CANTAB
PART OF GAM SYSTEMATIC

14 November 2018

Trading Markets with Online Convex Optimisation

Thomas Gillam, GAM Systematic | Cantab

Overview

- Asset management
- Online Convex Optimisation (OCO)
- Empirical application to an ‘asset selection’ problem

What is asset management?

- Very loosely: given cash, buy and sell assets such that:

  1. **Risk** is controlled.

  2. We have positive **expected return**.

- Positions must be adjusted over time to achieve this – this adds considerations regarding **trading costs**.

- Cantab is a *systematic* manager, since we use mathematical and statistical methods to devise algorithms, which in turn make all investment decisions.

Systematic investment

Since inception, Cantab has used statistical techniques for:

- Constructing models that can predict future asset returns, given some indicators of market conditions.
  [i.e. regression problems]

- Estimating asset volatility, tail risk, and correlations.
  [e.g. GARCH]

- Portfolio construction. [e.g. Risk parity]

- Designing algorithms for efficient trading when rebalancing.
  [e.g. multi-armed bandit]

Common important challenges

- Adapting to changing market conditions.
- Efficiently obtaining an accurate historical simulation of an algorithm's performance (i.e. not "future peeking")
- Working in a potentially adversarial market, especially on short timescales…

(Aside) S&P 500 futures price, January 2008

Normal market conditions:

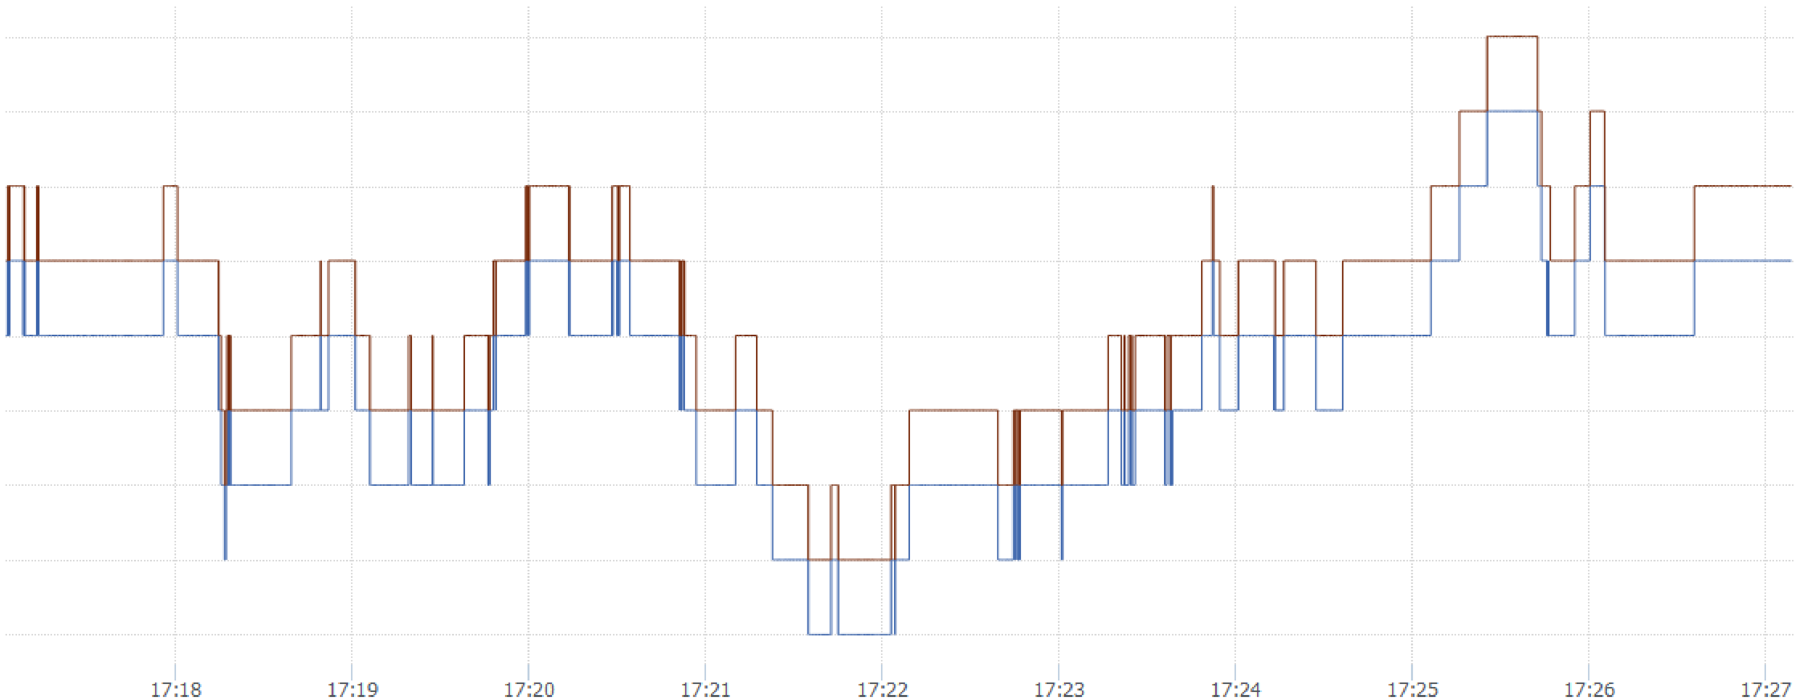

(Aside) S&P 500 futures price, January 2008

Abnormal market conditions:

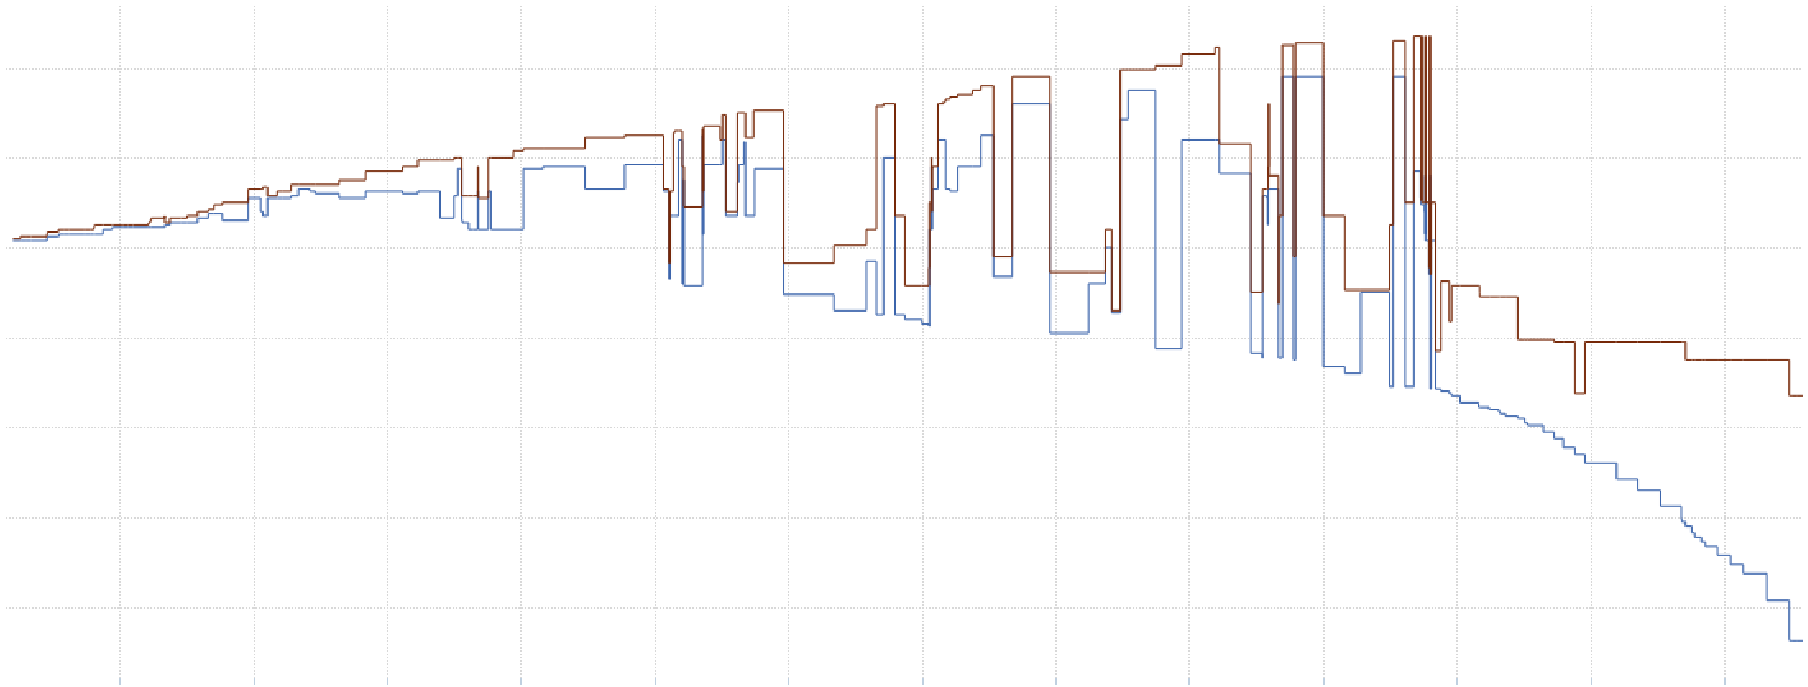

(Aside) S&P 500 futures price, January 2008

Abnormal market conditions:

