Expected utility maximization under incomplete information and with Cox-process observations

Wolfgang Runggaldier
(joint with K.Fujimoto and H.Nagai)

Dipartimento di Matematica, Università di Padova

Conference on Innovative Research in Mathematical Finance in Honor of Yuri Kabanov,
Luminy, September 2018

Outline

- Description of the model and the objective.
- Remarks on the problem setup and on the approach.
- Preliminaries for the main result.
- Main result:
  - i) An approximation result leading to a "value iteration"-type algorithm analogous to that for infinite horizon discounted Markov decision problems;
  - ii) a general dynamic programming principle

The model

- Given are m risky assets with prices S_t^i satisfying

$$dS_t^i = S_t^i\{r^i(\theta_t)dt + \sum_j \sigma_j^i(\theta_t)dB_t^j\}$$

  and let $X_t^i := \log S_t^i$.
- θ_t : is a hidden finite-state Markovian factor process

  $$d\theta_t = Q^*\theta_t dt + dM_t, \quad \theta_t \in E := \{e_1, \cdots, e_N\}, \ \theta_0 = \xi \in E$$

  Q : transition intensity matrix; M_t : jump martingale on $(\Omega, \mathcal{F}, \mathcal{F}_t, P)$.

$\rightarrow$ $p_t := (p_t^1, \cdots, p_t^N)$: state-probability vector, $p_t^i = P\{\theta_t = e_i\}$

The model

- Given is also a non-risky asset with price S_t^0 satisfying

$$dS_t^0 = r_0 S_t^0 dt$$

- Let $\tilde{S}_t^i := \frac{S_t^i}{S_t^0}$, with $\tilde{X}_t^i := \log \tilde{S}_t^i$ so that

$$d\tilde{X}_t^i = \{r^i(\theta_t) - r_0 - d(\sigma\sigma^*(\theta_t))^i\}dt + \sum_{j=1}^{m} \sigma_j^i(\theta_t)dB_t^j$$

with $d(\sigma\sigma^*(\theta)) = (\frac{1}{2}(\sigma\sigma^*)^{11}(\theta), \ldots, \frac{1}{2}(\sigma\sigma^*)^{mm}(\theta))$ *(column vector)*.

The model

Prices *(and thus also the logarithms of their discounted values)* are only observed at the random times $\tau_0, \tau_1, \tau_2, \cdots$ so that, putting $\tilde{X}_k^i := \tilde{X}_{\tau_k}^i$, the observations $(\tau_k, \tilde{X}_k)$ form a multivariate marked point process with counting measure

$$\mu(dt, dx) = \sum_k \mathbf{1}_{\{\tau_k < \infty\}} \delta_{\{\tau_k, \tilde{X}_k\}}(t, x) dtdx$$

→ The corresponding counting process

$$\Lambda_t := \int_0^t \int_{\mathbb{R}^m} \mu(dt, dx)$$

is supposed to be a Cox process with intensity $n(\theta_t)$, i.e.

$$\Lambda_t - \int_0^t n(\theta_s) ds \quad \text{is an} \quad (\mathcal{F}_t, P) - \text{martingale.}$$

The model

- Random time observations are more realistic in comparison with diffusion-type models, especially on small time scales: prices do not vary continuously but by tick-size at random times in reaction to *arrival of significant new information*.
- Restricting observations and trading to random times corresponds to the fact that portfolios cannot be re-balanced continuously: think of transaction costs and/or liquidity restrictions (*see Pham, Tankov 08/09 for a case of full observations*).
- The partial information setup allows for continuous updating of the underlying model.

Investment strategies, portfolios

- N_t^i : number of assets of type i in the portfolio at time t :

$$N_t^i = \sum_k \mathbf{1}_{[\tau_k, \tau_{k+1})}(t) N_{\tau_k}^i$$

- The wealth process at time t is then $V_t := \sum_{i=0}^m N_t^i S_t^i$. and the investment ratios

$$h_t^i := \frac{N_t^i S_t^i}{V_t}, \qquad (h_k^i := h_{\tau_k}^i)$$

are defined on

$$\bar{H}_m := \{(h^1, \ldots, h^m); h^1 + h^2 + \ldots + h^m \leq 1, 0 \leq h^i, i = 1, 2, \ldots, m\}$$

$\rightarrow$ *No shortselling is allowed and $\bar{H}_m$ is closed and bounded.*

