Problem 1

(a). Shiller famously argued in 1981 that there was *excess volatility* in stock prices in the sense that they moved too much to be justified by subsequent changes in dividends (fundamentals). Discuss this claim in the context of what you have learned in the course. For instance, should price changes necessarily be followed by equivalent dividend changes if the market were efficient? You can answer this question using models seen in the course, but you are welcome to bring in alternative explanations.

Solution: (There are many ways to answer the question – the following are some of the observations that a good solution *could* make). The course has emphasized different motives for trading and different motives for pricing, and very few of the models we have seen would predict that price changes should be one-to-one with expected discounted future dividends. For instance, almost all our models incorporate liquidity traders or traders who attach private value to a given stock. This is often motivated as traders who try to optimize their portfolio by choosing a certain risk-profile (or who must adhere to regulation/norms if they are institutional investors), or traders who have liquidity needs (for instance if they are facing margin calls). Such trading will move prices, but be unassociated with subsequent dividend changes. We have also seen that the prices are affected by market makers' holding cost and transaction costs. Such costs will lead to a greater price spread and therefore might contribute to increasing the volatility of price changes, but again, this will be unrelated to future dividends. Herding and speculative trading, as seen in the section on bubbles, can also contribute to price changes. Therefore, there seems to be many rational reasons why prices should move more than subsequent dividends.

(b). Consider the following sequence of orders.

(1). A limit sell order of 300 units at price 3.6.

(2). A limit sell order of 100 units at price 3.5.

(3). A limit buy order of 200 units at price 3.6.

(4). A limit sell order of 500 units at price 3.7.

(5). A limit buy order of 400 units at price 3.65.

(6). A market buy order of 200 units.

(7). A limit buy order of 400 units at price 3.5.

(8). A market sell order of 400 units.

We now investigate the execution of these orders under two different assumptions about the market structure.

(i). First, suppose that the market is a **continuous limit order book**, where orders are cleared sequentially in the manner described in chapter 1.2.1 of the textbook, pages 18 to 21.[1] Describe at each step in the sequence what happens with the incoming order: does it enter the book and if so on which side, or does it execute, and if so, how many units at what price. Find the total number of units traded.

(ii). Second, suppose instead that the market is cleared through a call auction, as described in chapter 1.2.1 of the textbook, pages 21 to 23. Find the market-clearing price and the number of units traded.

(iii). Use the results from (i) and (ii) to compare the number of units traded in each auction, and discuss their efficiency properties. Is one better than the other? Why?

Solution:

(i). The trades will execute in the following manner:

- Trader 1's order is entered into the book on the ask side.
- Trader 2's order is entered into the book on the ask side.
- Trader 3's order is fully executed: a 100 units are bought at price 3.5, a 100 units are bought at price 3.6.
- Trader 4's order is entered into the book on the ask side.
- Trader 5's order is partially executed: 200 units are bought at price 3.6, 200 units are entered into the book on the bid side.

[1] In particular, suppose that orders are filled according to time preference (not pro-rata), that limit orders which are partially marketable will execute the marketable part and enter the remaining amount in the limit order book, and that market orders that can be only partially filled execute against the available amount.

- Trader 6's order is fully executed: 200 units are bought at price 3.7.
- Trader 7's order is entered into the book on the bid side.
- Trader 8's order is fully executed: 200 units at price 3.65, 200 units at price 3.5.

In total, 1000 units are traded.

(ii). In the call auction market, the market-clearing price will be 3.6 and 800 units will be traded.

(iii). In the continuous LOB there is more trade than in the call auction market. Since trades are always mutually beneficial, this would seem to indicate that there is an efficiency gain in the continuous LOB. The difference in trades in the two markets is made up by the fact that in the continuous LOB, trader 4 who is a seller with limit price 3.7 is able to sell 200 units (to a market order) and trader 7 who is a buyer with a limit price of 3.5 is able to buy 200 units (against a market sell order), whereas neither of them trade in the call auction. But if these traders' valuations are reasonably close to the limit price they have stated, *ex post* trader 4 would want to buy from trader 7 the 200 units he has acquired. Thus, the market is not Pareto efficient in the sense that no profitable trades can be made after the market closes. The call auction market, on the other hand, has this property.

