

Improving the Stock Market

Philip Sterne
Frankfurt Institute for Advanced Studies

January 22, 2012

Abstract

This paper seeks to reduce the pricing noise inherent in a normal market structure, by requiring more pricing information to be communicated. Rather than individual buy or sell orders, the specification of an investor's entire demand function is provided. The demand function specifies how many shares the investor would like to own for all possible prices. The demand functions of all investors can then be aggregated and displayed to the market. This aggregate demand function is used to price securities efficiently. If an investor wishes to withdraw from the market his shares may be redistributed amongst the remaining investors, without any further orders being placed. This ensures liquidity is built into the market and volatility is reduced as a result.

To discourage short-term speculation we propose that a Tobin tax be levied whenever a demand function is changed. The Tobin tax encourages demand functions to estimate medium-term price movements and aid price-discovery. We show that the demand curve which maximizes profit is intimately related to the belief of the fair price.

1 Introduction

Since the financial crisis of 2008 there have been growing calls for financial reform. One avenue for potential reform that has been completely ignored so far is changing the rules of the stock market itself. Stock markets have existed since the 12th century in France, and very little has changed in the intervening centuries. Essentially all new technology in the market has just improved transaction processing times. But do we really need a market that settles trades in microseconds? This article argues that with current technology we now have the ability to create better markets which have far less volatility and are just as liquid.

Financial markets are an integral part of modern society. They provide the means to efficiently allocate capital through the correct pricing of companies. Companies can use primary markets to raise capital necessary to fund further growth, and the secondary markets can provide feedback on the company's future plans. However there is a growing body of evidence that markets are prone to irrationality (Shiller, 1981), and lack liquidity in times of uncertainty. The market structure proposed here aids in efficient, rational price discovery. In the current market structure there are market participants who make money by providing liquidity. Brokerages compete with one another to provide liquidity,

and are prepared to pay large sums of money to be physically close to the exchange so that their trades can be executed first¹ (Wissner-Gross and Freer, 2010). The proposed market structure has liquidity built into it, and does not require such liquidity providers.

Excess volatility in the stock market can be interpreted as evidence of insufficient information in the market. Market participants are aware that their pricing of a financial asset is based on incomplete information. A significant price movement can provide weak statistical evidence for the discovery of new information which is being spread through the market. However the agents are unable to distinguish from the equally-likely scenario that the price move is unrelated to any new significant information and instead presents an opportunity to provide liquidity by selling high or buying low. The prudent choice is to follow the trend until certain that there is no new information. However if significant numbers of market participants follow small trends then it is possible to profit from this herding behaviour by placing disruptive trades which move prices disproportionately and then moving against the trend once it has developed. All these strategies add noise to price discovery as well as adding noise to investors' performance². Moreover if any of these strategies are reliably profitable then the market will rapidly adopt the strategies until they are no longer profitable (Malkiel and McCue, 1985). When excess volatility is posed as a problem of insufficient information, it immediately suggests that a market structure which communicates more information might lead to lower volatility and better price discovery.

Proponents of the efficient-markets hypothesis argue that the random-walk behaviour of stock prices provides evidence that the market is informationally efficient. However there are many different types of random-walk behaviour, and it is possible a market will still be informationally-efficient with much reduced variance. This paper will look at a market structure for which profit is optimised by considering the medium- to long-term value of a share. As a result the new market structure should be far less volatile than current markets and we predict far less trading as a result. As an added bonus the market structure also provides far more information by giving the market's estimate of the distribution of the share price in the medium term.

We now describe the proposed market structure in detail. In the appendix we try to briefly cover minor implementation details.

2 Design

Currently an order in a normal stock market involves an amount of shares (positive or negative for buying or selling), and a price at which the order will take place. This paper proposes a different approach. One instead describes one's demand function. For a given price x the investor describes how many shares of company A they would like to own:

$$f_A(x) = \# \text{ shares demanded at price } x. \quad (1)$$

¹There is evidence that the speed of light is becoming a factor in response times

²A similar criticism can be made of technical analysis which tries to create profit from the analysis of the stock's history.

We assume that the demand function is monotonically non-increasing³:

$$x \leq y \Rightarrow f_A(x) \geq f_A(y) \quad (2)$$

The market for a given company is given by the aggregate demand function, which is the sum of all the investors' (i) demand functions:

$$d_A(x) = \sum_i f_{Ai}(x) \quad (3)$$

and we can guarantee that the aggregate demand function is also monotonically non-increasing. For pedagogical purposes we assume that money, prices and the number of shares a company can issue are real numbers. In appendix A the technical details of handling discrete amounts are covered.