Investment strategies, portfolios

The dynamics of a self-financing portfolio are ($h_t \in \bar{H}_m$)

$$dV_t = V_t \left\{ h_t' \{ r(\theta_t)\, dt + h_t' \sigma(\theta_t) dB_t \right\}$$

→ Defining $\gamma : \mathbb{R}^m \times \bar{H}_m \to \bar{H}_m$ by

$$\gamma^i(z,h) := \frac{h^i \exp(z^i)}{1 + \sum_{i=1}^{m} h^i(\exp(z^i) - 1)}, \qquad i = 1, \ldots, m$$

one has that, for $t \in [\tau_k, \tau_{k+1})$,

$$h_t^i = \gamma^i(\tilde{X}_t - \tilde{X}_k, h_k)$$

→ h_t is thus determined by $h_k, \tilde{X}_k, \tilde{X}_t$ where $\tilde{X}_t$ is unobserved for $t \in (\tau_k, \tau_{k+1})$.

Investment strategies, portfolios

- The set $\mathcal{A}$ of admissible strategies is

$$\begin{aligned} \mathcal{A} &:= \{\{h_k\}_{k=0}^{\infty} | h_k \in \bar{H}_m,\ \mathcal{G}_k \text{ measurable}\} \\ &\text{where} \\ \mathcal{G}_k &:= \mathcal{F}_0 \vee \sigma\{\tau_0, \tilde{X}_0, \tau_1, \tilde{X}_1, \tau_2, \tilde{X}_2, \ldots, \tau_k, \tilde{X}_k\} \end{aligned}$$

- For $n > 0$ let

$$\mathcal{A}^n := \{h \in \mathcal{A} | h_{n+i} = h_{\tau_{n+i}-} \text{ for all } i \geq 1\}$$

$\rightarrow$ For $h \in \mathcal{A}^n$ one has $N_{n+k} = N_{n+k-1} = N_n$

$\rightarrow$ $\mathcal{A}^0 \subset \mathcal{A}^1 \subset \cdots \mathcal{A}^n \subset \mathcal{A}^{n+1} \cdots \subset \mathcal{A}$.

log-utility

- Considering a log-utility and recalling

$$dV_t = V_t \left\{ h_t' \{ r(\theta_t)\, dt + h_t' \sigma(\theta_t) dB_t \right\},$$

for any given $T > 0$ one has

$$\begin{aligned} \log V_T &= \log v_0 + \int_0^T h_t' \sigma(\theta_t) dB_t \\ &\quad + \int_0^T h_t' r(\theta_t) - \tfrac{1}{2} h_t' \sigma\sigma'(\theta_t) h_t] dt \\ &= \log v_0 + \int_0^T h_t' \sigma(\theta_t) dB_t + \int_0^T f(\theta_t, h_t)\, dt \end{aligned}$$

having put $f(\theta, h) := h' - \frac{1}{2} h' \sigma\sigma'(\theta) h$

$\rightarrow$ *Our problem can now be formulated as follows*

The problem (log-utility)

Problem (Log-utility): Given a finite planning horizon $T > 0$, determine the optimal value

$$\sup_{h \in \mathcal{A}} E \{ \log V_T | \tau_0 = 0, p_0 = p \}$$

$$= \log v_0 + \sup_{h \in \mathcal{A}} E \left\{ \int_0^T f(\theta_t, h_t) dt | \tau_0 = 0, p_0 = p \right\}$$

as well as an optimal maximizing strategy

$$\hat{h} \in \mathcal{A}$$

→ The optimal strategy maximizes the expected log-value obtained at a fixed terminal time (T *is not considered a rebalancing/liquidation time*).

Remarks on problem setup

- → A stochastic control problem under incomplete information.
- → Standard approach: transform them into a complete information problem, the so-called ”separated problem”, where instead of the unobservable quantities one considers their distributions, conditional on the observations.

This requires:

i) solving the associated filtering problem;

ii) formulating the separated problem so that its solution is indeed a solution of the original incomplete information problem.

Remarks on problem setup

- The associated filtering problem has been solved in work by Cvitanic, Liptser, Rozovskii and it was found that *"the given problem does not fit into a standard diffusion or point process filtering framework"*.
- Not only the filtering problem, but also the control part of the problem does not fit into any standard framework and so there remained the task to find an approach also for the control part.

