On Tracking Portfolios with Certainty Equivalents on a Generalization of Markowitz Model: the Fool, the Wise and the Adaptive

Richard Nock RNOCK@MARTINIQUE.UNIV-AG.FR
Brice Magdalou BRICE.MAGDALOU@MARTINIQUE.UNIV-AG.FR
Eric Briys ERIC.B@CYBERLIBRIS.COM
CEREGMIA — Univ. Antilles-Guyane, B.P. 7209, 97275 Schoelcher Cedex, Martinique, France

Frank Nielsen NIELSEN@CSL.SONY.CO.JP
Sony Computer Science Laboratories, Inc., 3-14-13 Higashi Gotanda. Shinagawa-Ku. Tokyo 141-0022, Japan

Abstract

Portfolio allocation theory has been heavily influenced by a major contribution of Harry Markowitz in the early fifties: the mean-variance approach. While there has been a continuous line of works in on-line learning portfolios over the past decades, very few works have really tried to cope with Markowitz model. A major drawback of the mean-variance approach is that it is approximation-free only when stock returns obey a Gaussian distribution, an assumption known not to hold in real data. In this paper, we first alleviate this assumption, and rigorously lift the mean-variance model to a more general mean-*divergence* model in which stock returns are allowed to obey any exponential family of distributions. We then devise a general on-line learning algorithm in this setting. We prove for this algorithm the first lower bounds on the most relevant quantity to be optimized in the framework of Markowitz model: the *certainty equivalents*. Experiments on four real-world stock markets display its ability to track portfolios whose cumulated returns exceed those of the best stock by orders of magnitude.

1. Introduction

In PUDD'NHEAD WILSON, Mark Twain once quoted the wise man: "Put all your eggs in the one basket and — *watch that basket!*", against the fool who argues to rather scatter money (and attention). The large majority of works on on-line learning portfolios *watch* portfolios using their expected returns (Even-Dar et al., 2006). Very few works have started to look at the problem with a refined lens, relying on *risk premiums* instead of returns (Warmuth & Kuzmin, 2006), inspired by a theory born more than fifty years ago (Markowitz, 1952). As Markowitz has shown, investors know that they cannot achieve stock returns greater than the risk-free rate without having to carry some risk. The famed mean-variance approach was born, in which the variance term models the investor's aversion to risk. Under the assumptions that the investor obeys exponential utility and the stocks returns have Gaussian distribution, the optimal portfolio is that which maximizes the *difference* between expected returns and half the variance times the Arrow-Pratt risk aversion parameter (Pratt, 1964). This latter term in the difference quantifies the *risk premium* of the portfolio, while the difference — hence, the quantity which completely defines the optimal portfolio — is the *certainty equivalent*.

There are prominent limitations to both the model and the previous approaches that learn portfolios on-line. First, it is a well-known observation that empirical data do not obey Gaussian distribution, thus impairing the safe application of Markowitz' model to real domains. Second, all previous attempts to cast on-line learning in this model relied on approximations of the actual quantity to be maximized, the certainty equivalent (Even-Dar et al., 2006; Warmuth & Kuzmin, 2006).

In this paper, we alleviate these two limitations. We first replace the Gaussian distribution assumption about returns by the more realistic assumption that they obey general exponential families: we prove that

Appearing in *Proceedings of the 28th International Conference on Machine Learning*, Bellevue, WA, USA, 2011.

the mean-variance approach of Markowitz is generalized by a mean-*divergence* model, in which the divergence part heavily relies on a class of distortion popular in machine learning: Bregman divergences (Banerjee et al., 2005; Nock & Nielsen, 2009). We then provide, in this mean-divergence portfolio choice model, a general algorithm for on-line learning reference portfolios that are allowed to drift, based on a generalization of Amari's famed natural gradient (Amari, 1998). We show a lower bound on the certainty equivalent of this algorithm which depends on the certainty equivalents of the reference portfolios. No such bound was previously known, even in the restricted case of Markowitz' model. Our contribution is also experimental, as we provide results on four major stock markets (DJIA, NYSE, S&P500, TSE) that display (i) the interest in lifting the mean-variance to the mean-divergence model, as the mean-variance model appears to be suboptimal, (ii) the performance of the algorithm on real data, with its ability to adapt its allocation and simultaneously beat by orders of magnitude market contenders from both the "fool" and the "wise" families in Twain's acception (resp. uniform cost rebalanced portfolio and best stock).

The remaining of the paper is organized as follows: Section 2 presents the mean-divergence model. Section 3 presents our algorithm and its properties. Section 4 details the experiments, and Section 5 concludes.

Notations Italicized bold letters like $\boldsymbol{v}$ denote vectors and v_i their coordinates. Blackboard notations like $\mathbb{S}$ denote subsets of (tuples of) reals, and $|.|$ their cardinal. Calligraphic letters like $\mathcal{A}$ are reserved for algorithms. Economic concepts are distinguished with small capitals: for example, the certainty equivalent is denoted $\textsc{c}$, and utility functions are denoted $\textsc{u}$. We define $\mathbf{0}$, the null vector, $\mathbf{1}$, the all-1 vector and $\mathbf{1}_j$ the vector with "1" in coordinate j and zero elsewhere. Because of size constraints, parts of the technical and experimental material of this paper are available in a supplementary material file[1].

2. The mean-divergence model

We consider an (investor, market) pair setting, in which the investor is characterized by a vector $\boldsymbol{\alpha} \in \mathbb{P}$, a portfolio allocation vector over d assets, where $\mathbb{P}$ denotes the d-dimensional probability simplex. These d assets characterize the market, on which we compute a vector of returns $\boldsymbol{w} \in [-1, +\infty)^d$. Quantity

$$\omega_{\text{inv}} \doteq \boldsymbol{w}^\top \boldsymbol{\alpha} \quad (1)$$

models the investor's *wealth* brought by his/her portfolio. We assume that $\boldsymbol{w}$ is drawn at random from some density p_ψ which belongs to the exponential families of distributions (Banerjee et al., 2005):

$$\begin{aligned} p_\psi(\boldsymbol{w} : \boldsymbol{\theta}) &\doteq \exp\left(\boldsymbol{w}^\top \boldsymbol{\theta} - \psi(\boldsymbol{\theta})\right) b(\boldsymbol{w}) \ , \quad (2) \\ &= \exp\left(-D_{\psi^\star}(\boldsymbol{w} \| \boldsymbol{\nabla}_\psi(\boldsymbol{\theta})) + \psi^\star(\boldsymbol{w})\right) b(\boldsymbol{w}) \ , \end{aligned}$$

