Financial Markets Microstructure

Lecture 7

Depth determinants

Chapter 4 of FPR

Egor Starkov

Københavns Unversitet

Spring 2025

Previously on...

1. The spread is not only driven by adverse selection: order costs and inventory risk have an effect as well

Homework from last time

We said today that inventory risk is priced when the dealer is risk-averse. Risk-aversion is one explanation, but other factors can also contribute to inventory risk. The two following cases explore this issue:

- A big trader was punted off the Nordic power market after failing to meet margin calls (two articles on absalon).
  - How does inventory risk manifest in this story?
  - Explain why such inventory risk can be priced even by risk-neutral agents.
- Negative oil futures prices were registered last year (blog post on absalon or here).
  - Why did it happen? How do negative prices make sense?
  - How does inventory risk manifest in this story?
  - Explain why such inventory risk can be priced even by risk-neutral agents.

Inventory concerns more broadly

- We have discussed how dealers' risk-aversion can drive the spread, depth, and make prices deviate from the efficient level
- Two comments on that, from the cases you read:
- **Point 1**: risk-aversion in markets may stem from market risks, rather than inherent risk-aversion in preferences.
  - Standard story: $u(w)$ is concave in future wealth (e.g., MeanVar/CARA/CRRA prefs), and $w \sim z \cdot v + \ldots$ (position $\times$ asset value), hence $u(v)$ is concave in v, more so for higher z.
  - Alternative: $u(w)$ is linear (risk-neutrality), but low v creates higher risk of margin calls, which are costly: $w \sim z \cdot (v - c \cdot \mathbb{I}\{v < \underline{v}, z > 0\}) + \ldots \Rightarrow u(v)$ is again concave in v.
  - Either story leads to dealer's inventory affecting their willingness to buy/sell

Inventory concerns more broadly

- **Point 2**: what if traders in the market are risk-averse, and not just the dealer?
  - If traders provide liquidity (e.g., we are in LOB market and not a dealer market) – same inventory risks
  - If general market populace is risk-averse: in a similar way, traders' valuation for the asset would depend on how far their current position is from their ideal position.
  - So if there is some aggregate imbalance – i.e., current aggregate holdings (many traders long on oil futures) are different from aggregated ideal positions (everyone wants to dump their futures) – then market price might deviate from the fundamental value
  - (Although the question to ask is: why did such discrepancy in positions arise in the first place, and is it by itself informative about the fundamentals)

Today

- **Trade size**
  - How does trade size affect prices?
  - I.e., what determines market depth?
  - (Spoiler: mostly the same factors as with liquidity)
  - Will look at Kyle (1985) model – an alternative to GM that allows flexible trade size

Prices and trade size

- How does trade size affect prices?
  - Spread larger for large trades, price moves further from efficient level
  - I.e., market has limited depth
- Why?

Prices and trade size

- How does trade size affect prices?
  - Spread larger for large trades, price moves further from efficient level
  - I.e., market has limited depth
- Why?
  1. Adverse selection: larger trades indicate more/stronger news

Prices and trade size

- How does trade size affect prices?
  - Spread larger for large trades, price moves further from efficient level
  - I.e., market has limited depth
- Why?
  1. Adverse selection: larger trades indicate more/stronger news
  2. Inventory risk: large positions are risky and take dealers longer to unwind, hence require larger premiums

Prices and trade size

- How does trade size affect prices?
  - Spread larger for large trades, price moves further from efficient level
  - I.e., market has limited depth
- Why?
  1. Adverse selection: larger trades indicate more/stronger news
  2. Inventory risk: large positions are risky and take dealers longer to unwind, hence require larger premiums
  3. Imperfectly competitive dealers: market power allows dealers to set wider spread and steeper or flatter pricing schedules

Prices and trade size

- How does trade size affect prices?
  - Spread larger for large trades, price moves further from efficient level
  - I.e., market has limited depth
- Why?
  1. Adverse selection: larger trades indicate more/stronger news
  2. Inventory risk: large positions are risky and take dealers longer to unwind, hence require larger premiums
  3. Imperfectly competitive dealers: market power allows dealers to set wider spread and steeper or flatter pricing schedules
  4. Order processing costs: may increase or decrease (per stock) in total order size

Kyle model

- We will look at Kyle (1985) model which links market depth to adverse selection
- It can be extended to accomodate imperfect competition among dealers (see 4.2.4) and inventory risk (4.3)
    - the inventory risk version is broadly similar to the Stoll model that we skipped
    - trading costs are very difficult to incorporate in this model