Remarks on problem setup

- Our problem is defined over a finite horizon, but the number of transitions is random, possibly infinite and so it becomes intuitive to look for an *algorithm analogous to those for infinite horizon Markovian decision problems (e.g. Value Iteration)*.
- We show that also in our setup, which is intermediate between continuous and discrete time, one can obtain results that are analogous to the classical ones, in particular, we also obtain myopic optimal policies.

This can however not be shown directly as in the classical cases (*the number of observation/rebalancing times* τ_k *up to the horizon* T *is a.s. finite, but their number depends on* ω *and is not bounded from above*.

Filtering

To summarize the filtering results in Cvitanic, Liptser, Rozovskii (2006), denote by $\pi_t(f) = E[f(\theta_t)|\mathcal{G}_t]$ the filter of $f(\theta_t)$ given $\mathcal{G}_t$ with $\mathcal{G}_t := \mathcal{F}_0 \vee \sigma\{\mu((0,s] \times B) : s \leq t, B \in \mathcal{B}(\mathbb{R}^m)\}$.

- $\rightarrow$ Being $\theta_t \in \{e_1, \cdots, e_N\}$, we have $f(\theta_t) = \sum_i f(e_i)\mathbf{1}_{e_i}(\theta_t)$. It thus suffices to consider $\pi_t^i = \pi_t\left(\mathbf{1}_{e_i}(\theta_t)\right)$
- $\rightarrow$ Since the observations take place only along $\tau_1, \tau_2, \cdots$, useful information also arrives only along that sequence and we have

$$\pi^i_{\tau_{k+1}} = M^i\left(\tau_{k+1} - \tau_k, \tilde{X}_{\tau_{k+1}} - \tilde{X}_{\tau_k}, \pi_{\tau_k}\right)$$

for suitable functions $M^i(\cdot)$ and with $\pi_{\tau_k} := \left(\pi^1_{\tau_k}, \cdots, \pi^N_{\tau_k}\right)$

Filtering

Putting $\pi_k = \pi_{\tau_k}$, we obtain the Markov process $\left\{\tau_k, \pi_{\tau_k}, \tilde{X}_{\tau_k}\right\}_{k=1}^{\infty}$ with respect to $\mathcal{G}_k$ that will turn out to be the state process for the "separated" (completely observed) control problem.

Preliminaries to value iteration

- Considering the expected log-utility at a generic time $t \in [0, T]$ with $p = p_t$ we may write

$$E\{\log V_T|\tau_0 = t, p_0 = p\}$$
$$= \log V_t + E\left\{\int_t^T f(\theta_s, h_s)ds|\tau_0 = t, p_0 = p\right\}$$

$\rightarrow$ What matters is

$$E\left\{\int_t^T f(\theta_s, h_s)ds|\tau_0 = t, \pi_0 = \pi\right\}$$

(*notice the π instead of p*)

Preliminaries to value iteration

A standard approach to Value Iteration is to consider, for an admissible strategy $h. = (h_t)$, the value function

$$W(t,\pi,h.) := E\left\{\int_t^T f(\theta_s,h_s)ds | \tau_0 = t,\ \pi_0 = \pi\right\}$$

and put

$$W(t,\pi) \quad := \quad \sup_{h\in\mathcal{A}} W(t,\pi,h)$$

$$W^n(t,\pi) \quad := \quad \sup_{h\in\mathcal{A}^n} W(t,\pi,h)$$

→ Working directly with the above leads to various difficulties and so we consider a modified approach via auxiliary value functions.

A contraction operator

- Let

$$\mathcal{S}_N := \left\{ p \in \mathbb{R}^N \mid \sum_{i=1}^{N} p^i = 1\,;\ 0 \le p^i\ ,\ i = 1, \cdots, N \right\}$$

so that also $\pi_t = (\pi_t^1, \cdots, \pi_t^N) \in \mathcal{S}_N$.

- On $\mathcal{S}_N$ consider the Hilbert metric

$$d_H(\pi, \bar{\pi}) := \log \left(\sup_{\bar{\pi}(A)>0, A \subset E} \frac{\pi(A)}{\bar{\pi}(A)} \sup_{\pi(A)>0, A \subset E} \frac{\bar{\pi}(A)}{\pi(A)} \right)$$

and put $\Sigma := [0, \infty) \times \mathcal{S}_N$.

