Some Stylized Facts On Transaction Costs And Their Impact On Investors

Charles-Albert Lehalle (CFM, Paris and Imperial College, London)

Conseil Scientifique de l'AMF – Paris, janvier 2019

Some stylized fact on transaction costs and its impact on investors

The positionning of this talk

I have been studying optimal trading with different co authors for 14 years. With the sophistication of optimizers and mathematical frameworks, we had to investigate the nature of transaction costs as much as possible, since the nature of optimality we try to obtain emerges from the balance between urgency to trade (to be as close as possible to the traded information), and costs. This talk mainly relies on recent works

- Chapter 3 of *Market Microstructure in Practice*, 2nd Edition (Lehalle et al., 2018)
- *Market Impacts and the Life Cycle of Investors Orders*, (Bacry et al., 2015)
- *Mean Field Game of Controls and An Application To Trade Crowding*, (Cardaliaguet and Lehalle, 2017)
- *Incorporating Signals into Optimal Trading*, (Lehalle and Neuman, 2017) (accepted in Finance and Stochastics)
- *Stock Market Liquidity and the Trading Costs of Factor Based Investments*, on-going work with Amine Raboun, Marie Brière, and Tamara Nefedova (Université Paris Dauphine)
- *Stylized Facts on Price Formation on Corporate Bonds and Best Execution Analysis*, on-going work with Xin Guo and Renyuan Xu (UC Berkeley)

They exploit different databases: ANcerno, enhanced TRACE, NASDAQ OMX trades, RTH quotes, and a proprietary database from Crédit Agricole Cheuvreux.

Outline

Outline

Outline

Defining Transaction Costs

The position of transaction costs analysis in the investment process
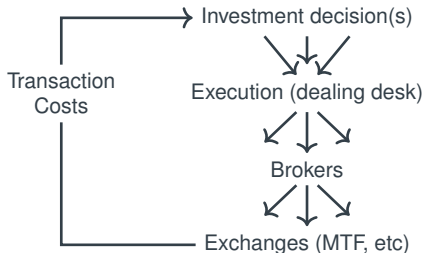
The investment process evolved the last 25 years towards specialization of the tasks:

- ► The world of asset management concentrated in the hand of few, very large companies,
- ► They created dealing desks, specialized in the execution of large orders ("instructions", "tickets", "*metaorders*"). They turned to be multi-assets (Equities and FIM) few years ago.
- ► Pushed by the regulation (Reg NMS, MiFID 1-2), the concept of *best execution* emerged, spreading the responsibility of all the process over all the involved layers, including brokers and vendors (data and software).
- ► Smallest asset managers largely rely on external companies and services.

Transaction Costs → Investment decision(s) → Execution (dealing desk) → Brokers → Exchanges (MTF, etc)

This "*culture of best execution*" is a step forward, nevertheless it did not propagated by now to the process of investment as a whole.

Defining Transaction Costs

The role of Transaction Costs (TC) in the investment process

We can easily list domains where Transaction Costs (TC) could be better taken into account

- ► Index Investing: an **investible index should be self-financed**... including transaction costs? (the optimal cut of a cap. weighted index could be characterized thanks to TC).
  The practice of building "style ETF" according to an ad hoc implementation of an academic factor (Harvey et al., 2013), and using market capitalization, it meant to take TC into account.
- ► Today a common practice is to **use a ranking of instruments** in the investible universe. Could follow the capitalization, the Average Daily Volume (ADV) or the average bid-ask spread (see next slide).
- ► The rebalancing of a portfolio: from the current portfolio to the desired one, instrument by instrument, could be **a balance between the tracking error between the two portfolios and the anticipated TC**.

Preliminary remark

Transaction costs are the cost to pay to convert uncertainty (i.e. tracking error to a desired portfolio), to dollars. These costs are partly deterministic (like *fees*), but are also the result of a *liquidity seeking process*. The latter is partly due to the *nature of the instrument* and partly to the *market context*.

Defining Transaction Costs

A static vision of usual "liquidity variables": 300 US stocks in 2014 (one point is one stock one day)
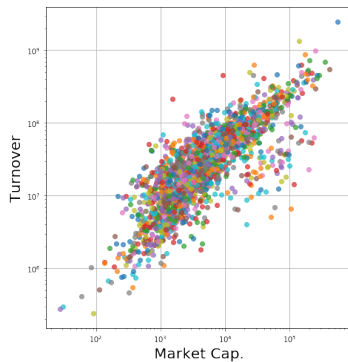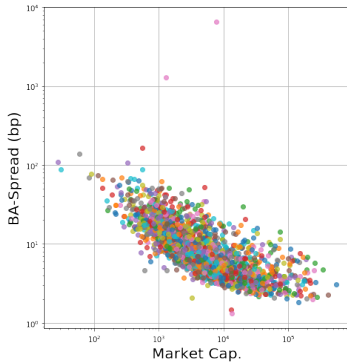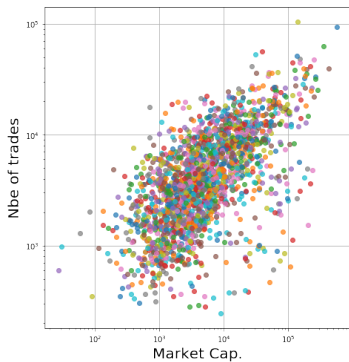
Defining Transaction Costs

A collection of metrics and their drawbacks

Qualitative indicators
Easy to compute ex-ante

Capitalization
Average Daily Volume
Bid-Ask Spread

why not "*Volume at first limit*"?
or the "*Avg. number of trades per day*"?

Quantitative metrics
Easy to compute ex-post

Implementation Shortfall:
(avg. price - decision price)$\times$ sign(ticket)
Slippage to WVAP:
Volume Weighted Average Price
(avg. price - VWAP) $\times$ sign(ticket)
Slippage to Close:
(avg. price - Close) $\times$ sign(ticket)

The question of **transparency and availability of datasets** is crucial. As soon as a volume or quotes (best bid and ask) are involved, the reference markets are of importance (MTF, Block matching, OTC, etc).

The **estimation of quantitative metrics** is key to go from a qualitative framework to a quantitative one, allowing comparisons and decision support.

Defining Transaction Costs

How could we use averages?

Eiffage, 01/04/2011 to 01/04/2012
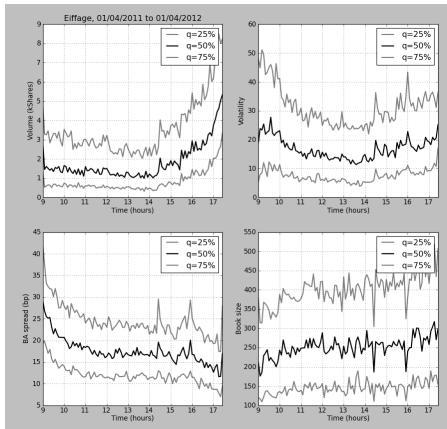
Here is an example of intra-day seasonality for four variables of interest:

- ► traded volumes
- ► volatility, using the GK estimates (Garman and Klass, 1980) on 5 min slices
- ► bid-ask spread
- ► average size at first limit.

Shapes can be easily seen, they reflect

- ► uncertainty of prices at the beginning of the day
- ► increase of activity at the end of the day
- ► a very different relation between volatility, bid-ask spread (and quantity at first limits) on the morning vs. on the afternoon.

What quantity to be used? if an asset manager trades at the open, at the close or all over the day?

Defining Transaction Costs
Who computes what?

- Usually asset managers have only access to daily data, "high frequency" (i.e. intraday detailed) datasets are not always available to their dealing desks. There is a clear difference between equities (some data are easy to access) and fixed income products (like corporate bonds).