where $\boldsymbol{\theta}$ defines the natural parameter of the family, and $b(.)$ normalizes the density. $\psi : \mathbb{S} \rightarrow \mathbb{R}$ ($\mathbb{S} \subseteq \mathbb{R}^d$) is strictly convex differentiable, and $\psi^\star$ is its *convex conjugate*, defined as $\psi^\star(\boldsymbol{z}) \doteq \sup_{\boldsymbol{t} \in \text{dom}(\psi)} \{\boldsymbol{z}^\top \boldsymbol{t} - \psi(\boldsymbol{t})\} = \boldsymbol{z}^\top \boldsymbol{\nabla}_\psi^{-1}(\boldsymbol{z}) - \psi(\boldsymbol{\nabla}_\psi^{-1}(\boldsymbol{z}))$ (Banerjee et al., 2005). We define the *Bregman divergence* D_ψ with *generator* ψ as (Banerjee et al., 2005):

$$D_\psi(\boldsymbol{x} \| \boldsymbol{y}) \doteq \psi(\boldsymbol{x}) - \psi(\boldsymbol{y}) - (\boldsymbol{x} - \boldsymbol{y})^\top \boldsymbol{\nabla}_\psi(\boldsymbol{y}) \ , \quad (3)$$

where $\boldsymbol{\nabla}_\psi$ denotes the gradient of ψ.

It is not hard to show that the gradients of ψ and $\psi^\star$ are inverse of each other ($\boldsymbol{\nabla}_\psi = \boldsymbol{\nabla}_{\psi^\star}^{-1}$), and furthermore the fundamental relationship holds:

$$D_\psi(\boldsymbol{x} \| \boldsymbol{y}) = D_{\psi^\star}(\boldsymbol{\nabla}_\psi(\boldsymbol{y}) \| \boldsymbol{\nabla}_\psi(\boldsymbol{x})) \ . \quad (4)$$

Exponential families contain popular members, such as the Gaussian, exponential, Poisson, multinomial, beta, gamma, Rayleigh distributions, and many others.

A quite counterintuitive observation about the investor is that he/she would typically *not* choose $\boldsymbol{\alpha}$ based on the maximization of the expected returns. This is the famed St. Petersburg paradox, which states that the expected return alone lacks crucial informations about the way $\boldsymbol{\alpha}$ is chosen, such as investor's being not unconscious to the fact that investments cannot be achieved without carrying out some *risk* (Chavas, 2004). A popular normative approach alleviates this paradox (von Neumann & Morgenstern, 1944): five assumptions about the way people build preferences among allocation vectors are enough to show that portfolios are ordered based on an expected *utility* of returns, $\mathrm{E}_{\boldsymbol{w} \sim p_\psi}[\textsc{u}(\boldsymbol{w}^\top \boldsymbol{\alpha})]$, where $\textsc{u}(.)$ denotes a real-valued *utility* function. It can be shown that this expectation, which is computed over numerous markets, equals the utility of a *single* equivalent case ("sure market") in which the expected wealth is minored by a *risk premium* (Chavas, 2004):

$$\mathrm{E}_{\boldsymbol{w} \sim p_\psi}[\textsc{u}(\omega_{\text{inv}})] = \textsc{u}\left(\mathrm{E}_{\boldsymbol{w} \sim p_\psi}[\omega_{\text{inv}}] - \textsc{p}(\boldsymbol{\alpha}; \boldsymbol{\theta})\right) \quad (5)$$

Because this case represents a sure money-metric equivalent of the left-hand side's numerous markets, the quantity $\textsc{c}(\boldsymbol{\alpha}; \boldsymbol{\theta}) \doteq \mathrm{E}_{\boldsymbol{w} \sim p_\psi}[\omega_{\text{inv}}] - \textsc{p}(\boldsymbol{\alpha}; \boldsymbol{\theta})$ is called the *certainty equivalent*. Markowitz has shown that the certainty equivalent may be derived exactly when

[1] http://www1.univ-ag.fr/~rnock/Articles/ICML11/

Table 1. Bregman divergences used in this paper; $\|\boldsymbol{x}\|_q \doteq (\sum_i |x_i|^q)^{1/q}$ denotes the q-norm.

$\varphi(\boldsymbol{x})$	$D_\varphi(\boldsymbol{x}\|\boldsymbol{y})$	Comments
$\frac{1}{2}\|\boldsymbol{x}\|_q^2$	$\frac{1}{2}\|\boldsymbol{x}\|_q^2 - \frac{1}{2}\|\boldsymbol{y}\|_q^2 - (\boldsymbol{x}-\boldsymbol{y})^\top \boldsymbol{\nabla}_\varphi(\boldsymbol{y})$	q-norm divergence, D_{L_q}; $(\nabla_\varphi(\boldsymbol{y}))_i = \frac{\mathrm{sign}(y_i)\|y_i\|^{q-1}}{\|\boldsymbol{y}\|_q^{q-2}}$
$\sum_i x_i \ln x_i - x_i$	$\sum_i (x_i \ln(x_i/y_i) - (x_i - y_i))$	Kullback-Leibler divergence, D_{KL}
$-\sum_i \ln x_i$	$\sum_i ((x_i/y_i) - \ln(x_i/y_i) - 1)$	Itakura-Saito divergence, D_{IS}
$\sum_i \exp x_i$	$\sum_i (\exp(x_i) - (x_i - y_i + 1)\exp(y_i))$	Exponential divergence, D_{EXP}

p_ψ is Gaussian. Applying the mean-variance model in the general case without caring for the Gaussian assumption incurs an approximation to the premium part in (5) which can be devastating (Chavas, 2004).

To summarize, alleviating the Gaussian assumption implies to find U, P and C with which (5) holds under the more general setting of exponential families. Finding U is in fact easy even when $d > 1$. We rely on Arrow-Pratt measure of absolute risk aversion (Chavas, 2004; Pratt, 1964), which can be computed for each stock as:

$$\mathrm{R}_i(\omega_{\mathrm{inv}}) \doteq -\frac{\partial^2}{\partial w_i^2}\mathrm{U}(\omega_{\mathrm{inv}})\left(\frac{\partial}{\partial w_i}\mathrm{U}(\omega_{\mathrm{inv}})\right)^{-1}, \forall i = 1,2,...,d \ .$$

We say that there is *constant absolute risk aversion* (CARA) whenever $\mathrm{R}_i(\omega_{\mathrm{inv}}) = a, \forall i = 1,2,...,d$, for some risk aversion parameter $a \in \mathbb{R}$. The following Lemma easily follows from (Chavas, 2004).

Lemma 1 *$\mathrm{R}(\omega_{inv}) = a$ for some $a \in \mathbb{R}$ iff $\mathrm{U}(x) = x$ (if $a = 0$) or $\mathrm{U}(x) = -\exp(-ax)$ (otherwise).*

Assuming that the investor is risk averse, we have $a > 0$. We can now provide the expressions of $\mathrm{C}(\boldsymbol{\alpha};\boldsymbol{\theta})$ and $\mathrm{P}(\boldsymbol{\alpha};\boldsymbol{\theta})$, which we now rename $\mathrm{C}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta})$ and $\mathrm{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta})$, since they depend on ψ, the *premium generator*.

