

Financial Markets Microstructure

Final Exam with Solutions

Københavns Universitet

June 2020

Problem 1

In the Glosten-Milgrom model, as well as many other models we had in the class (e.g., Glosten, Parlour, Duffie-Garleanu-Pedersen, etc), strategic traders are restricted to trading at most one unit of the asset. Without this restriction, they would be willing to buy or sell infinite amounts.

1. What factors preclude or disincentivize such behavior in the real world?
2. Is it reasonable to model these factors as an exogenous constraint on trade size? If not, how would you incorporate them in the Glosten-Milgrom model?

Solution:

The possible factors include the following:

- **Limited liquidity/positive price impact.** The models under consideration effectively assume traders to be competitive, i.e., (1) infinitesimally small, and for that reason (2) not internalizing their impact on the market price. In the real world any trade of nontrivial quantity of the asset will have a non-zero price impact. This is not captured in the GM model, but it is the main driving force behind Kyle model, which can be seen as “GM model with price impact”.
- **Risk aversion.** If traders are risk-averse, they would not be willing to assume infinite positions, since those entail infinite risks. If we assume that informed traders in the GM model are risk-averse then we can dispose with the assumption of fixed trade size, although this gives rise to the next issue.
- **Need to blend in with noise traders.** The informed traders obtain a better price on their trades by mimicking the behavior of uninformed traders. If uninformed traders have particular trade patterns (e.g., always trade a single unit at a time) then informed

traders have incentives to mimic these patterns (restrict their order sizes). This aspect can be introduced into the model without affecting its conclusions: it can be shown that if informed traders can submit larger orders, they will never do so as long as noise traders operate in unit orders only.

- **Funding liquidity/leverage/short-selling constraints.** In order to take a large long position in an asset that is currently underpriced, the trader must have the liquid funds to do so. Similarly, short sales typically require a margin (collateral). All investors have some limits on the funds at their disposal, which yields a natural upper bound on the size of positions they can take. Relatedly, borrowing money in order to trade is typically subject to leverage constraints. Capturing all of these constraints as an exogenous limit on trade size is reasonable, but one should keep in mind that these limits may differ across traders.

Problem 2

Consider an asymmetric Glosten-Milgrom model in which the fundamental value v is $v = 10$ with probability $\gamma < 1/2$ and $v = 5$ w.p. $1 - \gamma$. The arriving trader can submit a buy or a sell order for one unit of the asset. The trader is informed w.p. π and is a noise trader w.p. $1 - \pi$. In the latter case the trader submits a buy order w.p. $\rho > 1/2$ and a sell order w.p. $1 - \rho$ independently of v . The informed trader knows v and trades so as to maximize profit. The dealer is risk-neutral and competitive.

1. Derive the ask and bid quotes set by the dealer.
2. Calculate the bid-ask spread. How would the spread react to an increase in ρ ? How does your answer depend on π ?
3. What is the intuition behind this dependence on π ?

Solution:

1. Dealer is competitive, so the quotes will capture the expected value of the asset conditional

on buy and sell orders respectively.

$$\begin{aligned}
a &= \mathbb{E}[v|\text{buy}] \\
&= \mathbb{E}[v|\text{buy,noise}] \mathbb{P}(\text{noise}|\text{buy}) + \mathbb{E}[v|\text{buy,informed}] \mathbb{P}(\text{informed}|\text{buy}) \\
&= (10\gamma + 5(1 - \gamma)) \frac{(1 - \pi)\rho}{\pi\gamma + (1 - \pi)\rho} + 10 \frac{\pi\gamma}{\pi\gamma + (1 - \pi)\rho} \\
&= \mu + 5(1 - \gamma) \frac{\pi\gamma}{\pi\gamma + (1 - \pi)\rho},
\end{aligned}$$

where $\mu = \mathbb{E}[v] = 5 + 5\gamma$. Similarly,

$$b = \mathbb{E}[v|\text{sell}] = \mu - 5\gamma \frac{\pi(1 - \gamma)}{\pi(1 - \gamma) + (1 - \pi)(1 - \rho)}.$$

2. The spread is given by

$$S = a - b = \frac{5\pi\gamma(1 - \gamma)}{[\pi\gamma + (1 - \pi)\rho][\pi(1 - \gamma) + (1 - \pi)(1 - \rho)]}.$$

