A Mean Field Game of Controls: Closing The Loop of Optimal Trading

Charles-Albert Lehalle (Capital Fund Management, Paris and Imperial College, London)

Joint work with Pierre Cardaliaguet (CEREMADE, Université Paris Dauphine)

FIME – May 2017

Figure 5.5: Electronic market development by asset class, 2012

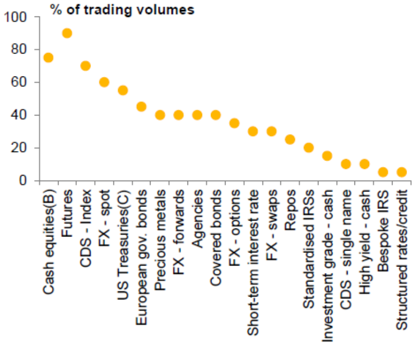

Source: FEMR[245]
Note: includes multi-dealer RFQ

Optimal Trading [Lehalle et al., 2013, Chapitre 3]

- Investors use trading algorithms to buy and sell large amounts of shares or contracts
- It meets the demand of regulators: more tractability, less complex products
- Intermediaries themselves use trading algorithms.

No more isolating an agent

- Up to now, the literature focused on one large investor facing a *background noise* (with the exception of [Jaimungal and Nourian, 2015], modeling one large risk-averse agent vs. small agents sensitive to their expected gain only).
- Here instead of having one isolated mean-variance agent [Almgren and Chriss, 2000],
- We will model all agents conducting simultaneously the same kind of strategies à la [Cartea and Jaimungal, 2015].

Generic considerations on MFG

- Mean Field Game is about a **continuum of agents**, characterized by their **distribution**,
- Each agent is fully identified by its position in the state space (from the viewpoint of one specific agent, others can be reordered),
- Each agent is sensitive to others via a **Mean Field**, and each agent contributes to this mean field (think about the pressure in a room where agents are particules),
- Each agent solves a (backward) optimization problem (his cost function can be a functional of the distribution at t),
- The distribution of agents is transported (via the controls) a forward way.

The natural mean field of financial markets

- Endogenous liquidity is often missing in the cost function of each agent (think about replicating bank's risk),
- Each bank is facing a **mean field**, i.e. the aggregation of others' actions is meant to be martingale,
- In reality **banks do communicate** via the global state of liquidity.
- **Liquidity is the natural mean field** to inject mathematical finance models in a game theoretical framework (slow: [Carmona et al., 2013] and HF: [Lachapelle et al., 2016], now *instantaneous*).

On our database of 300,000 large orders [Bacry et al., 2015]

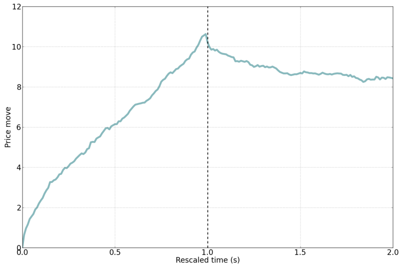

Optimal Trading is about

- Trading **slow enough** to avoid market impact
- and **fast enough** so that the price is close to the decision.

Investors

- tacke decisions based on private information and portfolio construction methods,
- concentrate their decisions on their dealing desk,
- who study the liquidity of the portfolios to buy and sell,
- and use brokers to execute an automated way these decisions.

Trading Algorithms: Typical Features

Benchmark	Type of stock	Type of trade	Main feature
PoV	Medium to large market depth	(1) Long duration position	(1) Follows current market flow, (2) Very reactive, can be very aggressive, (3) More price opportunity driven if the range between the max percent and min percent is large
VWAP / TWAP	Any market depth	(1) Hedging order, (2) Long duration position, (3) Unwind tracking error (delta hedging of a fast evolving inventory)	(1) Follows the “usual” market flow, (2) Good if market moves with unexpected volumes in the same direction as the order (up for a buy order), (3) Can be passive
Implementation Shortfall (IS)	Medium liquidity depth	(1) Alpha extraction, (2) Hedge of a non-linear position (typically Gamma hedging), (3) Inventory-driven trade	(1) Will finish very fast if the price is good and enough liquidity is available, (2) Will “cut losses” if the price goes too far away
Liquidity Seeker	Poor a fragmented market depth	(1) Alpha extraction, (2) Opportunistic position mounting, (3) Already split / scheduled order	(1) Relative price oriented (from one liquidity pool to another, or from one security to another), (2) Capture liquidity everywhere, (3) Stealth (minimum information leakage using fragmentation)

