

MANCHESTER
1824

The University
of Manchester

Economics Discussion Paper

EDP-0508

**Portfolio Management without Probabilities or
Statistics**

by

Sjur Didrik Flåm

April 2005

School of Social Sciences,
The University of Manchester
Oxford Road
Manchester M13 9PL
United Kingdom

Portfolio Management without Probabilities or Statistics

SJUR DIDRIK FLÅM*

April 27, 2005

ABSTRACT. Considered here is on-line financial management aimed at maximizing the long-run growth of wealth. The portfolio is repeatedly rebalanced in response to observed returns on diverse assets. Suppose statistical information and related methods are not available - or deemed either inappropriate or too difficult. On that assumption this paper explores how an adaptive procedure, which totally dispenses with statistics and associated competence, nonetheless may solve the problem over time.

Keywords: log-utility, portfolio selection, growth of wealth, evolutionary stability, replicator dynamics, stochastic approximation.

JEL Classification: C61, C62, G11.

1. INTRODUCTION

Imagine an investor who steadily aims at maximal, long-term growth of his financial savings, but always hesitates - or quite simply remains untrained - in using probabilistic reasoning and statistical methods. Thus, at any stage, he neither accesses nor produces empirical estimates of dividend distributions or associated parameters. In short, although empirically oriented, he never acts like a probabilist or statistician. *Could he, by some reasonable rule of thumb, improve his portfolio stepwise - and, all the same, maximize growth in the long run?*

Assuming a stochastic but stationary environment, this paper provides a constructive, positive and quite simple answer. Not surprisingly, the answer comes in terms of investment behavior that reallocates savings towards assets most recently observed to generate above average relative return. Such reallocation has several noteworthy features, relating to many and fairly diverse strands of literature. A first relation is to stochastic approximation methods [2], [3]; a second goes to evolutionary game theory [15], [21]. Other links point towards dynamical systems, adaptive algorithms, and reinforcement learning [3], [20]. Finally, and most important, it all revolves around Kelly's investment criterion [6], [18]. Proper presence of all these aspects suffice for convergence to optimal portfolio choice in the long run.

To bring this out Section 2 specifies the setting and the adaptive process. Section 3 links the latter to replicator dynamics and evolutionary stability in two-person games. Stability of that sort ensures global convergence, as proven in Section 4.

*Visiting professor, School of Economics Studies, University of Manchester, Oxford Rd. M13 9PL UK, on leave from University of Bergen; sjur.flaam@econ.uib.no. Support from Finansmarkedsfondet is gratefully acknowledged.

2. THE SETTING AND THE INVESTMENT STRATEGY

Suppose wealth grows multiplicatively as follows: At discrete times $t = 0, 1, \dots$ the investor puts the proportion p_t^a of his total holding into asset $a \in A$. That asset gives gross return $X_{t+1}^a \geq 0$ per unit, obtainable at the end of time interval $(t, t + 1]$. Hence his accumulated financial savings S_t develop by the dynamic

$$S_{t+1} := \langle p_t, X_{t+1} \rangle S_t := \left(\sum_{a \in A} p_t^a X_{t+1}^a \right) S_t,$$

where $\langle \cdot, \cdot \rangle$ designates the usual inner product on \mathbb{R}^A . As customary, the latter space consists of all functions $a \in A \mapsto x(a) =: x^a \in \mathbb{R}$. Without loss of generality posit $S_0 = 1$. The asset list A is fixed and finite. The restrictions $p_t^a \geq 0$ and $\sum_{a \in A} p_t^a = 1$, mean that $p_t := (p_t^a)$ always belongs to the standard simplex

$$\mathbb{P} := \left\{ p = (p^a) \in \mathbb{R}_+^A : \sum_{a \in A} p^a = 1 \right\}.$$

By tacit assumption p_t can be, and indeed is, chosen so that $\langle p_t, X_{t+1} \rangle > 0$ almost surely (a.s.).