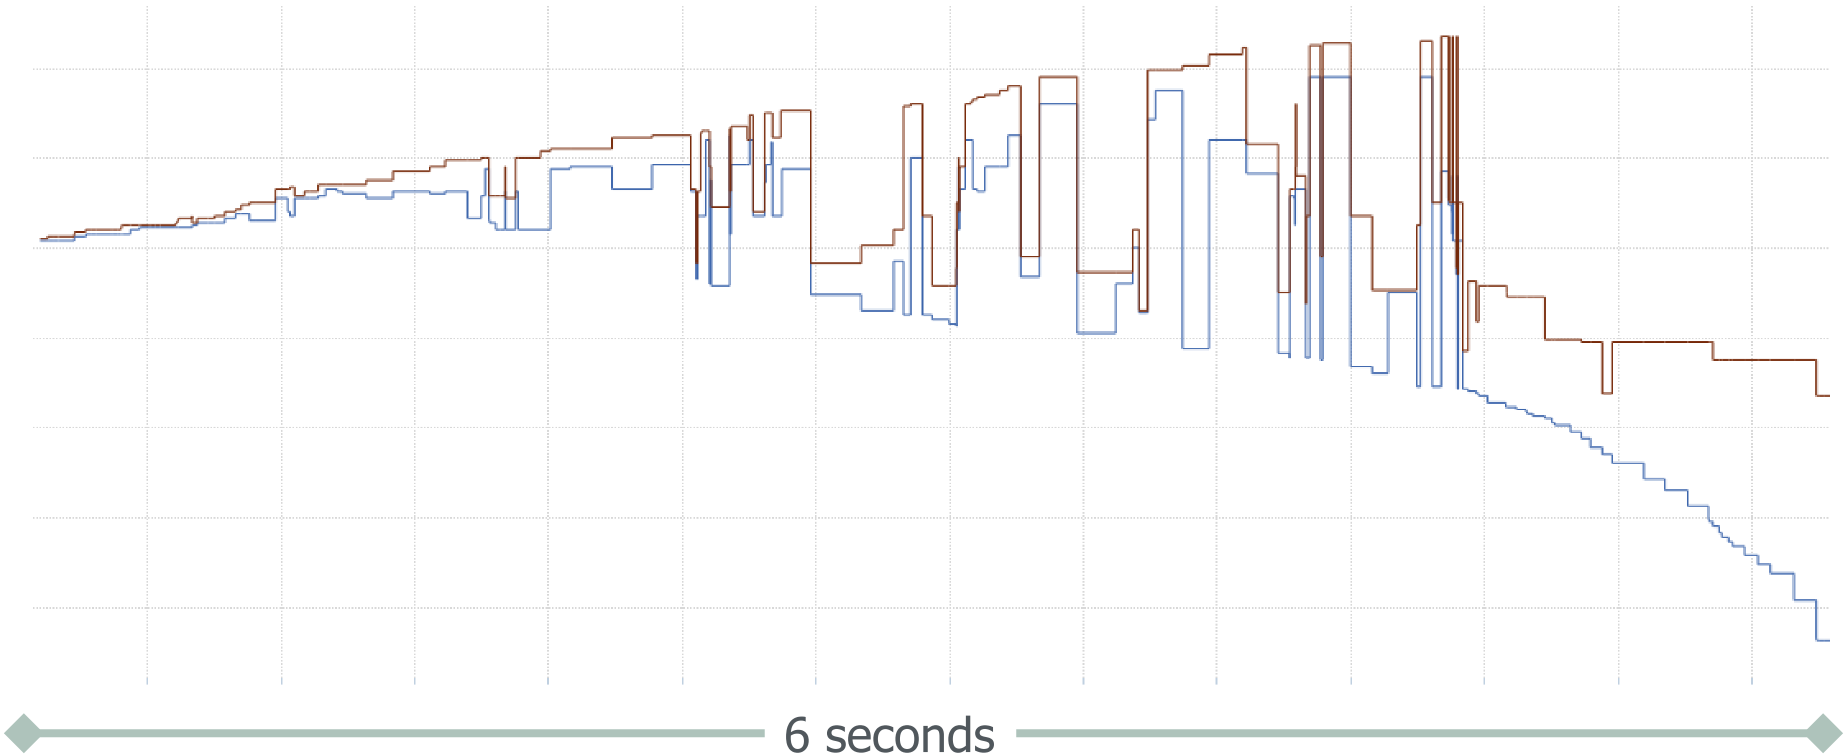

Common important challenges

- Adapting to changing market conditions.
- Efficiently obtaining an accurate historical simulation of an algorithm's performance (i.e. not "future peeking")
- Working in a potentially adversarial market, especially on short timescales…

Common important challenges

- Adapting to changing market conditions.
- Efficiently obtaining an accurate historical simulation of an algorithm’s performance (i.e. not “future peeking”)
- Working in a potentially adversarial market, especially on short timescales…

OCO can provide benefits in all of these respects.