Trading Algorithms: Typical Uses

Benchmark	Region of preference	Order characteristics	Market context	Type of hedged risk
PoV	Asia	Large order size (more than 10% of ADV: Average daily consolidated volume)	Possible negative news	Do not miss the rapid propagation of an unexpected news event (especially if I have the information)
VWAP / TWAP	Asia and Europe	Medium size (from 5 to 15% of ADV)	Any “unusual” volume is negligible	Do not miss the slow propagation of information in the market
Implementation Shortfall (IS)	Europe and US	Small size (0 to 6% of ADV)	Possible price opportunities	Do not miss an unexpected price move in the stock
Liquidity Seeker	US (Europe)	Any size	The stock is expected to “oscillate” around its “fair value”	Do not miss a liquidity burst or a relative price move on the stock

More on all this in the three “reference books” for practitioners:

- Market Microstructure in Practice [Lehalle et al., 2013]
- The Financial Mathematics of Market Liquidity [Guéant, 2016]
- Algorithmic and High-Frequency Trading [Cartea et al., 2015]
- Quantitative Trading: Algorithms, Analytics, Data, Models, Optimization [Guo et al., 2016]

Optimal Trading Rate

The first papers [Almgren and Chriss, 2000], [Bertsimas and Lo, 1998], focussed on the optimal trading rate, or trading speed (i.e. how many shares to buy or sell every 5 minutes) for long metaorders.

- it does not deal with microscopic orderbook dynamics,
- it is a convenient way to take into account any information or constraint at this time scale.

It is very useful for asset managers, brokers, or hedgers. I.e. especially when the decision step is separated from the execution step.
Nevertheless it can be used for opportunistic trading too, when risk management at an intraday scale is important.

The Cartea and Jaimungal Version

The usual (simplistic) example of (continuous time) optimal trading (for a large sell order)

1. Write the Markovian dynamics or the price P, the quantity to trade Q and the cash account X for a sell of Q_0 shares before $t = T$ (control is the –negative– trading speed ν)

$$dQ = \nu\, dt,\ dX = -\nu\,(P + \kappa\cdot\nu)dt,\ dP = \mu\, dt + \sigma\, dW.$$

2. Write the cost function to maximize

$$V(t,p,q,x,\nu) = \mathbb{E}\left(X_T + Q_T(P_T - A\cdot Q_T) - \phi\int_{\tau=t}^{T} Q_\tau^2\, d\tau \middle| \mathcal{F}_t\right).$$

3. it gives the HJB and its terminal condition $V(T_, \ldots) = x + q(p - Aq)$

$$-\mu\partial_P V = \partial_t V + \frac{\sigma^2}{2}\partial_P^2 V - \phi\, q^2 + max_\nu\left\{\nu\partial_Q V\, dt - \nu(p+\kappa\cdot\nu)\partial_X V\right\}.$$

4. After the change of variable $V(t,p,q,x) = x + q\,p + v(t,q)$, you have

$$-\mu\partial_P V = \partial_t v - \phi\, q^2 + \max_{\nu}\left\{\nu\partial_Q v - \kappa\nu^2\right\}.$$

5. The optimal control is $\nu^* = \partial_Q v/(2\kappa)$, and the PDE $-\mu\partial_P v = \partial_t v - \phi\, q^2 + \kappa(\partial_Q v)^2/(4\kappa)$.
6. When the value function is quadratic: $v(t,q) = h_0(t) + q\,h_1(t) - q^2\,h_2(t)/2$, you can separate the PDE in three:

$$\left\{\begin{array}{rcrcr} 2\kappa\phi & = & -2\kappa h_2' & + & h_2^2 \\ -\mu & = & h_1' & - & 2h_1 h_2 \\ 0 & = & h_0' & + & h_1^2 \end{array}\right.$$

And terminal conditions $h_0(T) = h_1(T) = 0$ and $h_2(T) = -2A$: **backward dynamics**.