- Often asset managers rely on brokers to **execute and report**. In such a case a good practice is to share the execution between several brokers and compare their "performances". This is how **Transaction Costs Analysis (TCA)** was born as a *monthly comparison* of brokers prices around these reference papers:
  - *A methodology for measuring transaction costs* (Collins and Fabozzi, 1991)
  - *A practical framework for estimating transaction costs and developing optimal trading strategies to achieve best execution* (Kissell et al., 2004)
  - *Performance of institutional trading desks: An analysis of persistence in trading costs* (Anand et al., 2011)

- It is important to note that, as long as brokers fees are paying "research" (fundamental analysis), the "optimal" balance between sending flows to pay for research, and allocating flows to the best brokers, is not obvious... There is hope that **MiFID 2 would help to allocate flows on brokers minimizing transaction costs**.

Defining Transaction Costs

Where should Transaction Costs be taken into account, and what could be the consequences

Modern Portfolio Theory (Markowitz, 1952) is based on *independent and identically distributed* assumptions (of the returns and of the risk). Some methods, like Black-Litterman (see (Bertsimas et al., 2012)), provide a forward-looking (but i.i.d. again, and "at the equilibrium") framework. **It is difficult to take transaction costs into account without introducing path dependency** (i.e. the cost is a function of the portfolio at t_0).

We have to go back to (Merton, 1969) to obtain an time-dependent framework, that has been solved in (Shreve and Soner, 1994) for **proportional transaction costs** (i.e. buying v shares at price p suffers from $c \cdot v$ costs: the total price to buy is $v \cdot (p + c)$):

▶ in such a case there is a **no-trade zone**: the portfolio manager should wait for a given tracking error to be reached, and trade only when this threshold is crossed.

Another approach has been initiated by (Garleanu and Pedersen, 2013), using quadratic costs (i.e price to buy is $v^T \cdot p + v^T C v$):

▶ in such a case the solution is very similar to Kalman filtering: **target the expected optimal portfolio in d days, and make a small step in its direction**; reassess your choice each day.

Defining Transaction Costs

In practice, there are two layers of decision

In practice, the time scale of decision and the time scale of execution are separated:

1. Given the current portfolio and the expected costs, define a "target portfolio"
2. The difference between the target and the current portfolio are given as "instructions" to the dealing desk
3. When possible, each line (i.e. *metaorder*) is labelled as "urgent", "discretion" or "neutral";
4. The dealing desk combine all the instructions, takes the current and short term expected states of liquidity into account to time the execution of the metaorders.
5. Different brokers are used, and some direct (block) trading mechanisms. Today, on equity markets, 90% of the flows are executed (at the end) by algorithms.
6. Brokers are meant to provide feedback thanks to **pre-trade analysis** and **post-trade analysis**.
7. Dealing desks are conducting **TCA** to assess brokers' quality and provide feedback to portfolio managers.

We will see later in this talk how the trading algorithms are built.

Measuring Implicit Costs of Large Orders

The Market Impact of large orders

in (Moro et al., 2009)
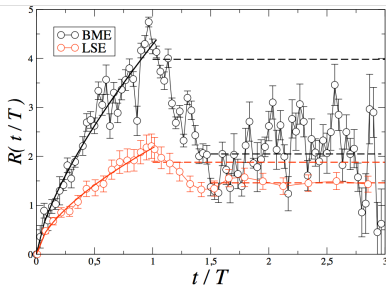
Market Impact takes place in different phases

- the transient impact, concave in time,
- reaches its maximum, the temporary impact, at the end of the metaorder,
- then it decays,
- up to a stationary level; the price moved by a permanent shift.

Measuring Implicit Costs of Large Orders

The Market Impact of large orders

On our database of 300,000 large orders
(Bacry et al., 2015)
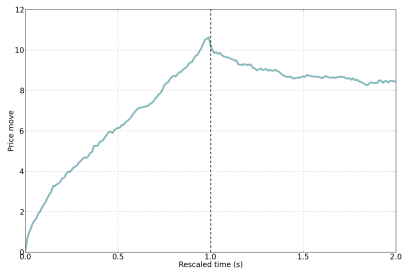
Market Impact takes place in different phases

- ► the transient impact, concave in time,
- ► reaches its maximum, the temporary impact, at the end of the metaorder,
- ► then it decays,
- ► up to a stationary level; the price moved by a permanent shift.

Measuring Implicit Costs of Large Orders

The Market Impact of Large Orders

To be more than anecdotical, it is needed to make statistics, that for we need a not of occurrences of " **metaorders** ". Some paper documented the "square root impact": the temporary impact of your flow is proportional to the its square root. But the three phases has been studied in fewer papers:

- ► *The Non-Linear Market Impact of Large Trades: Evidence from Buy-Side Order Flow* (Bershova and Rakhlin, 2013) – intraday impact
- ► *Is Market Impact a Measure of the Information Value of Trades? Market Response to Liquidity vs. Informed Trades* (Waelbroeck and Gomes, 2013) – daily impact of cash trades
- ► *Slow decay of impact in equity markets* (Brokmann et al., 2015) – daily impact of informed trades (hedge fund)
- ► *Market Impacts and the Life Cycle of Investors Orders* (Bacry et al., 2015) – intraday and daily impact of informed trades (bank)

Measuring Implicit Costs of Large Orders

The Market Impact of Large Orders: The Square Root Effect

Regression	Parameter	Coef. (log-log)	Coef. (L2)	Coef. (L1)
(R.1)				
	Daily participation	0.54	0.45	0.40
(R.2)				
	Daily participation	0.59	0.54	0.59
	Duration	−0.23	−0.35	−0.23
(R.3)				
	Daily participation	0.44	–	–
	Bid-ask spread	0.28	–	–
(R.4)				
	Daily participation	0.53	–	–
	Volatility	0.96	–	–
(R′.1)				
	Trading rate	0.43	0.33	0.43
(R′.2)				
	Trading rate	0.37	0.56	0.45
	Duration	0.15	0.24	0.23
(R′.3)				
	Trading rate	0.32	–	–
	Bid-ask spread	0.57	–	–
(R′.4)				
	Trading rate	0.32	–	–
	Volatility	0.88	–	–

Source: (Bacry et al., 2015)

The Formula should be close to

$$\text{MI} \propto \sigma \cdot \sqrt{\frac{\text{Traded volume}}{\text{Daily volume}}} \cdot T^{-0.2}$$

The term in duration is very difficult to estimate because you have a lot of conditioning everywhere:

► did you trading process reacted to market conditions?

► are you alone?

► etc.

We used different methods.

Measuring Implicit Costs of Large Orders

The Market Impact of Large Orders: The Square Root Effect

Source: on-going work with Amine Raboun, Marie Brière, and Tamara Nefedova using the ANcerno database (90 millions of metaorders, see Appendix)

The Formula should be close to

$$\text{MI} \propto \sigma \cdot \sqrt{\frac{\text{Traded volume}}{\text{Daily volume}}} \cdot T^{-0.2}$$

The term in duration is very difficult to estimate because you have a lot of conditioning everywhere:

- ▶ did you trading process reacted to market conditions?
- ▶ are you alone?
- ▶ etc.

We used different methods.

Measuring Implicit Costs of Large Orders

Parameters of the market impact for trading costs model

We regress the implementation shortfall on the bid-ask spread ψ and the volatility-participation one ($\sigma\sqrt{V/ADV}$); a linear regression on our 90 millions of metaorders. We bootstrap on samples of 10,000 metaorders to confirm. Over the whole period: $\alpha = 0.4$ and $\beta = 0.7$ (with respective t-stats of 17 and 37 and a R^2 of around 90%). Here you can see the variations of these coefficients:
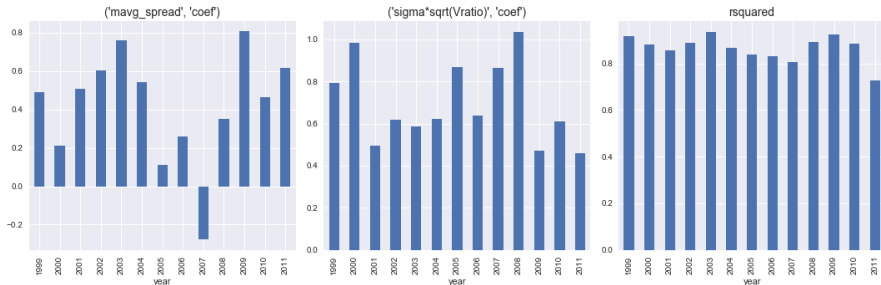
Measuring Implicit Costs of Large Orders

The non stationarity of trading costs.

