Risk-Adjusted On-line Portfolio Selection

Robert Dochow, Esther Mohr and Günter Schmidt

Abstract The objective of on-line portfolio selection is to design provably good algorithms with respect to some on-line or offline benchmark. Existing algorithms do not consider 'trading risk'. We present a novel risk-adjusted portfolio selection algorithm (RAPS). RAPS incorporates the 'trading risk' in terms of the maximum possible loss. We show that RAPS performs *provably* 'as well as' the Universal Portfolio (UP) [4] in the worst-case. We *empirically* evaluate RAPS on historical NYSE data. Results show that RAPS is able to beat BCRP as well as several 'follow-the-winner' algorithms from the literature, including UP. We conclude that RAPS outperforms in case the assets in the portfolio follow a positive trend.

1 On-line Portfolio Selection

Let P denote the on-line portfolio selection algorithm, and let OPT denote the optimal offline algorithm. The input sequence becomes available to P over time, while OPT knows the whole input sequence in advance. The performance of P is evaluated by means of worst-case competitive analysis without making any statistical assumptions, e.g., on the nature of the stock market. The outcome is the ratio between the value obtained by OPT and P on a worst-case instance, called *regret*.

More formally, on-line portfolio selection aims to determine practical P for investing wealth among a set of m assets ($i = 1, \ldots, m$) over T trading periods ($t = 1, \ldots, T$). The finance community mainly addresses the problem of

R. Dochow (✉)
Saarland University, Saarbrücken, Germany
e-mail: rd@orbi.uni-saarland.de

E. Mohr
University of Mannheim, Mannheim, Germany
e-mail: mohr@bwl.uni-mannheim.de

G. Schmidt
Department of Statistical Sciences, University of Cape Town, Cape Town, South Africa

D. Huisman et al. (eds.), *Operations Research Proceedings 2013*,
Operations Research Proceedings, DOI: 10.1007/978-3-319-07001-8_16,

maximizing the risk-adjusted return, while the information theory and machine learning community aims to maximize the terminal wealth $W_T(P)$ of P. The output of any P is a sequence of allocation vectors $\mathbf{b} = b_1, \ldots, b_T$ for the m assets, with $b_t = (b_{t1}, \ldots, b_{tm})$. The elements b_{ti} represent the proportion of wealth to be invested in the ith asset at the beginning of the tth period ($\sum_{i=1}^{m} b_{ti} = 1$). Let q_{ti} be the price of asset i at time t, and let $\mathbf{x} = x_1, \ldots, x_T$ denote an arbitrary sequence of m-dimensional price relative vectors x_t of the m assets over T. Then the elements of x_t are positive price relatives $x_{ti} = q_{ti}/q_{t-1i}$ of the ith asset at the end of the tth period. In other words, within the tth period the *portfolio return* in-/decreases by the factor $b_t^{\mathrm{T}} x_t = \sum_{i=1}^{m} b_{ti} x_{ti}$. Thus, after T trading periods, the terminal wealth achieved by P equals $W_T(P) = W_o \prod_{t=1}^{T} b_t^{\mathrm{T}} x_t$, where W_o denotes the initial wealth and is set to \$1 for convenience in this work.

In general, any P usually *learns* to compete with a target set of N reference algorithms ($j = 1, \ldots, N$). Let $\mathtt{Q} = \{Q^1, \ldots, Q^N\}$ denote this set. Following the concept of competitive analysis, the performance of P is measured by the *worst-case logarithmic wealth ratio* [3, p. 278]

$$\mathbf{W}_T(P, \mathtt{Q}) = \sup_{\mathbf{x}} \sup_{Q \in \mathtt{Q}} \ln \frac{W_T(Q)}{W_T(P)}, \tag{1}$$

where $\mathtt{Q}$ can be chosen arbitrarily. Most common is the class of constant-rebalanced portfolio (CRP) algorithms, or a mixture of different classes of algorithms.