With apologies for slight abuse of language, since proportions are at center stage, call any $p \in \mathbb{P}$ a *portfolio*. Note that short positions are precluded. Further note that money is never injected or withdrawn. And no transaction cost does ever incur. Then, *what investment strategy will maximize the growth rate of stock?*

For a start, and for simplicity, suppose the return vectors $X_t = (X_t^a) \in \mathbb{R}^A, t = 0, 1, \dots$ are independent, and all distributed as a generic X . Then plainly, one should hold the portfolio $p \in \mathbb{P}$ constant to enjoy stock $S_T = \prod_{t=0}^{T-1} \langle p, X_{t+1} \rangle$ at any future time T regarded worthy of special attention. In such settings, assuming a finite expectation $E \log \langle p, X \rangle$, the *Strong Law of Large Numbers* [4] says that

$$\log(S_T^{1/T}) = \frac{1}{T} \sum_{t=1}^T \log \langle p, X_{t+1} \rangle \rightarrow E \log \langle p, X \rangle \text{ a.s. when } T \rightarrow +\infty.$$

Thus an *unconditioned log-optimal portfolio* p , meaning one which will

$$\text{maximize } E \log \langle p, X \rangle \text{ s.t. } p \in \mathbb{P}, \tag{1}$$

ought appeal to the agent at hand. The optimal value $d := \max_p E \log \langle p, X \rangle$ in (1) describes the dominant behavior of S_T in that $S_T \sim 2^{Td}$ when $\log = \log_2$.

More generally, if the return process X_t displays dependencies over time, the portfolio p_t , chosen at time t , should respond to the history (X_0, \dots, X_t) and other relevant information observed so far. Let all such information be codified by means of a sigma-field \mathcal{F}_t . In essence, what is required is that p_t stays constant on each atom of \mathcal{F}_t .

If the investor worships maximization of the objective $E \log S_T$, T denoting his planning horizon, he would, at any interim time t , be well advised to

$$\text{maximize } E_t \log \langle p_t, X_{t+1} \rangle \quad \text{s.t. } p_t \in \mathbb{P}, \quad (2)$$

using then conditional expectation $E_t := E[\cdot | \mathcal{F}_t]$. Such maximization generates an expected growth rate $d_T := \max E(\log S_T)/T$ over the time span $[0, T]$. Define $d_\infty := \lim_{T \rightarrow \infty} d_T$, provided this limit exists. Employing 2 as base for the logarithm, the relation $S_T \sim 2^{Td_\infty}$ tells that d_∞ qualifies as *doubling rate* - and $1/d_\infty$ as the expected time until doubling of wealth.¹

These observations, and many studies, speak for iterative resolution of problem (2).² But then the investor had better identify the underlying probability distributions - and also be able to compute corresponding expectations. Quite often however, at least one of these two requirements is not met. Knowledge might be imperfect - or expectations too hard to compute. If so, *how can a scantily informed, statistically unskilled investor, while seeking maximal long-run growth of own savings, sequentially adapt his portfolio in response to observed returns?*

To come to grips with this question consider the Kuhn-Tucker conditions of problem (2):

$$\left. \begin{array}{l} E_t [X_{t+1}^a / \langle p_t, X_{t+1} \rangle] = 1 \quad \text{whenever } p_t^a > 0, \\ E_t [X_{t+1}^a / \langle p, X_{t+1} \rangle] \leq 1 \quad \text{otherwise.} \end{array} \right\} \quad (3)$$

Conditions (3) are necessary and sufficient for optimality. They tell that the relative returns $X_{t+1}^a / \langle p_t, X_{t+1} \rangle$ on assets $a \in A$ worthy of investment must, in conditional expectation, all equal the maximal value 1.

Optimality conditions often guide the design of adaptive processes, and so they'll do here as well. For the sake of argument imagine operator E_t were absent in (3). Its hypothetical absence permits a shift of focus towards the *realized*, relative returns $X_{t+1}^a / \langle p, X_{t+1} \rangle$. Trivially, these entities have weighted mean $\sum_{a \in A} p^a X_{t+1}^a / \langle p, X_{t+1} \rangle = 1$. This indicates that a positive position p_t^a in asset a should be reduced iff

$$X_{t+1}^a / \langle p, X_{t+1} \rangle < 1.$$

Specifically, it seems prudent that portfolio p_t be updated at the end of time period $(t, t + 1]$ as follows: For each $a \in A$ let

$$p_{t+1}^a = p_t^a + s_t p_t^a [X_{t+1}^a / \langle p_t, X_{t+1} \rangle - 1]. \quad (4)$$

¹When the return process is stationary, Theorem 15.5.1 in [6] says that

$$d_\infty = \lim_{t \rightarrow \infty} \max \{ E_t \log \langle p_t, X_{t+1} \rangle : p_t \in \mathbb{P} \}.$$

If moreover, returns are independent across time, $d_\infty = d :=$ the optimal value in (1).