A contraction operator

- Let $C_b(\Sigma)$ be the set of bounded continuous functions $g : \Sigma \to \mathbb{R}$ with norm $\|g\| := \max_{x\in\Sigma} | g(x) |$.
- Let $C_{b,lip}(\Sigma)$ be the set of bounded and Lipschitz continuous functions $g : \Sigma \to \mathbb{R}$ with norm

$$N^\lambda(g) := \lambda\|g\| + [g]_{lip}$$

where

$$[g]_{lip} := \sup_{\tau,\bar\tau,\,\pi,\bar\pi\in\mathcal{S}_N} \frac{|g(\tau,\pi) - g(\bar\tau,\bar\pi)|}{|\tau-\bar\tau| + d_H(\pi,\bar\pi)}$$

$\rightarrow$ $C_{b,lip}(\Sigma)$ is a Banach space with norm $N^\lambda(g)$, $\forall\lambda > 0$.

A contraction operator

Definition: Let $J: C_b(\Sigma) \to C_b(\Sigma)$ be the operator

$$Jg(\tau, \pi) = E\left\{g(\tau_1 \wedge T, \pi_{\tau_1} \wedge T) \mid \tau_0 = \tau, \pi_0 = \pi\right\}$$

Lemma 1: J is a contraction operator on $C_b(\Sigma)$ with contraction constant $c := 1 - e^{-\bar{n}T} < 1$, where $\bar{n} := \max n(\theta) = \max_i n(e_i)$.

Lemma 2: J is a contraction operator on $C_{b,lip}(\Sigma)$ having contraction constant $c' := (c + \max(\bar{n}, \frac{2}{\log 3})\frac{1}{\lambda})$ with λ large enough so that $c' < 1$.

Preliminaries to the optimal strategy

Recalling

$$E\{\log V_T|\tau_0 = t, \pi_0 = \pi\}$$
$$= \log V_t + E\left\{\int_t^T f(\theta_s, h_s)ds|\tau_0 = t, \pi_0 = \pi\right\}$$

let

Definition:

$$\hat{C}(\tau, \pi, h) = E\left\{\int_\tau^{T\wedge\tau_1} f(\theta_s, h_s)ds|\tau_0 = \tau, \pi_0 = \pi\right\}$$

Preliminaries to the optimal strategy

Lemma: We have

i) $E\left\{\int_t^T f(\theta_s, h_s)ds|\tau_0 = t, \pi_0 = \pi\right\}$
$$= E\left\{\sum_k \hat{C}(\tau_k, \pi_k, h_k)1_{\{\tau_k < T\}}|\tau_0 = t, \pi_0 = \pi\right\}$$

ii) $\hat{C}$ is bounded and continuous on $[0, T] \times \mathcal{S}_N \times \bar{H}_m$ for the metric $|t - \bar{t}| + d_H(\pi, \bar{\pi}) + \sum_{i=1}^m |h^i - \bar{h}^i|$

iii) $\exists\, \hat{h}(\tau, \pi)$ measurable s.t. $\sup_{h \in \bar{H}_m} \hat{C}(\tau, \pi, h) =$
$= \hat{C}(\tau, \pi, \hat{h}(\tau, \pi)) := C(t, \pi)$

iv) $C(t, \pi)$ is Lipschitz for the metric on $[0, T] \times \mathcal{S}_N$.

Preliminaries to the optimal strategy

The result in the previous point i) is rather crucial: the various h_k are chosen at the various τ_k but the objective function depends on h_t also between observation times, which in turn depends on the unobservable θ_t and X_t between observation times.

$\rightarrow$ *The optimal strategy will turn out to be myopic and given by a maximizer of the individual terms in the sum on the RHS in i). Due to the infinite sum, this however does not follow directly.*

$\rightarrow$ Next we describe our procedure and the results.

Preliminaries to value function

Definition: Based on the contraction property of J on $C_{b,lip}(\Sigma)$ with norm $N^\lambda(\cdot)$, let

$$\bar{W}(t,\pi) := \sum_{k=0}^{\infty} J^k C(t,\pi)$$

Lemma: $\bar{W} \in C_{b,lip}(\Sigma)$ and

$$\bar{W}(t,\pi) = C(t,\pi) + J\bar{W}(t,\pi)$$

"Value iteration" for the auxiliary value function

- Recall that, for $t \in [\tau_k, \tau_{k+1})$ we have $h^i_t = \gamma^i(\tilde{X}_t - \tilde{X}_k, h_k)$

Definition: For $h \in \bar{H}_m$ let (*no rebalancing*)

$$\bar{W}^0(t, \pi, h) := E\left\{ \int_t^T f(\theta_s, \gamma(\tilde{X}_s - \tilde{X}_t, h)) ds | \tau_0 = t, \ \pi_0 = \pi \right\}$$

which is bounded and continuous.