The market determines the current price of a share by finding the point at which the aggregate demand function equals the total number of shares issued. For the purposes of this paper we will assume that this procedure is run at discrete points in time (e.g. once a minute). It would also be possible to run in continuous time, by running the discrete procedure whenever a new order is received (it is possible that significant computational savings are possible when it is known that only a single new order has been received, but this is not explored further). Therefore at time step t the price is determined to be $x^{(t)}$, as this is the point for which $d(x^{(t)}) = N$.

Once the price of the shares has been determined then it is possible to allocate shares to each individual investor. This results in investor i receiving $f_i(x^{(t)})$ shares.

The share price would be in equilibrium until an investor changes their demand function. At this step the algorithm could be run again (or at specified time intervals if trading is frequent). This time the price is determined to be $x^{(t+1)}$, and investor i will now own $f_i(x^{(t+1)})$ shares. Investor i 's trading account will be credited with

$$-x^{(t)}(f_i(x^{(t)}) - f_i(x^{(t-1)})). \quad (4)$$

It is useful to consider a simple example in detail to see the profits that will result from such a strategy. Let company A float in the market with 75 available shares. Investor 1 has a demand function for company A given by:

$$f_1(x) = \max(100 - x^2, 0). \quad (5)$$

Initially Investor 1 owns the entire company, and the price per share is \$5 (since $f_1(5) = 75$). Furthermore Investor 1 currently has no cash in their trading

³It is possible that a rational investor would not have a monotonically non-increasing demand-curve. This situation could arise if the investor believes that an abnormally low price, indicates some unforeseen danger to the health of the company of which he is not aware. The low price then provides statistical evidence for increased uncertainty. In light of this increased uncertainty, an investor might want to hold less of the company. However to ensure stability we remove this possibility. A pragmatic approach in this case would be to halt trading in moments of extreme volatility. Enforcing monotonicity rules out the possibility of stop-losses in the market. Arguably stop-losses are bad for financial markets as they do not contribute to price stability or help price discovery. Later we will also assume that the demand functions are bounded; while this is restrictive in mathematical terms, it will not affect any practical implementation.

account. Investor 2 would like to purchase some of company A and has a demand function as follows:

$$f_2(x) = \max(65 - x, 0). \quad (6)$$

Now the aggregate demand function is:

$$d_A(x) = \max(100 - x^2, 0) + \max(65 - x, 0) \quad (7)$$

and the new price is now \$9 (since $d_A(9) = 75$). Investor 1 now owns 19 shares ($f_1(9)$) and \$504 has been paid into their trading account, which is the exact amount it cost Investor 2 to purchase the 56 remaining shares ($f_2(9)$). Some time passes and for whatever reason, Investor 2 now decides to sell their shares, in which case the price will revert to \$5 per share. Investor 2 receives \$280 at the new price (and hence Investor 2 suffers a loss of \$224). Conversely Investor 1 has a profit of \$224 (although all these calculations are without any of the potential tax incentives discussed later in the paper which encourage medium-term investing).

2.1 Generalisation of the current system

It is easy to show that the proposed new system is an extension of the existing system. If one wanted to buy and hold a set number n of shares in company A , then the trading order would simply be: $f_A(x) = n$, for all price levels x . When deciding to sell the shares one would issue a new trading order $f_A(x) = 0$, and then would have sold all shares in the next iteration. This would also be a useful strategy for owners of the company who do not want to be bought out.

It is also possible to include short-selling in this system. As an example, if one believes that company A is overpriced at \$100, then one can short sell 10 shares by ensuring the demand function includes $f_A(\$100) = -10$.

2.2 Applying a Tobin tax

The aggregate demand function would provide useful pricing information, and can be displayed as a relevant summary of the company's worth by the market. In order for investors not to try and communicate false information (bluffing) through their influence on the demand function, we recommend a small Tobin tax to be applied whenever an investor wishes to change their demand function. As the tax rate increases this forces investors to consider a longer time-frame for their investments. It must be stressed that once a demand function has been specified there should be no further trading costs as the demand function trades in and out of positions with price fluctuations.

While we believe a Tobin-style tax should be applied to discourage any false information in the short term there should be an incentive to maintain accurate pricing information in the medium term. This could be achieved through the lack of tax on updates to the demand function which are older than 3 months.

It is not clear how the Tobin tax should be applied. It is possible that a simple tax should be levied on the entire portfolio whenever changes to the portfolio are made. This would discourage the holding of a diversified portfolio with many shares though. Another alternative would be the use of a distance function which measures the difference between two demand functions. The

tax would then be proportional to both the distance between the old and new demand functions, and the amount of money invested. This would not penalize fine-tuning of the demand function and allows better pricing information. However it is not clear what the optimal distance function should be.

We will avoid further discussion of this issue and simply assume that a pragmatic solution (which has the benefit of simplicity) would be to tax all invested assets at a 1% rate for any update to the demand function that is less than 3 months old⁴.