²See [6] and references therein. Related material is also found in [5], [7], [9], [10], [14], [16], [17], [18].

Discrete-time, stochastic process (4) becomes the main object below. It differs from the procedures studied in [5], [11] where E and E_t , respectively, operate on $X_{t+1}^a / \langle p_t, X_{t+1} \rangle$, and where $s_t \equiv 1$. In contrast, the sequence $s_t \in (0, 1)$ of *step sizes* is here selected a priori subject to

$$\left. \begin{array}{l} \text{(I)} \quad \sum_t s_t = +\infty, \quad \text{and} \\ \text{(II)} \quad \sum_t s_t^2 < +\infty. \end{array} \right\} \quad (5)$$

In terms of behavior, divergence condition (I) reflects that adaptation remains quite effective from any stage onwards. Mathematically, in terms of *intrinsic time* $\tau_k := s_0 + \dots + s_k$ at stage k , condition (I) embodies that (4) - as a sort of numerical integration - must continue until time $\tau_\infty = \infty$. On the other hand, convergence condition (II) entails that $s_t \rightarrow 0$, thereby reducing the response to randomness over time. Taken together (I) and (II) require that s_t , as a measure of adaptiveness, should dwindle but not too fast.

Note that (4) is *not* a stochastic gradient method [8]. That method would here assume the form

$$p_{t+1} = \mathcal{P} [p_t + s_t X_{t+1} / \langle p_t, X_{t+1} \rangle],$$

\mathcal{P} denoting the orthogonal projection onto \mathbb{P} . The latter operation demands however, some competence and effort [19]. And why should an investor, who is unable to calculate expectations, be perfect in orthogonal projections?

A bonus of (4) is that neither operation is needed. In (4) it is prudent though, to choose initial proportions $p_0^a > 0$ for all a . Any such choice ensures that $p_t^a > 0$ for all t and each a . Thus full diversification always obtains - a feature that appears natural on two accounts: First, if $p_0^a = 0$, then $p_t^a = 0$ for all t to the effect that good investments might possibly be foregone. Second, for reasons not modelled, if direct observation of the entire return vector X_t is deemed desirable at every stage t , then something should always be held of each asset. In any case - provided $s_t \in [0, 1]$, $p_0 \in \mathbb{P}$, and $X_{t+1}^a / \langle p_t, X_{t+1} \rangle \geq 0$ - iteration (4) ensures that $p_t \in \mathbb{P}$ for all t . Further, to maintain p_t in the relative interior of \mathbb{P} it suffices to select p_0 there and pick every $s_t \in (0, 1)$.

3. LOG-OPTIMAL AND EVOLUTIONARY STABLE STRATEGIES

Does (4) converge? To elucidate that question it helps to digress briefly into noncooperative game theory. For the moment suppose *the return process is stationary*. In particular, take all X_t here to be distributed as one and the same X . Assume problem (1) has finite optimal value and at least one optimal solution. Recall that portfolio $\bar{p} \in \mathbb{P}$ is unconditioned log-optimal iff

$$\left. \begin{array}{l} E [X^a / \langle \bar{p}, X \rangle] = 1 \quad \text{whenever } \bar{p}^a > 0, \quad \text{and} \\ E [X^a / \langle \bar{p}, X \rangle] \leq 1 \quad \text{otherwise.} \end{array} \right\} \quad (6)$$

As announced, to appreciate such investment strategies from a different angle, consider now a noncooperative, two-person, symmetric game, featuring finite action space

A. Payoffs in that game are defined as follows: When player 1 and 2 use mixed strategies p and $\bar{p} \in \mathbb{P}$, respectively, the first of them takes home profit

$$\pi(p, \bar{p}) := \sum_{a \in A} p^a \cdot E[X^a / \langle \bar{p}, X \rangle]. \quad (7)$$

As customary, a pair $(\bar{p}, \bar{p}) \in \mathbb{P} \times \mathbb{P}$ is declared a *symmetric Nash equilibrium* iff

$$\pi(p, \bar{p}) \leq \pi(\bar{p}, \bar{p}) \quad \text{for all } p \in \mathbb{P}. \quad (8)$$

Since $\pi(\bar{p}, \bar{p}) = 1$, conditions (6) and (8) are equivalent, and it follows forthwith

Proposition 3.1. (Log-optimality and equilibrium) *A portfolio \bar{p} is unconditioned log-optimal iff it constitutes a Nash equilibrium (\bar{p}, \bar{p}) in the two-person, symmetric game with payoff (7). \square*