"Value iteration" for the auxiliary value function

Define, then, recursively,

$$\begin{aligned}\bar{W}^0(t,\pi) &:= \max_{h\in\bar{H}_m} \bar{W}^0(t,\pi,h)\\ \bar{W}^n(t,\pi) &:= C(t,\pi) + J\bar{W}^{n-1}(t,\pi)\\ &= \sum_{k=0}^{n-1} J^k C(t,\pi) + J^n \bar{W}^0(t,\pi)\end{aligned}$$

thereby recalling that

$$C(t,\pi) = \sup_{h\in\bar{H}_m} \hat{C}(\tau,\pi,h)$$

$$Jg(\tau,\pi) = E\left\{g(\tau_1,\pi_1)1_{\{\tau_1<T\}} \middle| \tau_0 = \tau, \pi_0 = \pi\right\}$$

(notice that $J^0 C(t,\pi) = C(t,\pi)$).

Main theorem

- "Approximation theorem". Given $\epsilon > 0$, let $n_\epsilon := (\log(1 - c') + \log \varepsilon - \log N^\lambda(\bar{W}^1 - \bar{W}^0)) / \log c'$. Then

$$N^\lambda(W - \bar{W}^n) < \epsilon \quad \forall\, n \geq n_\epsilon$$

i.e. the recursive algorithm for computing $\bar{W}^n$ is a "value iteration algorithm" for the actual optimal value function W.

→ Being $N^\lambda(W - \bar{W}^n) = \lambda \|W - \bar{W}^n\| + [W - \bar{W}^n]_{lip}$ the above implies that $W - \bar{W}^n$ is small for all (t, π) and does not vary abruptly.

Main theorem

- "Dynamic Programming Principle" (*concerns the actual optimal value function*). For any $n > 0$

$$W(t,\pi) = \sup_{h\in\mathcal{A}^n} E\left\{\sum_{k=0}^{n} \hat{C}(\tau_k,\pi_k,h_k)1_{\{\tau_k<T\}} + W(\tau_{n+1},\pi_{n+1})1_{\{\tau_{n+1}<T\}} \middle| \tau_0 = t, \pi_0 = \pi\right\}$$

Main theorem (contd.)

3. Optimal value and optimal strategy

- Given $V_0 = v_0, \tau_0 = 0, \pi_0 = \pi$ we have

$$\sup_{h \in \mathcal{A}} E\left\{\log V_T | \tau_0 = 0, \pi_0 = \pi\right\}$$
$$= \log v_0 + \sup_{h \in \mathcal{A}} E\left\{\int_0^T f(\theta_t, h_t)dt | \tau_0 = 0, \pi_0 = \pi]\right\}$$
$$= \log v_0 + C(0, \pi) + \sum_{k=1}^{\infty} E\left\{C(\tau_k, \pi_k) 1_{\{\tau_k < T\}} | \tau_0 = 0, \pi_0 = \pi\right\}$$

Main theorem (contd.)

- The optimal strategy is given by

i) for $t = \tau_k$: $\hat{h}_k = \hat{h}(\tau_k, \pi_{\tau_k})$ such that

$$C(t,\pi) = \sup_{h \in \bar{H}_m} \hat{C}(\tau,\pi,h) = \hat{C}(\tau,\pi,\hat{h}(\tau,\pi))$$

ii) for $t \in [\tau_k, \tau_{k+1})$: $\hat{h}^i_t = \gamma^i(\tilde{X}_t - \tilde{X}_k, \hat{h}_k)$

$\rightarrow$ The optimal strategy is derived directly on the basis of the local dynamics of the asset prices and not on the basis of the value function (*no corresponding regularity is thus required on the value function*).

$\rightarrow$ The value function has of course its own interest.

References

- K. Fujimoto, H. Nagai, W.J.Runggaldier, *Expected log-utility maximization under incomplete information and with Cox-process observations.* Asia-Pacific Financial Markets. 21 (2014), No 1, pp 35-66.

- K. Fujimoto, H. Nagai, W.J.Runggaldier, *Expected power-utility maximization under incomplete information and with Cox-process observations.* Applied Mathematics and Optimization, 67 (2013) No.1, pp. 33-72.

Happy Birthday, Yura!