Theorem 1 *Assume CARA and p_ψ as in (2). Then:*

$$\mathrm{C}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = \frac{1}{a}\left(\psi(\boldsymbol{\theta}) - \psi(\boldsymbol{\theta} - a\boldsymbol{\alpha})\right) \ , \quad (6)$$

$$\mathrm{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = \frac{1}{a}D_\psi\left(\boldsymbol{\theta} - a\boldsymbol{\alpha}\|\boldsymbol{\theta}\right) \ . \quad (7)$$

Proof: We have:

$$\begin{aligned}\mathrm{E}_{\boldsymbol{w}\sim p_\psi}[\mathrm{U}(\omega_{\mathrm{inv}})] &= \int -\exp\left(\boldsymbol{w}^\top(\boldsymbol{\theta} - a\boldsymbol{\alpha}) - \psi(\boldsymbol{\theta})\right) b(\boldsymbol{w})\mathrm{d}\boldsymbol{w} \\ &= -\exp\left(\psi(\boldsymbol{\theta} - a\boldsymbol{\alpha}) - \psi(\boldsymbol{\theta})\right) \times \\ &\quad \underbrace{\int \exp\left(\boldsymbol{w}^\top(\boldsymbol{\theta} - a\boldsymbol{\alpha}) - \psi(\boldsymbol{\theta} - a\boldsymbol{\alpha})\right) b(\boldsymbol{w})\mathrm{d}\boldsymbol{w}}_{=1} \\ &= -\exp\left(-a\mathrm{C}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta})\right) \ , \quad (8)\end{aligned}$$

where we have used in (8) Lemma 1 and (5). The definition of the certainty equivalent yields $\mathrm{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = \mathrm{E}_{\boldsymbol{w}\sim p_\psi}[\boldsymbol{w}^\top\boldsymbol{\alpha}] - \mathrm{C}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = \boldsymbol{\alpha}^\top\boldsymbol{\nabla}_\psi(\boldsymbol{\theta}) - \mathrm{C}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = \frac{1}{a}\left(\psi(\boldsymbol{\theta} - a\boldsymbol{\alpha}) - \psi(\boldsymbol{\theta}) + a\boldsymbol{\alpha}^\top\boldsymbol{\nabla}_\psi(\boldsymbol{\theta})\right) = \frac{1}{a}D_\psi\left(\boldsymbol{\theta} - a\boldsymbol{\alpha}\|\boldsymbol{\theta}\right)$, as claimed. □

Various safe checks, explained in the following Lemma, show that the risk premium behaves consistently (proof omitted).

Lemma 2 *(i) $\lim_{a\to 0}\mathrm{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = 0$, (ii) $\mathrm{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta})$ is strictly increasing in a, (iii) $\lim_{\boldsymbol{\alpha}\to\boldsymbol{0}}\mathrm{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) = 0$ (holds under any vector norm convergence); (iv) assuming p_ψ Gaussian allows to recover the variance premium of the mean-variance model:*

$$p_\psi = \mathcal{N}(\boldsymbol{\mu},\boldsymbol{\Sigma}) \quad \Rightarrow \quad \mathrm{P}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\boldsymbol{\alpha};\boldsymbol{\theta}) = (a/2)\boldsymbol{\alpha}^\top\boldsymbol{\Sigma}\boldsymbol{\alpha} \ . (9)$$

The proof of (9) involves considering the vector-matrix encoding of the Gaussian (Nielsen & Nock, 2009), with the matrix part of the allocation being the null matrix. The following Lemma provides simple illustrative examples of upperbounds on P_ψ for some popular exponential families.

Lemma 3 *Denote respectively $\mathrm{P}_{d,\boldsymbol{q}}(\boldsymbol{\alpha};\boldsymbol{\theta})$, $\mathrm{P}_\lambda(\boldsymbol{\alpha};\boldsymbol{\theta})$, $\mathrm{P}_{\lambda'}(\boldsymbol{\alpha};\boldsymbol{\theta})$ the premiums associated to the d-dimensional multinomial (parameter $\boldsymbol{q} \in \mathbb{P}$), Poisson (parameter $\lambda > 0$) and exponential (parameter $\lambda' > 0$) distributions. Then (D_{KL} is defined in Table 1):*

$$\mathrm{P}_{d,\boldsymbol{q}}(\boldsymbol{\alpha};\boldsymbol{\theta}) \leq dD_{\mathrm{KL}}\left(\frac{1}{a}\middle\|\frac{1}{1-\exp(-a)}\right) \ , \quad (10)$$

$$\mathrm{P}_\lambda(1;\theta) \leq a\lambda \ , \quad (11)$$

$$\mathrm{P}_{\lambda'}(1;\theta) \leq \frac{1}{\lambda'} - \frac{1}{\lambda' + a} \ . \quad (12)$$

(proof omitted) Poisson and exponential distributions have a single natural parameter, which explains the "1" in lieu of $\boldsymbol{\alpha}$ in (11-12). The bounds in (10-12) are all increasing in a; those of (11-12) are also increasing with the variance of the distribution, showing that variance minimization as in the mean-variance model may be an approximate primer to control P_ψ.

General comments There is a striking parallel between $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}$ in (1) and (2). Everything is like if the natural parameter $\boldsymbol{\theta}$ were acting as a *natural market allocation.* The corresponding *natural investor* is optimal in the sense that its allocation is based on the market's expected behavior (Banerjee et al., 2005): indeed, exponential families satisfy $\boldsymbol{\theta} = \boldsymbol{\nabla}_{\psi^\star}(\mathrm{E}_{\boldsymbol{w}\sim p_\psi}[\boldsymbol{w}])$. For p_ψ Gaussian, it was previously known that the optimal allocation is proportional to $\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$ (Markowitz, 1952): this is precisely the vector part of the Gaussian's natural parameters (Nielsen & Nock, 2009).

3. Tracking portfolios

We wish to build a portfolio with guarantees (*e.g.* lower bounds) on its certainty equivalents in the mean-divergence model. As usual in on-line learning, we update this portfolio, say $\boldsymbol{\alpha}_0$, $\boldsymbol{\alpha}_1$, ..., with the will to track sufficiently closely a *reference* portfolio allowed to drift over iterations: $\boldsymbol{r}_0, \boldsymbol{r}_1, ...$ Intuitively, the drifting reference is assumed to bring large certainty equivalents. There is a third parameter allowed to drift, the natural market allocation: $\boldsymbol{\theta}_0, \boldsymbol{\theta}_1, ...$. Naturally, we could suppose that $\boldsymbol{r}_t = \boldsymbol{\theta}_t, \forall t$, which would amount to tracking directly the best possible allocation, but this setting would be too restrictive because it may be easier to track some $\boldsymbol{r}_t$ close to $\boldsymbol{\theta}_t$ but having specific properties that $\boldsymbol{\theta}_t$ does not have (*e.g.* sparsity). In order not to laden the analysis, the reference portfolio enjoys the same risk aversion parameter a as ours.