(c). We have noted on several occasions that traders may not receive information at the same time. Suppose we are in a setting with 2 periods where one trader arrives in each period. There are three cases: (i) both traders are noise traders, (ii) both traders are informed and each trader observes an independent signal,[2] or (iii) both traders are informed, but one trader is a *fast trader* and the other is a *slow trader*, who both observe the same piece of private information, but the fast trader gets to trade before the slow trader.[3] The market maker cannot distinguish between these three settings.

Write up a model of this in the style of the Glosten-Milgrom model. Without solving

[2] In particular, suppose each trader i observes a signal s_i that is informative but imperfectly correlated with the asset value, but that s_1 and s_2 are drawn independently, conditional on the true value of the asset.

[3] Notice that the difference to case (ii) is that the traders observe the same signal.

the model, answer the following: Compared to a standard model where we are either in case (i) or (ii), but never in case (iii), how do you think prices will be different? That is to say, what is the effect of potentially having traders who trade at different times with the same information?

Solution: (There are many ways to answer the question – the following are some of the observations that a good solution *could* make). Suppose we are in a two-period model where at each time $t = 1, 2$ a new trader arrives, trades and subsequently disappears. There is a single asset with value V uniformly distributed on $\{0, 1\}$ is being traded each period. Furthermore,

- With probability α, the traders at $t = 1, 2$ observe an independent signal s_t where $s_t = V$ with probability $\pi > 1/2$, and $s_t = 1 - V$ otherwise.
- With probability β, the traders at $t = 1, 2$ observe the same signal s where $s = V$ with probability $\pi > 1/2$, and $s = 1 - V$ otherwise.
- With probability $1 - \alpha - \beta$, the traders at $t = 1, 2$ are noise traders.

The market maker sets prices competitively, and each trader can buy or sell one unit, or abstain. Rational traders maximize expected profits. Noise traders buy or sell with equal probability.

Suppose there is an equilibrium in which rational traders with positive (negative) information always buy (sell). The fact that some traders trade on the same information will bring the period-2 price more toward the a priori expected value (1/2) compared to the case in which all traders have independent information, whenever there are two consecutive trades in the same direction. The reason is that the market maker will interpret consecutive trades this as potentially arising from traders trading on the same information, rather than new information.

Problem 2

In this exercise, we consider a model of the Glosten-Milgrom type, but we add the possibility of trading either one or two units each period.

Suppose there is a single asset with value

$$V \in \{v^H, v^L\},$$

where $\mathbb{P}(V = v^H) = \alpha$. Let $\overline{v} = \alpha v^H + (1-\alpha)v_L$. Market makers do not know the true value, but it is observed by traders. There is a single trader who is either a *rational* trader, with probability $\pi > 0$, or a *noise* trader, with probability $1 - \pi$.

- If the trader is rational, he maximizes his expected payoff (risk neutral).
- If he is a noise trader, with probability β he trades one unit, and with probability $1 - \beta$ he trades two units. In either case, he buys or sells with probability $\frac{1}{2}$ each.

The market maker is competitive and makes zero expected profits. He sets ask prices $a(1)$ and $a(2)$ for the cases where one and two units are bought, respectively, and similarly bid prices $b(1)$ and $b(2)$. Suppose that $a(k)$ is the per-unit price for k units, such that the price of buying one unit is $1 \cdot a(1)$ and the price of buying two units is $2 \cdot a(2)$, and similarly for $b(k)$. Finally, assume that the rational trader always buys when he observes $V = v^H$ and sells when he observes $V = v^L$.

(a). Suppose first that the rational trader always trades two units. What is $a(1)$ and $b(1)$?

Solution: Since there is no informed trading at these 'ticks' then $a(1) = b(1) = \overline{v}$.

(b). Suppose still that the rational trader always trades two units. What is $a(2)$ and $b(2)$?

Solution: Now, we have to take into account the probability of trading with an informed trader. Applying Bayes' Rule yields

$$
\begin{aligned}
a(2) &= \frac{\pi\alpha}{\pi\alpha + (1-\pi)(1-\beta)\frac{1}{2}} v^H + \frac{(1-\pi)(1-\beta)\frac{1}{2}}{\pi\alpha + (1-\pi)(1-\beta)\frac{1}{2}} \overline{v}, \\
b(2) &= \frac{\pi(1-\alpha)}{\pi(1-\alpha) + (1-\pi)(1-\beta)\frac{1}{2}} v^H + \frac{(1-\pi)(1-\beta)\frac{1}{2}}{\pi(1-\alpha) + (1-\pi)(1-\beta)\frac{1}{2}} \overline{v}.
\end{aligned}
$$

(c). Given the prices you have derived in (a) and (b), show a condition for when it is optimal for the rational trader to trade two units. Comment on your result: why is it always/not always optimal to trade two units?