Unless a log-optimal \bar{p} selects only *one* asset, which seems rather unlikely, many diverse p maximize (7). So, one can rarely expect an equilibrium strategy \bar{p} to be *strict*. Yet all equilibria enjoy a sort of stability that goes beyond (8) - as explained next. A set $\bar{\mathbb{P}} \subseteq \mathbb{P}$ of equilibria is declared *evolutionary stable* [15] iff for each $\bar{p} \in \bar{\mathbb{P}}$ there is some $\delta > 0$ such that

$$p \in \mathbb{P} \setminus \bar{\mathbb{P}}, \quad \|p - \bar{p}\| \leq \delta \ \& \ \pi(p, \bar{p}) = \pi(\bar{p}, \bar{p}) \quad \Rightarrow \quad \pi(\bar{p}, p) > \pi(p, p).$$

The last inequality says that \bar{p} fares strictly better against any non-optimal, sufficiently close, best response p than does the latter against it self. Typically $\bar{\mathbb{P}}$ is a singleton. Indeed, this happens if the possible return vectors are confined to no single hyperplane [1], [5]. In other words: when X has support of full dimension, \bar{p} becomes unique.

Proposition 3.2. (Log-optimality and evolutionary stability) *The set $\bar{\mathbb{P}}$ of unconditioned log-optimal portfolios is closed convex and evolutionary stable. In fact, it holds for any $p \in \mathbb{P} \setminus \bar{\mathbb{P}}$ that $\pi(\bar{p}, p) > \pi(p, p)$.*

Proof. The function $p \mapsto L(p) := E \log \langle p, X \rangle$ is concave and upper semicontinuous. Therefore the set $\bar{\mathbb{P}} := \arg \max L$ is closed convex. For any non-optimal p and optimal \bar{p} it holds of course that $L(p) < L(\bar{p})$. Further, because of concavity, $L(\bar{p}) \leq L(p) + \langle L'(p), \bar{p} - p \rangle$. Adding the left hand sides and the right hand sides of the last two inequalities we get $\langle L'(p), \bar{p} - p \rangle > 0$. Finally, since $\langle L'(p), \bar{p} - p \rangle = \pi(\bar{p}, p) - \pi(p, p)$, the conclusion follows. \square

Still pursuing the game theoretic perspective, let $\pi(a, p) := E[X^a / \langle p, X \rangle]$ denote the payoff to a player who opts for pure strategy $a \in A$ against the mixed $p \in \mathbb{P}$. For the argument, let us subscribe next to the idea that pure strategies that perform

above average, will be replicated. That is, assume they are at evolutionary advantage is so far as

$$\frac{dp^a}{d\tau} := \dot{p}^a = p^a [\pi(a, p) - \pi(p, p)] \quad (9)$$

for each $a \in A$. We immediately record that system (9) is remarkably stable:

Proposition 3.3. (Asymptotic stability) *System (9) is globally asymptotically stable on the relative interior of the simplex. That is, for any initial p_0 , having all $p_0^a > 0$, it holds that each accumulation point of the resulting trajectory is log-optimal.*

Proof. Recall that the *relative entropy* (alias *Kullback-Leibler distance*)

$$K(\bar{p}, p) := \sum_{a \in A} \bar{p}^a \log \frac{\bar{p}^a}{p^a} = - \sum_{a \in A} \bar{p}^a \{ \log p^a - \log \bar{p}^a \} \quad (10)$$

between any two probability distributions $\bar{p}, p \in \mathbb{P}$ is jointly convex, nonnegative, and vanishes only when $\bar{p} = p$. Continuity considerations lead to the convention that $0 \log 0 = 0$ and $r \log(r/0) = +\infty$ when $r > 0$. Introduce a "distance function"

$$\lambda(p) := \min \{ K(\bar{p}, p) : \bar{p} \text{ log-optimal} \}. \quad (11)$$

Note that $\lambda(p)$ so defined is (convex and) differentiable. Indeed, for each $p \in \mathbb{P}$ there is a unique minimizing $\bar{p}(p) \in \bar{\mathbb{P}}$. Uniqueness derives from $K(\cdot, p)$ being strictly convex for each p . Consequently, by Danskin's envelope theorem, when all $p^a > 0$,

$$\nabla \lambda(p) = \frac{\partial}{\partial p} K(\bar{p}, p)|_{\bar{p}=\bar{p}(p)} = - [\bar{p}^a / p^a] \in \mathbb{R}^A.$$

