ln(x_i T_i) &= \sum_{i=1}^{N} p_i \ln\left(\frac{x_i}{q_i} \times Q^{-1}\right) \\
&= \sum_{i=1}^{N} p_i \ln\left(\frac{p_i}{q_i} \times \frac{x_i}{p_i} \times Q^{-1}\right) \\
&= \sum_{i=1}^{N} p_i \ln \frac{p_i}{q_i} - \sum_{i=1}^{N} p_i \ln \frac{p_i}{x_i} + \sum_{i=1}^{N} p_i \ln Q^{-1} \\
&= D\left(\mathbf{p} \;||\; \mathbf{q}\right) - D\left(\mathbf{p} \;||\; \mathbf{x}\right) - \ln Q\,,
\end{aligned}$$

Since $D\left(\mathbf{p} \;||\; \mathbf{x}\right) \geq 0$, the best we can do is to make it zero by taking $\mathbf{x} = \mathbf{p}$, in agreement with our earlier calculation.

In this formula we can regard the quantity $D\left(\mathbf{p} \;||\; \mathbf{q}\right)$ as a measure of the inefficiency of the betting market: unless $D\left(\mathbf{p} \;||\; \mathbf{q}\right) \geq \ln Q$ the gambler will have a negative expected growth rate for any choice of $\mathbf{x}$.

For example, in a *pari-mutuel* or totalisator betting market, the returns T_i are set by distributing some fixed fraction $\rho < 1$ of total receipts among the winning bets. Suppose a total of N_i units is bet on horse i to win, and

$N = \sum_i N_i$ is the total number of units bet. Then the value T_i of a winning unit bet on i is $\rho N/N_i$ and

$$\sum_{i=1}^{N} T_i^{-1} = \sum_{i=1}^{N} \frac{N_i}{\rho N} = \rho^{-1} .$$

In this situation $\ln Q = \ln \rho^{-1} > 0$, quantifying the additional obstacle in the way of positive expected growth rate for this strategy.

In practice the gambler can always choose not to bet, so the choice $T_0 = 1$ gives a more realistic model. In fact any value $T_0 > 0$ leads to a similar calculation, and we now consider this general case. Then

$$F_i(z) = \frac{p_i}{z} \qquad F_0(z) = \sum_{i=1}^{N} \min\left(\frac{p_i}{z}, T_i^{-1}\right) .$$

The condition for $x_0 > 0$ is then

$$F_0(0) = \sum_{i=1}^{N} T_i^{-1} > T_0^{-1} .$$

As noted above $\sum_i T_i^{-1} > 1$ for a typical pari-mutuel market so with $T_0 = 1$ this condition is satisfied. For the calculation of x_0, we start with $A_i = p_i$, $B_i = 0$, so

$$\frac{A_i T_i}{1 - B_i T_i} = p_i T_i$$

and i belongs to $C_2(0)$, implying inclusion of runner i in the portfolio, when $F_0(p_i T_i) < T_0^{-1}$. Then

$$A_0 = \sum_{i \in C_1(0)} p_i = 1 - \sum_{i \in C_2(0)} p_i \, , \quad B_0 = \sum_{i \in C_2(0)} T_i^{-1} .$$

Hence

$$X_0 = x_0 T_0 = \frac{A_0}{T_0^{-1} - B_0} = \frac{1 - \sum_{i \in C_2(0)} p_i}{T_0^{-1} - \sum_{i \in C_2(0)} T_i^{-1}} ,$$

in agreement with Kelly's formula in case $T_0 = 1$. For $i \in C_2(0)$ the allocation x_i is given by solving $F_i(X_i) = T_i^{-1}$. This requires $x_0 T_0 + x_i T_i = p_i T_i$, so

$$x_i = \frac{p_i T_i - x_0 T_0}{T_i} .$$

Here is a numerical example which illustrates the main points of the process. Table 1 shows the probabilities and returns for a race with six runners. The entries in the table are placed in order of decreasing values of the quantity $p_k T_k$, which gives the expected value of a unit bet on k. In this example $T_0 = 1$ and $\sum_{i=1}^{n} T_i^{-1} \approx 1.1$. The final column shows $F_0(p_k T_k)$, so values less than 1 indicate inclusion in the portfolio. The gambler should therefore bet on the first three and ignore the rest of the field. The formula for x_i gives

$$x_0 = 0.86486 \, , \quad x_1 = 0.06351 \, , \quad x_2 = 0.04838 \, , \quad x_3 = 0.02324$$

k	p_k	T_k	p_kT_k	$F_0(p_kT_k)$
1	0.15	10.0000	1.50000	0.66667
2	0.23	4.7619	1.09524	0.87609
3	0.30	3.1250	0.93750	0.97133
4	0.18	4.0000	0.72000	1.07444
5	0.09	7.1429	0.64286	1.09778
6	0.05	12.5000	0.62500	1.10000

Table 1. Optimal Portfolio for Win Bets

and a value for the expected logarithm at this point of approximately 0.016. This represents an expected return of about 1.6% on the gambler's total funds. However, the gambler has only bet 13.5% of total wealth, so the return on funds actually in play is much higher. In addition, a typical weekend may offer several dozen opportunities for betting, so an expected gain of one or two percent per race may still be worthwhile.