(If you can not show the result generally, try substituting a specific value for (α, β, π) and see what happens.)

Solution: Suppose $V = v^H$. We are assuming that the trader always busy, so the question is whether it is optimal to buy one or two units, given the prices. The profits are the following

$$
\begin{aligned}
u_a^H(1) &= v^H - a(1) = (1-\alpha)(v^H - v^L), \\
u_a^H(2) &= 2[v^H - a(2)] \\
&= 2\frac{(1-\pi)(1-\beta)\frac{1}{2}}{\pi(1-\alpha) + (1-\pi)(1-\beta)\frac{1}{2}} (v^H - \overline{v}) \\
&= \frac{(1-\pi)(1-\beta)}{\pi(1-\alpha) + (1-\pi)(1-\beta)\frac{1}{2}} (1-\alpha)(v^H - v^L).
\end{aligned}
$$

Thus, $u_a^H(2) \geq u_a^H(1)$ is equivalent to

$$\begin{aligned}
\frac{(1-\pi)(1-\beta)}{\pi(1-\alpha)+(1-\pi)(1-\beta)\frac{1}{2}} &\geq 1 \\
\Leftrightarrow (1-\pi)(1-\beta) &\geq \pi(1-\alpha)+(1-\pi)(1-\beta)\frac{1}{2} \\
\Leftrightarrow (1-\pi)(1-\beta) &\geq 2\pi(1-\alpha) \\
\Leftrightarrow 1-\beta &\geq \pi[2(1-\alpha)+1-\beta].
\end{aligned}$$

This can be rewritten as

$$\pi \leq \frac{1-\beta}{2(1-\alpha)+1-\beta}.$$

Notice that the right-hand side is in $(0,1)$. Similarly, on the bid side.

$$\begin{aligned}
u_b^H(1) &= b(1) - v^L = \alpha(v^H - v^L), \\
u_b^H(2) &= 2[b(2) - v^L] = \frac{(1-\pi)(1-\beta)}{\pi\alpha+(1-\pi)(1-\beta)\frac{1}{2}}\alpha(v^H - v^L).
\end{aligned}$$

Thus, $u_b^H(2) \geq u_b^H(1)$ is equivalent to

$$\frac{(1-\pi)(1-\beta)}{\pi\alpha+(1-\pi)(1-\beta)\frac{1}{2}} \geq 1$$

which can be found to be

$$\pi \leq \frac{1-\beta}{2\alpha+1-\beta}.$$

Thus, for it to be optimal to trade two units on both the ask and bid side, we need

$$\pi \leq \min\left\{\frac{1-\beta}{2\alpha+1-\beta}, \frac{1-\beta}{2(1-\alpha)+1-\beta}\right\}.$$

Intuition: Ask side. The rational trader trades two units when there is a sufficient amount of noise trading to hide behind. But now, he can only hide behind noise trader who trade two units. If π is too high, then there is not enough noise/trading to hide behind. For instance, on the ask side, the higher the probability of noise trading of two units (the lower β) and the higher the probability that V is high (the higher α), the higher the threshold. The effect of β is straightforward, but the effect of α is not, since higher α implies higher $a(2)$. However, it also implies higher $a(1)$, but because more is traded at $a(2)$, this makes it more attractive to trade two units. Bid side is similar.

(d). Now suppose that $\alpha = 1/2$. Show that for certain values of π it is also possible to have an equilibrium in mixed strategies where the rational trader sometimes trades one unit and sometimes two units. You can focus on the ask side for this question, i.e. you need to show the existence of ask prices $a(1)$ and $a(2)$, as well as a mixed buying strategy for the rational trader, that form an equilibrium. Find the probability σ with which the trader buys two units in equilibrium.