CRP maintains a constant fraction of the total wealth in each of the underlying m assets. In an *i.i.d.* market if T is large, then *OPT* is the Best CRP (BCRP) [2]. Thus, on-line portfolio selection always chooses a target set $\mathtt{B} = \{B^1, \ldots, B^N\}$ of N CRP reference algorithms, known as 'experts'. If P is compared with *any* possible 'expert' in the simplex domain $\Delta_m = \{b_t : b_t \in \mathbb{R}_+^m, \sum_{i=1}^{m} b_{ti} = 1\}$ then (1) becomes the so-called *regret* [4]

$$r(P) = \mathbf{W}_T(P, \mathtt{B}) = \sup_{\mathbf{x}} \sup_{B \in \Delta_m} \ln \frac{W_T(B)}{W_T(P)}, \tag{2}$$

where $\sup_{B \in \Delta_m} W_T(B) = W_T^*(B)$ is the wealth achieved by BCRP. Note that P outperforms BCRP if $r(P) < 0$.

Further, P is called *universal* if it achieves asymptotically *no regret* on average for T periods and arbitrary bounded $\mathbf{x}$ with respect to BCRP [4, (1.7)]

$$\frac{1}{T} r(P) = \frac{1}{T} \ln W_T(P) - \frac{1}{T} \ln W_T^*(B) \to 0. \tag{3}$$

In the recent years there has been a growing interest and skepticism concerning the value of competitive theory to analyze on-line portfolio selection algorithms. In particular, competitive analysis is inconsistent with the more widely accepted analyses and theories based on statistical assumptions. The main criticisms are: (i) $r(P)$ gives a theoretical upper bound on the loss of P relative to BCRP but omits to analyze its applicability in practice [2], and (ii) existing CRP based algorithms do not consider 'trading risk'. We address both drawbacks.

In short, our risk-adjusted on-line portfolio selection algorithm (RAPS) incorporates the 'trading risk' in terms of the maximum observed fluctuation of the period wealth up to time t. A systematic higher $W_T(P)$ can only be achieved by accepting a higher risk [7], i.e., a higher fluctuation. Addressing (i) we *empirically* evaluate the practical applicability of RAPS on historical NYSE data.[1] Results show that RAPS is able to beat BCRP as well as several known 'follow-the-winner' algorithms, including UP of [4]. Addressing (ii) we show that RAPS performs *provably* 'as well as' UP in the worst-case.

The rest of the paper is organized as follows. In the next section we give the necessary theoretical background. We formally present and analyze RAPS. Section 3 shows the benefits of RAPS on a numerical example. Section 4 concludes.

2 Algorithm RAPS

Without making any statistical assumptions on the nature of the stock market, [4] proves that certain P are *universal*. Cover's algorithm, UP, achieves asymptotically *no regret*.

UP: The idea is to start with the Uniform CRP (UCRP) in period $t = 1$, i.e., $b_1 = (\frac{1}{m}, \ldots, \frac{1}{m})$. For $t \geq 2$, the $\mathbf{b}$ is approximated by the past performance of the N 'experts' [4, (1.3), p. 2]

$$\hat{b}_{t+1i} = \frac{\sum_{j=1}^{N} b_{ti}^{j} \cdot W_t(B^j)}{\sum_{j=1}^{N} W_t(B^j)}, \tag{4}$$

where $W_t(B^j) = W_{t-1}(B^j) \cdot b_t^{\mathrm{T}} x_t$ denotes the *compound period wealth* of the jth 'expert' in the tth period. Thus, in hindsight, $\hat{b}_{t+1}$ is the weighted average over all 'experts' in target set B [4, p. 3].