Average Residuals of Trading Costs Model
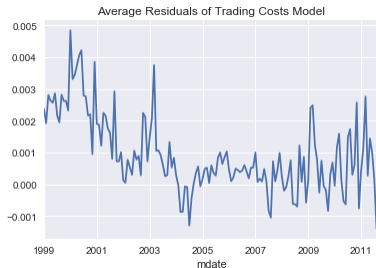
► To estimate the trading costs, we use a 2-years rolling model, nevertheless, to qualify this modelling, we performed a model on the whole dataset and draw the average of the residuals (monthly) over the year: the **transaction costs would be under estimated during the pre-fragmentation** (i.e. before 2005) period and overestimated afterwards.

► With such a model, we take into account the "market context" (via volatility) and the "characteristics of the instrument" (via bid-ask spread and ADV).

► Nevertheless there is another effect: the flows around the considered metaorder. Do I trade an isolated way? In the direction of the crowd or opposite to the crowd?

Measuring Implicit Costs of Large Orders

Metaorders are correlated, and the relative sign of your metaorder largely drives your costs

Define **metaorder imbalance** the (normalized) difference between buy and sell metaorders. It has been recently shown on the ANcerno database that this imbalance has a large role in the transaction costs (Bucci et al., 2018). We recover and document this feature.

	Volume Imbalance							
	≤ -0.75	-0.5	-0.25	0.0	0.25	0.5	0.75	≥ 0.75
Av Trading cost (bps)	32.87	8.18	5.98	4.23	3.49	4.86	6.90	21.03
Occurrences (≥ 10) (%)	19.66	10.92	9.40	9.07	9.06	9.49	11.00	21.40

Measuring Implicit Costs of Large Orders

Metaorders are correlated, and the relative sign of your metaorder largely drives your costs

Define **metaorder imbalance** the (normalized) difference between buy and sell metaorders. It has been recently shown on the ANcerno database that this imbalance has a large role in the transaction costs (Bucci et al., 2018). We recover and document this feature.

Measuring Implicit Costs of Large Orders

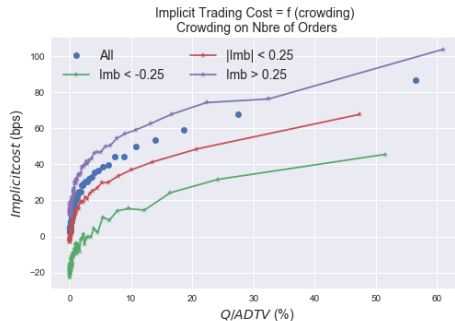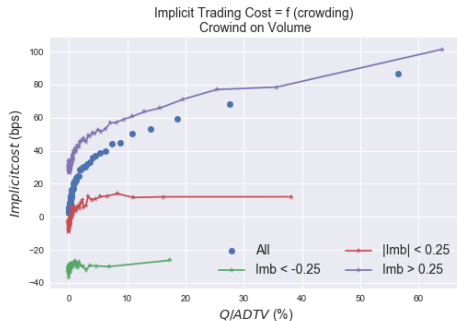Metaorders are correlated, and the relative sign of your metaorder largely drives your costs

Define **metaorder imbalance** the (normalized) difference between buy and sell metaorders. It has been recently shown on the ANcerno database that this imbalance has a large role in the transaction costs (Bucci et al., 2018). We recover and document this feature.

Parametric vs Non Parametric Estimation

This "imbalance effect" is why when we draw the non parametric estimate (that should be independent to the size of the metaorder) on the parametric chart, it is not flat: "**conditionally to the fact that metaorders on the market are in my direction, their imbalance of metaorders is rather in the same direction**". It is enough to influence the (red) curve.

Measuring Implicit Costs of Large Orders

The Market Impact of Large Orders: From Temporary To Permanent Market Impact

In (Bacry et al., 2015) we had enough data to investigate long term impact, exploring the relationships between permanent impact and traded information.

Daily price moves

▶ If you plot the long term price moves (x-axis in days), you see an regular increase;

▶ But the same stock is traded today, tomorrow, the day after, etc.

▶ Once you remove the market impact of "future" trades (similarly to Waelbroeck and Gomes (2013)), you obtain a different shape.

▶ If you look each curve: the yellow one contains the CAPM β (the metaorders are trading market-wide moves), the green curve contains the idiosyncratic moves, this shape can be read as the daily decay of metaorders impact.

Measuring Implicit Costs of Large Orders

The Market Impact of Large Orders: From Temporary To Permanent Market Impact

In (Bacry et al., 2015) we had enough data to investigate long term impact, exploring the relationships between permanent impact and traded information.

Daily price moves

- ▶ If you plot the long term price moves (x-axis in days), you see an regular increase;
- ▶ But the same stock is traded today, tomorrow, the day after, etc.
- ▶ Once you remove the market impact of "future" trades (similarly to Waelbroeck and Gomes (2013)), you obtain a different shape.
- ▶ If you look each curve: the yellow one contains the CAPM β (the metaorders are trading market-wide moves), the green curve contains the idiosyncratic moves, this shape can be read as the daily decay of metaorders impact.

Outline

Transaction Costs of Corporate Bonds

1. Focus on bid-ask spread costs and brokers evaluation

In a soon available paper with Renyuan Zu and Xin Guo, we investigate ways to provide best execution analysis for corporate bonds. Since the database TRACE is available since 2002 in the US and records all transactions on corporate bonds, we decided to work in two directions

- ▶ Provide a proxy for the bid-ask spread of a corporate bond that is rarely traded. Some practitioners argue "each bond is so specific that it is impossible to use some bonds to understand what is happening on others". We wanted to test this claim. We started from existing regressions (cf. Appendix) and used ML penalization techniques to provide a **bid-ask spread to be expected** to be used as a benchmark to compare brokers efficiency.
  This should allow to perform TCA in the sense of the seminal papers like (Collins and Fabozzi, 1991).

- ▶ Once this benchmark is available, it is possible to conduct a systematic review of *best execution* (in the MiFID 2 sense), and identify "suspicious trades". Then we wanted to provide a way to review these trades. We identified a buy/sell asymmetry on our (2015-2016) data: **buying corporate bonds is more expensive than selling them** (to broker-dealers).

Transaction Costs of Corporate Bonds

2. Modelling the bid-ask spread on rarely traded bonds

An example on US375558BG78
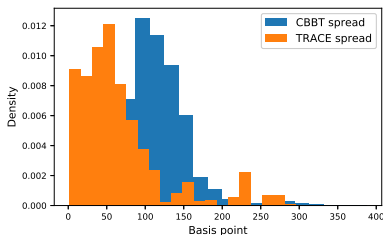
The first steps are to

▶ assign a *side* to each trade. We took the standard convention: *client* orders are assumed to the be the liquidity consuming ones;

▶ having a proxy for the bid-ask spread while we only have transactions. We used the methodology proposed in (Harris, 2015): when we witness two consecutive trades in opposite directions, we assume the bid-ask spread is the difference between the two prices.

To validate this methodology, we use 12 months of CBBT *indicative quotes* published by Bloomberg and compared them with our estimates.

Stationarity tests validated that our estimates are consistent with the CBBT, but smaller (mean ~ 0.9 and median ~ 0.7).