The algorithm we propose is named $\mathcal{OMD}_{\phi,\psi}$, for "On-line learning in the Mean-Divergence model". To state $\mathcal{OMD}_{\phi,\psi}$, we abbreviate the gradient (in $\boldsymbol{\alpha}$) of the risk premium as: $\boldsymbol{\nabla}_{\mathrm{P}}(\boldsymbol{\alpha};\boldsymbol{\theta}) \doteq \boldsymbol{\nabla}_\psi(\boldsymbol{\theta}) - \boldsymbol{\nabla}_\psi(\boldsymbol{\theta} - a\boldsymbol{\alpha})$ (a, ψ implicit in the notation). $\mathcal{OMD}_{\phi,\psi}$ initializes $\boldsymbol{\alpha}_0 = (1/d)\mathbf{1}$, *learning rate* parameter $\eta > 0$, and then iterate the following update, for $t = 0, 1, ..., T-1$:

$$\boldsymbol{\alpha}_{t+1} \quad \leftarrow \quad \boldsymbol{\nabla}_\phi^{-1}(\boldsymbol{\nabla}_\phi(\boldsymbol{\alpha}_t) - \eta\boldsymbol{\nabla}_{\mathrm{P}}(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t) - z_t\mathbf{1}) \quad (13)$$

where z_t is chosen so that $\boldsymbol{\alpha}_{t+1} \in \mathbb{P}$[2]. There are several quantities of interest to state our main result:

$$\varsigma \quad \doteq \quad \max_{t\geq 0}\max_{i\neq j}(\mathbf{1}_i - \mathbf{1}_j)^\top \boldsymbol{\nabla}_{\mathrm{P}}(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t) \ , \quad (14)$$

$$\nu \quad \doteq \quad \max_{t\geq 0}\|\boldsymbol{\nabla}_\psi(\boldsymbol{\theta}_t)\|_\infty \ , \quad (15)$$

$$\alpha \quad \doteq \quad \min_{t\geq 0}\min_i \alpha_{t,i} \ . \quad (16)$$

ς is the maximal scope of the premium gradient, ν is the maximal market return in absolute value, and α is the minimal allocation made by $\mathcal{OMD}_{\phi,\psi}$. We finally denote as λ the minimal eigenvalue, over all iterations, of the Hessian of ψ which fits a Taylor-Lagrange expansion of P_ψ's Bregman divergence (see *e.g.* (Nock et al., 2008), Lemma 2). $\lambda > 0$ since ψ is strictly convex.

[2]When $\mathrm{dom}(\phi) \not\subseteq \mathbb{R}_+$, we scale and renormalize $\boldsymbol{\alpha}_{t+1}$ when necessary to ensure that $\boldsymbol{\alpha}_{t+1} \in \mathbb{P}$.

Theorem 2 *Let $\upsilon > 0$ be user-fixed. Let $\mathbb{T} \subseteq \{0, 1, ..., T-1\}$ group iterations s. t. $\boldsymbol{\alpha}_t \neq \boldsymbol{r}_t$. Fix*

$$a \quad = \quad \frac{(\upsilon + 2\nu)}{\lambda \min_{t\in\mathbb{T}}\|\boldsymbol{\alpha}_t - \boldsymbol{r}_t\|_2^2} \ . \quad (17)$$

*Then, for **any** $\eta > 0$, the certainty equivalent of $\mathcal{OMD}_{\mathrm{KL},\psi}$ can be lower bounded as follows, $\forall T > 0, \forall p, q \geq 1, (1/p) + (1/q) = 1$:*

$$\sum_{t=0}^{T-1}\mathrm{C}_\psi(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t) \geq \sum_{t=0}^{T-1}\mathrm{C}_\psi(\boldsymbol{r}_t;\boldsymbol{\theta}_t) - d^{\frac{1}{q}}\ln\left(\frac{1}{\alpha}\right)\sum_{t=0}^{T-1}\|\boldsymbol{r}_{t+1} - \boldsymbol{r}_t\|_p + |\mathbb{T}|\upsilon - T\varsigma - \frac{1-\alpha}{\eta}\ln\left(\frac{1}{\alpha(1-\alpha)}\right) - \ln d \ . \quad (18)$$

Proof: The proof exploits a popular high-level trick consisting in crafting a (lower) bound to the progress to the shifting reference:

$$\delta_t \doteq D_{\mathrm{KL}}(\boldsymbol{r}_t\|\boldsymbol{\alpha}_t) - D_{\mathrm{KL}}(\boldsymbol{r}_{t+1}\|\boldsymbol{\alpha}_{t+1}) = \delta_{t,1} + \delta_{t,2} \quad (19)$$

with

$$\delta_{t,1} \quad \doteq \quad D_{\mathrm{KL}}(\boldsymbol{r}_t\|\boldsymbol{\alpha}_t) - D_{\mathrm{KL}}(\boldsymbol{r}_t\|\boldsymbol{\alpha}_{t+1}) \ ,$$
$$\delta_{t,2} \quad \doteq \quad D_{\mathrm{KL}}(\boldsymbol{r}_t\|\boldsymbol{\alpha}_{t+1}) - D_{\mathrm{KL}}(\boldsymbol{r}_{t+1}\|\boldsymbol{\alpha}_{t+1}) \ .$$

We bound separately the two terms, starting with $\delta_{t,1}$. Using (13), the definition of $\boldsymbol{\nabla}_{\mathrm{P}}(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t)$ and the fact that $\boldsymbol{r}_t \in \mathbb{P}$ and $\boldsymbol{\alpha}_t \in \mathbb{P}$, we have:

$$\delta_{t,1} \quad = \quad (\eta/a)\tau_t - D_{\mathrm{KL}}(\boldsymbol{\alpha}_t\|\boldsymbol{\alpha}_{t+1}) \ , \quad (20)$$

with $\tau_t \doteq ((\boldsymbol{\theta}_t - a\boldsymbol{\alpha}_t) - (\boldsymbol{\theta}_t - a\boldsymbol{r}_t))^\top(\boldsymbol{\nabla}_\psi(\boldsymbol{\theta}_t - a\boldsymbol{\alpha}_t) - \boldsymbol{\nabla}_\psi(\boldsymbol{\theta}_t))$. We now bound the two terms in (20).