Cartea and Jaimungal (with misc. co-authors) developed this framework for plenty versions: with a (slightly) different objective function (VWAP, PoV), with permanent market impact $\mu \to \mu + \nu$, with μ_t any (adapted) process, etc.

1. Standard Algorithmic Trading
2. Closing The Loop
3. An Explicit Solution For Identical Preferences

INSIGHT.DATA.CLARITY.

Different Levels of Loops to Close

Two areas are not explored enough

- for **practitioners** : statistical learning; how to adapt online to regime switches (remember what we said about liquidity game vs. price game)? How to be robust to transitory phases? “Closing the loop” with learning is mixing exploration and exploitation.
- for **regulators** : game theory; what is the result of putting rational agents together? The more quants will read the 3 books, the more it will be needed to understand such interactions, and how changing “meta parameters” (ie rules) will modify the outcome of this game?

For game theory on financial market:

- few agents usually leads to principal - agent problems,
- a lot of agents usually leads to mean field games.

Moreover, game theory is a way to obtain **robust control** .

- the number of players needs to be "large enough"
- all players contribute to a "mean field" (i.e. a global variable: available shares, volatility, resource, etc)
- a function of this mean field (at least its mean, may be its standard deviation, etc) appear in this utility function of the players

→ the name on the player cannot be used, but they can have a parameter (like a time horizon or risk aversion) of their own

The methodology is similar to the one to solve static Nash games:

- express the solution (for one agent) and find the solution as if the mean field was known
- you obtain a backward pde
- combine what you know about the mean field to find its forward pde

Liquidity is typically a mean field: the state of the inventory of participants influence their costs and can lead to fire sales [Carmona et al., 2013]. What practitioners call "velocity" of the liquidity (the flows) is a mean field too, it probably forms the prices along with market impact.

A continuum of agents trade optimally "à la Cartea-Jaimungal".

$$dS_t = \alpha\mu_t\, dt + \sigma\, dW_t.$$

$$dQ^a_t = \nu^a_t\, dt, \tag{1}$$

now for a seller, $Q^a_0 > 0$ (the associated control ν^a will be mostly negative) and the wealth suffers from linear trading costs driven by κ (or *temporary*, or *immediate market impact*):

$$dX^a_t = -\nu^a_t(S_t + \kappa \cdot \nu^a_t)\, dt.$$

Same equations as for the standard framework, except **the trend is made of the permanent impact of all agents**:

$$\mu = \int_{a\in\mathfrak{A}} \nu^a\, df(a),$$

where $f(a)$ is the density of the agents in a feature space $\mathfrak{A}$.

The cost function of investor a selling from $t = 0$ and T is similar to the ones used in [Cartea et al., 2015]: the terminal inventory is penalized and a quadratic running cost is subtracted:

$$V_t^a := \sup_{\nu} \mathbb{E}\left(X_T^a + Q_T^a (S_T - A^a \cdot Q_T^a) - \phi^a \int_{s=t}^{T} (Q_s^a)^2 \, ds \middle| \mathcal{F}_t \right).$$

Here we took T common to all investors, i.e. **the end of the trading day**.

Our framework is then

- Each agent a has an initial quantity Q_0^a to buy ($Q_0^a < 0$) or to sell ($Q_0^a > 0$) we can even have purely opportunistic agents ($Q_0^a = 0$).
- They all start at the open of the trading session $t = 0$ and end at the close $t = T$.
- Each of them maximizes the value of his trades for the day: cash + penalized remaning quantity (by A^a) - cost of risk (with his own risk aversion ϕ^a).