Setup: Broad strokes

- A call auction; orders come from a "large" speculator and a population of noise traders; market cleared by a dealer.
- Speculator/informed trader: has private information
  - Trades using a 'large' speculative market order
  - Strategically moderates order size to reduce price impact
  - 'Hides' behind noise traders who submit a random size order
- Representative market maker (MM)/dealer
  - Risk neutral and competitive (zero profits)
  - Clears orders in *batches* (as opposed to one-by-one in Glosten & Milgrom)
  - Cannot distinguish speculative orders from noise orders in a batch

Setup

- Asset: Trade in one risky asset with value $v \sim \mathcal{N}(\mu, \sigma_v^2)$
- Speculator: Observes true value v (perfect information)
  - Places market order x
  - If the order clears at price p: gain is $x(v - p)$
  - Does not know p when choosing x: maximizes *expected* gain (risk neutral) given $\mathbb{E}[p|x]$
- Noise trader: Has random demand $u \sim \mathcal{N}(0, \sigma_u^2)$
- MM: Submits a supply schedule of (q, p) combinations:
  - "If the order imbalance is $q = x + u$, I will absorb it at price p"
  - Observes aggregate flow $q = x + u$, but not x and u
  - Competitive (zero profit): $p = \mathbb{E}[v|q]$
- Assumption: u and v are jointly normal and independent

Setup: Timing

- To be explicit, the timing is as follows:

1. at the beginning of the period:
    - speculator chooses order size x
    - noise traders submit their order u
    - dealer submits price schedule (q, p)
2. then market price $p(q)$ is determined given total order $q = x + u$
3. at the end of the period payoffs are realized

Linear equilibrium: outline

- The equilibrium is described by the speculator's strategy $x(v)$ and the dealer's pricing schedule $p(q)$.

Linear equilibrium: outline

- The equilibrium is described by the speculator's strategy $x(v)$ and the dealer's pricing schedule $p(q)$.
- Look for equilibrium where speculator's strategy is linear: $x = \beta(v - \mu)$
  - Note: β is endogenously determined by the equilibrium, we'll derive it
  - $\beta > 0$ measures speculator *aggression*
- MM's pricing is driven by the zero-profit condition: $p = \mathbb{E}[v|q]$
  - In eqm, MM knows the speculator's strategy $(x = \beta(v - \mu))$
  - So MM observes $q = x + u = \beta(v - \mu) + u$, and wants to estimate v

Aside: Hitchhiker's guide to normal updating

Updating normal beliefs with normal signals might seem daunting at first, but here's how it works.

Aside: Hitchhiker's guide to normal updating

Updating normal beliefs with normal signals might seem daunting at first, but here's how it works.

- Suppose there is some uncertain variable x and our prior belief is that $x \sim \mathcal{N}(\mu_x, \sigma_x^2)$
- This prior belief is said to have precision $\tau_x \equiv \frac{1}{\sigma_x^2}$.

Aside: Hitchhiker's guide to normal updating

Updating normal beliefs with normal signals might seem daunting at first, but here's how it works.

- Suppose there is some uncertain variable x and our prior belief is that $x \sim \mathcal{N}(\mu_x, \sigma_x^2)$
- This prior belief is said to have precision $\tau_x \equiv \frac{1}{\sigma_x^2}$.
- Suppose we get signal y about x with precision τ_ϵ.
  I.e., we observe $y = x + \epsilon$, where $\epsilon \sim \mathcal{N}\left(0, \frac{1}{\tau_\epsilon}\right)$ (so $\tau_\epsilon \equiv \frac{1}{\sigma_\epsilon^2}$).

Aside: Hitchhiker's guide to normal updating

Updating normal beliefs with normal signals might seem daunting at first, but here's how it works.