Online convex optimisation

Online convex optimisation

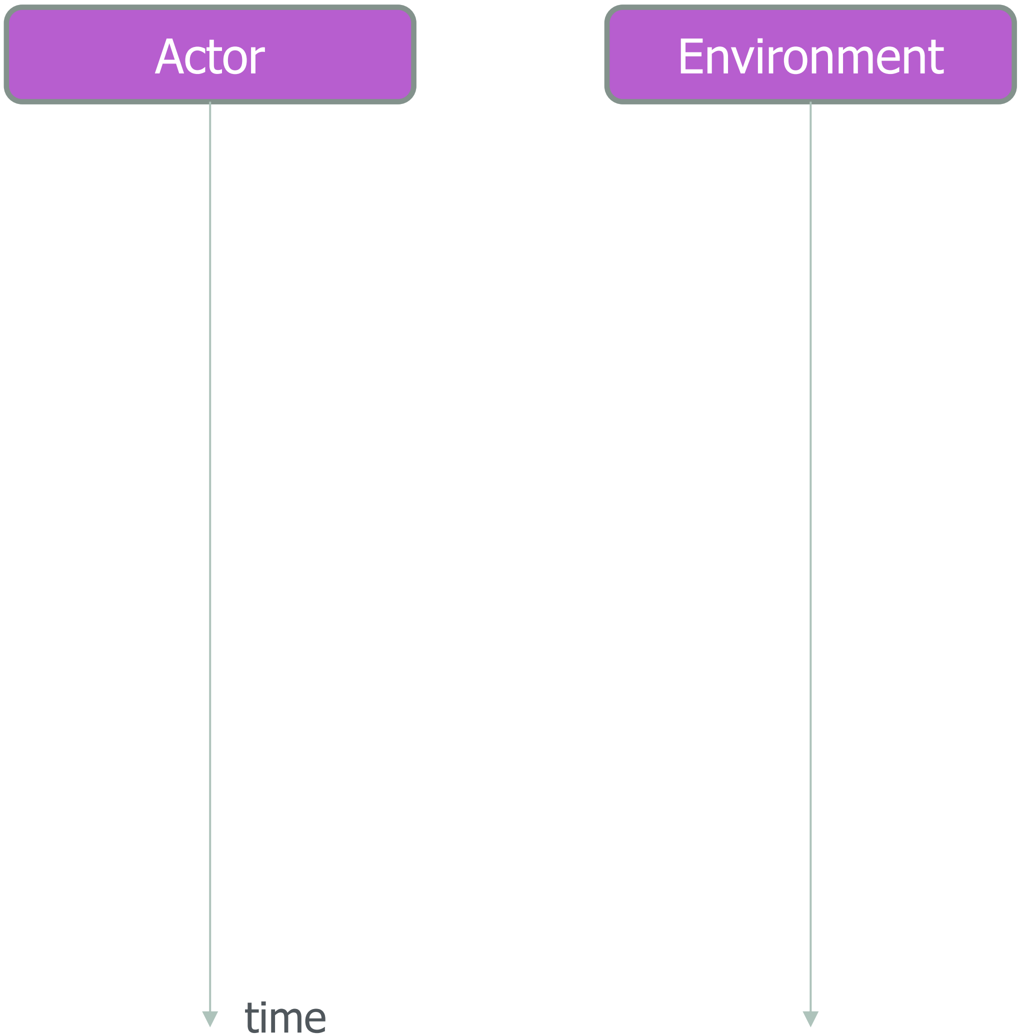

Online convex optimisation

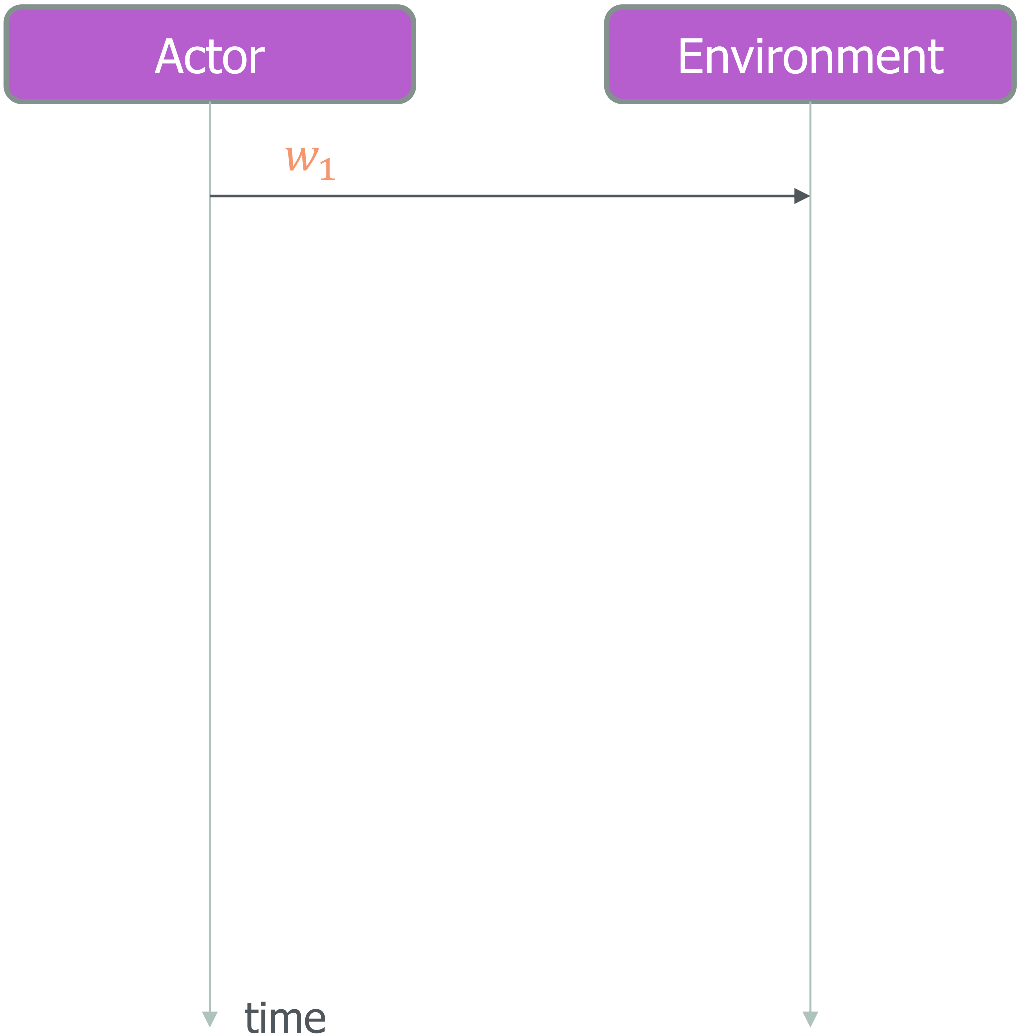

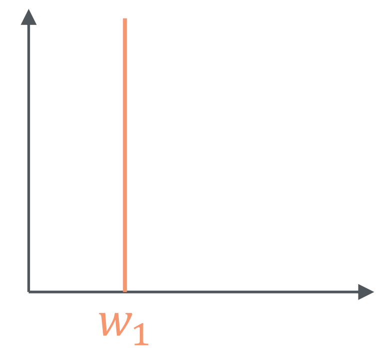

Online convex optimisation

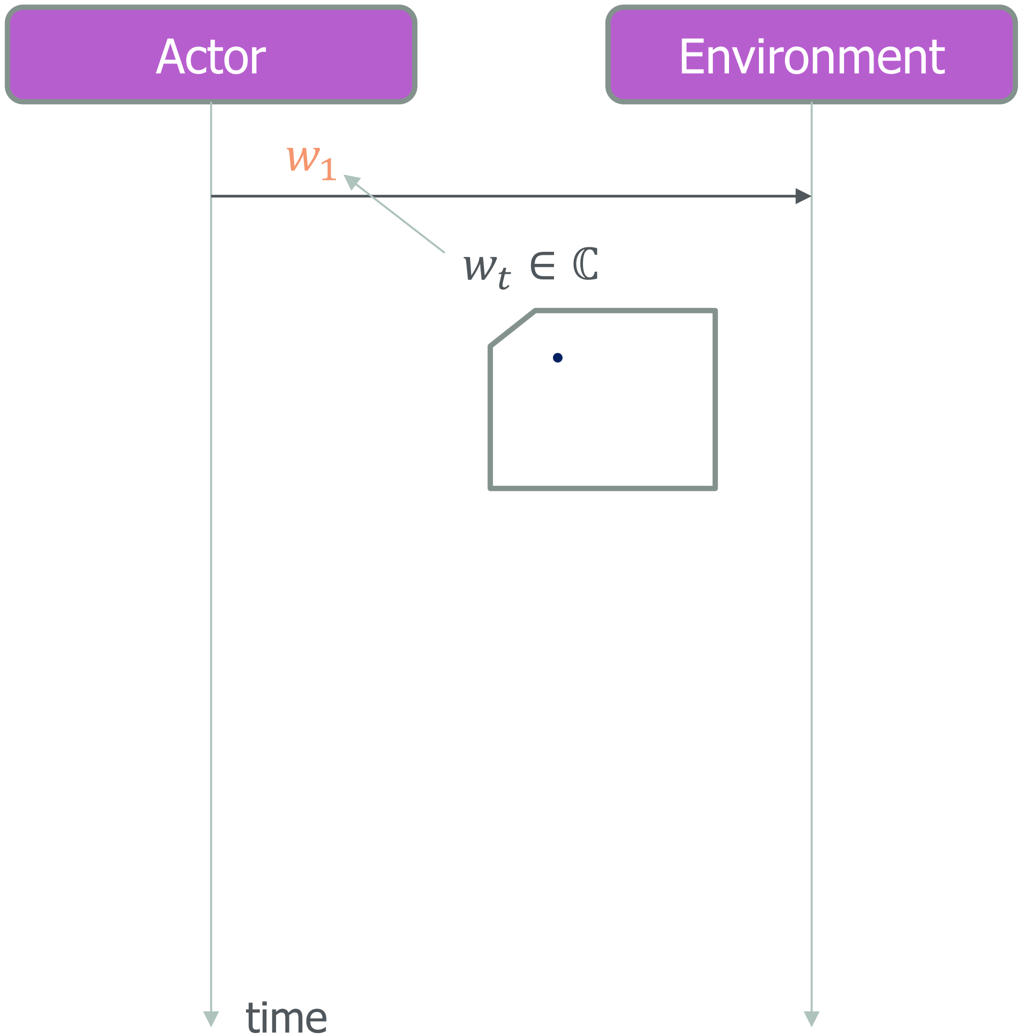

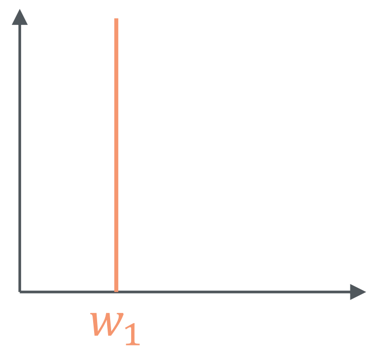

Online convex optimisation

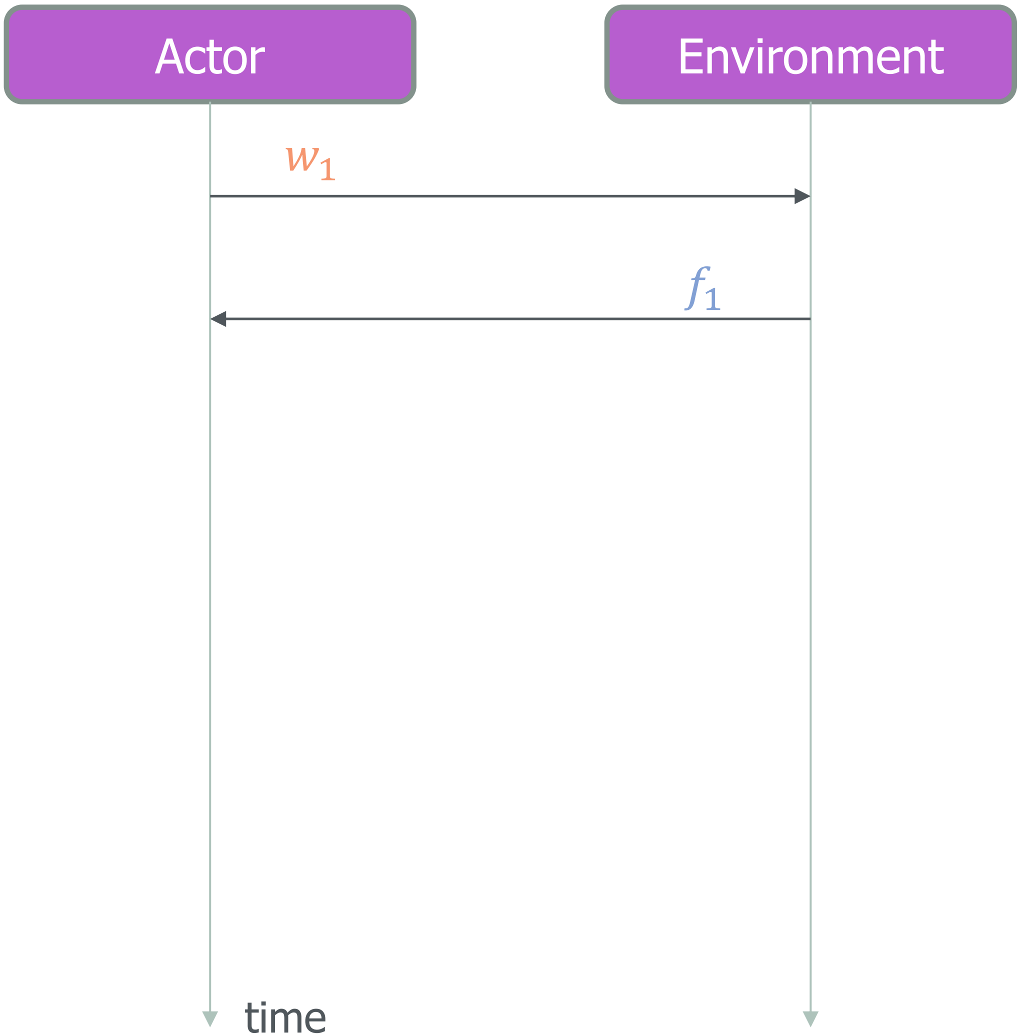

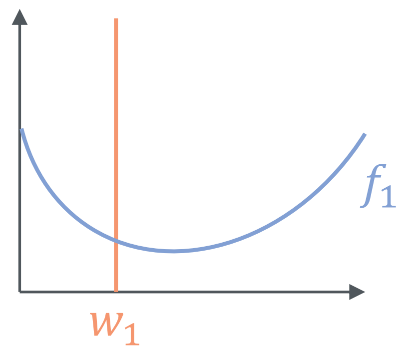

Online convex optimisation

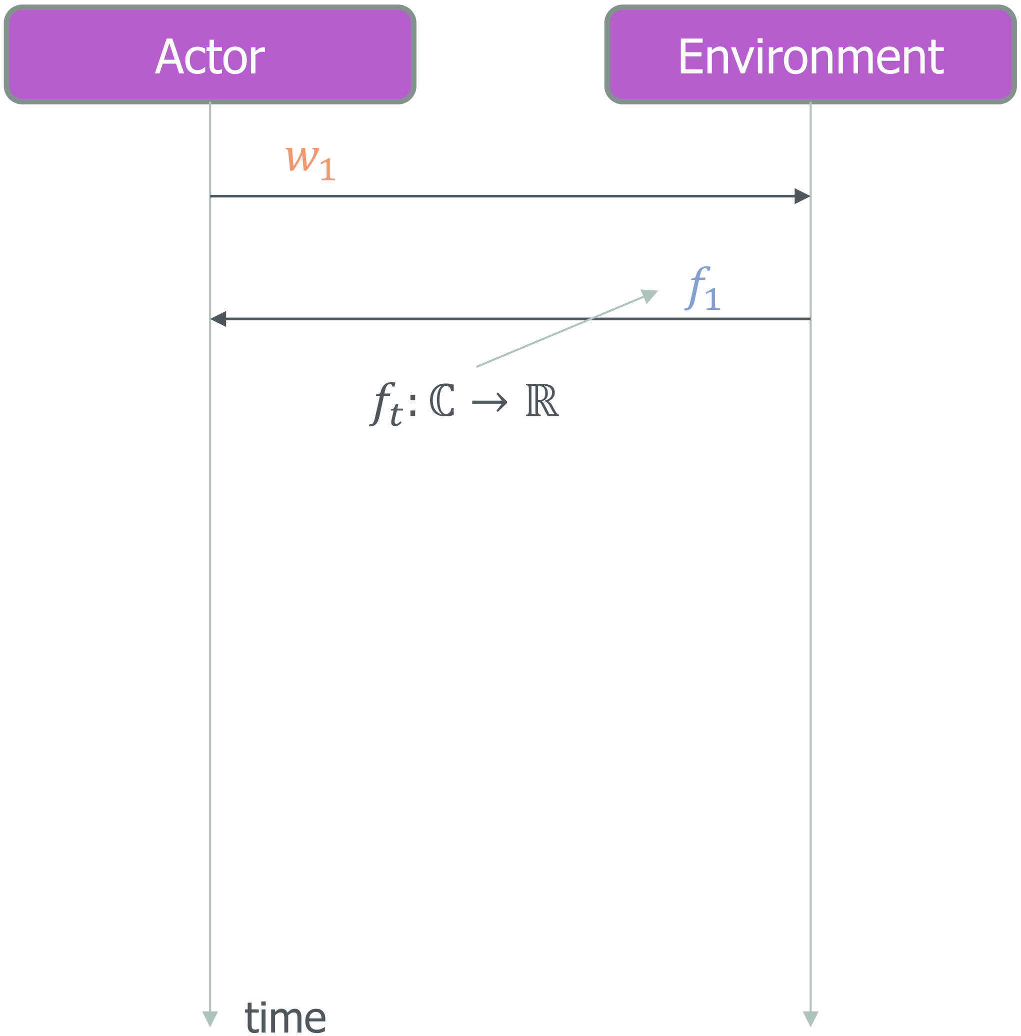

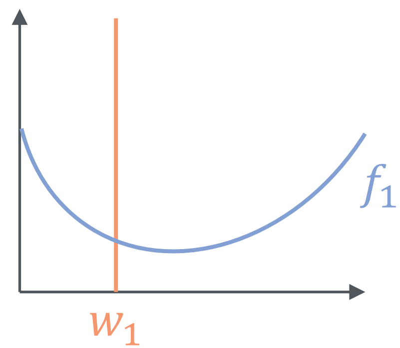

Online convex optimisation

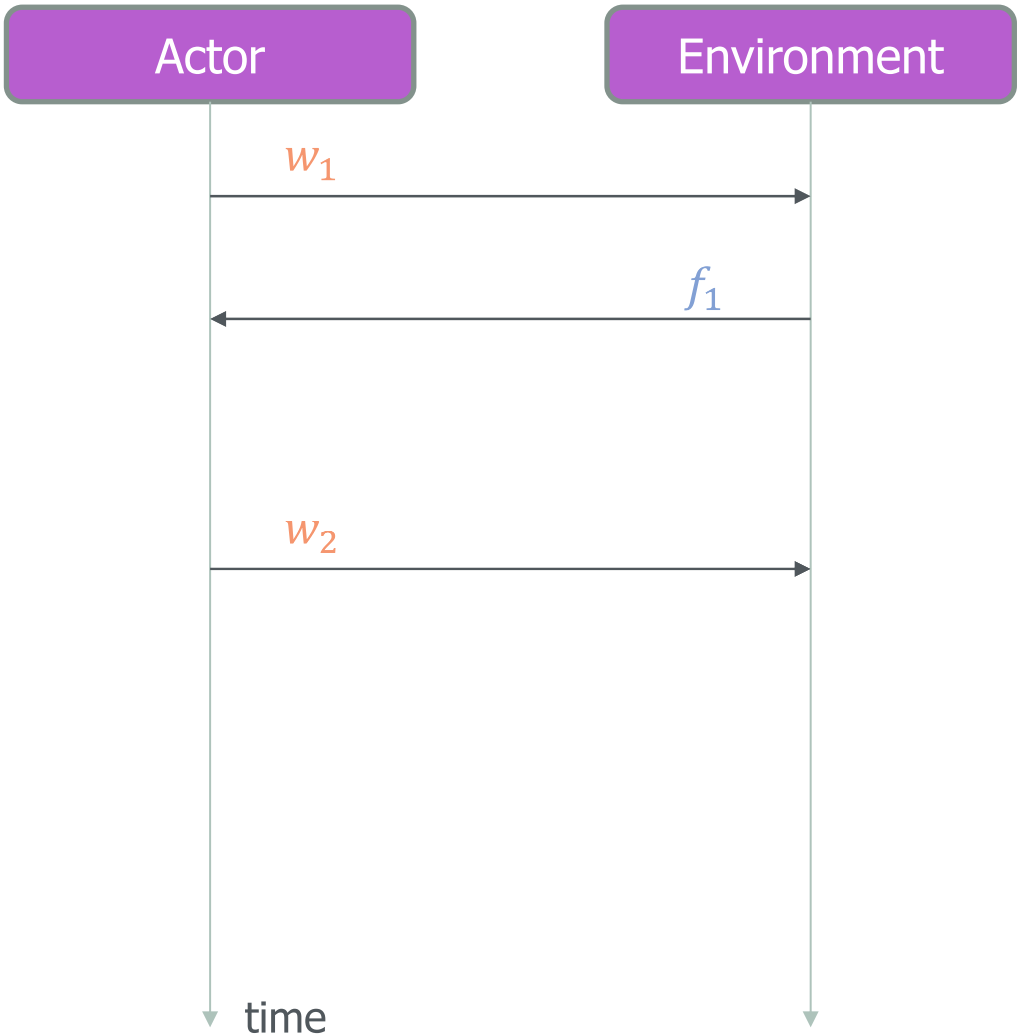

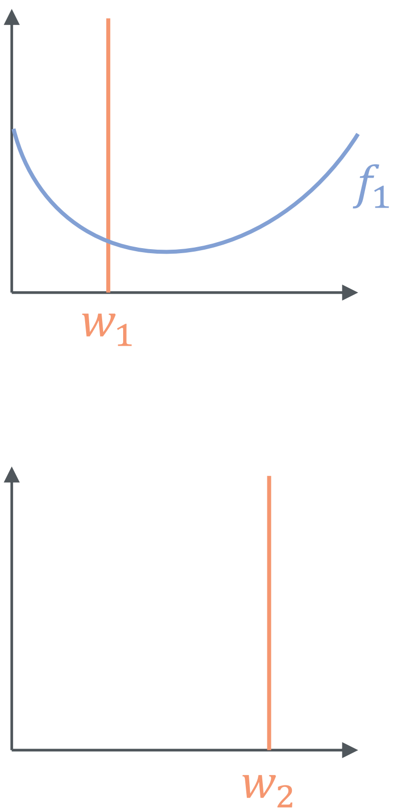

Online convex optimisation

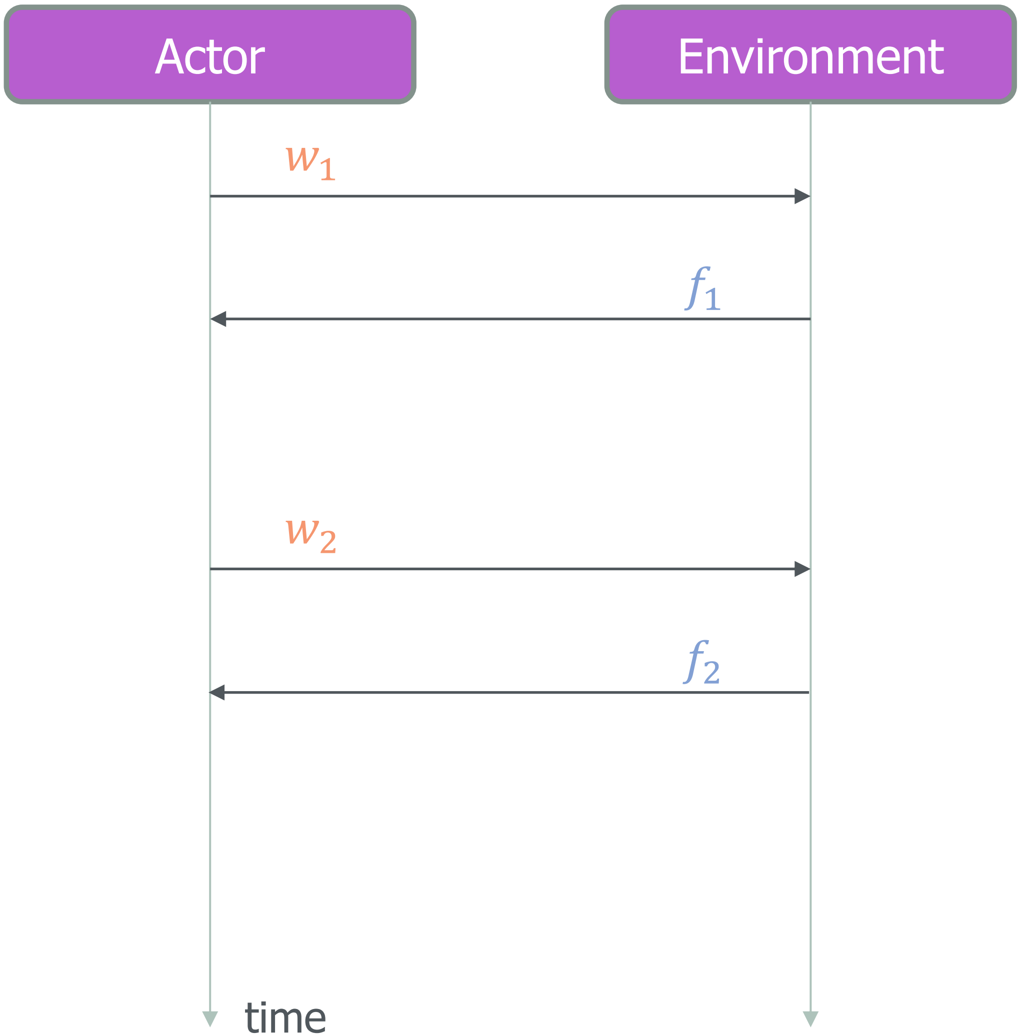

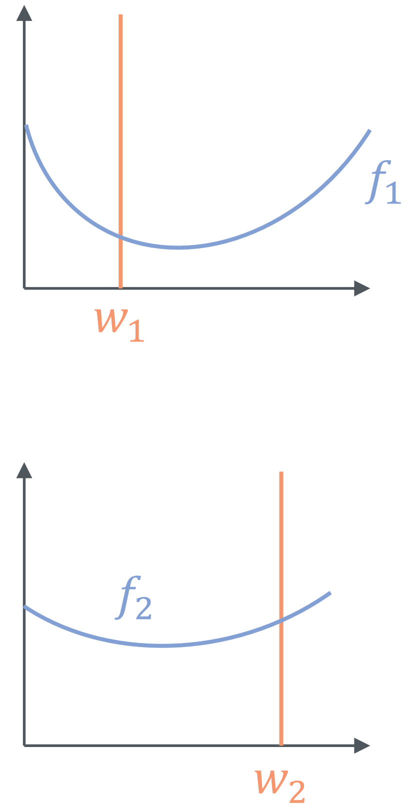

Why is this formulation relevant?

- Naturally suited for sequence prediction.

- Minimal assumptions – returns can come from any distribution, that need not be stationary.

- Environment can be adversarial!

Classes of algorithms

- How much about f_t are we told?
  - Zeroth order: $f_t(w_t)$. Bandit problem.
  - First order: $f_t(w_t), \nabla f_t(w_t)$. Most common.
  - Second order: $f_t(w_t), \nabla f_t(w_t), \nabla^2 f_t(w_t)$. etc…