Lemma 4 *$\tau_t \geq a\left(\mathrm{C}_\psi(\boldsymbol{r}_t;\boldsymbol{\theta}_t) - \mathrm{C}_\psi(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t) + \upsilon\right)$ if $t \in \mathbb{T}$, and $\tau_t = a\left(\mathrm{C}_\psi(\boldsymbol{r}_t;\boldsymbol{\theta}_t) - \mathrm{C}_\psi(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t)\right)$ otherwise.*

(proof given in the supplementary material[1])

Lemma 5 *$D_{\mathrm{KL}}(\boldsymbol{\alpha}_t\|\boldsymbol{\alpha}_{t+1}) \leq \eta\varsigma$.*

(proof given in the supplementary material[1])

Putting altogether Lemmata 4 and 5 in (20), we obtain the following lower bound on the sum of $\delta_{t,1}$:

$$\sum_{t=0}^{T-1}\delta_{t,1} \quad \geq \quad \eta\left(\sum_{t=0}^{T-1}\mathrm{C}_\psi(\boldsymbol{r}_t;\boldsymbol{\theta}_t) - \sum_{t=0}^{T-1}\mathrm{C}_\psi(\boldsymbol{\alpha}_t;\boldsymbol{\theta}_t)\right) + \eta\left(|\mathbb{T}|\upsilon - T\varsigma\right) \ . \quad (21)$$

Table 2. Experimental market domains. Returns are daily (DJIA, NYSE and TSE) or weekly (S&P500).

name	d	T	start date	end date
DJIA	30	506	01/14/01	01/14/03
NYSE	36	5650	07/03/62	12/31/84
S&P500	324	618	01/08/98	11/12/09
TSE	88	1257	01/04/94	12/31/98

Working on a lowerbound for $\delta_{t,2}$ is easier, as $\delta_{t,2}$ simplifies to:

$$\begin{aligned}\delta_{t,2} &= \phi(\boldsymbol{r}_t) - \phi(\boldsymbol{r}_{t+1}) + (\boldsymbol{r}_{t+1} - \boldsymbol{r}_t)^\top \boldsymbol{\nabla}_{\text{KL}}(\boldsymbol{\alpha}_{t+1}) \\ &\geq \phi(\boldsymbol{r}_t) - \phi(\boldsymbol{r}_{t+1}) - \|\boldsymbol{r}_{t+1} - \boldsymbol{r}_t\|_p d^{\frac{1}{q}} \ln \frac{1}{\alpha} \quad (22)\end{aligned}$$

where (22) follows from Hölder inequality ($p, q \geq 1, (1/p) + (1/q) = 1$). There remains to sum (19) for $t = 0, 1, ..., T-1$, use (21) and (22), rearrange and use the facts $D_{\text{KL}}(\boldsymbol{r}_0\|\boldsymbol{\alpha}_0) = \phi(\boldsymbol{r}_0) + \ln d$, $D_{\text{KL}}(\boldsymbol{r}_T\|\boldsymbol{\alpha}_T) - \phi(\boldsymbol{r}_T) \geq (1-\alpha)\ln(\alpha(1-\alpha))$ to get (18). □

Comments on $\mathcal{OMD}_{\phi,\psi}$ and Theorem 2 The choice $\phi = \text{KL}$ in Theorem 2 was made in part to fuel experimental observations (See Section 4). Notice also the absence of constraint on η: previous theoretical results on on-line algorithms tend to put very tight constraints on η for efficient learning (Borodin et al., 2004; Kivinen & Warmuth, 1997). $\mathcal{OMD}_{\phi,\psi}$ explicitly relies on the optimization of the premiums, yet it implicitly works on maximizing the certainty equivalents as well, as indeed (6) implies $\boldsymbol{\nabla}_{\text{P}}(\boldsymbol{\alpha};\boldsymbol{\theta}) = \boldsymbol{\nabla}_\psi(\boldsymbol{\theta}) - \boldsymbol{\nabla}_{\text{C}}(\boldsymbol{\alpha};\boldsymbol{\theta})$, where $\boldsymbol{\nabla}_{\text{C}}(\boldsymbol{\alpha};\boldsymbol{\theta})$ is the gradient in $\boldsymbol{\alpha}$ of the certainty equivalent. It is thus not surprising that $\mathcal{OMD}_{\phi,\psi}$ meets guarantees on the certainty equivalents. From the information geometric standpoint, $\mathcal{OMD}_{\phi,\psi}$ turns out to approximate a generalization of Amari's natural gradient (Amari, 1998), to progress towards the optimization of a cost function using a geometry induced by a Bregman divergence (D_ϕ).

Lemma 6 *The solution to* $\boldsymbol{\alpha}' = \arg\min_{\boldsymbol{\alpha}\in\mathbb{A}} D_\phi(\boldsymbol{\alpha}\|\boldsymbol{\alpha}_t)$, *where* $\mathbb{A} = \{\boldsymbol{\alpha} \in \mathbb{R} : (\boldsymbol{\alpha}^\top \mathbf{1} = 1) \wedge (\text{P}_\psi(\boldsymbol{\alpha};\boldsymbol{\theta}) \leq k)\}$, *satisfies the following set of non-linear inequalities:*

$$\boldsymbol{\alpha}' = \boldsymbol{\nabla}_\phi^{-1}(\boldsymbol{\nabla}_\phi(\boldsymbol{\alpha}_t) - \eta\boldsymbol{\nabla}_{\text{P}}(\boldsymbol{\alpha}';\boldsymbol{\theta}_t) - z_t\mathbf{1}) \;. \quad (23)$$

(proof omitted) Notice that, to enforce $\boldsymbol{\alpha}' \in \mathbb{P}$ in (23), it is enough to ensure that $\text{dom}(\phi) \subseteq \mathbb{R}_+$. One may easily check that fixing $\phi(\boldsymbol{x}) = \boldsymbol{x}^\top \mathbf{G}\boldsymbol{x}$ ($\mathbf{G}$ symmetric positive definite) in (23) and removing the constraint $\boldsymbol{\alpha}^\top\mathbf{1} = 1$ ($z_t = 0$) allows to retrieve exactly Theorem 1 in (Amari, 1998). The update (13) in $\mathcal{OMD}_{\phi,\psi}$ appears as a tractable approximation to (23) — all the better as $D_\phi(\boldsymbol{\alpha}'\|\boldsymbol{\alpha}_t)$ is small — in which $\boldsymbol{\alpha}_t$ replaces $\boldsymbol{\alpha}'$ in the premium gradient. Since $\boldsymbol{\alpha}_t, \boldsymbol{\alpha}' \in \mathbb{P}$, a most natural choice for D_ϕ suggested by Lemma 6 is Kullback-Leibler divergence (Table 1), in which case $\mathcal{OMD}_{\phi,\psi}$ resembles $\mathcal{EG}$ algorithms (Kivinen & Warmuth, 1997).