Thus, on the interior of the simplex, $\lambda(\cdot)$ is a Lyapunov function because along a solution trajectory of (9) it has time derivative

$$\frac{d}{d\tau} \lambda(p(\tau)) = - \sum_{a \in A} \bar{p}^a [\pi(a, p) - \pi(p, p)] = - [\pi(\bar{p}, p) - \pi(p, p)] < 0,$$

\bar{p} denoting the unique minimizer in (11). Asymptotic stability follows from this because $\lambda(p)$, being bounded below, must converge. Its limit value must equal 0, and the conclusion follows. \square

4. CONVERGENCE

Given so much stability of (9), it is tempting next to discretize time in that system, using (tick or) step size s_t . Numerical integration of (9), using Euler's direct method, then takes the form

$$p_{t+1}^a = p_t^a + s_t p_t^a [E [X^a / \langle p_t, X \rangle] - 1].$$

Upon dropping the expectation operator in the last equation, format (4) comes up anew. In short, convergence of (4) revolves around the following question: *Does the stability of continuous-time, deterministic system (9) transfer to its discrete-time, stochastic counterpart (4)?* That question, quite standard in stochastic approximation theory, gets a positive answer here.

It hinges however, on *two* hypotheses stated next. *First*, the return process must have a particular Markovian feature, called the *Robbins-Monro assumption*, namely: *The conditional distribution of X_{t+1} given \mathcal{F}_t depends at most on p_t .* This means, in essence, that the return process X_t be stationary. *Second*, and quite reasonably, assume that

$$\Gamma := \sup \left\{ E \sum_a \{X^a / \langle p, X \rangle - 1\}^2 : p \in \mathbb{P} \right\} < \infty. \quad (12)$$

Clearly, (12) holds if all random variables $X^a / \langle p, X \rangle$ for $a \in A, p \in \mathbb{P}$ are uniformly bounded.

Theorem 3.1. (Global convergence to log-optimality) *Under the Robbins-Monro assumption and (12) each cluster point of iteration (4) is almost surely a log-optimal portfolio. In particular, when the latter is unique, convergence to that portfolio obtains a.s.*

Proof. The demonstration derives from established theory, using the results in [2] or Chapter 5 of [3]. For completeness main arguments are included. But the following auxiliary result is just stated:

Lemma (See [3], Section 5.2.1) *Suppose $\mathcal{A}_t, B_t, C_t, D_t, t = 0, 1, \dots$ are finite-valued, non-negative random variables, all measurable with respect to a sigma-field $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$, which satisfy*

$$E[\mathcal{A}_{t+1} | \mathcal{F}_t] \leq (1 + B_t)\mathcal{A}_t + C_t - D_t. \quad (13)$$

Then, in the event $\{\sum_t B_t < +\infty, \sum_t C_t < +\infty\}$ it holds that

$$\mathcal{A}_t \rightarrow \mathcal{A} < \infty \quad \text{and} \quad \sum_t D_t < +\infty \quad \text{a.s.} \quad (14)$$

As said, in the present case, let \mathcal{F}_t be generated by the return vectors X_0, \dots, X_t and other relevant information unveiled up to time t . Plainly, $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$. Posit

$$\mathcal{A}_t := \min \{ \|p_t - \bar{p}\|^2 / 2 : \bar{p} \text{ log-optimal} \}$$

Note that $\mathcal{A}_t = \|p_t - \bar{p}_t\|^2 / 2 \geq 0$ for a unique \mathcal{F}_t -measurable, log-optimal \bar{p}_t . Also,

$$\begin{aligned} \mathcal{A}_{t+1} &\leq \|p_{t+1} - \bar{p}_t\|^2 / 2 = \|(p_t - \bar{p}_t) + (p_{t+1} - p_t)\|^2 / 2 \\ &\leq \mathcal{A}_t + s_t \sum_{a \in A} (p_t^a - \bar{p}_t^a) [X_{t+1}^a / \langle p_t, X_{t+1} \rangle - 1] + s_t^2 \sum_{a \in A} [X_{t+1}^a / \langle p_t, X_{t+1} \rangle - 1]^2 / 2 \\ &= \mathcal{A}_t + s_t \sum_{a \in A} (p_t^a - \bar{p}_t^a) X_{t+1}^a / \langle p_t, X_{t+1} \rangle + s_t^2 \sum_{a \in A} [X_{t+1}^a / \langle p_t, X_{t+1} \rangle - 1]^2 / 2 \end{aligned}$$