- What does $\mathbb{C}$ look like? (simplex, unit ball, …)

Asset selection problem

An asset i has price p_t^i at time t.

From time $t \to t+1$, define the 'return' of the asset $r_{t+1}^i = \frac{p_{t+1}^i}{p_t^i}$.

A (long only) portfolio $\boldsymbol{x}_t \in \Delta^n$, an element of the n dimensional simplex.

The cumulative return of a sequence of such portfolios is:

$$R_T^P = \prod_{t=2}^{T} \boldsymbol{x}_{t-1} \cdot \boldsymbol{r}_t = \exp\left(\sum_{t=2}^{T} \log \boldsymbol{x}_{t-1} \cdot \boldsymbol{r}_t\right)$$

We take $-\log R_T^P$ to be the cumulative loss of the algorithm, hence at each step we incur loss

$$l_t = -\log \boldsymbol{x}_{t-1} \cdot \boldsymbol{r}_t$$

Loss and regret

Picking the ‘best’ $\boldsymbol{x}_t \forall t$ (in terms of loss minimisation) is an unrealistic goal!

Rather we define a constant optimal point, $\boldsymbol{x}^*$, in the concept of regret:

$$\mathrm{Regret}_T = \sum_{t=2}^{T} -\log \boldsymbol{x}_{t-1} \cdot \boldsymbol{r}_t - \min_{\boldsymbol{x}^*} \sum_{t=2}^{T} -\log \boldsymbol{x}^* \cdot \boldsymbol{r}_t$$

i.e. benchmark against the best *constant* portfolio with hindsight.

Algorithms can achieve guaranteed sublinear regret: typically $O\left(\sqrt{T}\right)$ or even $O(\log T)$.