Corporate Bonds I

3. Drivers of the bid-ask spread of corporate bonds

These regressions allow to compute an "expected bid-ask spread" for a given week, to be used as a benchmark cost for TCA (Transaction Cost Analysis). Main results are:

- ▶ **The volatility** is an important feature, as expected by empirical observations and the theory: the larger the volatility, the larger the bid-ask. Practically we observe that an increase of 5% of the volatility (that is $1/2$ of its standard deviation on our dataset) corresponds to an increase of the bid-ask spread by 25bp (that is around one third of its standard deviation).

- ▶ **The number of trades per day** N and *the traded volume V* (in dollars) are both important variables (in log units), with coefficient suggesting that $N/\sqrt{V}$ is the feature impacting the bid-ask spread in bp. This implies that:
  - for a given **trading activity** N, the larger the traded volume, the smaller the bid-ask spread (in bp);
  - for a given traded volume in dollars, the lower average trade size (i.e. the more trades), the large the bid-ask spread.

  It is compatible with the documented stylized fact that for corporate bonds, small trade size obtain a worst bid-ask spread than large trades (Fermanian et al., 2015).

- ▶ The value of the coupon and the duration of the corporate bond play a small role in the formation of the bid-ask spread (both a via positive coefficients).

Corporate Bonds II

3. Drivers of the bid-ask spread of corporate bonds

- ▶ The **number of years to maturity** and the *years since issuance* are selected by our robust regressions.

  These two variables are linked, via the maturity of the bond, thanks to the relation: *Year to maturity = Maturity - years since issuance*. Naturally, the coefficient of *year to maturity* is negative while the one of *years since issuance* is positive: the further away from the maturity, the smallest bid-ask spread (in bp). This could support the claim of some market participants that there is only a short period after the issuance during which it is not too expensive to trade them on secondary markets.

Other variables appearing in the OLS are not robust enough to be selected by penalized regressions. Removing these 17 variables from the regression only reduces the R^2 from around 0.55 to around 0.50. It is a cheap cost to pay for the gained robustness.

Note that the R^2 of these regressions are around 50%, that is in line with the best results obtained in the academic literature: Dick-Nielsen et al. (2012) obtains R^2 between 0.50 and 0.80, while the R^2 of other papers are by far below 0.50.

Transaction Costs of Corporate Bonds

4. An buy/sell asymmetry of importance

We wanted to provide a way to investigate at the level of "abnormal" trades (in terms of bid-ask spread). Restricting the study to not isolated trades (i.e. surrounded by other trades at +/-2min; 3M of points), we could draw these "price profiles". They exhibit:

- the usual "price impact" phases: impulse followed by a decay (in accordance with the usual *propagator models*, see (Taranto et al., 2016))
- a buy/sell asymmetry, **specific to corporate bonds**.
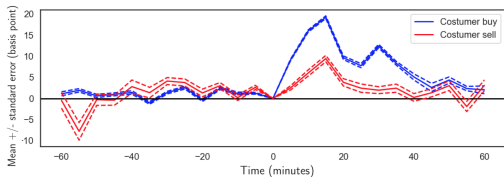
This feature has been already noticed in: the regression of Table IV of Hendershott and Madhavan (2015);Figure 15 of Mizrach (2015) plots the yearly average price change after 5 trades from 2003 to 2015 (the impact of buys is 25% more than of sells); and Table 1 of Ruzza (2016) exhibits an average price deviation between the price of a transaction and the average price of the day is on average of 56bp to 33bp for institutional buyers and of -25bp to -21bp for institutional sellers on TRACE data from 2004 to 2012.

Common Practices of Asset Management

Different components of the trading process

Asset managers delegate the trading to a dealing desk (internal or external).

► The good practice should be to label the *metaorders* as "urgent", or "discretion" or "neutral", so that the correct risk aversion (i.e. level of urgency) can be applied.

► Dealing desks rely on electronic tools, but continue to "seek liquidity" on small caps (equities) and some corporate bonds.

► They are meant to conduct a **pre-trade analysis** based on expected transaction costs (this expectation is based on external and internal databases...).

► **Trade scheduling** emerged 20 years ago to help algo designers to provide decision support tools. They are (or should be...) made of

— liquidity predictors (volumes, bid-ask spread, short term moves, volatility, etc), at least based on calendars

— an optimizer, taking a utility function as input (with few parameters), and providing a guideline for high frequency and opportunistic interactions with liquidity

— a collection of "short term robots", providing and consuming liquidity under the envelope provided by the optimizer.

► At the end of the day, operators are meant to provide a **post-trade analysis**, to check its consistency with the pre-trade analysis (and possible trigger reassessments of TC models)

► On a regular basis (weekly or monthly), the asset manager should take care of TCA report to take macroscopic decisions,

► The dealing desk should report to portfolio managers metrics to be taken into account at the allocation level.

Common Practices of Asset Management

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)

Common Practices of Asset Management

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)

Common Practices of Asset Management

Optimal Trading Rate

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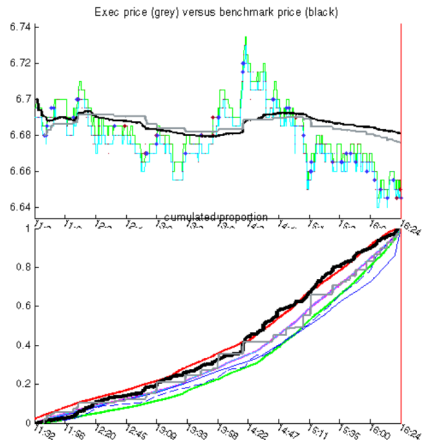
The first papers Almgren and Chriss (2000), Bertsimas and Lo (1998), focussed on the optimal trading rate, or trading speed (i.e. how many shares to buy or sell every 5 minutes) for long metaorders.

- ▶ 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.

Common Practices of Asset Management

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 $V(T, \ldots) = x + q(p - Aq)$

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

Common Practices of Asset Management

The Quadratic Solution is Nice

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 $-\mu x = \partial_t v - \phi\,q^2 + \kappa(\partial_Q v)^2/(4\kappa)$.

6. 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:

$$\begin{cases} 2\kappa\phi & = & -2\kappa h_2' & + & h_2^2 \\ -\mu & = & h_1' & - & 2h_1 h_2 \\ 0 & = & h_0' & + & h_1^2 \end{cases}$$

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.

Outline

Information to be Taken Into Account

"Fast" price moves and brokers' flows can be taken into account

---

Up to now, the trading schedulers (i.e. algorithms) have been built using two main components

- ▶ an "urgency" term, captured thanks to a "risk aversion coefficient";
- ▶ a trading-cost term, giving an incentive to slow down the trading.

In reality, we can go back to the **rationals for urgency**:

- ▶ trading "fast information" (like news)
- ▶ need for an "hedged portfolio" (beta-neutral, cash-neutral, or Tracking-Error constraints).

Information about the action of other participants could be added in the framework (see the MFG part of this talk for a model). Brokers always published on flows[1], and now they are selling information on their past flows...

The regulatory pressure and the available technology pushes dealing desks of asset managers to take control on the balance between the private information and the aggregated (more or less anonymous) information available. Brokers usually had this role, but they usually failed to not fall in a conflict of interest. See the "back-running" organized by brokers identified in (Di Maggio et al., 2016), thanks to ANcerno database.

---

[1] For instance, JP Morgan's corporate and investment bank has a twenty pages recurrent publication titled "*Flows and Liquidity*", with charts and tables describing different metrics like *money flows*, turnovers, and other similar metrics by asset class (equity, bonds, options) and countries. Almost each investment bank has such a publication for its clients; Barclay's one is titled "*TIC Monthly Flows*". As an example its August 2016 issue (9 pages) subtitle was "*Private foreign investors remain buyers*".