Its derivative w.r.t. ρ is

$$\frac{\partial S}{\partial \rho} = -S \frac{(1 - \pi)[\pi(1 - 2\gamma) + (1 - \pi)(1 - 2\rho)]}{[\pi\gamma + (1 - \pi)\rho][\pi(1 - \gamma) + (1 - \pi)(1 - \rho)]}.$$

The denominator is positive, hence the spread reacts positively to an increase in ρ if and only if

$$\pi(1 - 2\gamma) + (1 - \pi)(1 - 2\rho) < 0 \Leftrightarrow \pi < \frac{\rho - 1/2}{\rho - \gamma}.$$

(The assumption $\gamma < 1/2 < \rho$ is important for the sign of the final inequality.)

3. In this problem the spread is asymmetric across buy and sell orders. When ρ increases, there is relatively less informed trading on the buy side (so the buy-side half-spread $a - \mu$ shrinks), but more on the sell side. The final effect on S is thus determined by which of these two components dominates.

Each half-spread is most sensitive to parameters when the probability of informed trading in the respective direction is close to 50% – this is the state of maximal uncertainty, – while otherwise the order flow is very likely to be either informative, or not, and comparable parameter value changes do not affect this probability as strongly. Given our parameter assumptions, if $\pi < \frac{\rho - 1/2}{\rho - \gamma}$ then $\mathbb{P}(\text{informed}|\text{sell})$ is closer to 1/2 than $\mathbb{P}(\text{informed}|\text{buy})$, hence the sell-side effect dominates, and vice versa.

Problem 3

This question is based on the Kondor model. Suppose that the fundamental value of the asset is given by the sum of L individual components:

$$\theta = \theta_1 + \theta_2 + \dots + \theta_L,$$

where each component $\theta_l \sim \text{i.i.d.}\mathcal{N}(0, \sigma^2)$. There are two strategic traders, label them i and j . Trader i observes the first I components of θ (denote their sum as $x_i = \theta_1 + \dots + \theta_I$), while trader j observes the last J components (denote their sum as $x_j = \theta_{L-J+1} + \dots + \theta_L$). Assume $I + J < L$, so there are no components observed by both traders, but there are some which are not observed by either one (denote their sum as $x_k = \theta_{I+1} + \dots + \theta_{L-J}$). In addition, both traders observe the same public signal $y = \theta + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$.

In answering the questions below, you can use the following fact:

If $q = v + u$, where $v \sim \mathcal{N}(\mu_v, \sigma_v^2)$ and $u \sim \mathcal{N}(\mu_u, \sigma_u^2)$, and the two are independent, then: (1) $q \sim \mathcal{N}(\mu_v + \mu_u, \sigma_v^2 + \sigma_u^2)$, and (2) $\mathbb{E}[v|q] = \mathbb{E}[v] + (q - \mathbb{E}[q]) \frac{\mathbb{C}(v,q)}{\mathbb{V}(q)}$.

1. Calculate $\mathbb{E}[\theta|x_j, y]$, i.e., trader j 's asset valuation conditional on the information available to him.¹
2. Calculate the second-order expectation $\mathbb{E}[\mathbb{E}[\theta|x_j, y] | x_i, y]$, i.e., trader i 's expectation of trader j 's valuation, conditional on i 's information.

HINT: you should get an expression of the form $\mathbb{E}[\mathbb{E}[\theta|x_j, y] | x_i, y] = \alpha y + \beta(y - x_i)$. You need to find α and β .

3. Take the coefficient β you have obtained in part 2 (if you did everything correctly, it should be positive). It captures the degree of divergence of second-order beliefs: the higher is trader I 's private signal of asset value, the lower he expects J 's valuation to be, and the higher is β , the stronger this effect is. How does β depend on I , the number of components that trader i observes? Explain the intuition behind this.
4. How does β depend on J , the number of components that trader j observes? Explain the intuition behind this.

NOTE: if you could not solve part 2, you can still try to provide an educated guess for the directions and the reasons of the effects in parts 3 and 4.

¹In general, conditioning on x_j is not the same as conditioning on $(\theta_{L-J+1}, \dots, \theta_L)$, since the latter contains more information. But the two are equivalent in this problem.