In this string take conditional expectation $E_t := E[\cdot | \mathcal{F}_t]$ to get inequality (13) with

$$B_t := 0, \quad C_t := s_t^2 \Gamma, \quad \text{and } D_t := -s_t \left\{ \sum_{a \in A} (p_t^a - \bar{p}_t^a) E_t [X_{t+1}^a / \langle p_t, X_{t+1} \rangle] \right\} \geq 0.$$

Now, via condition (II) in (5), since the event $\{\sum_t B_t < +\infty, \sum_t C_t < +\infty\}$ carries full probability, (14) follows. Whenever $\mathcal{A} > 0$, condition (I) in (5) implies $\sum D_t = +\infty$. Consequently, $\mathcal{A} = 0$ a.s. and this completes the proof. \square

5. CONCLUDING REMARKS

This paper emphasized three views on portfolio management. One related to how an imperfectly informed, statistically non-trained financial manager might learn growth-optimal investment step by step. Another looked merely for a tractable algorithm to compute a log-optimal portfolio. And finally, one perspective was on how equilibrium might emerge in a symmetric game of finance. Whatever viewpoint or vantage-ground one chooses, eventually to find a good portfolio is likely to take some time and some adaptation - as reflected in process (4).

REFERENCES

- [1] P. H. Algoet and T. M. Cover, Asymptotic optimality and asymptotic equipartition properties of log-optimum investment, *The Annals of Probability* 16, 2, 876-898 (1988).
- [2] M. Benaim, A dynamical system approach to stochastic approximation, *SIAM J. of Control and Optimization* 34, 437-472 (1996).
- [3] A. Benveniste, M. Métivier and P. Priouret, *Adaptive Algorithms and Stochastic Approximations*, Springer, Berlin (1990).
- [4] K. L Chung, *A Course in Probability Theory* (2. ed.), Academic Press, San Diego (1974).
- [5] T. M. Cover, An algorithm for maximizing expected log investment return, *IEEE Trans. Inform. Th.* 30, 2, 369-373 (1984).
- [6] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, J. Wiley, New York (1991).

- [7] M. A. H. Dempster, I. V. Evstigneev and K. R. Schenk-Hoppé, Exponential growth of fixed-mix strategies in stationary asset markets, *Finance and Stochastics* 7, 263-276 (2003).
- [8] Y. Ermoliev, *Methods of Stochastic Programming*, Nauka, Moscow (in Russian) (1976).
- [9] I. V. Evstigneev and K. R. Schenk-Hoppé, From rags to riches: on constant proportions investment strategies, *Inter. J. Theoretical and Applied Finance* 5, 6, 563-573 (2002).
- [10] I. V. Evstigneev and K. R. Schenk-Hoppé, The von Neumann-Gale growth model and its stochastic extensions, to appear in *Handbook of Economic Growth Theory*.
- [11] S. D. Flåm, Pursuing maximal growth of financial wealth, Typescript (2005).
- [12] A. A. Gaivoronski and F. Stella, On-line portfolio selection using stochastic programming, *J. Economic Dynamics and Control* 27, 1013-1043 (2003).
- [13] A. A. Gaivoronski and F. Stella, Stochastic nonstationary optimization for finding universal portfolios, *Annals of Operations Research* 100, 165-188 (2000).
- [14] T. Hens and K. R. Schenk-Hoppé (eds.) *Special Issue on Evolutionary Finance*, *J. Math. Econ.* 1-2, 41 (2005).
- [15] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge University Press (1998).
- [16] G. N. Iyengar and T. M. Cover, Growth optimal investment in horse race markets with costs, *IEEE Trans. Information Theory* 46, 7, 2675-2683 (2000).
- [17] J. B. Long, The numeraire portfolio, *J. Financial Economics* 26, 29-69 (1990).
- [18] D. G. Luenberger, *Investment Science*, Oxford University Press, Oxford (1998).
- [19] C. Michelot, A finite algorithm for finding the projection of a point onto the canonical simplex of \mathbb{R}^n , *J. Optimization Theory and Applications* 50., 1, 195-200 (1986).
- [20] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge MA (1998).
- [21] J. Weibull, *Evolutionary Game Theory*, MIT Press, Cambridge Mass. (1995).