- Suppose there is some uncertain variable x and our prior belief is that $x \sim \mathcal{N}(\mu_x, \sigma_x^2)$
- This prior belief is said to have precision $\tau_x \equiv \frac{1}{\sigma_x^2}$.
- Suppose we get signal y about x with precision τ_ϵ.
  I.e., we observe $y = x + \epsilon$, where $\epsilon \sim \mathcal{N}\left(0, \frac{1}{\tau_\epsilon}\right)$ (so $\tau_\epsilon \equiv \frac{1}{\sigma_\epsilon^2}$).
- Then our posterior belief about x has precision $\tau_{x|\epsilon} = \tau_x + \tau_\epsilon$, and the posterior mean weighs the prior mean μ_x and the signal y according to their precisions:

$$x|y \sim \mathcal{N}\left(\frac{\tau_x}{\tau_x + \tau_\epsilon}\mu_x + \frac{\tau_\epsilon}{\tau_x + \tau_\epsilon}y,\ \frac{1}{\tau_{x|\epsilon}}\right)$$

(you can verify this by directly calculating the conditional pdf)

Deriving the distribution of $v|q$

In our case:

- $v \sim \mathcal{N}(\mu, \sigma_v^2)$,
- $q = \beta(v - \mu) + u$, where $u \sim \mathcal{N}(0, \sigma_u^2)$, $u \perp v$.

Deriving the distribution of $v|q$

In our case:

- $v \sim \mathcal{N}(\mu, \sigma_v^2)$,
- $q = \beta(v - \mu) + u$, where $u \sim \mathcal{N}(0, \sigma_u^2)$, $u \perp v$.
- So the prior precision (of the belief about v) is $\tau_v \equiv \frac{1}{\sigma_v^2}$,
- and the signal is $\tilde{q} \equiv \frac{q}{\beta} + \mu = v + \frac{u}{\beta}$, which has precision $\tau_{\tilde{q}} = \frac{\beta^2}{\sigma_u^2}$

Deriving the distribution of $v|q$

In our case:

- $v \sim \mathcal{N}(\mu, \sigma_v^2)$,
- $q = \beta(v-\mu) + u$, where $u \sim \mathcal{N}(0, \sigma_u^2)$, $u \perp v$.
- So the prior precision (of the belief about v) is $\tau_v \equiv \frac{1}{\sigma_v^2}$,
- and the signal is $\tilde{q} \equiv \frac{q}{\beta} + \mu = v + \frac{u}{\beta}$, which has precision $\tau_{\tilde{q}} = \frac{\beta^2}{\sigma_u^2}$
- Then

$$
\begin{aligned}
v|q &\sim \mathcal{N}\left(\frac{\tau_v}{\tau_v + \tau_{\tilde{q}}}\mu + \frac{\tau_{\tilde{q}}}{\tau_v + \tau_{\tilde{q}}}\tilde{q},\ \frac{1}{\tau_v + \tau_{\tilde{q}}}\right) \\
&\sim \mathcal{N}\left(\mu + \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2 + \sigma_u^2} q,\ \left(\frac{1}{\sigma_v^2} + \frac{\beta^2}{\sigma_u^2}\right)^{-1}\right)
\end{aligned}
$$

Dealer's strategy

Going back to the zero profit condition:

$$
\begin{aligned}
& p = \mathbb{E}[v|q] \\
\iff & p = \mu + \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2 + \sigma_u^2} q \\
\iff & p = \mu + \lambda q, \\
\text{where } \lambda \equiv & \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2 + \sigma_u^2} \; \left(= \frac{\mathbb{C}(v, q)}{\mathbb{V}(q)} \right)
\end{aligned}
$$

Here λ is the price impact coefficient.
Conversely, $1/\lambda$ is a measure of market depth

Speculator's strategy

Back to the speculator's problem. The speculator takes for granted the MM's pricing rule $p = \mu + \lambda q$

- Speculator's profit is $\Pi(x) = x(v - p) = x(v - \mu - \lambda x - \lambda u)$
- Expected profit is $\mathbb{E}[\Pi(x)] = x(v - \mu - \lambda x)$, since $\mathbb{E}[u] = 0$
- Speculator chooses x to maximize $\mathbb{E}[\Pi(x)]$. Using the first-order condition:

$$
\begin{aligned}
v - \mu - 2\lambda x &= 0 \\
&\Rightarrow x = \beta(v - \mu),
\end{aligned}
$$

where $\beta = 1/(2\lambda)$

- Confirmed that it is optimal for the speculator to use a linear strategy!
- Note analogy to monopoly problem:

Speculator's strategy

Back to the speculator's problem. The speculator takes for granted the MM's pricing rule $p = \mu + \lambda q$