Loss and regret

Finally, note that we can equivalently write:

$$\text{Regret}_T = \sum_{t=2}^{T} -\log \boldsymbol{x}_{t-1} \cdot \tilde{\boldsymbol{r}}_t - \min_{\boldsymbol{x}^*} \sum_{t=2}^{T} -\log \boldsymbol{x}^* \cdot \tilde{\boldsymbol{r}}_t$$

Where:

$$\tilde{r}_t^i = \frac{r_t^i}{\max_j r_t^j}$$

This has the useful property that $0 < \tilde{r}_t^i \leq 1 \ \forall i$.

Empirical case study

Basket of 6 energy futures markets (crude oil (x2), natural gas, gas oil, heating oil, reformulated blendstock). 25 years of market-close prices.

Each market represents a tradable *sequence* of futures contracts.

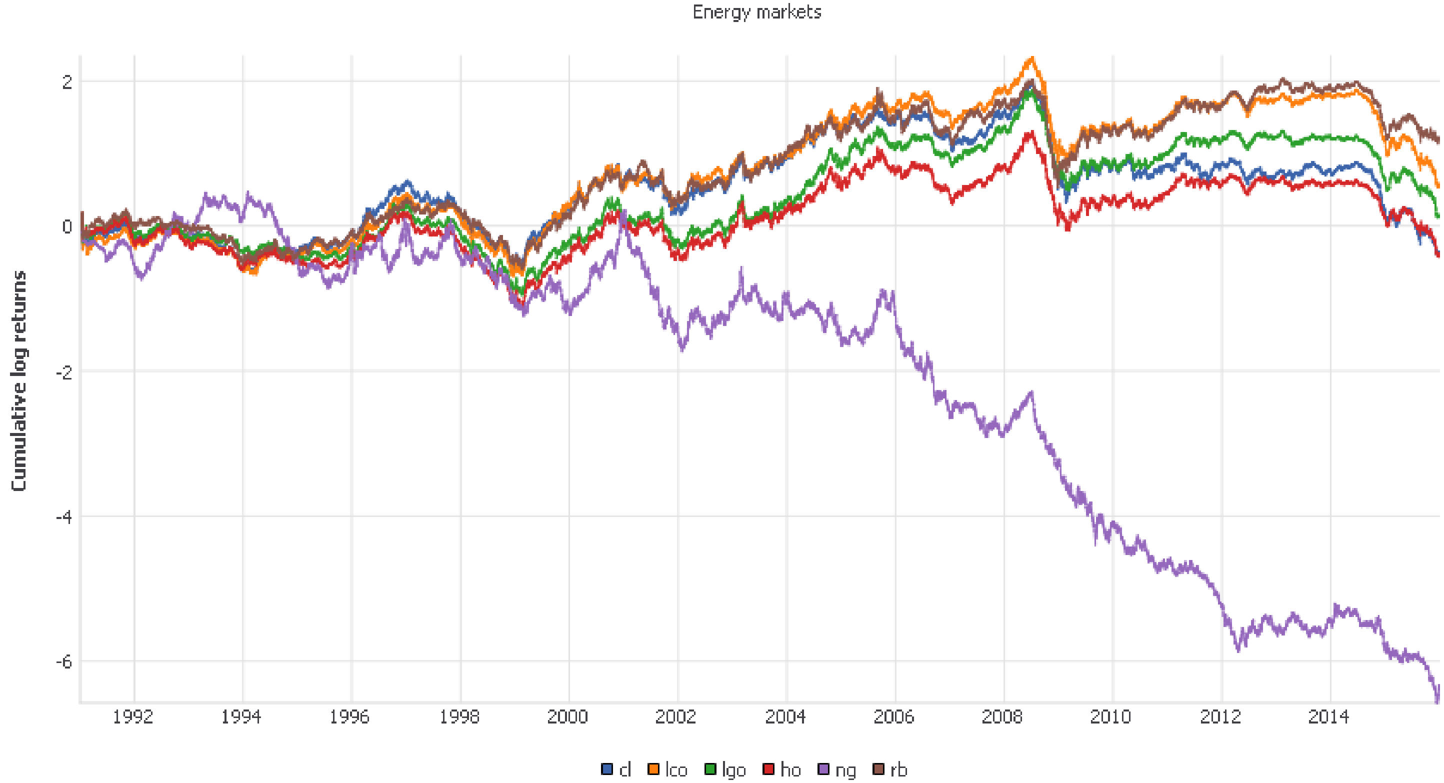

Empirical case study

Task: allocate a portfolio $x_t \in \Delta^n$ at every time step between the markets.

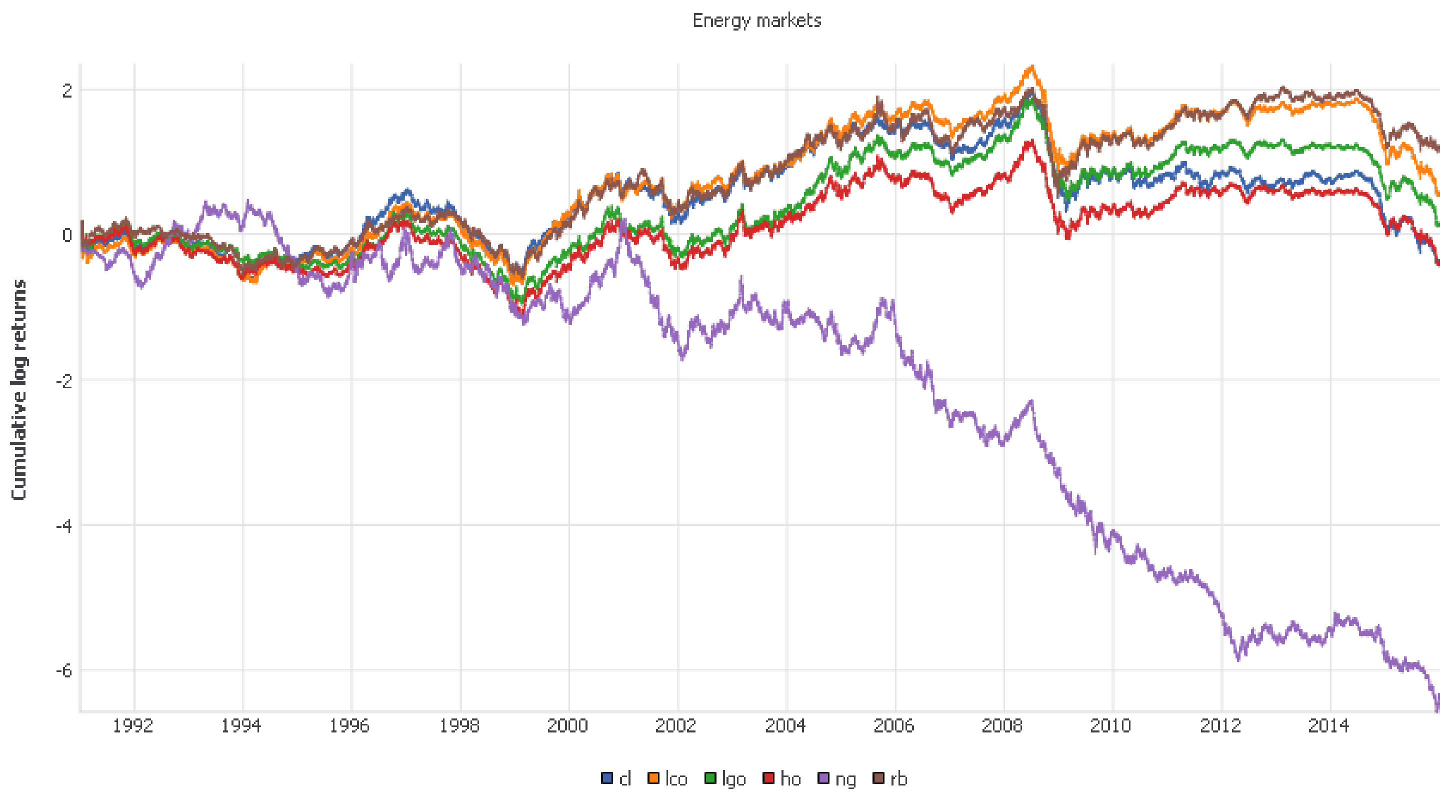

Online Gradient Descent (OGD)

Closely related to stochastic gradient descent.

Online Gradient Descent. (Zinkevich's online version of Stochastic Gradient Descent)
Inputs: convex set $K \subset \mathbb{R}^n$, step sizes $\eta_1, \eta_2, \ldots \geq 0$.

- On period 1, play an arbitrary $x_1 \in K$.
- On period $t > 1$: play

$$x_t = \Pi_K(x_{t-1} - \eta_t \nabla f_{t-1}(x_{t-1}))$$

Here, Π_K denotes the *projection* onto nearest point in K, $\Pi_K(y) = \arg\min_{x \in K} \|x - y\|$.

Source Logarithmic Regret Algorithms for Online Convex Optimization. Hazan, Agarwal, Kale. 2007.

OGD's regret bound

If we follow this algorithm with $\eta_t = \frac{1}{Ht}$, our regret is bounded by

$$\frac{G^2}{2H}(1 + \log T)$$

With the definitions for G and H:

$$\sup_{\boldsymbol{x},t} \|\nabla(-\log \boldsymbol{x} \cdot \boldsymbol{r}_t)\|_2 \leq G$$

$$\forall \boldsymbol{x}, t\colon\ (\nabla^2(-\log \boldsymbol{x} \cdot \boldsymbol{r}_t) - HI_n) \text{ is P.S.D.}$$

These show how the realised values of $\boldsymbol{r}_t$ will influence the tightness of the bound, and hence how well the algorithm is likely to perform.