Solution:

1. We have

$$\begin{aligned}
\mathbb{E}[\theta|x_j, y] &= \mathbb{E}\left[\sum_{l=1}^L \theta_l \mid x_j, y\right] \\
&= \mathbb{E}[x_j + x_i + x_k \mid x_j, y] \\
&= x_j + \mathbb{E}[x_i + x_k \mid x_j, y]
\end{aligned} \tag{1}$$

To calculate the latter expectation we notice that $y = x_j + x_i + x_k + \epsilon$, where x_j is known (to j). Applying the fact given in the problem, we infer that

$$\mathbb{E}[x_i + x_k \mid x_j, y] = \alpha(y - x_j),$$

where $\alpha = \frac{(L-J)\sigma^2}{(L-J)\sigma^2 + \sigma_\epsilon^2}$. Plugging it back into (1), we get

$$\mathbb{E}[\theta|x_j, y] = \alpha y + (1 - \alpha)x_j. \tag{2}$$

2. Plugging (2) into the target expectation, we get

$$\begin{aligned}
\mathbb{E}[\mathbb{E}[\theta|x_j, y] \mid x_i, y] &= \mathbb{E}[\alpha y + (1 - \alpha)x_j \mid x_i, y] \\
&= \alpha y + (1 - \alpha)\mathbb{E}[x_j \mid x_i, y]
\end{aligned} \tag{3}$$

Recalling again that $y = x_i + x_j + x_k + \epsilon$ and applying the fact we obtain

$$\mathbb{E}[x_j \mid x_i, y] = (y - x_i) \frac{J\sigma^2}{(L - I)\sigma^2 + \sigma_\epsilon^2}$$

Plugging this back into (3), we finally get the desired expression

$$\begin{aligned}
\mathbb{E}[\mathbb{E}[\theta|x_j, y] \mid x_i, y] &= \alpha y + \beta(y - x_i), \\
\text{where } \beta &= (1 - \alpha) \frac{J\sigma^2}{(L - I)\sigma^2 + \sigma_\epsilon^2} \\
&= \frac{J\sigma^2\sigma_\epsilon^2}{[(L - J)\sigma^2 + \sigma_\epsilon^2][(L - I)\sigma^2 + \sigma_\epsilon^2]} > 0
\end{aligned}$$

3. Weight β is increasing in I , i.e., as trader i gets to observe more components, the divergence amplifies. For example, suppose $\theta_1 > 0$ and $\theta_2 = 0$. Then increasing I from one to two will lead trader i to amplify the inference he makes from θ_1 and decrease his opinion about j 's valuation – even though θ_2 by itself is uninformative. However, even though θ_2 does

not contain any information about the actual asset value in this example, it leads trader i to believe that information in $y - x_i$ is more likely to be observed by trader j . This is because i knows that j observes share $J/(L - I + 1)$ of components that constitute $y - x_i$.

4. Weight β is increasing in J , i.e., as trader j gets to observe more components, the divergence amplifies:

$$\frac{\partial \beta_2}{\partial J} = \beta_2 \left(\frac{1}{J} + \frac{\sigma^2}{(L - J)\sigma^2 + \sigma_\epsilon^2} \right) < 0$$

The intuition is the same as in part 3: as J increases, trader i perceives it to be more probable that j observes whatever components are responsible for the information in $y - x_i$.

Problem 4

Below you can find an Economist article from July 17, 2019 on the liquidity of the corporate bond market.² The article mentions that the actual liquidity in that market is well below what the investors seem to expect it to be. What are the possible consequences of such misalignment? Should we attempt to alleviate it by making the market more transparent? How does it relate to the discussion of market transparency we had in class?

Solution: A surface-level consequence is that the investors are exposed to more risk than they account for, since they may be unable to unwind their positions when they decide to do so. This is, in itself, an inefficiency that we would typically want to correct.

The counterargument is also described in the article: this false belief in liquidity leads investors to hold bonds, which they would or could not invest in otherwise due to liquidity risk. This by itself creates some liquidity, supplied by such traders. An attempt to correct the investors' belief would lead to the market's suffering liquidity dropping even further, completely paralyzing the bond market and leading to a crash, which is not the most desirable outcome.