Lemma 1 *Assume that UP competes against target set* B. *UP divides* W_o *in* N *equal parts and invests according to* B^j $(j = 1, \ldots, N)$. *Then the* terminal wealth *of UP equals* $W_T(UP) = \frac{1}{N}\sum_{j=1}^{N} W_T(B^j)$, *and its* worst-case logarithmic wealth ratio *is bounded as* [*cf.* (1)] [3, Example 10.3, p. 278]

$$\mathbf{W}_T(UP, \mathrm{B}) = \sup_{\boldsymbol{x}} \sup_{B^j \in B} \ln \frac{W_T(B^j)}{W_T(UP)} \leq \ln N. \tag{5}$$

Lemma 2 *If* μ *is the uniform density on* Δ_m, *then UP of* [4] *satisfies* [*cf.* (2)] [3, Theorem 10.3, p. 283]

[1] http://www.cs.bme.hu/~oti/portfolio/data/nyseold.zip

$$r(UP) = \sup_{x} \sup_{B \in \Delta_m} \ln \frac{W_T(B)}{W_T(UP)} \leq (m-1)\ln(T+1). \tag{6}$$

UP exploits the 'follow-the-winner' principle, and performs *provably* 'almost as well' as BCRP [4, Theorem 7.1].

RAPS: Let $m_t^j = \min\{W_o, \ldots, W_t(B^j)\}$ and $M_t^j = \max\{W_o, \ldots, W_t(B^j)\}$ be the minimum and maximum compound period wealth of the jth 'expert' up to time t. Then $\phi_t(B^j) = \frac{M_t^j}{m_t^j}$ equals the maximum observed fluctuation of the period wealth up to time t, and the inverse $\phi_t(B^j)^{-1}$ quantifies an experts' possible maximum *loss* up to time t. To compute $\hat{b}_{t+1}$, UP uses the experts compound period wealth. Instead, the idea of RAPS is to replace the $W_t(B^j)$ in (4) by $\phi_t(B^j)^{-1}$. Like UP, RAPS starts with UCRP in $t = 1$. For the subsequent $t \geq 2$ periods, **b** is approximated by

$$\hat{b}_{t+1i} = \frac{\sum_{j=1}^{N} b_{ti}^j \cdot \phi_t(B^j)^{-1}}{\sum_{j=1}^{N} \phi_t(B^j)^{-1}}. \tag{7}$$

Lemma 3 *Assume that all* $x_{ti} \leq 1$*, and that RAPS competes against a target set of* $\mathbb{B}$ *algorithms. RAPS divides* W_o *in* N *equal parts and invests according to* B^j*. Then the* terminal wealth *of RAPS equals* $W_T(RAPS) = \frac{1}{N}\sum_{j=1}^{N} \phi_t(B^j)^{-1}$*, and its* worst-case logarithmic wealth ratio *is bounded as* [*cf.* (1)]

$$\boldsymbol{W}_T(RAPS, \mathbb{B}) \leq \ln N. \tag{8}$$

Proof The proof is based on Lemma 1. We know that *iff* the assets in the portfolio do not follow a positive trend, then $x_{ti} \leq 1$. It follows $\phi_t(B^j)^{-1} = m_t^j = W_t(B^j)$ for $t = 1, \ldots, T$ and $j = 1, \ldots, N$. Thus

$$\begin{aligned}
\mathbf{W}_T(RAPS, \mathbb{B}) = \sup_{\mathbf{x}} \ln \frac{\max_{j=1,\ldots,N} \phi_T(B^j)^{-1}}{\frac{1}{N}\sum_{j=1}^{N} \phi_T(B^j)^{-1}} &\leq \sup_{\mathbf{x}} \ln \frac{\max_{j=1,\ldots,N} \phi_T(B^j)^{-1}}{\frac{1}{N}\max_{j=1,\ldots,N} \phi_T(B^j)^{-1}} \\
&= \sup_{\mathbf{x}} \ln \frac{\max_{j=1,\ldots,N} W_T(B^j)}{\frac{1}{N}\max_{j=1,\ldots,N} W_T(B^j)} \\
&= \ln N. \qquad \square
\end{aligned} \tag{9}$$

Under worst-case assumptions $\mathbf{W}_T(RAPS, \mathbb{B}) = \mathbf{W}_T(UP, \mathbb{B})$, cf. (5). Consequently, $r(RAPS) \leq (m-1)\ln(T+1)$ also equals UP, cf. (6). The worst-case performance of UP is basically unimprovable [3, p. 285], but UP has some practical disadvantages which are addressed by [5, 6]. We aim to answer the question if RAPS is able to outperform (some of) these algorithms from the literature in case $x_{ti} > 1$.