HJB For One Player (Backward Value Function)

The associated Hamilton-Jacobi-Bellman is

$$0 = \partial_t V^a - \phi^a q^2 + \frac{1}{2}\sigma^2 \partial_S^2 V^a + \alpha\mu\partial_S V^a + \sup_\nu \{\nu\partial_Q V^a - \nu(s+\kappa\,\nu)\partial_X V^a\},$$

with the terminal condition $V^a(T, x, s, q; \mu) = x + q(s - A^a q)$.

The usual solution: Following the Cartea and Jaimungal's approach, we will use the following ersatz: $V^a = x + qs + v^a(t, q; \mu)$. Thus the HJB on v is

$$-\alpha\mu\, q = \partial_t v^a - \phi^a q^2 + \sup_\nu \left\{\nu\partial_Q v^a - \kappa\,\nu^2\right\},$$

with the teminal condition $v^a(T, q; \mu) = -A^a q^2$.

The associated optimal feedback / control is straightforward to find:

$$\nu^a(t, q) = \frac{\partial_Q v^a(t, q)}{2\kappa}. \tag{2}$$

$\Rightarrow$ We know that if we have the value function of an agent v, we can deduce the associated optimal control.

Transport of the Mass of the Players (Forward)

Distribution of agents is mainly defined by the joint distribution $m(t, dq, da)$ of

- the inventory Q_t^a, with known initial values.
- the preferences of the agent: the risk aversion ϕ^a, and the terminal penalization A^a.

The net trading flow μ driving the trend of the public price at time t reads:

$$\mu_t = \int_{(q,a)} \nu_t^a(q)\, m(t, dq, da) = \int_{q,a} \frac{\partial_Q v^a(t, q)}{2\kappa}\, m(t, dq, da).$$

$\Rightarrow$ v^a is an implicit function of μ (look at the HJB), meaning **we will have a fixed point problem to solve** in μ.

By the dynamics of Q_t^a, the transport of the measure $m(t, dq, da)$ has to follow (continuity equation)

$$\partial_t m + \partial_q \left(m \frac{\partial_Q v^a}{2\kappa} \right) = 0 \text{ with initial condition } m_0 = m_0(dq, da).$$

Obtaining The Backward-Forward Dynamics

Now we can have side to side:

- the HJB (backward) PDE where we plug the value of μ;
- the (Forward) transport of the mass of agents m, driven by the aggregation of their instantaneous decisions.

$$\left\{\begin{array}{rcl} \underbrace{-\alpha q \int_{(q',a')} \frac{\partial_Q v^{a'}(t,q')}{2\kappa} m(t,dq',da')}_{\text{aggregate of all agents}} & = & \underbrace{\partial_t v^a - \phi^a q^2 + \frac{(\partial_Q v^a)^2}{4\kappa}}_{\text{optimal for one agent}} \\ \partial_t m + \partial_q \left(m \frac{\partial_Q v^a}{2\kappa} \right) & = & 0 \end{array}\right.$$

Under boundary (resp. initial and terminal) conditions:

$$\left\{\begin{array}{rcll} m(0,dq,da) & = & m_0(dq,da) & , \\ v^a(T,q;\mu) & = & -A^a q^2 & , \quad \forall a. \end{array}\right.$$

Explicit Solution For a Special Case

We will need a notation for the aggregated (i.e. net) position of all agents $E(t) = \mathbb{E}[Q_t] = \int_q q\, m(t, dq)$.

Then we can write:

$$\begin{aligned} E'(t) &= \int_q q \partial_t m(t, dq) && \leftarrow \text{definition} \\ &= -\int_q q \partial_q \left(m(t,q) \frac{\partial_Q v(t,q)}{2\kappa} \right) dq && \leftarrow \text{forward dynamics (transport)} \\ &= \int_q \frac{\partial_Q v(t,q)}{2\kappa} m(t, dq) && \leftarrow \text{integration by parts.} \end{aligned}$$

Moreover, $v(t,q)$ can be expressed as a quadratic function of q: $v(t,q) = h_0(t) + q\, h_1(t) - q^2 \frac{h_2(t)}{2}$, leading to:

$$E'(t) = \int_q m(t,q) \left(\frac{h_1(t)}{2\kappa} - \frac{h_2(t)}{2\kappa} q \right) dq = \frac{h_1(t)}{2\kappa} - \frac{h_2(t)}{2\kappa} E(t).$$