Information to be Taken Into Account

Even very short term flows (and inventories) can be exploited

With Eyal Neuman, we proposed recently a framework to take high frequency flows into account. Following the reasoning of (Foucault et al., 2012), part of the "slow, macroscopic" flows should be seen in orderbooks. Moreover, the "orderbook imbalance signal" is now well identified as the "*worst kept secret of HFT*".
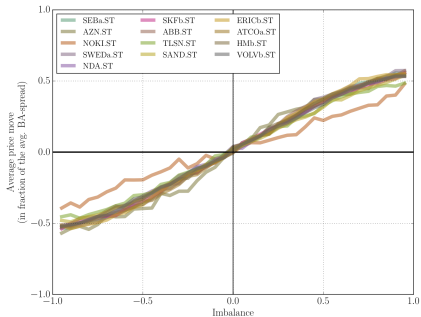
We expect the information contained in high frequency flows to be of two different natures(see (Huang et al., 2015) and subsequent papers for details):

- ► **imbalance of liquidity provision** (i.e. bid size vs. ask size, aggregated a "correct" way),
- ► **imbalance of liquidity consumption** (i.e. recent trades, seen via a propagator model or Hawkes processes).

We used the NASDAQ OMX trade database (before 2014) and RTH to match quotes, because the ID of the owners of the trades is disclosed; see (Van Kervel and Menkveld, forthcoming).

Information to be Taken Into Account

When do market participants accept to trade?

Avg. Imbalance Limit (left) and Market (right) Orders — Pct. of Limit Orders Sent

Legend: SEBa.ST, AZN.ST, NOKI.ST, SWEDa.ST, NDA.ST, SKFb.ST, ABB.ST, TLSN.ST, SAND.ST, ERICb.ST, ATCOa.ST, HMb.ST, VOLVb.ST
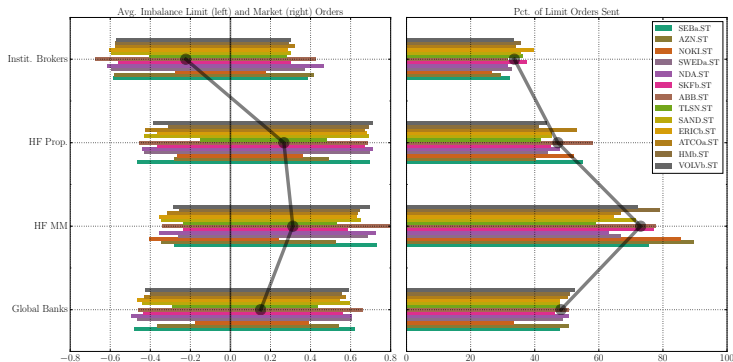
Use of limit and market orders and state of the imbalance before a trade, for each market participant. On the left panel: average imbalance just before a trade obtained with a limit order (left part, negative), and average imbalance just before a trade obtained with a market order (right part, positive). On the right panel: percentage of trades with a limit orders. The dark lines with large dots are averages.

Information to be Taken Into Account

How does it influences their trading speed?

We see that each participant has his own "signature", that is a characteristic of its constraints:

► Executing brokers have a large (blind) risk aversion: they cannot afford to wait for the "best opportunities" seen from the orderbook (De March and Lehalle, 2018)

► HF MM try to trade when the price will "not move", to earn the bid-ask spread by alternating buys and sells;

► HF Prop follow the signals with as few constraints as possible;

► International investment banks have a mix of activities.

Renormalized average trading rate in the direction of the imbalance $\hat{r}_+$ (solid line) and in the opposite direction $\hat{r}_-$ (dotted line), during 10 consecutive minutes, for each type of participant.

Summary Of This Section

Optimal decisions is meant to balance between information and costs

- ▶ The **execution layer** is now sophisticated enough to take a lot of effects into account;
- ▶ It focussed on a "simple" risk aversion vs. cost balance;
- ▶ But in practice there is a competition between algo designers to use as much information as possible: news, cross-sectional behaviours, and flows.
- ▶ The *allocation process* is (or should be) the result of an optimization, taking *expected costs* and not simply filter according to proxies;
- ☞ The two layers should share information, so that **investors can act as liquidity providers when needed**. "*Why no natural buyer was in the market during the US Flash Crash of May 2010?*"

Liquidity as a "Mean Field"

It is now possible to study the consequences of "liquidity games"

As intermediaries, investment banks are using different tools to hedge their exposure on markets. All these tools are inspired from the seminal work of (Bachelier, 1900). They all rely on an intensive use of *stochastic control.* This includes market making, optimal trading, derivative hedging, yield curve exposure, etc.

These "ultimate hedging trades" consume liquidity (it is often possible to first stop into different liquidity pools inside the bank; here I only consider the remaining trades, going outside of the bank).

► If they match natural liquidity providers (investors or corporates), they operate the expected matching of financial risk vs. real (economic) risk;

► If thy match similar hedging trades coming from other banks, they simply operate a netting inside the financial system;

► But if they meet no such counterparts, they pushe the price an unnoticed way (because their books are mark-to-market at close, i.e. impacted prices).

Thanks to recent mathematical tools (i.e. Mean Field Games, see Lasry and Lions (2006a,b, 2007) and Huang et al. (2006) for seminal papers), it is possible to model with an high accuracy the outcome of a time continuous game where players implement stochastic control. It opens a lot of perspectives. With Pierre Cardaliaguet, **we investigated how to use this framework for optimal trading, the simplest hedging game**.

Liquidity as a "Mean Field"

Mean Field Games

- ► 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.

Liquidity as a "Mean Field"

MFG of Controls and An Application To Trade Crowding

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

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

$$dQ_t^a = \nu_t^a \, dt, \tag{1}$$

now for a seller, $Q_0^a > 0$ (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*):

$$dX_t^a = -\nu_t^a (S_t + \kappa \cdot \nu_t^a) \, dt.$$

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 $\mathfrak{A}$.

Liquidity as a "Mean Field"

From Cost Function to HJB

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:

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

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).

Liquidity as a "Mean Field"

HJB For One Player (Backward Value Function)

The associated Hamilton-Jacobi-Bellman is

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

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

$$-\alpha\mu\, q = \partial_t v^a - \phi^a q^2 + \sup_\nu \left\{ \nu\partial_Q v^a - \kappa\nu^2 \right\}, \text{ with the teminal condition } v^a(T, q; \mu) = -A^a q^2.$$

The associated optimal feedback / control is straightforward to find:

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

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

Liquidity as a "Mean Field"

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)

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

Liquidity as a "Mean Field"

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.

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

Under boundary (resp. initial and terminal) conditions:

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

Liquidity as a "Mean Field"

Explicit Solution For a Special Case: same preferences for all agents: $\phi^a \equiv \phi$, $A^a \simeq A$

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:

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

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:

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

Liquidity as a "Mean Field": Qualitative Meaning of All This

(Back To Identical Preferences To Ease The Writing)

Dependence of the Solution to the Mean Field

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

- ▶ 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** : $h_1(t) = 2\kappa \cdot E'(t) + h_2(t) \cdot E(t)$.

- ⇒ locally you adapt your behaviour to the mean field via h_1,

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

Liquidity as a "Mean Field": Qualitative Meaning of All This

The Standard Case

$E(t)$, $E_0 = 10.00$, $\alpha = 0.400$

$T = 5.00$, $A = 2.50$, $\kappa = 0.20$, $\phi = 0.10$
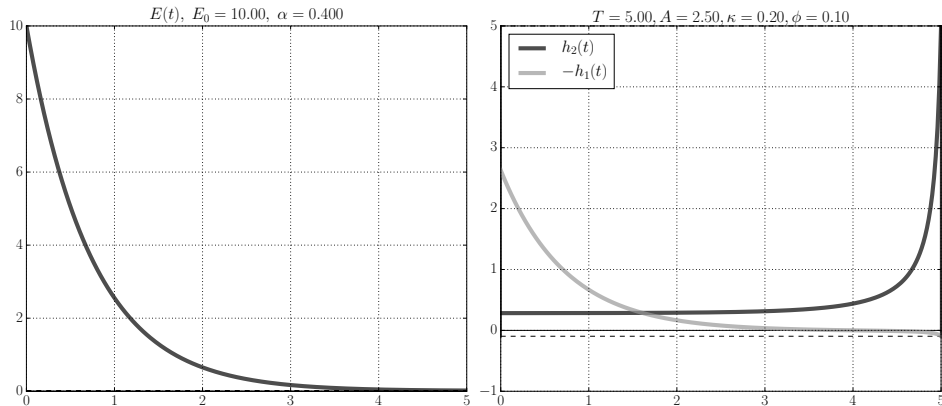
Dynamics of E (left) and $-h_1$ and h_2 (right) for a standard set of parameters: $\alpha = 0.4$, $\kappa = 0.2$, $\phi = 0.1$, $A = 2.5$, $T = 5$, $E_0 = 10$.