OGD – empirical results

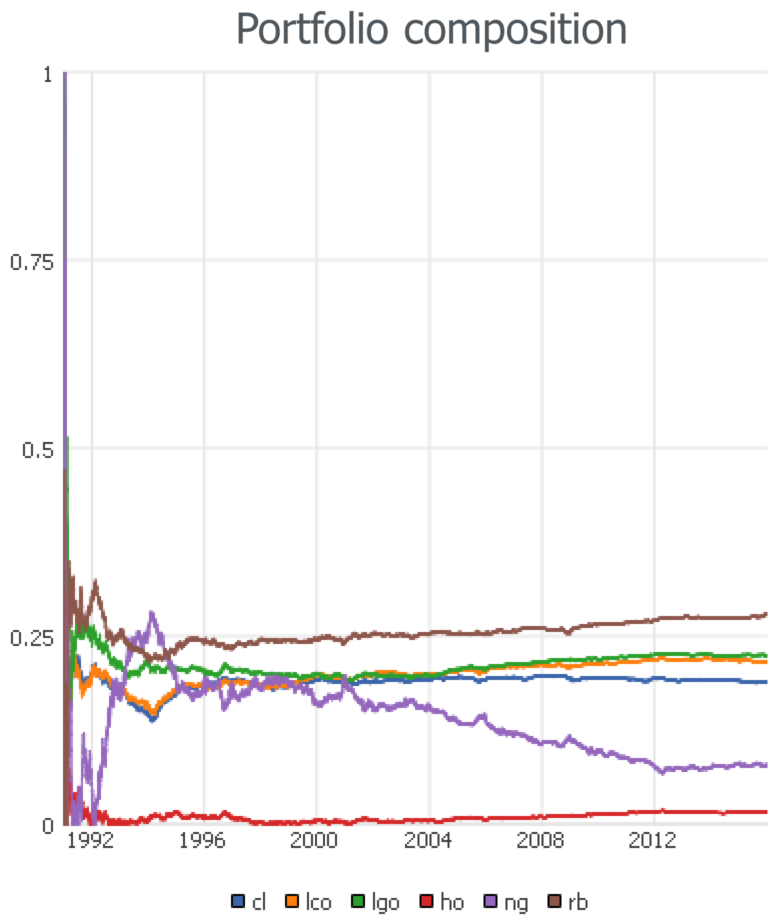

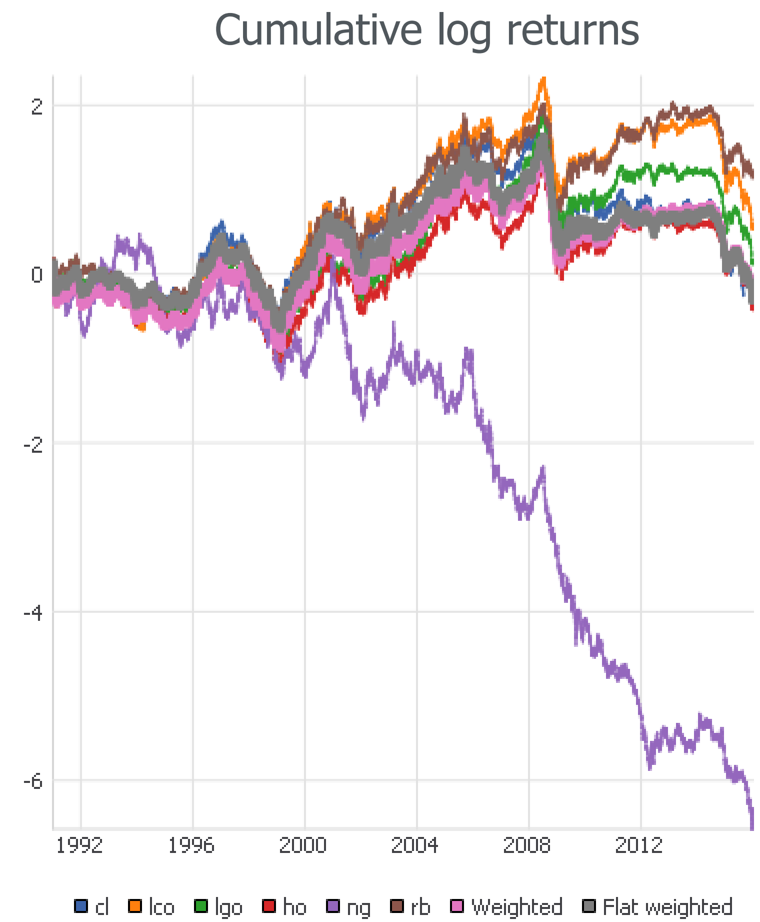

OGD – empirical results

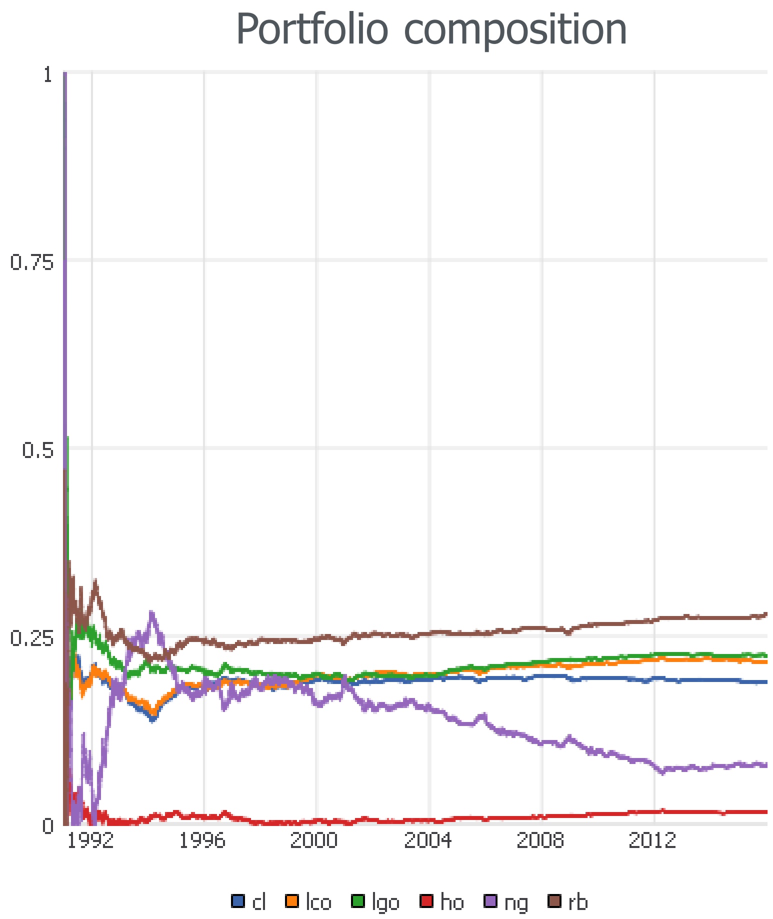

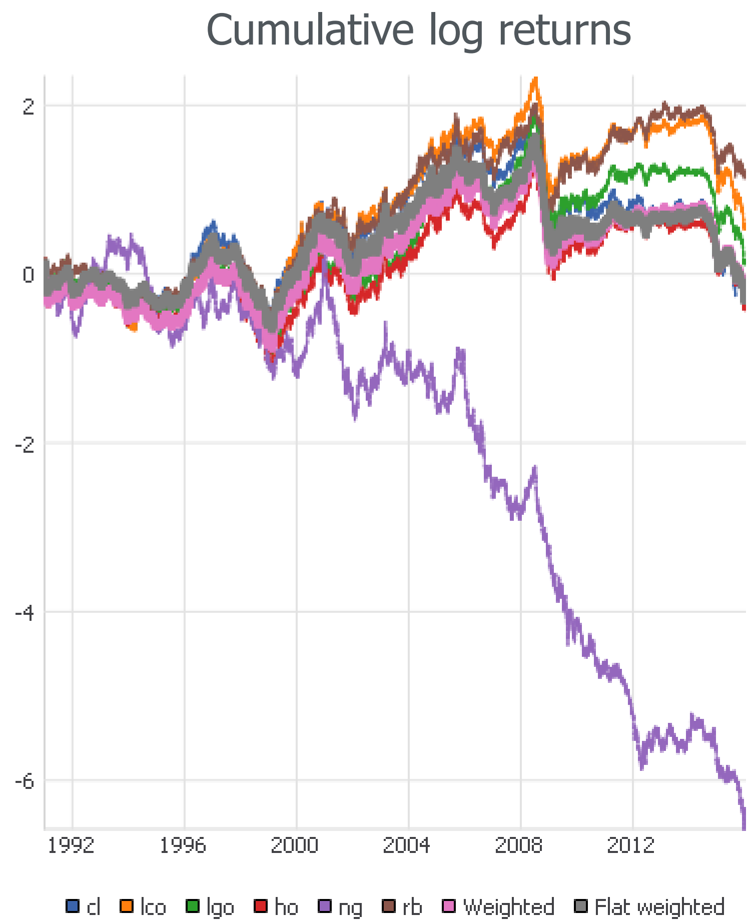

Decaying learning rate results in some lack of adaptivity.

Online Newton Step (ONS)

ONS(η, β, δ)

- On period 1, use the uniform portfolio $\mathbf{p}_1 = \frac{1}{n}\mathbf{1}$.
- On period $t > 1$: Play strategy $\tilde{\mathbf{p}}_t \triangleq (1-\eta)\mathbf{p}_t + \eta \cdot \frac{1}{n}\mathbf{1}$, such that:

$$\mathbf{p}_t = \Pi_{S_n}^{\mathbf{A}_{t-1}}\left(\delta \mathbf{A}_{t-1}^{-1}\mathbf{b}_{t-1}\right)$$

where $\mathbf{b}_{t-1} = (1 + \frac{1}{\beta}) \sum_{\tau=1}^{t-1} \nabla[\log_\tau(\mathbf{p}_\tau \cdot \mathbf{r}_\tau)]$, $\mathbf{A}_{t-1} = \sum_{\tau=1}^{t-1} -\nabla^2[\log(\mathbf{p}_\tau \cdot \mathbf{r}_\tau)] + \mathbf{I}_n$, and $\Pi_{S_n}^{\mathbf{A}_{t-1}}$ is the projection in the norm induced by $\mathbf{A}_{t-1}$, viz.,

$$\Pi_{S_n}^{\mathbf{A}_{t-1}}(\mathbf{q}) = \arg\min_{\mathbf{p}\in S_n} (\mathbf{q}-\mathbf{p})^\top \mathbf{A}_{t-1}(\mathbf{q}-\mathbf{p})$$

Figure 1. The Online Newton Step algorithm.

Source Algorithms for Portfolio Management based on the Newton Method. Agarwal, Hazan, Kale, Schapire. 2006.

ONS's regret bound

With suitable parameter choice, the regret bound is shown to be:

$$\mathrm{Regret}_T \leq \frac{10n^{1.5}}{\alpha} \log\left[\frac{nT}{\alpha^2}\right]$$

When applied to n assets with 'market variability parameter' α – this is the *minimum scaled return* one might encounter.

Regret scaling is *super-linear* in *n* .

Source Algorithms for Portfolio Management based on the Newton Method. Agarwal, Hazan, Kale, Schapire. 2006.