In a more compact form:

$$2\kappa E'(t) = h_1(t) - E(t) \cdot h_2(t).$$

We now collect all the equations:

$$
\begin{cases}
4\kappa\phi = -2\kappa h_2'(t) + (h_2(t))^2, & \text{(3a)} \\
\alpha h_2(t)E(t) = 2\kappa h_1'(t) + h_1(t)\left(\alpha - h_2(t)\right), & \text{(3b)} \\
-\left(h_1(t)\right)^2 = 4\kappa h_0'(t), & \text{(3c)} \\
2\kappa E'(t) = h_1(t) - h_2(t)E(t). & \text{(3d)}
\end{cases}
$$

with the boundary conditions $h_0(T) = h_1(T) = 0,\ h_2(T) = 2A,\ E(0) = E_0$, where $E_0 = \int_q q m_0(q) dq$ is the net initial inventory of market participants (i.e. the expectation of the initial density m).

The Master Equation For Identical Preferences

The previous system of ordinary differential equations implies

$$0 = 2\kappa E''(t) + \alpha E'(t) - 2\phi E(t) \quad (4)$$

with boundary conditions $E(0) = E_0$ and $\kappa E'(T) + AE(T) = 0$.

Closed form for the net inventory dynamics $E(t)$

For any $\alpha \in \mathbb{R}$, the problem (4) has a unique solution E, given by

$$E(t) = E_0 a (\exp\{r_+ t\} - \exp\{r_- t\}) + E_0 \exp\{r_- t\}$$

where a is given by

$$a = \frac{(\alpha/4 + \kappa\theta - A)\exp\{-\theta T\}}{-\frac{\alpha}{2}\mathrm{sh}\{\theta T\} + 2\kappa\theta \mathrm{ch}\{\theta T\} + 2A\mathrm{sh}\{\theta T\}},$$

the denominator being positive and the constants $r_\alpha^\pm$ and θ being given by

$$r_\pm := -\frac{\alpha}{4\kappa} \pm \theta, \qquad \theta := \frac{1}{\kappa}\sqrt{\kappa\phi + \frac{\alpha^2}{16}}.$$

Solving $h_2(t)$

h_2 solves the following backward ordinary differential equation (3a): $0 = 2\kappa \cdot h_2'(t) + 4\kappa \cdot \phi - (h_2(t))^2$ under $h_2(T) = 2A$. It is easy to check the solution is

$$h_2(t) = 2\sqrt{\kappa\phi}\frac{1 + c_2 e^{rt}}{1 - c_2 e^{rt}},$$

where $r = 2\sqrt{\phi/\kappa}$ and c_2 solves the terminal condition. Hence

$$c_2 = \frac{1 - A/\sqrt{\kappa\phi}}{1 + A/\sqrt{\kappa\phi}} \cdot e^{-rT}.$$

Keep in mind the optimal control is

$$\nu^* = \frac{\partial_Q v(t,q)}{2\kappa} = \frac{h_1(t)}{2\kappa} - q \cdot \frac{h_2(t)}{2\kappa},$$

Solving $h_1(t)$

The affine component of the control can be easily deduced from $h_2(t)$ and $E(t)$:

$$h_1(t) = 2\kappa \cdot E'(t) + h_2(t) \cdot E(t).$$

Dependence of the Solution to the Mean Field

The optimal control is

$$\nu^* = \frac{\partial_Q v(t,q)}{2\kappa} = \underbrace{\frac{h_1(t)}{2\kappa}}_{\text{reaction to the mean field}} - \underbrace{q \cdot \frac{h_2(t)}{2\kappa}}_{\text{inventory control}}.$$

- The second term is proportional to your **inventory**, i.e; the *remaning quantity to buy/sell*, it is **independent of E** ;
- The first term embeds **the dependence to the mean field** : $h_1(t) = 2\kappa \cdot E'(t) + h_2(t) \cdot E(t)$.