On-line Benchmarks: Motivated by the 'follow-the-winner' principle we limit to P which increase the b_{ti} of more successful assets. Rather than targeting BH_{best},

Table 1 Portfolio comparison in terms of the $W_T(P)$ achieved for $N = 101$

#	Assets	BCRP	BH$_{best}$	UP	EG (0.01)	UCRP	SCRP	RAPS	$r(RAPS)$
1	Comm. Metals and Kin Arc	**144.01**	52.02	78.47	117.15	118.69	26.36	127.96	+0.12
2	IBM and Coca Cola	15.07	13.36	14.18	15.00	15.02	5.48	**15.36**	−0.02
3	Comm. Metals and Mei Corp.	102.96	52.02	72.63	97.94	98.89	28.14	**109.57**	−0.06
4	$\bar{W}_T(P) = \frac{1}{630}\sum_{p=1}^{630} W_T(P)$	**26.57**	20.72	18.89	21.73	21.84	12.13	23.07	+0.14

these algorithms mainly track BCRP. Besides (i) UP and (ii) UCRP, we consider: (iii) Exponential Gradient (EG(η)) of [6] which aims to reduce the computational costs of UP from exponential to linear. The key parameter of EG(η) is the learning rate $\eta > 0$. For $\eta \to 0$ EG(η) reduces to UCRP [6, p. 35:11]. (iv) Successive CRP (SCRP) of [5, p. 170] which directly adopts BCRP up to the tth period, i.e., b_{t+1} equals the subsequent BCRP allocation vector (b_t^*). Note that, compared to UP and RAPS, the worst-case performance guarantees of EG(η) and SCRP are inferior (not as tight).

Offline Benchmarks: In the financial community the optimal offline benchmark is to buy-and-hold the best-performing asset of the portfolio, denoted by BH_{best} [2]. In contrast, the information theory and machine learning community considers BCRP. [4, Proposition 2.1] proved that BCRP exceeds BH_{best}. Obviously, BH_{best} and BCRP can only be computed in hindsight.

3 Numerical Results

The NYSE data set includes daily closing prices of 36 assets for 22 years ($T = 5,651$). We only consider portfolios containing $m = 2$ assets, resulting in $\binom{36}{2} = 630$ possible portfolio combinations, and limit to three pairs of assets, cf. Table 1. We selected these pairs in order to make our results comparable, cf. [4, p. 23], [6, p. 340], and [5, p. 181]. Portfolios #1 and #2 can be found in [4–6], and #3 in [4, 6]. In addition, #4 gives the average $\bar{W}_T(P)$; column $r(RAPS)$ indicates whether RAPS outperforms BCRP (< 0) or not (> 0).

#1: From [4, p. 26] we know that the outperformance of BCRP is due to the leverage effect in the posteriori computed BCRP. Thus, Cover compared UP to a randomly generated portfolio (98.4). Contrary to UP, RAPS clearly outperforms the random sample ($127.96 > 98.4$), and all P.

#2: The assets show a lockstep performance (12.21 and 13.36). Though, like UP, RAPS barely outperforms them (cf. [4, p. 23]). Further, RAPS outperforms BCRP and all P.

#3: Volatile uncorrelated assets (52.02 and 22.92) lead to great gains compared to BH_{best}. This also holds for #1 (52.02 and 4.13). RAPS clearly outperforms BCRP and all P.