The bound of Theorem 2 is not directly applicable, like most bounds in on-line learning (Kivinen et al., 2006), yet it provides intuitive clues about the dependencies between the parameters, and their choices to efficiently tune $\mathcal{OMD}_{\text{KL},\psi}$. If we except the term $|\mathbb{T}|\upsilon - T\varsigma$, the remaining part of the penalty in (18) is in fact familiar to on-line learning (Kivinen et al., 2006), and says that tracking the reference may indeed be more efficient as it gets sparse. The term $|\mathbb{T}|\upsilon - T\varsigma$ is interesting for the premium choice: ς actually depends on a, yet a appears in the gradient of ψ. Hence, premiums with a slowly increasing gradient, *e.g.* concave like for $\psi = \text{KL}$ or $\psi = \text{IS}$, dampen the penalty $-T\varsigma$ in (18), thus potentially leading to improved performances.

4. Experiments

We have considered four market domains, summarized in Table 2. They cover overall a wide period, from the early sixties to the last financial crisis. Experiments were devised to assess various objectives, including in particular (i) whether tracking portfolios on the basis of their risk premiums or certainty equivalents allows to find portfolios with good returns; (ii) whether lifting the mean-variance model to the more general mean-divergence model allows to cope more efficiently with different markets, in particular against two popular market opponents: the uniform cost rebalanced portfolio, $\mathcal{UCRP}$, which represents the average market's performance, and the best stock, $\mathcal{BEST}$, which is the stock giving the largest cumulative returns over all market iterations; (iii) whether the mean-divergence model improves the acuteness to spot, with new premiums, events at the market scale that would otherwise be missed — or at least dampened — in the mean-variance model.

General results: on each domain, $\mathcal{OMD}_{\phi,\psi}$ was run with every possible combination of the following parameters: $a \in \{0.01, 1, 100\}, \eta \in \{0.01, 1, 100\}, \psi \in \{\text{M}, \text{KL}, \text{IS}\}, \phi \in \{\text{L}_q, \text{KL}, \text{IS}\}$ (Table 1: $q \in \{2.001, 3, 4\}$ for the q-norm). Finally, in order to assess whether the update (13) can be made more efficient using more than just the last returns, we test the possibility of using, in the premium gradient update, a window average of the last r iterations, for $r \in \{1, 2, 4\}$. The results, integrating the cumulated returns of $\mathcal{BEST}$ and $\mathcal{UCRP}$, are given in Table 3. Due to the lack of space, we only provide the results for $\mathcal{OMD}_{\text{KL},\psi}$, but the interested reader may check the supplementary material[1]

Table 3. Cumulated returns (left table) and cumulated premiums (right table, y-scales are logscales) for $\mathcal{OMD}_{\text{KL},\psi}$ on the four domains, using three different premium generators ψ (leftmost column: see Table 1; M is Markowitz' variance premium). On each plot, $\mathcal{OMD}_{\text{KL},\psi}$'s synthetic results are given as follows: the *light grey* part covers the interval of the [25%, 75%] quantiles of $\mathcal{OMD}_{\text{KL},\psi}$, the *red* curve displays $\mathcal{OMD}_{\text{KL},\psi}$'s median results, the lower and upper *green* curves display respectively $\mathcal{OMD}_{\text{KL},\psi}$'s min and max results. The results of $\mathcal{BEST}$ are in *purple*, and those of $\mathcal{UCRP}$ are in *cyan*.

	Cumulated returns				Cumulated premiums			
	DJIA	NYSE	S&P500	TSE	DJIA	NYSE	S&P500	TSE
M, KL, IS	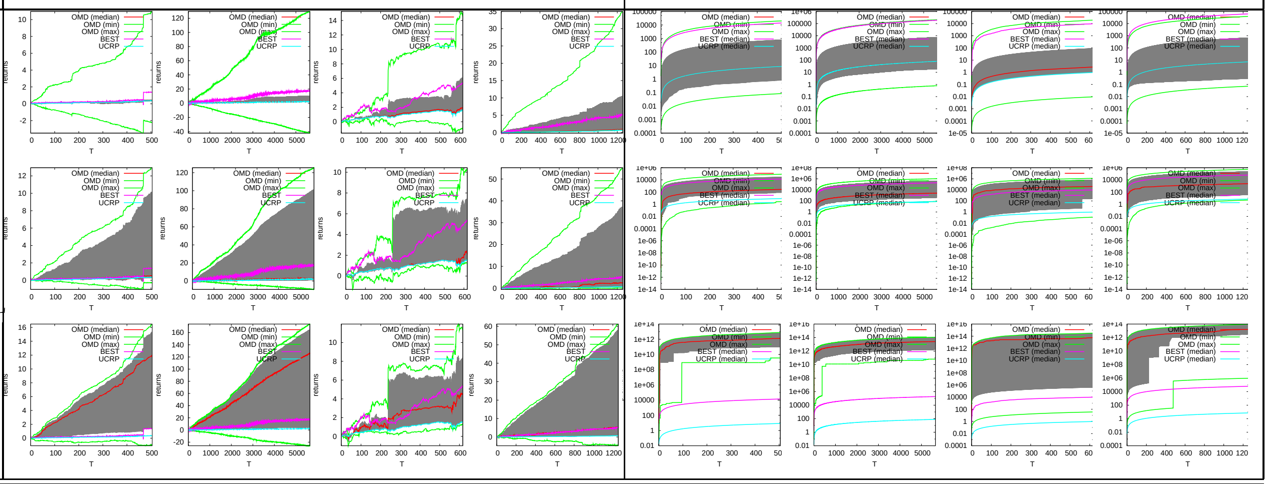							

for the results of the other choices of ϕ. The following conclusions can be drawn from these experiments: the better the cumulated returns for $\mathcal{OMD}_{\text{KL},\psi}$, the larger its premiums; in some sense, the paying strategies are noted as riskiest in the mean-divergence model. The poorest results according to cumulated returns are obtained for Markowitz' variance premium (M), with premiums almost always smaller than $\mathcal{BEST}$'s by orders of magnitude. Compared to $\mathcal{BEST}$'s, the premiums for KL are quite comparable at least for the median values, while those for IS are clearly huge. But the returns are up to the task: on the DJIA, $\mathcal{OMD}_{\text{KL,IS}}$'s median return with IS is more than *six* times that of $\mathcal{BEST}$, while more than *75%* of the possible combinations of parameters of $\mathcal{OMD}_{\text{KL,IS}}$ give better results than $\mathcal{BEST}$. On the NYSE, $\mathcal{OMD}_{\text{KL,IS}}$'s median returns are this time more than *ten* times those of $\mathcal{BEST}$. Recall that premiums are not honored by investors (unlike *e.g.* in insurance), hence one can judge results on the basis of returns only: with respect to this standpoint, $\mathcal{OMD}_{\text{KL,IS}}$ gives by far the best results, the second best being clearly $\mathcal{OMD}_{\text{KL,KL}}$. This is quite in accordance with the comments of Section 3, and comes as a strong advocacy to lift the mean-variance model to the mean-divergence model. Finally, we spotted no significant difference when varying window size r.