Liquidity as a "Mean Field": Qualitative Meaning of All This

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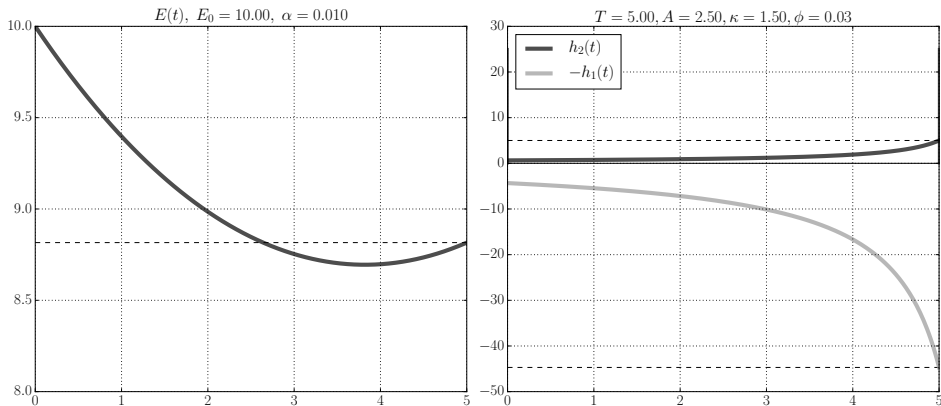
A specific case for which *E* is not monotonous: $\alpha = 0.01$, $\kappa = 1.5$, $\phi = 0.03$, $A = 2.5$, $T = 5$ and $E_0 = 10$.

Outline

Conclusions

Main take-away

1. Ex-ante measure of transaction costs are needed to allocated capital accordingly.
2. The transaction costs are not only a function of the "liquidity", but of the "market context", and the "investment flows" too.
3. Market Impact is an important component of transaction costs and should be taken into account. Recent empirical evidences on cross-impact (Benzaquen et al., 2017) have to be considered.
4. On corporate bonds, it is possible to conduct in depth TCA. Next step should be a cross-instrument analysis (1: taking bonds on the same –or correlated– underlying into account, 2: Bond $\simeq$ Govies + CDS).
5. In practice, two "decisions" are taken: allocation and trade scheduling, they are well understood and should be coupled some way. Splitting the two scale is good for operational risk, nevertheless aggregated information on TC should feed allocation.
6. The "high frequency" decisions/splitting should take low and high frequency flows into account to allow investors to compete with MM and HFT.
☞ Crowding and game-related issues can be studied thanks to MFG.

Regulatory Challenges I

Measuring the effect of regulation on investors (not on traders), and be sure traders implement risk-control features

Transaction costs transparency and an adequate capital allocation (i.e. taking the real illiquidity premium into account), is of paramount importance from a regulatory viewpoint. During the "10 years after Lehman" conference, a representative of the ECB said "*when liquidity is not priced properly, the financial system takes more risk than we think*".

▶ First of all, the Transaction Costs (TC) are more than an "illiquidity premium", they vary with market conditions (including flows of the day). Hence **regulators should be able to compute realized transaction costs for each type of investors**, and monitor investment flows. This means to have an ANcerno-like database of all metaorders for investors. And other databases for other market participants (brokers, banks, market makers, HFT, etc).

▶ The concept of "self-financed" is crucial for indexing, but not present in the Benchmark Regulation... TC should be more seriously taken into account than "filtering" or weights proportional to capitalization. At least "TC cost reviews" could take place. The main issue would be that TC are a function of the Asset under Management, hence such a review should be partly model-driven.

Regulatory Challenges II

Measuring the effect of regulation on investors (not on traders), and be sure traders implement risk-control features

- ► **Simple averages are not enough**, "liquidity" that affects TC is dynamic (Binkowski and Lehalle, 2018), hence an average improvement of the bid-ask spread cannot be considered as an improvement for any market participant (Lachapelle et al., 2016).
  It means adequate measurements will be CPU and memory consuming... What about an ESMA-wide computing facility?

- ► MiFID 2 provided a clear cut between Brokers' Fundamental Research and Execution Services will put pressure and competition on both sides. On the execution side, some algo designers may be tempted to put too much weight on "signals" (especially with the raise of Machine Learning). The risk-control layer of algo have to be assessed: "*what will your algo do if it is too late, and how does it defines being 'too late'? Is it 'market-impact aware'?*"

- ► In this talk I did not considered the matching mechanism (CLOB, RFQ, etc), nevertheless, all the optimal trading literature is **in favour of continuous trading**. It can of course address fragmentation and dark pools (Laruelle et al., 2011), but from a welfare viewpoint it carries a simple message: with risk-controlled continuous trading, you increase the probability to find a match (if other players do the same). It is nevertheless clear that if some large players do not accept this practice, optimal trading will provide good liquidation strategies, but the TC will improve for all participants.

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

Amber Anand, Paul Irvine, Andy Puckett, and Kumar Venkataraman. Performance of institutional trading desks: An analysis of persistence in trading costs. *The Review of Financial Studies*, 25(2):557–598, 2011. URL http://www.sbuweb.tcu.edu/pirvine/Rev.%20Financ.%20Stud.-2012-Anand-557-98.pdf.

Paul Asquith, Andrea S. Au, Thomas Covert, and Parag A. Pathak. The market for borrowing corporate bonds. *Journal of Financial Economics*, 107(1):155–182, 2013.

Louis Bachelier. *Théorie de la spéculation*. PhD thesis, 1900. URL http://www.numdam.org/numdam-bin/fitem?id=ASENS_1901_3_18__143_0.

Emmanuel Bacry, Adrian Iuga, Matthieu Lasnier, and Charles-Albert Lehalle. Market Impacts and the Life Cycle of Investors Orders. *Market Microstructure and Liquidity*, 1(2), December 2015. URL http://ssrn.com/abstract=2532152.

M. Benzaquen, I. Mastromatteo, Z. Eisler, and J. P. Bouchaud. Dissecting cross-impact on stock markets: an empirical analysis. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(2):023406, 2017. URL http://stacks.iop.org/1742-5468/2017/i=2/a=023406.

Nataliya Bershova and Dmitry Rakhlin. The Non-Linear Market Impact of Large Trades: Evidence from Buy-Side Order Flow. *Social Science Research Network Working Paper Series*, January 2013. URL http://ssrn.com/abstract=2197534.

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

Dimitris Bertsimas, Vishal Gupta, and Ioannis C. Paschalidis. Inverse optimization: a new perspective on the Black-Litterman model. *Operations research*, 60(6):1389–1403, 2012.

Hendrik Bessembinder, William Maxwell, and Kumar Venkataraman. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics*, 82(2):251–288, 2006. URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.475.6120&rep=rep1&type=pdf.

Bruno Biais and Richard C. Green. The Microstructure of the Bond Market in the 20th Century. Technical report, Carnegie Mellon University, 2007.

Mikolaj Binkowski and Charles-Albert Lehalle. Endogenous Dynamics of Intraday Liquidity, November 2018. URL http://arxiv.org/abs/1811.03766.

X. Brokmann, E. Serie, J. Kockelkoren, and J. P. Bouchaud. Slow decay of impact in equity markets. *Market Microstructure and Liquidity*, 1(2), July 2015. URL http://arxiv.org/abs/1407.3390.