$\Rightarrow$ locally you adapt your behaviour to the mean field via h_1,

$\rightarrow$ then (you changed your inventory), you slowly (re)adapt to be ready for boundary conditions / costs.

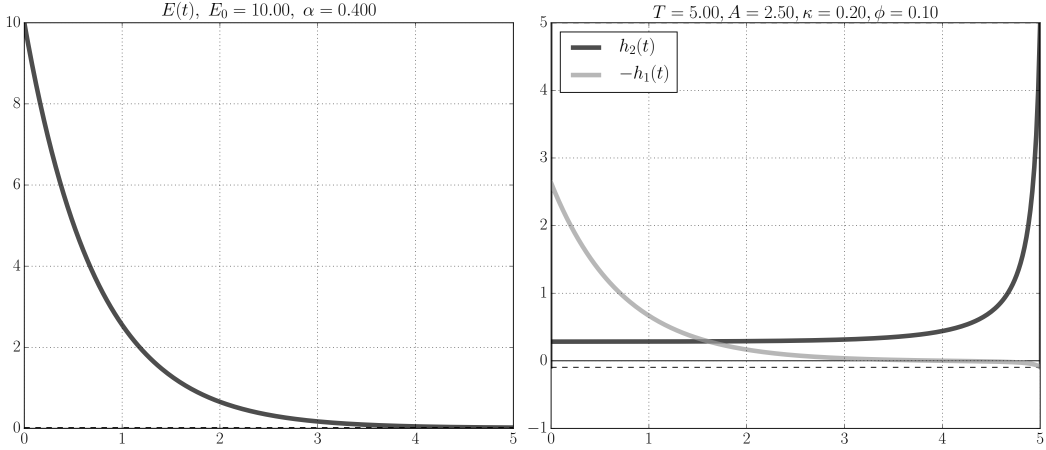

Dynamics of E (left) and $-h_1$ and h_2 (right) for a standard set of parameters: $\alpha = 0.4$, $\kappa = 0.2$, $\phi = 0.1$, $A = 2.5$, $T = 5$, $E_0 = 10$.

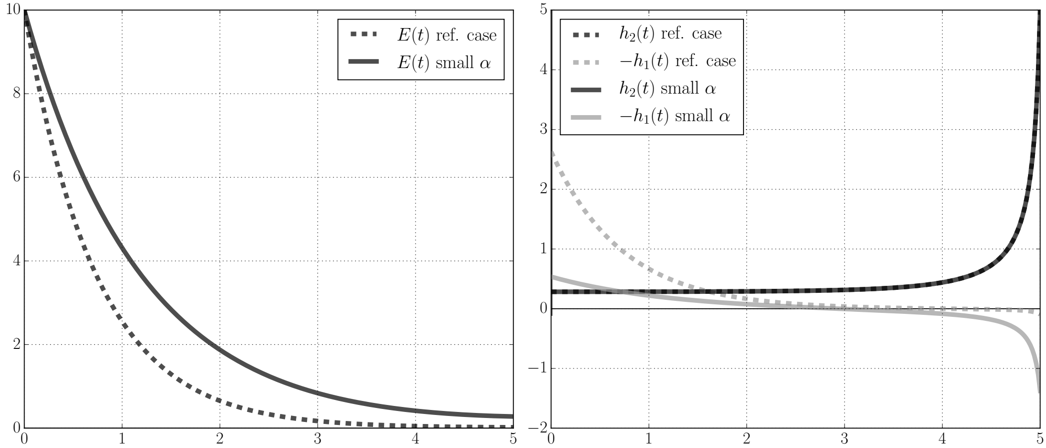

Comparison of the dynamics of E (left) and $-h_1$ and h_2 (right) between the "reference" parameters of Figure **??** and smaller α (i.e. $\alpha = 0.1$ instead of 0.4) such that $|h_1(0)|$ is smaller.