Figure 5 shows the behaviour of the function $F_0(z)$ over the range $0 \le z \le 1.5 = p_1T_1$. The solid vertical line shows the solution of the equation $F_0(z) = 1$. The positions of the broken lines show the values of the p_iT_i.

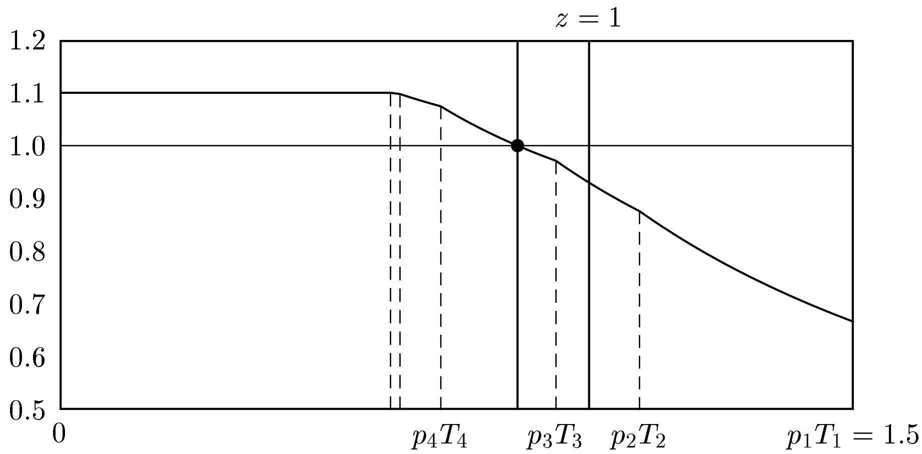

Figure 5. $F_0(z)$ for $0 \le z \le 1.5$

On the other hand, if

$$F_0(0) = \sum_{i=1} NT_i^{-1} \le T_0^{-1}$$

then $x_0 = 0$ and we can again compare $F_i(X_0)$ with T_i^{-1} to see which $x_i > 0$. But now $F_i(X_0) = F_i(0) = +\infty$ so all x_i are included. Then the relation

$$T_i^{-1} = F_i(X_i) = F_i(x_iT_i) = \frac{p_i}{x_iT_i}$$

gives $x_i = p_i$, as for the case $T_0 = 0$.

6.3. **Win and Exacta.** We conclude with another numerical example, this time for a horse race with three runners, where the gambler can bet on each horse to win, or on any pair to finish first and second. Table 2 shows the log-optimal portfolio, along with some intermediate results in the calculation. In this case the expected growth rate is approximately 0.068. This example

α	p_α	T_α	$X_{\pi(\alpha)}$	$F_\alpha(X_{\pi(\alpha)})$	T_α^{-1}	$p_\alpha T_\alpha$	x_α
0	1.00	1.0	0.0000	1.0428	1.0000	1.0000	0.4827
1	0.20	1.5	0.4827	0.3500	0.6666	0.3000	0.0000
2	0.50	2.5	0.4827	0.5333	0.4000	1.2500	0.2069
3	0.30	4.0	0.4827	0.4535	0.2500	1.2000	0.1793
12	0.10	4.0	0.4827	0.2071	0.2500	0.4000	0.0000
13	0.10	7.0	0.4827	0.2071	0.1428	0.7000	0.0310
21	0.30	5.0	1.0000	0.3000	0.2000	1.5000	0.1000
23	0.20	3.0	1.0000	0.2000	0.3333	0.6000	0.0000
31	0.15	3.0	1.2000	0.1250	0.3333	0.4500	0.0000
32	0.15	7.0	1.2000	0.1250	0.1428	1.0500	0.0000

TABLE 2. Optimal portfolio for Win and Exacta

in particular shows that the quantity $p_\alpha T_\alpha$, representing the expected return from a unit bet of type α, is not a useful indicator of inclusion. In fact the condition $p_\alpha T_\alpha > 1$, corresponding to expectation of profit, is neither necessary nor sufficient for inclusion of α in the log-optimal portfolio.

REFERENCES

[1] T.M. Cover. An algorithm for maximising expected log investment return. *IEEE Transactions on Information Theory*, IT-30:369–373, 1984.
[2] T.M. Cover and J.A. Thomas. *Elements of Information Theory.* Wiley Series in Telecommunications. Wiley-Interscience, 1991.
[3] D.C. Edelman and N.R. O'Brian. Tote arbitrage and lock opportunities in racetrack betting. Preprint, 2000.
[4] J.L. Kelly, Jr. A new interpretation of information rate. *Bell Sys. Tech. Journal*, 35:917–926, July 1956.
[5] A.L. Peressini, F.E. Sullivan, and J.J. Uhl, Jr. *The Mathematics of Nonlinear Programming.* Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1988.
[6] Roger Webster. *Convexity.* Oxford University Press, Oxford ; New York, 1994.

SCHOOL OF MATHEMATICS AND STATISTICS, UNIVERSITY OF SYDNEY, SYDNEY, NSW 2006, AUSTRALIA