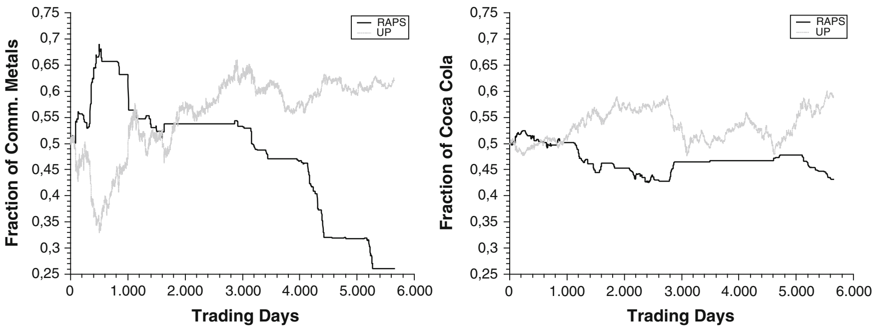

Fig. 1 Proportion of wealth ($\hat{b}_{ti}$) RAPS and UP invested in BH_{best}: #1 (*left*) and #2 (*right*)

#4: We run experiments on all 630 portfolio combinations. For each of the 36 assets $\bar{x}_{ti} > 1$ holds, where $\bar{x}_{ti} = \frac{1}{T}\sum_{t=1}^{T} x_{ti}$. On average, RAPS outperforms all P and BH_{best} but not BCRP. We claim that RAPS outperforms the online benchmarks in case the assets in the portfolio follow a *positive trend*, i.e., $\bar{x}_{ti} > 1 \; \forall \; m$ assets.

Summing up, RAPS outperforms BH_{best} and all P in all cases, and is superior to BCRP in two of three cases. Hence, Fig. 1 shows that targeting BCRP is superior to targeting BH_{best} (Comm. Metals (#1; #3) and Coca Cola (#2)) as RAPS stepwise reduces the $\hat{b}_{ti}$ invested in BH_{best}.

4 Conclusions

To the best of our knowledge, existing 'follow-the-winner' algorithms do not consider 'trading risk' when computing **b**. In contrast to existing P, RAPS targets the expert with the lowest possible loss $(\phi_t(B^j)^{-1})$. We *prove* that RAPS performs 'as well as' UP in the worst-case, and its computational costs are also exponential. Our numerical results (Table 1) are encouraging that RAPS performs well in practice. The constituent assets and all benchmark algorithms from the literature (UP, EG(0.01), UCRP, SCRP) are outperformed. In general, RAPS outperforms in case the assets in the portfolio follow a *positive trend*. Volatile uncorrelated stocks (like in #1 and #3) lead to great gains over BH_{best}. Figure 1 shows that targeting BCRP is superior to targeting BH_{best}. However, ponderous stocks (like in #2) show only modest improvements. This result is consistent with [4]. An open question is the universality of RAPS.

References

1. Borodin, A., El-Yaniv, R., & Gogan, V. (2000). On the competitive theory and practice of portfolio selection. In G. Gonnet & A. Viola (Eds.), *LATIN 2000: Theoretical informatics*. Springer, Berlin.
2. Borodin, A., El-Yaniv, R., & Gogan, V. (2004). Can we learn to beat the best stock. *Journal of Artificial Intelligence Research*, *21*(1), 579–594.

3. Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, learning, and games*. New York: Cambridge University Press.
4. Cover, T. (1991). Universal portfolios. *Mathematical Finance*, *1*(1), 1–29.
5. Gaivoronski, A. A., & Stella, F. (2000). Stochastic nonstationary optimization for finding universal portfolios. *Annals of Operations Research*, *100*, 165–188.
6. Helmbold, D., Schapire, R., Singer, Y., & Warmuth, M. (1998). On-line portfolio selection using multiplicative updates. *Mathematical Finance*, *8*(4), 325–347.
7. Modigliani, F., & Modigliani, L. (1997). Risk-adjusted performance. *Journal of Portfolio Management*, *23*(2), 45–54.