ONS – empirical results

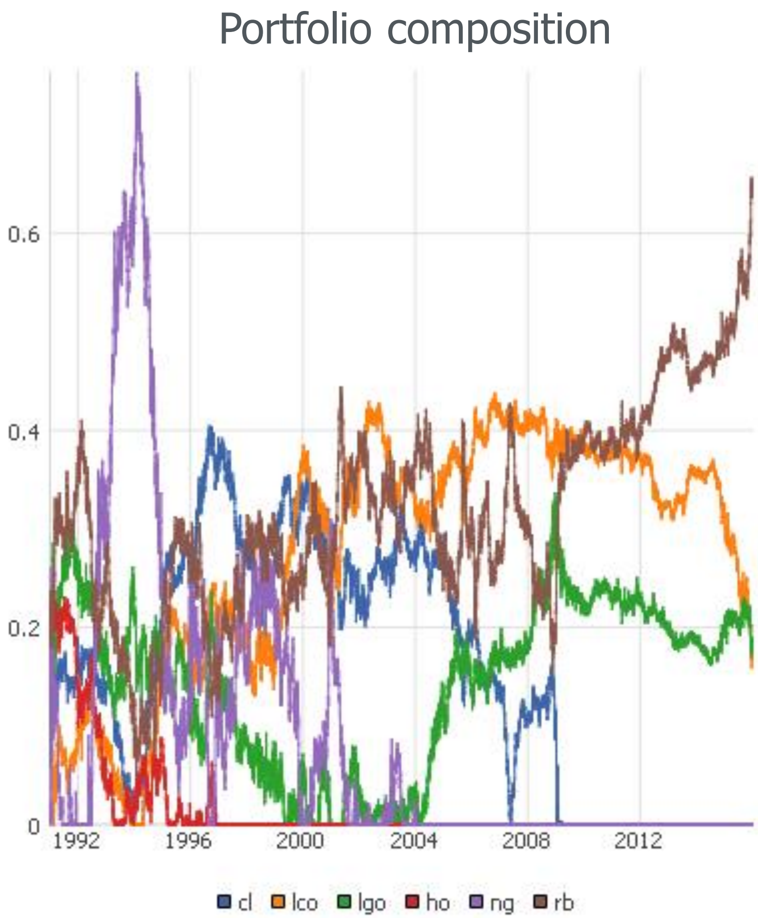

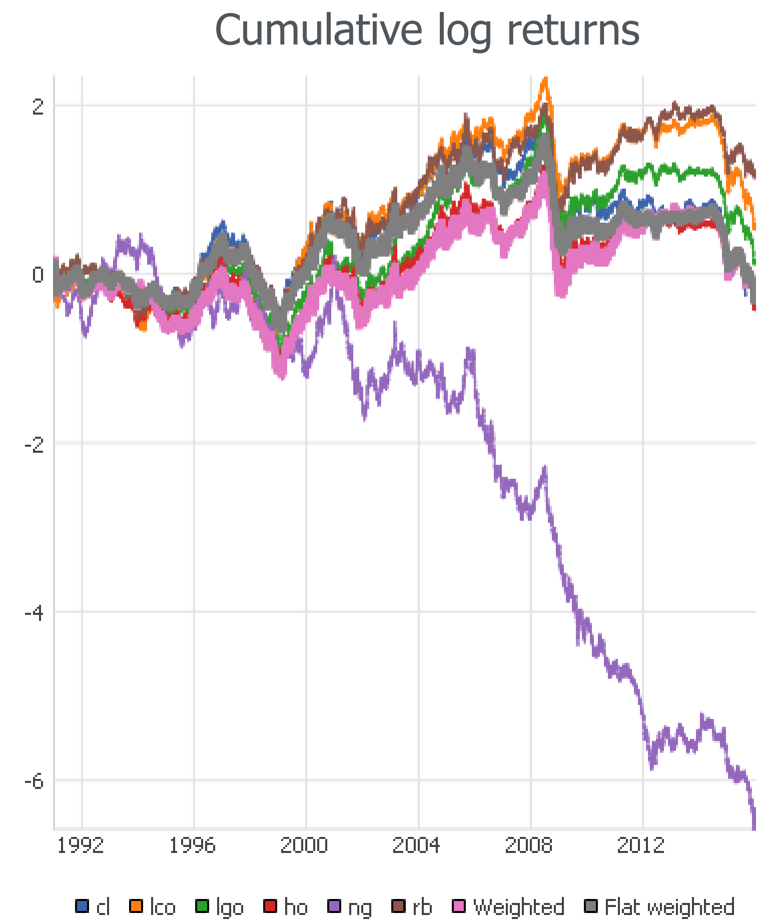

Very reactive compared to OGD. But, heavily dependent on choice of β.

Expert learning problem

Expert learning problem

- Action set is probability simplex, $\mathbb{C} = \Delta_n$.
- Loss function is inner product of action & a per-expert loss, $f_t(w_t) = w_t \cdot l_t$.
- Often restrict $l_t \in [0,1]^n$ - this is known as the "hedge setting".
- Algorithms: Multiplicative Weights, NormalHedge, AdaNormalHedge, Squint, …

Squint

- Squint is a recent algorithm (2015) to solve the Hedge setting:
  - **Parameter free**: aggregates over learning rates.
  - **Second-order bound**: adapts well to stochastic case.
  - **Quantile bound**: regret against subsets of experts, not just single experts.

Squint algorithm

Instantaneous regret vs. expert k: $r_t^k = \boldsymbol{w}_t^\intercal \boldsymbol{\ell}_t - \ell_t^k$

$$\boldsymbol{w}_{T+1} = \frac{\mathbb{E}_{\pi(k)\gamma(\eta)}\left[e^{\eta R_T^k - \eta^2 V_T^k} \eta \boldsymbol{e}_k\right]}{\mathbb{E}_{\pi(k)\gamma(\eta)}\left[e^{\eta R_T^k - \eta^2 V_T^k} \eta\right]}$$

Source Second-order Quantile Methods for Experts and Combinatorial Games. Koolen, van Erven. 2015.

Squint algorithm

Instantaneous regret vs. expert k: $r_t^k = \boldsymbol{w}_t^\intercal \boldsymbol{\ell}_t - \ell_t^k$

Regret vs. expert k

$R_T^k = \sum_{t=1}^T r_t^k$

Learning rate

$$\boldsymbol{w}_{T+1} = \frac{\mathbb{E}_{\pi(k)\gamma(\eta)}\left[e^{\eta R_T^k - \eta^2 V_T^k}\eta \boldsymbol{e}_k\right]}{\mathbb{E}_{\pi(k)\gamma(\eta)}\left[e^{\eta R_T^k - \eta^2 V_T^k}\eta\right]}$$

Prior over experts

Prior over learning rate

Cumulative variance

$V_T^k = \sum_{t=1}^T (r_t^k)^2$

Source Second-order Quantile Methods for Experts and Combinatorial Games. Koolen, van Erven. 2015.

Squint

An improper prior, $\gamma(\eta) \propto 1/\eta$, is found to give good bounds.

Weight can be written in closed form:

$$w_{T+1}^k \propto \pi(k) \int_0^{1/2} e^{\eta R_T^k - \eta^2 V_T^k} \, \mathrm{d}\eta = \pi(k) \frac{\sqrt{\pi} e^{\frac{(R_T^k)^2}{4V_T^k}} \left(\operatorname{erf}\left(\frac{R_T^k}{2\sqrt{V_T^k}} \right) - \operatorname{erf}\left(\frac{R_T^k - V_T^k}{2\sqrt{V_T^k}} \right) \right)}{2\sqrt{V_T^k}}$$

Authors kindly provide numerically stable implementation of this!

Source Second-order Quantile Methods for Experts and Combinatorial Games. Koolen, van Erven. 2015.

Squint's regret

Regret of Squint:

$$R_T^{\mathcal{K}} \leq \sqrt{2V_T^{\mathcal{K}}}\left(1+\sqrt{2\ln\left(\frac{\frac{1}{2}+\ln(T+1)}{\pi(\mathcal{K})}\right)}\right)+5\ln\left(1+\frac{1+2\ln(T+1)}{\pi(\mathcal{K})}\right)$$

$$V_T^{\mathcal{K}} = \mathbb{E}_{\pi(k|\mathcal{K})} V_T^k$$

- $\mathcal{K}$ represents a *subset* of experts – the above holds for all.
- Consistently performing experts => lower regret.
- V_T^k implicitly scales like T.

Source Second-order Quantile Methods for Experts and Combinatorial Games. Koolen, van Erven. 2015.

Applying Squint to asset selection problem

Expert learning is not directly compatible with asset selection.

However note that:

$$\tilde{\boldsymbol{r}}_t = 1 - \boldsymbol{\delta}_t$$

where $0 < \delta_t^i < 1 \;\; \forall i$.

Observe that often $\delta_t^i \ll 1$ in interesting cases, and approximate:

$$\sum_{t=2}^{T} -\log \boldsymbol{x}_{t-1} \cdot \tilde{\boldsymbol{r}}_t \approx \sum_{t=2}^{T} \boldsymbol{x}_{t-1} \cdot \boldsymbol{\delta}_t$$

We then recover the hedge setting.

Squint – empirical results

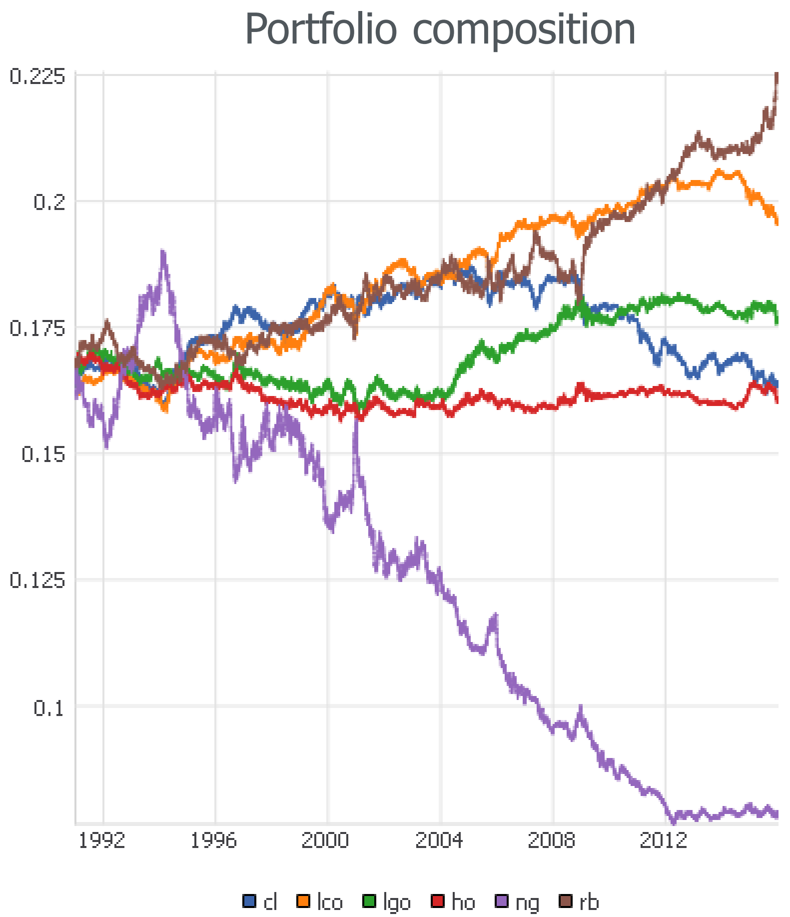

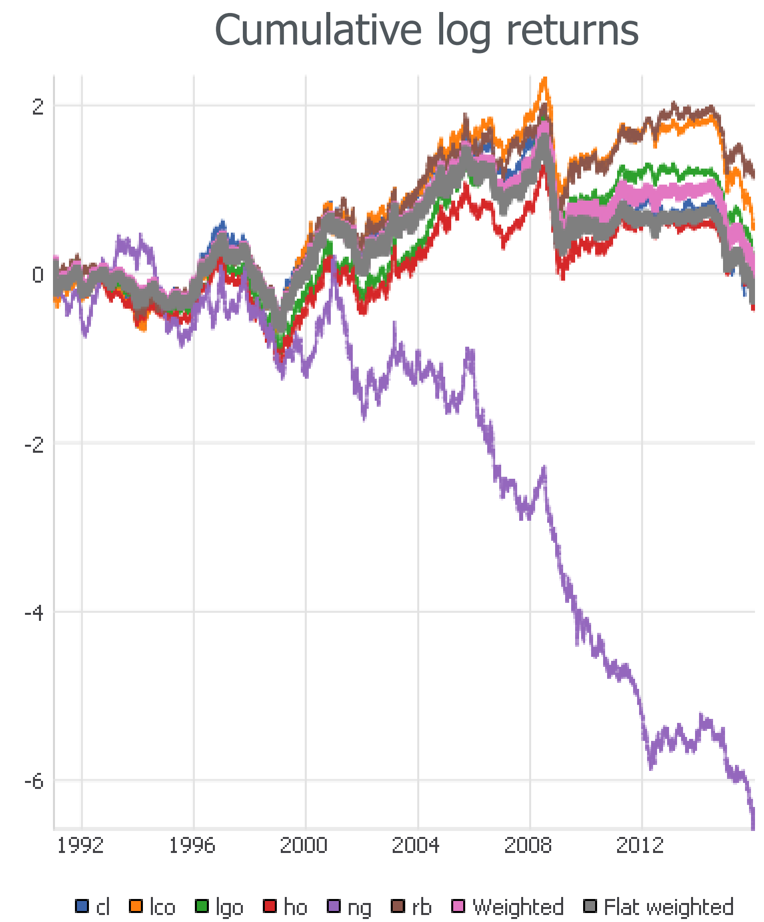

Relatively small deviations from equal weights.