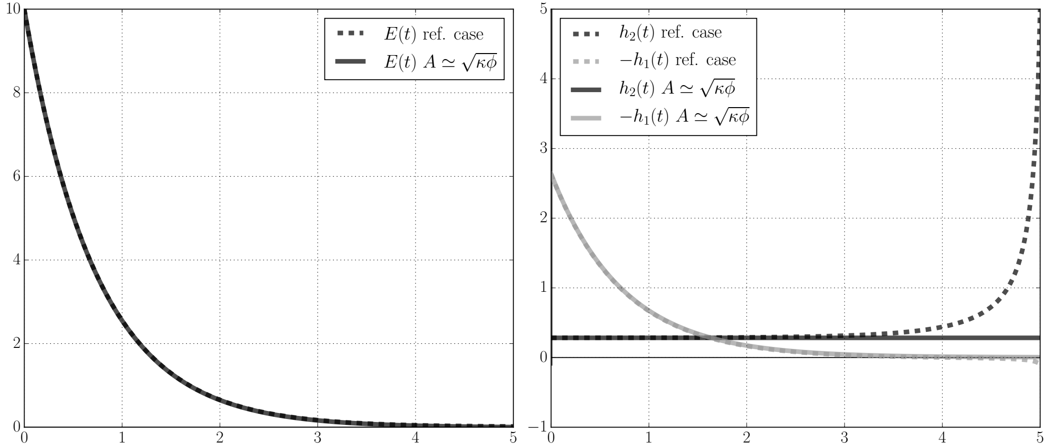

Comparison of the dynamics of E (left) and $-h_1$ and h_2 (right) between the "reference" parameters of Figure **??** and when $\sqrt{\kappa\phi} \simeq A$: in such a case h_2 is almost constant but E and h_1 are almost unchanged.

A Case With Not Monotonous E

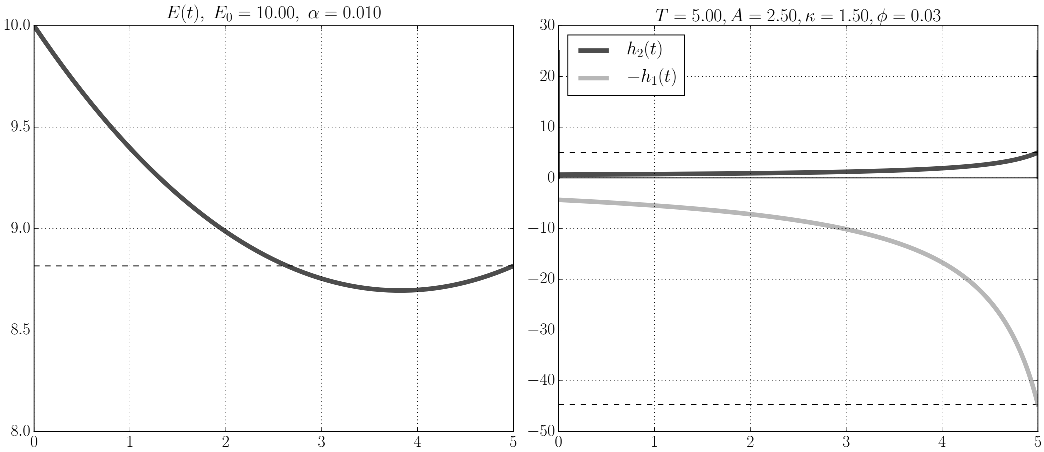

A specific case for which E is not monotonous: $\alpha = 0.01$, $\kappa = 1.5$, $\phi = 0.03$, $A = 2.5$, $T = 5$ and $E_0 = 10$.