Influences of a and η: Two major parameters in running $\mathcal{OMD}_{\phi,\psi}$ are a and η. To evaluate their influence, we filtered the general result, and plot in Table 4 the cumulated returns of $\mathcal{OMD}_{\text{KL,IS}}$ as a function of the values of a and η. The results for the other choices for ψ can be consulted in the supplementary material[1]. Table 4 clearly displays two opposite behaviors for the influence of a and η: while returns increase with a, they decrease with η. Results for $\mathcal{OMD}_{\text{KL,M}}$ tend to display that the *opposite* pattern holds for Markowitz' variance premium, as returns tend to decrease with a and increase with η. The case of $\mathcal{OMD}_{\text{KL,KL}}$ is also different, the median values ($a = \eta = 1$) seemingly being the best choice for all four domains. A plausible explanation to this phenomenon may lie in the second derivative of ψ, and thus in the convexity regime of the premium: for small returns, the second derivative values can roughly be ordered as IS $\gg$ KL $\gg$ M, and thus yield allocations that are much more spread before normalization for IS in (13). This perhaps provides a better acuteness to $\mathcal{OMD}_{\text{KL,IS}}$ through the risk premium, and to be used to its full potential, one does not have interest in fixing small values for a that would otherwise cloud the issue by reducing this premium. We thus see two opposite strategies through $\mathcal{OMD}_{\text{KL},\psi}$: the choice $\psi = $ M provides us with an algorithm which works at best when taking the less risks, giving in return portfolios with suboptimal returns, sometimes competing with the best stock. The "opposite" choice $\psi = $ IS gives a much more aggressive, high-premium / higher-return algorithm. For such aggressive strategies, the high premiums do not only act as signals to spot potential portfolios being subject to

Table 4. Cumulated returns of $\mathcal{OMD}_{\text{KL,IS}}$ as a function of $a \in \{0.01, 1, 100\}$ (left table) and $\eta \in \{0.01, 1, 100\}$ (right table). Each grey curve represents a run of $\mathcal{OMD}_{\text{KL,IS}}$. The results of $\mathcal{BEST}$ are in *purple*, and those of $\mathcal{UCRP}$ are in *cyan*.

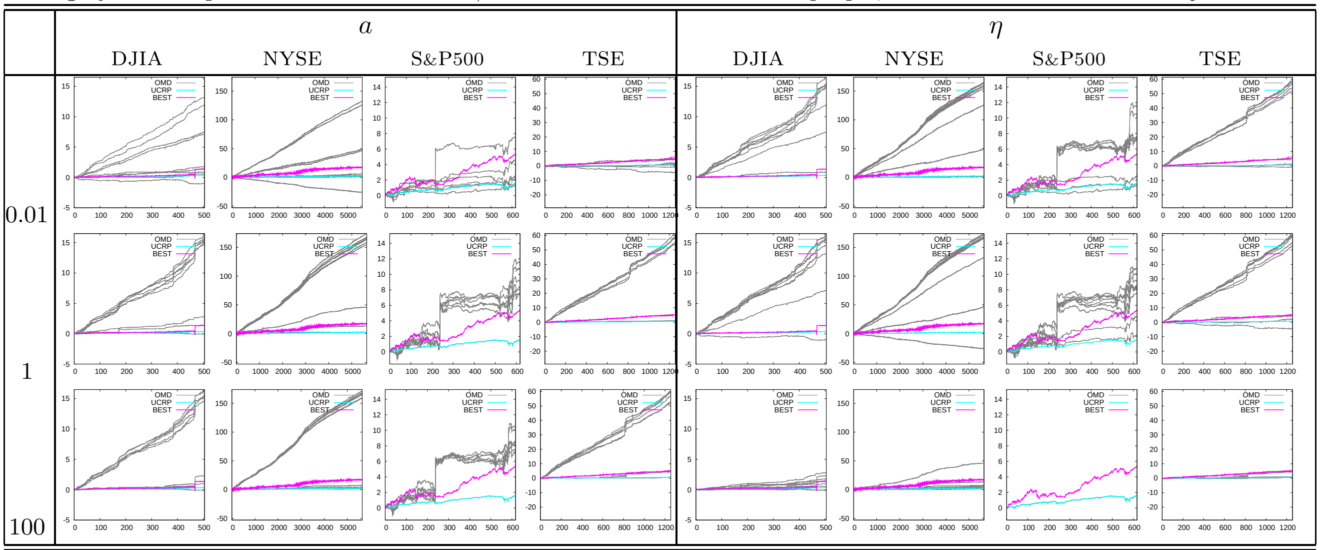

risk: they somehow act as parapets for $\mathcal{OMD}_{\text{KL,IS}}$ to "stay in line", and thus *need* to be high (a large) to really be efficient in this role. This being explained, the somehow "opposite" behavior observed with η may indicate that a and η act as offsets for each other in the update (13): small premium variations allow large learning rates for better results, while large premium variations enforce small learning rates.

$\mathcal{OMD}_{\phi,\psi}$ **watches its basket:** We have drilled down further into the portfolios of $\mathcal{OMD}_{\text{KL,IS}}$, to assess the way allocations are carried out. Table 5 provides some of the results obtained, the remaining of which appear in the supplementary material[1]. In each row, the right table gives the topmost stocks that represented more than 50% of $\mathcal{OMD}_{\text{KL,IS}}$'s portfolio, ordered according to the percentage of the iterations (shown) during which this occurred ("None"= no stock had absolute majority). A ($\star$) indicates $\mathcal{BEST}$. $\mathcal{OMD}_{\text{KL,IS}}$ has a prominent tendency to follow few stocks at a time, quite often catching $\mathcal{BEST}$, thus following Twain's "wise" behavior and playing efficiently against stocks' volatility; yet experiments demonstrate that some iterations tagged as "None" clearly favor a spreading of stocks, thus following Twain's "fool" behavior. Interestingly, the domain on which this spreading is the most frequent has also the most irregular average returns (See $\mathcal{UCRP}$): S&P500. Here, "None" is almost ten times more frequent than the following stock in the list. This fact, after comparison with DJIA and TSE, cannot be explained only by the increase in the number of stocks. In Table 5, the cumulated returns of stocks PHILIP MORRIS (DJIA), DUPONT (NYSE), PURE GOLD MINERALS INC. and INTERNATIONAL FOREST PRODUCTS LTD (TSE) display the ability of $\mathcal{OMD}_{\text{KL,IS}}$ to bet "just in time" on stocks, just before or during periods where they enjoy comparatively more important returns.