Summary of all empirical results

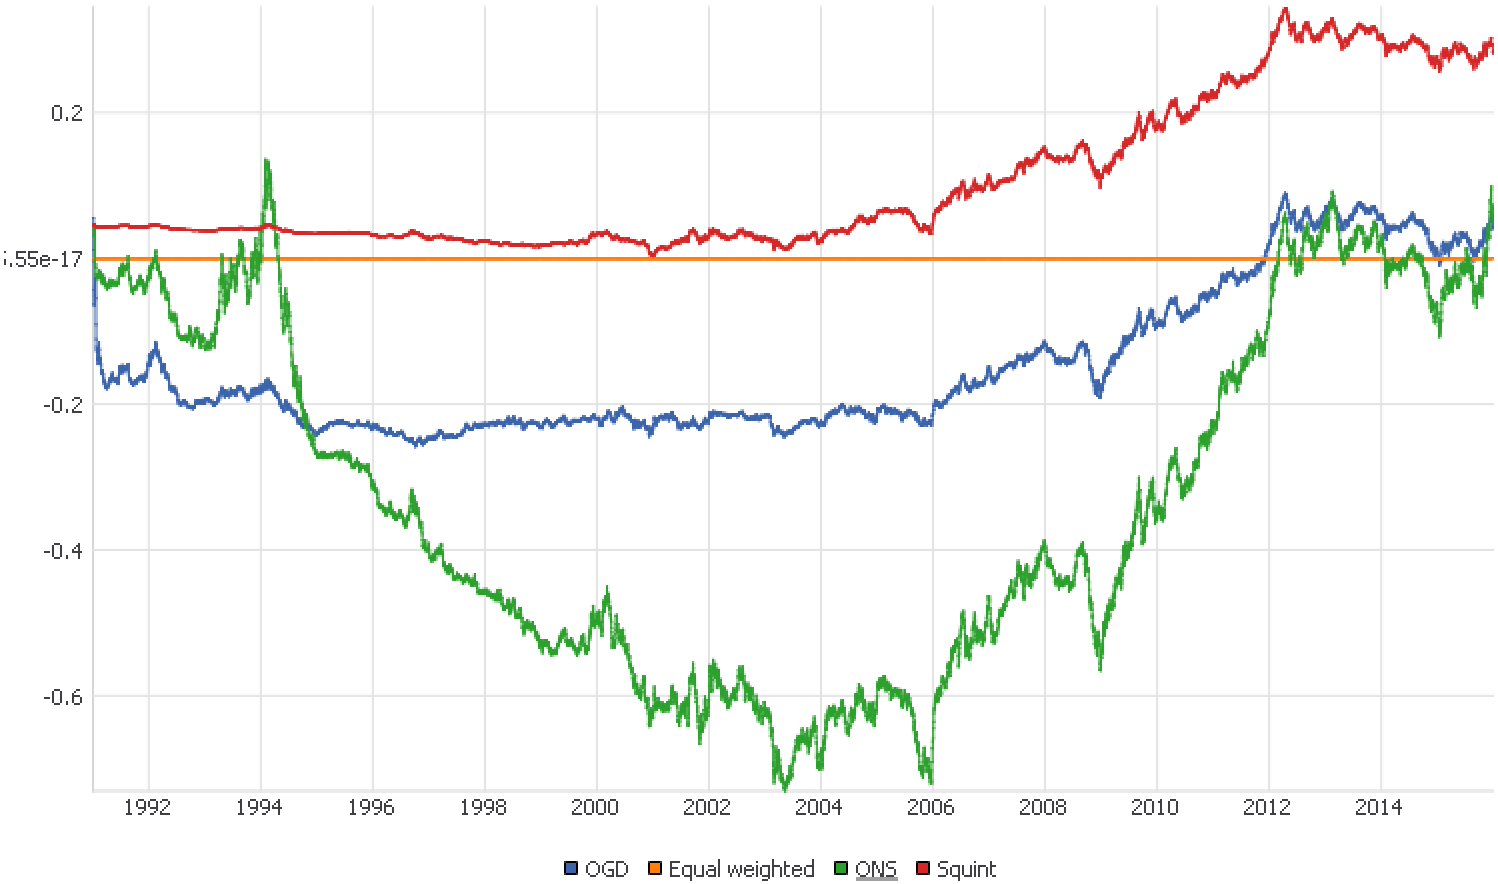

Difference of log returns between each algorithm and equal weights.

Conclusions

- Demonstrated the application of various OCO algorithms to the asset selection problem.

- Second order bounds are, in practice, useful for near-stochastic setting

- Many more useful areas:
  - Adaptive regret
  - Second-order bounds for general OCO (MetaGrad)
  - Including side information

Disclaimer

This document is issued by Cantab Capital Partners LLP ("CCP") which is authorised and regulated by the Financial Conduct Authority ("FCA"). The fund(s) under our management (the "Fund(s)") are defined as "Unregulated Collective Investment Schemes" ("UCIS") and the promotion of a UCIS either within the UK or from the UK is severely restricted by statute. Consequently, this document is only made available to professional clients and eligible counterparties as defined by the FCA and also to persons of a kind to whom the Fund may lawfully be promoted by an authorized person by virtue of Section 238(5) of the Financial Services and Markets Act 2000 and COBS 4.12.1R. Shares in the Funds ("Shares") should only be purchased by persons with experience of participating in unregulated schemes and any other person who receives this document should not rely upon it.

The value of the Shares, and any income from them, may go down as well as up and an investor may not receive back, on redemption of his Shares, the amount which was invested. Commodity trading involves substantial risk of loss. Changes in rates of exchange between the US Dollar and the currencies in which the investments of the Fund are denominated may cause the value of the Shares to go up or down. Past results are not necessarily indicative of future results. The Shares will not be dealt in on a recognised or designated investment exchange for the purposes of the Act, nor will there be a market maker in the Shares, and it may therefore be difficult for an investor to dispose of his Shares otherwise than by way of redemption or to obtain reliable information about the extent of the risks to which his investment is exposed.

This document does not constitute or form part of any offer to issue or sell, or any solicitation of any offer to subscribe or purchase, the Shares, nor shall it or the fact of its distribution form the basis of or be relied on in connection with, any contract therefore. Recipients of this document who intend to apply for Shares following the publication of the prospectus to be issued by the Fund are reminded that any such application must be made solely on the basis of the information and opinions contained in the prospectus which may be different from the information and opinions contained herein.

Neither CCP, nor its partners or employees warrant the accuracy, adequacy or completeness of the information contained herein and CCP expressly disclaims liability for errors or omissions in such information. No warranty of any kind implied, express or statutory is given by CCP or any of its directors or employees in connection with the information contained herein. Under no circumstances may this document, or any part thereof, be copied, reproduced or redistributed without the express permission of a partner of CCP.

RULE 4.7 DISCLOSURE: PURSUANT TO AN EXEMPTION FROM THE COMMODITY FUTURES TRADING COMMISSION IN CONNECTION WITH POOLS WHOSE PARTICIPANTS ARE LIMITED TO QUALIFIED ELIGIBLE PERSONS, AN OFFERING MEMORANDUM FOR THIS POOL IS NOT REQUIRED TO BE, AND HAS NOT BEEN, FILED WITH THE COMMISSION. THE COMMODITY FUTURES TRADING COMMISSION DOES NOT PASS UPON THE MERITS OF PARTICIPATING IN A POOL OR UPON THE ADEQUACY OR ACCURACY OF AN OFFERING MEMORANDUM. CONSEQUENTLY, THE COMMODITY FUTURES TRADING COMMISSION HAS NOT REVIEWED OR APPROVED THIS OFFERING OR ANY OFFERING MEMORANDUM FOR THIS POOL.

YOU SHOULD ALSO BE AWARE THAT THIS COMMODITY POOL MAY TRADE FOREIGN FUTURES OR OPTIONS CONTRACTS. TRANSACTIONS ON MARKETS LOCATED OUTSIDE THE UNITED STATES, INCLUDING MARKETS FORMALLY LINKED TO A UNITED STATES MARKET, MAY BE SUBJECT TO REGULATIONS WHICH OFFER DIFFERENT OR DIMINISHED PROTECTION TO THE POOL AND ITS PARTICIPANTS. FURTHER, UNITED STATES REGULATORY AUTHORITIES MAY BE UNABLE TO COMPEL THE ENFORCEMENT OF THE RULES OF REGULATORY AUTHORITIES OR MARKETS IN NON-UNITED STATES JURISDICTIONS WHERE TRANSACTIONS FOR THE POOL MAY BE EFFECTED.

This document and the information contained within it are highly confidential and intended solely for the use of the individual or entity to whom they are addressed or those who have registered with Cantab. You should not disseminate, distribute or copy this document or any of the information contained within it. If you are not the intended recipient, you are notified that disclosing, copying, distributing or taking any action in reliance on the contents of this information is strictly prohibited.

Registered in England No. OC317557. Registered office: City House, 126-130 Hills Road, Cambridge, CB2 1RE © Cantab Capital Partners LLP

Backup

Adaptive regret

Regret: $$R_T = \sum_{t=1}^{T} f_t(w_t) - \min_{u \epsilon \mathbb{C}} \sum_{t=1}^{T} f_t(u)$$

Adaptive regret: $$\sup_{I=[r,s]\subseteq[T]} \left\{ \sum_{t=r}^{s} f_t(w_t) - \min_{u_I \epsilon \mathbb{C}} \sum_{t=\mathrm{r}}^{s} f_t(u_I) \right\}$$

There are methods for adapting regret-minimising algorithms to target adaptive regret.

Some history

- **1991**: Cover. Universal portfolios.
- **2007**: Hazan. Online Newton Step.
- **2009**: Hazan. Adaptive regret.
- **2016**: van Erven, Koolen. MetaGrad, Squint.
- **2017**: Veness et. al: Gated linear networks.