Not Monotonous E

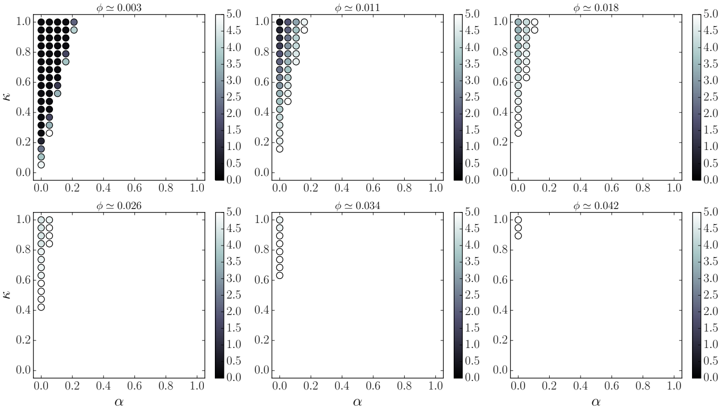

Numerical explorations of t^m for different values of ϕ (very small ϕ at the top left to small ϕ at the bottom right) on the $\alpha \times \kappa$ plane, when $T = 5$ and $A = 2.5$. The color circles codes the value of t^m: small values (dark color) when E changes its slope very early; large values (in light colors) when E changes its slope close to T.

Conclusion on MFG of Controls For Liquidation

It is a proof of maturity of the use if stochastic control in financial math:

- Four years ago, it was difficult to think about a game theoretical version of the Almgren and Chriss optimal liquidation problem (schied and jaimungal).
- Our understanding of the problem itself improved (see Guéant and Cartead and Jaimungal books)
- and some extensions of MFG have been needed (see the paper).
- but we now know how to handle it (and in a specific case it is fully solved)

Solving game theoretical versions of what we know is important (instead of sophisticating it in a mean field ~~game~~), because

- it is a way to obtain robust control
- it helps regulator to understand the system to adjust some meta parameters (κ is this example)

MFG is not the only way to answer to such questions. Nevertheless in general Mean Field Games can take into account interactions between different market participants **as soon as they interact via liquidity** (i.e. the mean field).
Moreover learning should not be forgot (done in our paper): what does change when information is not complete?

Thank You For Your Attention

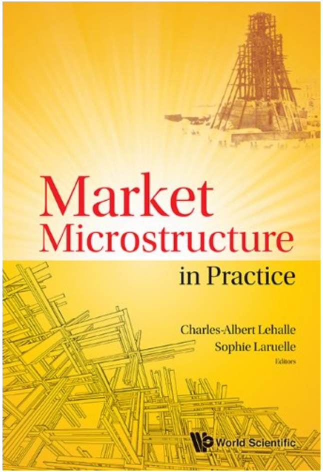

Almgren, R. F. and Chriss, N. (2000).
Optimal execution of portfolio transactions.
Journal of Risk, 3(2):5–39.

Bacry, E., Iuga, A., Lasnier, M., and Lehalle, C.-A. (2015).
Market Impacts and the Life Cycle of Investors Orders.
Market Microstructure and Liquidity, 1(2).

Bertsimas, D. and Lo, A. W. (1998).
Optimal control of execution costs.
Journal of Financial Markets, 1(1):1–50.

Carmona, R., Fouque, J.-P., and Sun, L.-H. (2013).
Mean Field Games and Systemic Risk.
Social Science Research Network Working Paper Series.

Cartea, A. and Jaimungal, S. (2015).
Incorporating Order-Flow into Optimal Execution.
Social Science Research Network Working Paper Series.

Cartea, A., Jaimungal, S., and Penalva, J. (2015).
Algorithmic and High-Frequency Trading (Mathematics, Finance and Risk).
Cambridge University Press, 1 edition.

Guéant, O. (2016).
The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making.
Chapman and Hall/CRC.

References II

Guo, X., Lai, T. L., Shek, H., and Wong, S. P. (2016).
Quantitative Trading: Algorithms, Analytics, Data, Models, Optimization.
Chapman and Hall/CRC.

Jaimungal, S. and Nourian, M. (2015).
Mean-Field Game Strategies for a Major-Minor Agent Optimal Execution Problem.
Social Science Research Network Working Paper Series.

Lachapelle, A., Lasry, J.-M., Lehalle, C.-A., and Lions, P.-L. (2016).
Efficiency of the Price Formation Process in Presence of High Frequency Participants: a Mean Field Game analysis.
Mathematics and Financial Economics, 10(3):223–262.

Lehalle, C.-A., Laruelle, S., Burgot, R., Pelin, S., and Lasnier, M. (2013).
Market Microstructure in Practice.
World Scientific publishing.