In the end, the blanket conclusion we made about opaque markets – that opaqueness hurts the uninformed traders while benefitting other market participants, – gets an extra layer in this setting, in that the aggregate effects of transparency may be so catastrophic that the uninformed traders will lose more than they gain if the market is made more transparent.

²Also available at: <https://www.economist.com/finance-and-economics/2019/07/11/why-everybody-is-concerned-about-corporate-bond-liquidity>

Buttonwood – Why everybody is concerned about corporate-bond liquidity

In September 2007 Britain suffered its first bank run in a century. Television pictures showed a long queue of depositors outside a branch of Northern Rock. Alistair Darling watched in dismay from Portugal, where he and his fellow European Union finance ministers were gathered. “They’re behaving perfectly rationally, you know,” Mervyn King, the governor of the Bank of England, said in the smarty-pants manner that economists are cherished for. Mr Darling was uncharmed. “It was not what I wanted to hear,” he recalled.

What Lord King probably had in mind was a well-thumbed textbook model. Banks have a liquidity mismatch. One side of the balance-sheet is hard-to-sell loans; the other side is deposits that can be withdrawn in a trice. If depositors believe that a bank is sound, there will be no runs on it. But if enough start to demand their deposits back, it makes sense for everybody to join the rush.

This model can also be applied in other areas. Take the corporate-bond market. Every policy body of stature, from the IMF to the European Central Bank (ECB), has worried about a growing mismatch between investors’ expectations that they can sell out at any moment and an underlying shortage of liquidity in the market. More investors are using corporate-bond funds as an alternative to cash. But fewer dealers are willing to trade bonds in size. A big scare could feasibly start a run.

The dynamics of capital-market runs are similar to those of bank runs. You see them in currency crises. Foreign-exchange reserves, say, are slim relative to the scale of local-currency assets held by flighty investors. Should enough of those investors sell out, others will soon follow. The result is a rout. There is a similar pattern with investment funds that promise speedy withdrawals but hold assets that cannot be sold quickly. Bad news prompts withdrawals. The speedy get paid. Other investors then try to get out, too. But the fund cannot sell assets fast enough. It is forced to suspend redemptions.

Such trouble is especially likely with corporate bonds, which are inherently illiquid. In contrast with trading in shares, where buy and sell orders are matched on electronic order books, corporate bonds are traded over-the-counter. Bonds are not as standardised as shares. A company may have bonds of several different maturities. If you want to buy or sell, you call a dealer.

The ease with which investors can trade bonds—the market’s liquidity—depends a lot, then, on the readiness of dealer banks to stockpile securities. Where there is heavy selling, dealers would ideally warehouse cheaper bonds for when people want to buy again. But since the financial crisis new rules have made it less cost-effective for banks to use capital for trading of any kind. The inventory of corporate bonds held by dealers has fallen sharply in the past decade.

As the role of dealers has shrunk, the thirst for instant liquidity has increased. Derisory yields on the safest government debt have drawn investors towards riskier securities, including corporate bonds. A cheap and convenient way to invest in them is to buy an exchange-traded fund, or ETF. These are low-cost investment funds that hold a basket of bonds, usually mirroring a benchmark index. They trade on stock exchanges just as listed shares do. The ease of buying and selling

bond ETFs is a big part of their appeal. They are also often used as depositories for spare cash. Studies are divided on whether ETFs make the underlying bonds more or less liquid. But there are concerns that in a stressed market, outflows from ETFs might make a bad situation worse. And it is not hard to make a case that the corporate-bond market has become more fragile. Many firms in America have issued lots of bonds to buy back their own shares. With extra leverage comes more risk. Half of all investment-grade bonds have a credit rating of BBB. In a recession a chunk of those bonds will be downgraded to junk. Many mutual funds and ETF s can hold only investment-grade bonds. If a lot of bonds have to change hands quickly, that could easily overwhelm the market's limited liquidity. Prices might fall a long way.

Just how messy the next big shake-out in the corporate-bond market is depends on many things: on how weak the economy gets; on how many BBB borrowers can avert a downgrade; on how quickly funds can be raised to buy at fire-sale prices. For now, it seems rational to hold bonds that afford a little extra yield. Smart-alecks say this will surely end badly. But who wants to hear that?