- Speculator's profit is $\Pi(x) = x(v - p) = x(v - \mu - \lambda x - \lambda u)$
- Expected profit is $\mathbb{E}[\Pi(x)] = x(v - \mu - \lambda x)$, since $\mathbb{E}[u] = 0$
- Speculator chooses x to maximize $\mathbb{E}[\Pi(x)]$. Using the first-order condition:

$$v - \mu - 2\lambda x = 0$$
$$\Rightarrow x = \beta(v - \mu),$$
$$\text{where } \beta = 1/(2\lambda)$$

- Confirmed that it is optimal for the speculator to use a linear strategy!
- Note analogy to monopoly problem: trade-off b/w trading more and trading at better price

Closing the equilibrium

- Finally, 'match' the coefficients:

$$\frac{1}{2\beta} = \lambda = \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2 + \sigma_u^2}$$

  i.e. $\beta^2\sigma_v^2 + \sigma_u^2 = 2\beta^2\sigma_v^2$ which yields

$$\beta = \frac{\sigma_u}{\sigma_v} \text{ and } \lambda = \frac{\sigma_v}{2\sigma_u}.$$

- Thus: the strategies are optimal given the prices, and the prices optimal given the strategies $\rightarrow$ **equilibrium**

Equilibrium properties

$$\beta = \frac{\sigma_u}{\sigma_v} \text{ and } \lambda = \frac{\sigma_v}{2\sigma_u}.$$

- Speculator is more **aggressive** (β higher) when:
  1. The informational advantage σ_v is smaller (why?)
  2. There's more noise σ_u to hide behind (why?)

Equilibrium properties

$$\beta = \frac{\sigma_u}{\sigma_v} \text{ and } \lambda = \frac{\sigma_v}{2\sigma_u}.$$

- Speculator is more **aggressive** (β higher) when:
  1. The informational advantage σ_v is smaller (why?)
  2. There's more noise σ_u to hide behind (why?)
- **Market depth**:

$$\frac{1}{\lambda} = 2\beta = 2\frac{\sigma_u}{\sigma_v}$$

The market is deeper when there is less insider trading and more noise trading

Equilibrium properties

- Insider's a priori (before observing v) **expected gain**:

$$\mathbb{E}[x(v-\mu-\lambda x)] = \mathbb{E}\left[\beta(v-\mu)\left(v-\mu-\frac{v-\mu}{2}\right)\right]$$
$$= \beta\frac{\sigma_v^2}{2} = \frac{\sigma_v\sigma_u}{2}$$

Comment: speculator expects a positive profit (could abstain). Competitive risk-neutral MM earns zero profits. Noise traders lose. Same as in GM.

Equilibrium properties

- Insider's a priori (before observing v) **expected gain**:

$$\mathbb{E}[x(v-\mu-\lambda x)] = \mathbb{E}\left[\beta(v-\mu)\left(v-\mu-\frac{v-\mu}{2}\right)\right]$$
$$= \beta\frac{\sigma_v^2}{2} = \frac{\sigma_v \sigma_u}{2}$$

Comment: speculator expects a positive profit (could abstain). Competitive risk-neutral MM earns zero profits. Noise traders lose. Same as in GM.

- Market maker's **posterior variance** of v is

$$\mathbb{V}(v|q) = \frac{1}{1/\sigma_v^2 + \beta^2/\sigma_u^2} = \frac{\sigma_v^2}{2}$$

Exactly half the prior variance: Insider reveals half his information

Kyle's model: summary

- **Dealer/market maker model**: Competitive, risk-neutral (zero profit) dealer chooses a supply schedule
- **Informed trader**: Observes signal about asset value and places market order
- **Market clearing**: Auction, dealer observes only total demand (informed + noise), total demand clears
- **Insights**: informed trading is a factor generating limited market depth, insider always reveals half his information
- **Advantage**: Richer trading opportunities, trader not price-taker
- **Shortcomings**: Still no resale

Kyle with inventory risk I

- Now let's look at how market maker's inventory risk can lead to limited depth.
- Assume no informed trading: $x = 0$.
- Asset value $v \sim \mathcal{N}(\mu, \sigma_v^2)$
- Market maker has mean-variance preferences over their next-period wealth:

$$U(w_{t+1}) = \mathbb{E}[w_{t+1}] - \frac{\rho}{2}\mathbb{V}(w_{t+1}),$$

  where w is composed of cash and asset holdings: $w_{t+1} = (z_t - q_t)v + q_t p_t$
- MM's initial asset position is z_t (initial cash is irrelevant, ignore it).