Solution: First we find the prices for general α. Whenever $V = v^H$ the rational trader buys two units with probability σ and one unit with probability $1-\sigma$. The ask prices will then be

$$\begin{aligned} a(1) &= \frac{\pi\alpha(1-\sigma)}{\pi\alpha(1-\sigma)+(1-\pi)\beta\frac{1}{2}}v^H + \frac{(1-\pi)\beta\frac{1}{2}}{\pi\alpha(1-\sigma)+(1-\pi)\beta\frac{1}{2}}\overline{v} \\ a(2) &= \frac{\pi\alpha\sigma}{\pi\alpha\sigma+(1-\pi)(1-\beta)\frac{1}{2}}v^H + \frac{(1-\pi)(1-\beta)\frac{1}{2}}{\pi\alpha\sigma+(1-\pi)(1-\beta)\frac{1}{2}}\overline{v} \end{aligned}$$

In order to play a mixed strategy, the rational trader has to be indifferent between buying one or two units. Thus we require $u_a^H(1) = u_a^H(2)$, which translates into

$$\begin{aligned} v^H - a(1) &= 2[v^H - a(2)] \\ \Leftrightarrow \frac{(1-\pi)\beta\frac{1}{2}}{\pi\alpha(1-\sigma)+(1-\pi)\beta\frac{1}{2}}(1-\alpha)(v^H-v^L) &= 2\frac{(1-\pi)(1-\beta)\frac{1}{2}}{\pi\alpha\sigma+(1-\pi)(1-\beta)\frac{1}{2}}(1-\alpha)(v^H-v^L) \\ \Leftrightarrow \frac{\beta}{\pi\alpha(1-\sigma)+(1-\pi)\beta\frac{1}{2}} &= 2\frac{1-\beta}{\pi\alpha\sigma+(1-\pi)(1-\beta)\frac{1}{2}} \\ \Leftrightarrow \sigma &= \frac{(1-\beta)(4\alpha\pi+\beta(1-\pi))}{2\alpha(2-\beta)\pi}. \end{aligned}$$

For $\alpha = 1/2$, this reduces to $\sigma = \frac{(1-\beta)(2\pi+\beta(1-\pi))}{\pi(2-\beta)}$. Clearly, this is greater than zero, and it is smaller than one whenever

$$\frac{(1-\beta)(2\pi+\beta(1-\pi))}{\pi(2-\beta)} \leq 1 \Leftrightarrow \pi \geq \frac{1-\beta}{2-\beta}$$

When this condition is satisfied, there is an equilibrium in mixed strategies. The intuition is the following. When there is a great deal of adverse selection (high π), trading at a single 'tick' is very revealing for the trader, and therefore he prefers to spread his trading over several ticks.

(e). Consider now a two-period model, $t = 1, 2$, where the two periods can be thought of as trading in the morning and in the afternoon on the same day. Suppose that at the beginning of the day, with probability $\gamma > 0$ there is an information event, and $V = -1$ or $V = 1$ with equal probability. With probability $1 - \gamma$ there is no information event, and $V = 0$. If there is an information event, then in each period, with probability π a rational trader, who knows V, arrives. With probability $1 - \pi$ a noise trader arrives. The noise trader behaves as in (a)-(d), i.e. he buys/sells with equal probability, and trades one/two units with probability $\beta/1 - \beta$. If there is no event, the trader is a noise trader in both periods.

The market maker sets prices $a_t(k)$ and $b_t(k)$, where $k \in \{1, 2\}$ is the trade size and $t \in \{1, 2\}$ is the period. Notice that now, the period-2 prices depend on the trade in period 1 as well:

$$a_2(k) = \mathbb{E}[V|d_1, d_2 = k],$$
$$b_2(k) = \mathbb{E}[V|d_1, d_2 = -k],$$

where d_t is the signed period-t trade (i.e. 1 if a buy of size 1, -2 if a sell of size 2, etc.).

Suppose a separating equilibrium is played in both periods, where the rational trader always trades two units (just as in (a) and (b)). Suppose further that there was a buy order of size 2 in first period and a buy order of one unit in period 2, that is, suppose

$$d_1 = 2 \text{ and } d_2 = 1.$$

Calculate the resulting prices, i.e. calculate $a_1(2)$ and $a_2(1)$ conditional on $d_1 = 2$.