Frédéric Bucci, Iacopo Mastromatteo, Zoltán Eisler, Fabrizio Lillo, Jean-Philippe Bouchaud, and Charles-Albert Lehalle. Co-impact: Crowding effects in institutional trading activity, July 2018. URL http://arxiv.org/abs/1804.09565.

Pierre Cardaliaguet and Charles-Albert Lehalle. Mean Field Game of Controls and An Application To Trade Crowding. *Mathematics and Financial Economics*, pages 1–29, December 2017. URL http://arxiv.org/abs/1610.09904.

Rene Carmona, Jean-Pierre Fouque, and Li-Hsien Sun. Mean Field Games and Systemic Risk. *Social Science Research Network Working Paper Series*, August 2013. URL http://arxiv.org/abs/1308.2172.

Álvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and High-Frequency Trading (Mathematics, Finance and Risk)*. Cambridge University Press, 1 edition, October 2015. ISBN 1107091144. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/1107091144.

Sugato Chakravarty and Asani Sarkar. Trading costs in three US bond markets. *The Journal of Fixed Income*, 13(1):39–48, 2003.

Bibliography III

Tarun Chordia, Asani Sarkar, and Avanidhar Subrahmanyam. An Empirical Analysis of Stock and Bond Market Liquidity. *Review of Financial Studies*, 18(1):85–129, March 2005. ISSN 1465-7368. doi: 10.1093/rfs/hhi010. URL http://www.econstor.eu/bitstream/10419/60642/1/363003657.pdf.

Bruce M. Collins and Frank J. Fabozzi. A methodology for measuring transaction costs. *Financial Analysts Journal*, 47(2):27–36, 1991. URL https://faculty.fuqua.duke.edu/~{}charvey/Teaching/BA453_2006/Collins_A_methodology_for_1991.pdf.

Hadrien De March and Charles-Albert Lehalle. Optimal trading using signals, November 2018. URL http://arxiv.org/abs/1811.03718.

Marco Di Maggio, Francesco Franzoni, Amir Kermani, and Carlo Sommavilla. The Relevance of Broker Networks for Information Diffusion in the Stock Market. Technical Report 63, Swiss Finance Institute Research Paper, October 2016. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2860118.

Jens Dick-Nielsen, Peter Feldhütter, and David Lando. Corporate bond liquidity before and after the onset of the subprime crisis. *Journal of Financial Economics*, 103(3):471–492, March 2012. ISSN 0304405X. doi: 10.1016/j.jfineco.2011.10.009. URL http://dx.doi.org/10.1016/j.jfineco.2011.10.009.

Jean-David Fermanian, Olivier Guéant, and Arnaud Rachez. Agents' Behavior on Multi-Dealer-to-Client Bond Trading Platforms, November 2015. URL http://arxiv.org/abs/1511.07773.

Thierry Foucault, Johan Hombert, and Ioanid Rosu. News Trading and Speed. *Social Science Research Network Working Paper Series*, December 2012. URL http://ssrn.com/abstract=2188822.

Nils Friewald and Florian Nagler. Dealer Inventory and the Cross-Section of Corporate Bond Returns. *Social Science Research Network Working Paper Series*, November 2014. URL https://editorialexpress.com/cgi-bin/conference/download.cgi?db_name=AFA2016&paper_id=163.

Nicolae B. Garleanu and Lasse H. Pedersen. Dynamic Trading with Predictable Returns and Transaction Costs. *Journal of Finance*, 68(6):2309–2340, December 2013. URL http://www.nber.org/papers/w15205.

Mark B. Garman and Michael J. Klass. On the Estimation of Security Price Volatility from Historical Data. *Journal of Business*, 53 (1):67–78, 1980.

Michael A. Goldstein, Edith S. Hotchkiss, and Erik R. Sirri. Transparency and Liquidity: A Controlled Experiment on Corporate Bonds. *Review of Financial Studies*, 20(2):235–273, March 2007. ISSN 1465-7368. doi: 10.1093/rfs/hhl020. URL http://dx.doi.org/10.1093/rfs/hhl020.

Olivier Guéant. *The Financial Mathematics of Market Liquidity: From Optimal Execution to Market Making*. Chapman and Hall/CRC, April 2016. ISBN 1498725473. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/1498725473.

Xin Guo, Tze L. Lai, Howard Shek, and Samuel P. Wong. *Quantitative Trading: Algorithms, Analytics, Data, Models, Optimization*. Chapman and Hall/CRC, December 2016. ISBN 1498706487. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/1498706487.

Lawrence Harris. Transaction Costs, Trade Throughs, and Riskless Principal Trading in Corporate Bond Markets. *Social Science Research Network Working Paper Series*, September 2015. URL http://ssrn.com/abstract=2661801.

Campbell R. Harvey, Yan Liu, and Heqing Zhu. ...and the Cross-Section of Expected Returns. *Social Science Research Network Working Paper Series*, April 2013. URL http://ssrn.com/abstract=2249314.

Terrence Hendershott and Ananth Madhavan. Click or Call? Auction Versus Search in the Over-the-Counter Market. *The Journal of Finance*, 70(1):419–447, November 2015. URL http://ssrn.com/abstract=1961350.

Minyi Huang, Roland P. Malhamé, Peter E. Caines, and Others. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3): 221–252, 2006.

Weibing Huang, Charles-Albert Lehalle, and Mathieu Rosenbaum. Simulating and analyzing order book data: The queue-reactive model. *Journal of the American Statistical Association*, 10(509), December 2015. doi: 10.1080/01621459.2014.982278. URL http://dx.doi.org/10.1080/01621459.2014.982278.

Robert Kissell, Morton Glantz, and Roberto Malamut. A practical framework for estimating transaction costs and developing optimal trading strategies to achieve best execution. *Finance Research Letters*, 1(1):35–46, 2004. URL http://isiarticles.com/bundles/Article/pre/pdf/19891.pdf.

Aimé Lachapelle, Jean-Michel Lasry, Charles-Albert Lehalle, and Pierre-Louis Lions. 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, June 2016. URL http://link.springer.com/article/10.1007/s11579-015-0157-1.

Sophie Laruelle, Charles-Albert Lehalle, and Gilles Pagès. Optimal split of orders across liquidity pools: a stochastic algorithm approach. *SIAM Journal on Financial Mathematics*, 2:1042–1076, 2011. URL http://epubs.siam.org/doi/abs/10.1137/090780596.

Jean-Michel Lasry and Pierre-Louis Lions. Jeux à champ moyen. I - Le cas stationnaire. *Comptes Rendus Mathematique*, 343 (9):619–625, November 2006a. ISSN 1631073X. doi: 10.1016/j.crma.2006.09.019. URL http://dx.doi.org/10.1016/j.crma.2006.09.019.

Jean-Michel Lasry and Pierre-Louis Lions. Jeux à champ moyen. II - Horizon fini et contrôle optimal. *Comptes Rendus Mathematique*, 343(10):679–684, November 2006b. ISSN 1631073X. doi: 10.1016/j.crma.2006.09.018. URL http://dx.doi.org/10.1016/j.crma.2006.09.018.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, March 2007. ISSN 0289-2316. doi: 10.1007/s11537-007-0657-8. URL http://dx.doi.org/10.1007/s11537-007-0657-8.

Charles-Albert Lehalle and Eyal Neuman. Incorporating Signals into Optimal Trading, April 2017. URL http://arxiv.org/abs/1704.00847.

Charles-Albert Lehalle, Sophie Laruelle, Romain Burgot, Stéphanie Pelin, and Matthieu Lasnier. *Market Microstructure in Practice*. World Scientific publishing, 2013. URL http://www.worldscientific.com/worldscibooks/10.1142/8967.

Charles-Albert Lehalle, Sophie Laruelle, Romain Burgot, Stéphanie Pelin, and Matthieu Lasnier. *Market Microstructure in Practice*. World Scientific publishing, 2018.