Kyle with inventory risk II

- To get the pricing schedule, follow the competitive logic:
  - The market-maker takes some market price p as given, chooses how much $q(p)$ to sell at this price:

$$\max_q \left\{ \underbrace{(z_t - q)\mathbb{E}[v] + qp}_{\mathbb{E}[w_{t+1}]} - \underbrace{\frac{\rho}{2}(z_t - q)^2 \mathbb{V}(v)}_{\mathbb{V}(w_{t+1})} \right\}$$

  - FOC: $p - \mu + \rho(z_t - q)\sigma_v^2 = 0 \iff q(p) = z_t + \frac{p-\mu}{\rho\sigma_v^2}$
- For market to clear, need $q(p) = u = q$ (dealer's supply $=$ total traders' market order), so inverting the pricing schedule we get:

$$p(q) = \mu + \rho\sigma_v^2(q - z_t).$$

Kyle with inventory risk III

$$p(q) = \mu + \rho\sigma_v^2(q - z_t)$$

Takeaways:

1. Depth (dictated by the dealer's willingness to trade at a given price) is limited
2. This is despite traders still being completely price-insensitive in this model!
3. Price impact depends on asset riskiness σ_v^2 and MM's risk aversion ρ.
4. Midquote depends on z_t

So really, all the same stuff as in GM with inventory risk.
The book also looks at versions with many MMs with heterogeneous ρs, and many imperfectly competitive MMs.

Extensions

Other extensions are possible:

1. **Dynamics**
   - In a fully dynamic model, the insider reveals less than half of the information in each period. Why?

Extensions

Other extensions are possible:

1. **Dynamics**
   - In a fully dynamic model, the insider reveals less than half of the information in each period. Why?
     In order to better benefit from informational advantage

Extensions

Other extensions are possible:

1. **Dynamics**
   - In a fully dynamic model, the insider reveals less than half of the information in each period. Why? In order to better benefit from informational advantage
   - In the limit where trade is continuous over $[0, T]$, then $\mathbb{V}(v|q_0, ..., q_t) \simeq (T - t)\frac{\sigma_v^2}{T}$: variance decreases linearly in time.

Extensions

Other extensions are possible:

1. **Dynamics**
   - In a fully dynamic model, the insider reveals less than half of the information in each period. Why? In order to better benefit from informational advantage
   - In the limit where trade is continuous over $[0, T]$, then $\mathbb{V}(v|q_0, ..., q_t) \simeq (T - t)\frac{\sigma_v^2}{T}$: variance decreases linearly in time. Model of how to split a large trade over time

Extensions

Other extensions are possible:

1. **Dynamics**
   - In a fully dynamic model, the insider reveals less than half of the information in each period. Why? In order to better benefit from informational advantage
   - In the limit where trade is continuous over $[0, T]$, then $\mathbb{V}(v|q_0, \ldots, q_t) \simeq (T - t)\frac{\sigma_v^2}{T}$: variance decreases linearly in time. Model of how to split a large trade over time
2. **More insiders**
   - More insiders are more competitive; more aggressive
   - The market is more liquid and more information revealed
   - In dynamic model with several insiders: rush to trade on common information from the beginning

3 **Imperfect market maker competition** (Cournot style)

- Finite number of market makers, $k = 1, ..., K$
- Market maker k supplies $y^k = \phi(p - \mu)$
- Market clears at price p with $\sum y^k = q$
- Strategic market maker takes into account effect of orders on profits
- Now: $p = \mu + \lambda q$ where $\lambda = \alpha(K-1)/(K-2) > \alpha$.

3. **Imperfect market maker competition** (Cournot style)
   - Finite number of market makers, $k = 1, ..., K$
   - Market maker k supplies $y^k = \phi(p - \mu)$
   - Market clears at price p with $\sum y^k = q$
   - Strategic market maker takes into account effect of orders on profits
   - Now: $p = \mu + \lambda q$ where $\lambda = \alpha(K-1)/(K-2) > \alpha$.
4. **Trading costs**
   - Trivial in GM. Very difficult here, both technically and conceptually.
   - Don't know how many trades there are, don't know the total volume (not q – some noise traders' orders could've cancelled each other out)
   - Even taking costs as a linear function of order imbalance $|q|$ makes things difficult

Homework

1. We will talk about empirical estimation of factors of illiquidity next time (ch.5) and begin talking about LOB markets (without dealers; ch.6)
2. Solve ex 3 in ch.4 (p.159): Kyle's model with competition among speculators.