Solution: The period-1 price can be calculated using the same method as above. In particular, now $v^H = 1$ whereas the expectation $\overline{v} = 0$, $\alpha = 1/2$ and we have to take account of γ.

$$\begin{aligned} a_1(2) &= \frac{\gamma\pi\frac{1}{2}(1) + [\gamma(1-\pi)(1-\beta)\frac{1}{2} + (1-\gamma)(1-\beta)\frac{1}{2}](0)}{\gamma\pi\frac{1}{2} + \gamma(1-\pi)(1-\beta)\frac{1}{2} + (1-\gamma)(1-\beta)\frac{1}{2}} \\ &= \frac{\gamma\pi}{\gamma\pi + \gamma(1-\pi)(1-\beta) + (1-\gamma)(1-\beta)}. \end{aligned}$$

In the period 2, on the other hand, the price must take into account the first trade as well. Notice that the second trade can only have come from a noise trader, and therefore we only need to consider whether the first trade came from a rational or a noise trader. In particular,

$$\begin{aligned} a_2(1) &= \frac{\gamma\pi\frac{1}{2}(1-\pi)\beta\frac{1}{2}(1) + [\gamma((1-\pi)\frac{1}{2})^2(1-\beta)\beta + (1-\gamma)(\frac{1}{2})^2(1-\beta)\beta](0)}{\gamma\pi\frac{1}{2}(1-\pi)\beta\frac{1}{2} + \gamma((1-\pi)\frac{1}{2})^2(1-\beta)\beta + (1-\gamma)(\frac{1}{2})^2(1-\beta)\beta} \\ &= \frac{\gamma\pi(1-\pi)}{\gamma\pi(1-\pi) + \gamma(1-\pi)^2(1-\beta) + (1-\gamma)(1-\beta)} \end{aligned}$$

(f). Let the realized prices in (e) be denoted p_1 and p_2, i.e. $p_1 = a_1(2)$ and $p_2 = a_2(1)$. Is it possible that $p_2 < p_1$? If not possible, explain why this is so. If possible, explain how the price can decrease following a buy order.

Solution: Rewrite $a_2(1) < a_1(2)$ as

$$\begin{aligned} \frac{\gamma\pi(1-\pi)}{\gamma\pi(1-\pi) + \gamma(1-\pi)^2(1-\beta) + (1-\gamma)(1-\beta)} &< \frac{\gamma\pi}{\gamma\pi + \gamma(1-\pi)(1-\beta) + (1-\gamma)(1-\beta)} \\ \frac{1-\pi}{\gamma\pi(1-\pi) + \gamma(1-\pi)^2(1-\beta) + (1-\gamma)(1-\beta)} &< \frac{1}{\gamma\pi + \gamma(1-\pi)(1-\beta) + (1-\gamma)(1-\beta)} \\ \frac{1}{\gamma\pi + \gamma(1-\pi)(1-\beta) + (1-\gamma)(1-\beta)\frac{1}{1-\pi}} &< \frac{1}{\gamma\pi + \gamma(1-\pi)(1-\beta) + (1-\gamma)(1-\beta)} \end{aligned}$$

Thus, the inequality will be satisfied if

$$(1-\gamma)(1-\beta)\frac{1}{1-\pi} > (1-\gamma)(1-\beta),$$

which is always true.

This occurs because a small trade can be seen as a signal indicating that it is less likely that an information event has occurred. This brings the price toward the no-event level, which is zero.

Problem 3

In the following pages you will find a 2009 article from The Economist on transparency in financial markets. Discuss the article using what you have learned in the course. Summarize the main arguments, evaluate them using theory and discuss points which have been omitted. You are welcome to bring in theories and facts from outside the course in the discussion.

Solution: Below I outline a list of points that a good answer should touch upon. This is not a complete list, and the discussion of these issues is as important as mentioning them. The conclusion does not have to be the same as the one I have given, but should be well-argued.

Summary. The main points of the article are

- The case for transparency is not clear-cut. Even though it is being lauded by experts, few have tried to assess its merits.
- Until recently, there was substantial backing for opacity, based on the arguments that traders need opacity to make money from correcting market inefficiencies.
- Increasing transparency is not just a question of increasing information, which may be irrelevant or incomprehensible.
- Trying to regulate transparency can introduce lopsided incentives.
- What matters for liquidity is symmetry, rather than quantity, of information.
- Politics make the implementation of transparency hard because (i) the costs are born by a few, but the benefits by many, and (ii) disclosure requirements tend to be a fast response to a recent crisis.
- Even with good information, people may still make bad choices...
- ...but in the end, transparency is better because opacity may create systemic risk – however, disclosure rules should be tailored individually to markets, and possibly delayed, to give incentives to those who profit from 'correcting' market inefficiencies.