Harry Markowitz. Portfolio Selection. *The Journal of Finance*, 7(1):77–91, March 1952. ISSN 00221082. doi: 10.2307/2975974. URL http://dx.doi.org/10.2307/2975974.

Robert C. Merton. Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *The Review of Economics and Statistics*, 51(3), 1969. URL http://www.jstor.org/stable/1926560.

Bruce Mizrach. Analysis of Corporate Bond Liquidity. Technical report, FINRA, December 2015.

Esteban Moro, Javier Vicente, Luis G. Moyano, Austin Gerig, Doyne J. Farmer, Gabriella Vaglica, Fabrizio Lillo, and Rosario N. Mantegna. Market impact and trading profile of large trading orders in stock markets, August 2009. URL http://arxiv.org/abs/0908.0202.

Alessio Ruzza. Agency Issues in Corporate Bond Trading. Technical report, SSRN, December 2016. URL https://ssrn.com/abstract=2759766.

Paul Schultz. Corporate Bond Trading Costs: A Peek Behind the Curtain. *The Journal of Finance*, 56(2):677–698, April 2001. doi: 10.1111/0022-1082.00341. URL https://www3.nd.edu/~{}pschultz/paper1.pdf.

S. E. Shreve and H. M. Soner. Optimal Investment and Consumption with Transaction Costs. *The Annals of Applied Probability*, 4(3):609–692, August 1994.

Damian E. Taranto, Giacomo Bormetti, Jean-Philippe Bouchaud, Fabrizio Lillo, and Bence Toth. Linear models for the impact of order flow on prices I. Propagators: Transient vs. History Dependent Impact, February 2016. URL http://arxiv.org/abs/1602.02735.

Vincent Van Kervel and Albert Menkveld. High-frequency trading around large institutional orders. *Journal of Finance*, forthcoming.

Henri Waelbroeck and Carla Gomes. Is Market Impact a Measure of the Information Value of Trades? Market Response to Liquidity vs. Informed Trades. *Social Science Research Network Working Paper Series*, July 2013. URL http://ssrn.com/abstract=2291720.

Outline

Outline

The ANcerno Database

The fragmentation can be easily read in these data: more trade scheduling and a pressure on fees

	Number of Tickets (Million)	Av Ticket Size ($ Million)	Av Partici- pation Rate (%)	Execution period (days)	Av Number of brokers	Commi- ssions (bps)
Full Sample	90.35	0.59	1.98	1.82	1.16	9.30
1999	1.92	1.12	6.42	2.85	1.30	11.52
2000	3.02	1.09	4.43	2.24	1.25	12.47
2001	3.71	0.87	3.89	2.03	1.28	14.62
2002	4.30	0.80	3.82	2.15	1.25	18.85
2003	4.99	0.62	3.53	1.96	1.22	17.08
2004	6.18	0.92	3.38	1.96	1.18	13.00
2005	6.73	0.80	2.74	1.79	1.18	9.99
2006	9.47	0.56	1.74	1.56	1.14	9.73
2007	10.23	0.55	1.48	1.50	1.15	7.07
2008	11.66	0.48	0.90	1.40	1.16	8.60
2009	10.31	0.39	0.94	1.47	1.11	8.05
2010	10.62	0.37	0.84	1.35	1.13	4.04
Jan-Sept 2011	7.20	0.36	0.71	1.40	1.09	2.94

Outline

Corporate Bonds

Literature review

Reference	Dataset(s) Name(s)	Period covered
Schultz (2001)	CAI	1995-1997
Chakravarty and Sarkar (2003)	NAIC	1995-1997
Chordia et al. (2005)	TAQ, ISSM (Nyse), GovPX	1991-1998
Bessembinder et al. (2006)	NAIC + TRACE	2001 ; 2002
Goldstein et al. (2007)	TRACE	2002-2004
Biais and Green (2007)	Nyse Archives	1926-1930; 1943-1948
Dick-Nielsen et al. (2012)	TRACE	2003-2005
Asquith et al. (2013)	TRACE	2004-2007
Friewald and Nagler (2014)	Labelled TRACE	2003-2013
Mizrach (2015)	TRACE	2002-2007
Hendershott and Madhavan (2015)	TRACE + MarketAxes	2011-2011
Ruzza (2016)	TRACE	2004-2012

Quick list of empirical papers on transaction costs of corporate bonds

Corporate Bonds

The TRACE database

	Removal (nbe)	Removal (pct)	Number left
Step 1 Keep settled trades	1,877,866	5.4%	32,931,539
Step 2 Keep trades reported by dealers	2,095,934	6.36%	30,835,605
Step 3 Keep business days	5,735	0.02%	30,829,853
Step 4 Keep opened hours	745,619	2.4%	30,084,233
Step 5 Keep regular trades	583,157	1.9%	29,501,076
Step 6 Keep compatible prices	217,321	0.074%	29,182,755
Step 7 Keep bonds only	563,942	1.94%	28,719,813
Selection LR For BA-spread regression	–	–	4,371,363
Selection PI For price impact curves	–	–	3,564,264

Two-step Lasso regression table: The impact on bid-ask spread (in bp)

	Model OLS All features	Model L1 $\lambda_l = 0.2$	Model L2 $\lambda_l = 0.6$	Model L3 $\lambda_l = 1.0$
Volatility	4.9123***	5.1053***	5.2836***	5.4709***
	(0.020)	(0.016)	(0.014)	(0.015)
Log(Number of trades)	45.6241***	42.1450***	39.0703***	
	(0.544)	(0.451)	(0.433)	
Log(Total Volume)	−21.3397***	−21.1246***	−19.4449***	
	(0.272)	(0.252)	(0.243)	
Coupon	−0.4616***	1.2071***		
	(0.120)	(0.089)		
Duration	1.4201***	0.2453***		
	(0.3546)	(0.035)		
Year-to-maturity	−0.0643	−0.1078		
	(0.076)	(0.075)		
Issued years	1.2576***	−0.0216	0.2336***	1.0396***
	(0.065)	(0.052)	(0.046)	(0.011)
Constant	46.055	86.1737	85.6982	−2.9270
N	152.408	152.408	152.408	152.408
R^2	55.4 %	52.8 %	49.2 %	43.22 %

Standard errors in parenthesis. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Two-tailed test.
Source: Enhanced TRACE(2015-2016).

Outline

Dynamics For Identical Preferences

We now collect all the equations:

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

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 Main Equation For Identical Preferences

The previous system of ordinary differential equations implies

$$
\text{(4)} \qquad\qquad 0 = 2\kappa E''(t) + \alpha E'(t) - 2\phi E(t)
$$

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

Solving The Mean Field

> ### 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
>
> $$E(t) = E_0 a \left(\exp\{r_+ t\} - \exp\{r_- t\} \right) + E_0 \exp\{r_- t\}$$
>
> where a is given by
>
> $$a = \frac{(\alpha/4 + \kappa\theta - A)\exp\{-\theta T\}}{-\frac{\alpha}{2}\text{sh}\{\theta T\} + 2\kappa\theta\text{ch}\{\theta T\} + 2A\text{sh}\{\theta T\}},$$
>
> 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 the Control

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

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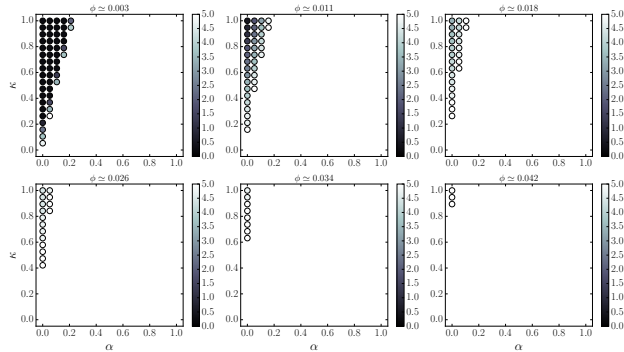
Solving $h_1(t)$

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

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

Liquidity as a "Mean Field": Qualitative Meaning of All This

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.