Premium values and market events: Finally, we drilled down into the values of premiums obtained, in particular to evaluate differences as a function of the premium P_{ψ}. Table 6 gives three examples of curves obtained on domain S&P500 ($a = 1$), chosen for its average behavior more irregular than the other markets. One can check that all premiums detect events during the last financial crisis (rightmost peaks), but relative variations are much smaller for P_{M}. On the other hand, P_{KL} peaks much more distinctively on these events, while P_{IS} yields very large premiums, as expectable from the theory and experiments developed above.

5. Conclusion

Carefully crafted heuristics have already demonstrated their capacities in beating $\mathcal{BEST}$ (Borodin et al., 2004), yet these are still crucially lacking theoretical foundations; to the best of our knowledge, our work may be the first attempt to show that such attainable performances may borne out a sound theory, moreover forged more than a decade ago (Amari, 1998; Kivinen & Warmuth, 1997) and popular ever since in machine learning. Our main objective is not in talking experimentally the big numbers with respect to other contenders: there are of course caveats to applying our algorithm, like for any other in the category (Borodin et al., 2004). Instead, even when we have not

Table 5. Allocations of $\mathcal{OMD}_{\text{KL,IS}}$ $(a = 100.0, \eta = 0.01)$. Each row relates to a domain (top to bottom: DJIA, NYSE, S&P500, TSE). In each row, the *right* table shows the most prominent stocks in $\mathcal{OMD}_{\text{KL,IS}}$'s portfolio (see text). The *left* plot displays the cumulated returns of the topmost stock of this list; vertical black bars indicate the iterations during which this stock had absolute majority in the portfolio (the KIN ARK plot may be misleading because of its size and the width of the vertical bars). The *center* plot displays the cumulated returns of another stock appearing in the list (convention for vertical black bars are the same).

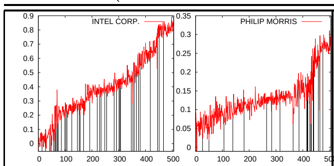

DJIA::	
None:	16.01%
INTEL CORP.:	7.91%
(⋆) AT&T CORP.:	7.31%
HP:	6.32%
JP MORGAN:	4.94%
PHILIP MORRIS:	4.35%
HONEYWELL:	4.35%

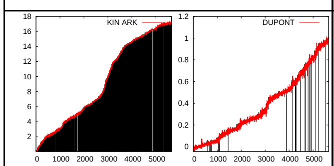

NYSE:	
(⋆) KIN ARK:	17.47%
None::	16.53%
IROQUOIS:	9.43%
ESPEY MAN.:	9.40%
MEI CORP.:	7.75%
COMM METALS:	5.17%
LUKENS:	3.94%

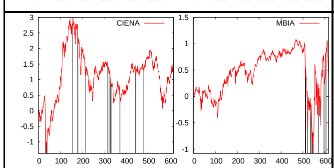

S&P500:	
None::	18.45%
CIENA:	1.94%
JDS UNIPHASE:	1.94%
MBIA:	1.62%
ADV. MIC. DEVC.:	1.46%
JABIL CIRCUIT:	1.29%
QWEST COMMS.:	1.29%

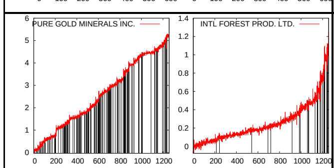

TSE:	
None::	17.33%
(⋆) PURE GOLD MIN.:	9.70%
BREAKWATER RES.:	8.27%
REPAP ENT. INC.:	5.72%
GENTRA INC.:	3.50%
COTT CORP.:	3.34%
MIRAMAR MIN.:	3.18%

found the golden eggs to put in our Twain's basket, we do believe that this possible bond between theory and such attainable experimental performances is as interesting as ordinary looking eggs with silver yolk to start filling this basket. In particular, our results show that the mean-divergence model may present new avenues for research on popular on-line learning algorithms like $\mathcal{EG}$ (Kivinen & Warmuth, 1997), such as the ways the parameters of the expected utility theory (Pratt, 1964) may be plugged in the algorithms and bounds. This also includes the experimental standpoint, as looking at the results in (Borodin et al., 2004) (DJIA and TSE in their Table 1: we used the same data) clearly displays that working with certainty equivalents or premiums, instead of returns like in the original $\mathcal{EG}$, skyrockets returns to the point that we become much more than a legal contender to $\mathcal{ANTICOR}$ (Borodin et al., 2004): we may beat it by orders of magnitude.

Table 6. Premiums on S&P500: P_{M}, P_{KL}, P_{IS} (left to right).

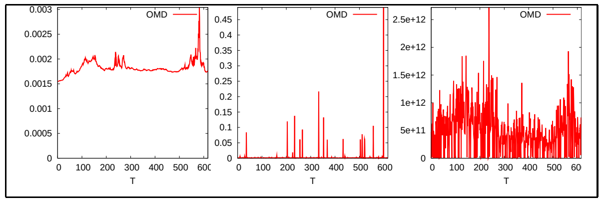

References

Amari, S.-I. Natural Gradient works efficiently in Learning. *Neural Computation*, 10:251–276, 1998.

Banerjee, A., Merugu, S., Dhillon, I., and Ghosh, J. Clustering with Bregman divergences. *J. of Mach. Learn. Res.*, 6:1705–1749, 2005.

Borodin, A., El-Yaniv, R., and Gogan, V. Can we learn to beat the best stock. *JAIR*, 21:579–594, 2004.

Chavas, J.-P. *Risk analysis in theory and practice.* Academic press advanced finance, 2004.

Even-Dar, E., Kearns, M., and Wortman, J. Risk-sensitive online learning. In *17th ALT*, pp. 199–213, 2006.

Kivinen, J. and Warmuth, M. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 132:1–63, 1997.

Kivinen, J., Warmuth, M., and Hassibi, B. The p-norm generalization of the LMS algorithm for adaptive filtering. *IEEE Trans. SP*, 54:1782–1793, 2006.

Markowitz, H. Portfolio selection. *Journal of Finance*, 6:77–91, 1952.

Nielsen, F. and Nock, R. Sided and symmetrized Bregman centroids. *IEEE Trans. IT*, 55:2882–2904, 2009.

Nock, R. and Nielsen, F. Bregman divergences and surrogates for learning. *IEEE Trans. PAMI*, 31(11): 2048–2059, 2009.

Nock, R., Luosto, P., and Kivinen, J. Mixed Bregman clustering with approximation guarantees. In *23 rd ECML*, pp. 154–169. Springer-Verlag, 2008.

Pratt, J.W. Risk aversion in the small and in the large. *Econometrica*, 32:122–136, 1964.

von Neumann, J. and Morgenstern, O. *Theory of games and economic behavior.* Princeton University Press, 1944.

Warmuth, M. and Kuzmin, D. Online variance minimization. In *19 th COLT*, pp. 514–528, 2006.