Theory. Assessing the article's arguments, the following theory can be used:

- *Incentives to acquire information.* The argument that traders need opacity to gain from correcting market inefficiencies can be linked to the incentive to acquire costly information. Recall the exercise from problem set 1, where we investigate endogenous information acquisition by traders. If we think of higher opacity as smaller π, then indeed it leads to a greater incentive to acquire information. Thus, in this sense the 'classic' argument is correct.

- *Symmetry, rather than quantity, of information is what matters.* This argument has a great deal of merit if we are concerned about liquidity, in the sense that eliminating information asymmetry would also eliminate the information driven part of the spread, since traders would only be trading for non-informational motives (Glosten and Milgrom or Kyle will tell you this). But it is not always clear how to separate asymmetry from quantity of information. Consider the Kyle model, where the price-impact parameter λ is equal to $\frac{\sigma_v}{2\sigma_u}$. In some sense, σ_u measures asymmetry in that it measures how much the market maker can deduce about the trader's incentives, and letting it go to zero gives extreme price sensitivity, which in standard parlance would correspond to an illiquid market. Similarly, releasing public information should reduce σ_v which would bring down price sensitivity. The model by Kurlat and Veldkamp also shows that mandatory disclosure can help solving an adverse selection problem.

- *Uneven distribution of costs and gains.* There are two types of costs the argument might refer to: the cost of procuring information, and lost profit due to transparency. The cost of acquiring information may be great in some cases (for instance, reporting according to regulation might require costly expert assistance), but a great deal of the information that could lead to more transparency does not have to be costly, i.e. trading data already in the dealers' possession. When it comes to winners and losers from greater transparency, we have seen in chapter 8 that transparency in many cases benefits liquidity traders and hurts informed traders. For the most part, we have not been able to analyze the effect on market makers, because we assume that they are perfectly competitive and make zero profits. To the extent that imperfect competition allows market makers to make money from acquiring information, the effect may be ambiguous since more transparency limits their losses against informed traders, but also their gains against liquidity traders.

The article by Kurlat and Veldkamp shows that informed traders may actually benefit from mandatory disclosure, even if it takes away their informational advantage. The reason is that if the market suffers from a great deal of adverse selection, it may be so illiquid that trading is 'expensive' and therefore the traders' private information yields little profit. But the illiquidity of the market also hurts the traders in another way: traders make a profit from holding risk, and if trading becomes more 'expensive', then they will make less profit from their risk-holding activity. The article shows that since more traders will become informed when information is cheap, the adverse selection problem, and therefore the gain to mandatory disclosure, is greater when information is cheap. Thus, contradictorily, it may be optimal to force information disclosure in markets where information is cheap to acquire and held by many, which goes somewhat against the argument of the article.

- *Giving incentives to arbitrageurs.* The article argues in the end that delayed disclosure would be one way of giving incentives to acquire information, whilst still avoiding many of the problems associated with opacity. This argument has some merit to it. We can think about this in terms of a herding model. Here, the inefficiency of the market results from many successive traders ignoring their private information and instead herding around previous actions. Delayed revelation of information would in many cases limit the extent of such herds, in that eventually the private information becomes available, leading prices to move toward the information-efficient level. However, this may not prevent short-term bubbles. As seen in Abreu and Brunnermeier, even when there is a known exogenous revelation point, traders may still speculate in 'riding the bubble'.

The Economist

Economics focus

Full disclosure

The case for transparency in financial markets is not clear-cut

Feb 19th 2009 | From the print edition

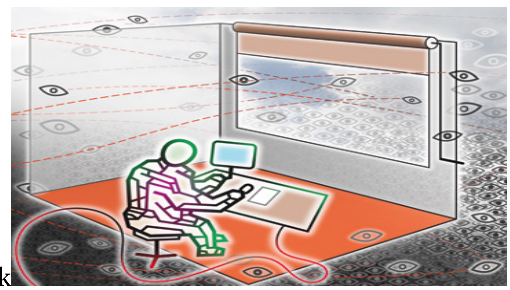

Illustration by Jac Depczyk

ITS promises are alluring, yet elusive; everyone, from politician to pundit, calls for more. In its recent report on financial reform, the Group of Thirty, a body of financial experts, mentioned it more than 30 times. Transparency is in vogue. Yet few ask whether it actually works.

Not long ago the cheerleaders of opacity were the loudest. Without privacy, they argued, financial entrepreneurs would be unable to capture the full value of their trading strategies and other ingenious intellectual property. Forcing them to disclose information would impair their incentive to uncover and correct market inefficiencies, to the detriment of all. And for years the so-called shadow banking system thrived, away from prying eyes. Then crisis hit, lending weight to the quip "What you see is what you get; what you don't see gets you." Few saw it coming, but if a lack of transparency was pervasive, how could they have?

As clear as mortgage-backed securities

"Sunlight is said to be the best of disinfectants," wrote Louis Brandeis, later a Supreme Court justice, in 1913, and almost a century later his words have become a maxim. Yet transparency is amorphous; it can, frustratingly, be anything but transparent and, implemented wrongly, may harm the very interests it is supposed to serve. In financial markets, the word is nearly always equated with information disclosure. The trouble is that the information is often incomplete, irrelevant or outright incomprehensible. Subprime-mortgage-backed securities are a case in point. These instruments—whose value remains shrouded in mystery—can have prospectuses of about 500-600 pages, most of which are devoted to intricate legalese. Yet, inexplicably, they do not contain the information about individual loans that is needed to detect default risk.

Nor is transparency free. The Sarbanes-Oxley act, which partly restored confidence after the scandals of Enron, WorldCom and others, came at a cost—not only in terms of the burden of compliance it imposed on companies. In order to shield small firms, those with a stockmarket value of less than $75m were initially exempted. This created a peculiar incentive: at least one study suggests that firms just below the threshold began disbursing unusual amounts of cash to shareholders and making fewer investments. The act has also been accused of stifling risk-taking and increasing directors' pay.

At its onset, the turmoil in financial markets was described as a liquidity crisis. And transparency and liquidity are close relatives. One enemy of liquidity is "asymmetric information". To illustrate this, look at a variation of the "Market for Lemons" identified by George Akerlof, a Nobel-prize-winning economist, in 1970. Suppose that a wine connoisseur and Joe Sixpack are haggling over the price of the 1998 Château Pétrus, which Joe recently inherited from his rich uncle. If Joe and the connoisseur only know that it is a red wine, they may strike a deal. They are equally uninformed. If vintage, region and grape are disclosed, Joe, fearing he will be taken for a ride, may refuse to sell. In financial markets, similarly, there are sophisticated and unsophisticated investors, and unless they have symmetrical information, liquidity can dry up. Unfortunately transparency may reduce liquidity. Symmetry, not the amount of information, matters.

The good news is that transparency can work. When information is relevant, standardised and public, it fosters intelligent decision-making. Lenders, for instance, are required to quote interest rates as annual percentage rates, making loans easy to compare. Some behavioural economists call this "simplified transparency", and think similar requirements should be imposed on complex financial products. Information must also be accurate as the credit-rating debacle shows: an AAA rating is harmful rather than helpful if it describes a CCC asset.

But politics impedes the ideal of transparency for at least two reasons. First, the benefits of transparency are widely dispersed among information users, whereas the costs are borne by few information disclosers; the disclosers therefore dominate the political process. Second, disclosure requirements are often drawn up after crises. They therefore tend to be hurried and haphazard, and support for them fades with memory of the hard times.

And even well-designed disclosure requirements may not suffice. People may make ill-informed choices, simplified transparency or not. In a recent study, two groups (made up of Harvard University staff) were asked to pick mutual funds. One group was given prospectuses which neatly summarised the funds' objectives, risk profiles, costs and past performance in a few pages. The other group received the standard long-winded and hard-to-understand prospectuses. They nonetheless made nearly identical choices, opting for funds with good past performance and largely neglecting fees. Academic research suggests that people should do precisely the opposite.

Still, for all its difficulties, transparency is usually better than the alternative. The opaque innovations of the recent past, rather than eliminating market inefficiencies, unintentionally created systemic risks. The important point is that financial markets are not created equal: they may require different levels of disclosure. Liquidity in the stockmarket, for example, thrives on differences of opinion about the value of a firm; information fuels the debate. The money markets rely more on trust than transparency because transactions are so quick that there is little time to assess information. The problem with hedge funds is that a lack of information hinders outsiders' ability to measure their contribution to systemic risk. A possible solution would be to impose delayed disclosure, which would allow the funds to profit from their strategies, provide data for experts to sift through, and allay fears about the legality of their activities. Transparency, like sunlight, needs to be looked at carefully.

From the print edition: Finance and economics