2.3 Maximizing profit

Here we show that committing to a demand-function style of trading will provide high-quality pricing information. To maximise profit from such a market requires the prediction of the price distribution over the next three months. The share price of this company in the medium term could depend on many things, eg. the possibility of winning a big contract, or the threat of suffering an extended labour dispute.

If we assume that the investor is small relative to the market then the investor's demand curve will not materially affect prices⁵. The trading profit from a demand curve for a given price trajectory will be:

$$- \sum_{t=1}^T x^{(t)} (f_i(x^{(t)}) - f_i(x^{(t-1)})) \quad (8)$$

and we will have a final holding $f_i(x^{(T)})$ shares. Focusing on insight we can think of the limit when demand function updates occur in continuous time and get the following integral:

$$- \int_{t=0}^{t=T} x(t) \frac{df_i}{dx} \frac{dx}{dt}. \quad (9)$$

This shows that profit is maximised by having a steep demand curve in areas where the price is likely to be volatile. In short one maximizes trading profit by being a liquidity provider for market-related prices.

It is interesting to note from equation 9 that the trading profit from a demand curve is unchanged if a constant value is added (while the trading profit is unchanged, the holdings and dividends will certainly change). If all demand curves are bounded, then without loss of generality consider a translation and linear scaling of the demand curve, such that:

$$f(0) = 1.0 \quad (10)$$

$$f(\infty) = 0.0 \quad (11)$$

then the (negative) derivative of the demand curve forms a well-defined probability distribution. This distribution can be interpreted as a prediction of the

⁴It is still possible to game this suggested system by allowing a single portfolio to be shared amongst multiple accounts and each account could be updated at an interval of 3 months, which would effectively update the entire portfolio at a rate much more frequently than 3 months.

⁵This section could be considerably expanded with more technical details, and an exploration of what the Nash equilibrium solution looks like when one's own demand curve does affect prices.

medium term price of a share, since profit is maximised when the derivative is large for market-related prices. Thus the aggregate demand function represents the market's believed distribution over the true (but unknown) medium-term value of the company.

The notion of market-related prices means that the pricing system does suffer from the Keynesian beauty contest problem (Keynes, 1936). Our profits are maximised by estimating what the market prices will be, rather than what a fair value of the stock is. This is ameliorated by the fact that every investor is required to commit to their supplied demand curves or pay a penalty if they do not. This commitment means that there is a risk of increased losses if prices are significantly different from a fair value.

If we let σ_A represent the short-term volatility ($\approx \frac{dx}{dt}$) of company A then the profits that accumulate from an investment in A will be proportional to σ_A . The same will apply for an investment in company B , which will provide an incentive to invest in volatile companies and reduce the volatility. Under this framework, volatile companies are difficult to price by the market, but bring increased profit when accurately priced.

3 Discussion

This market design allows all participants to explicitly state their demand function and trade effectively on it. Even if shares are infrequently traded there will be minimal liquidity risk. This market design offers the opportunity of reduced trading costs for investors who choose to invest on fundamentals and provide high-quality pricing information. This pricing information can be aggregated and displayed as common knowledge and is useful to both management and investor alike.

The proposed market works with a Tobin tax to discourage short-term speculation and increase the time-frame which investors should consider, without the negative effects of throwing "sand in the wheels of our excessively efficient... markets" (Tobin, 1978). In particular, applying a Tobin tax to the traditional financial markets reduces liquidity, but a Tobin tax to the proposed market leaves liquidity unaffected.

However it is not clear that any market participants would want to trade in such a market. The requirement of explicitly stating one's entire demand function is onerous when compared to the current system. Practically it might require legislation for market creation.

Finally while it is possible to structure this market to trade in real-time, it is a natural concern that this market will penalize more casual traders who are unable to continuously monitor the market. The casual trader can suffer from negative news as professional traders offload the stock at high prices (and the casual trader's unchanged demand function automatically buys the stock). Conversely the casual trader also suffers when extremely positive news is made public, and professional traders are able to buy the stock very cheaply. As a possible solution to this concern, one could instead allow the market to trade only at specified times (e.g. once an hour, or once a day), placing the casual trader at less of a disadvantage.

This article has argued that it is possible to construct a new market structure which will have less volatility and very little liquidity risk. This new market

structure requires a trader to provide more information to be successful, but this is good news: society uses financial markets as information aggregators, and now there is more information to aggregate. This market structure is impractical without significant technology, but again there is good news: we finally have sufficient computational resources to implement such a market on a large scale. Further research into this market structure could provide an efficient practical market that is less prone to gyrations than the current market and could help global financial stability.

A Discrete prices and shares

This paper has been greatly simplified by assuming that prices and shares are continuous real variables. In this appendix we relax that constraint and explore how a system would be implemented with discrete prices and shares.

Let N be the number of issued shares. For each discrete price level x the aggregate demand function will also be a discrete number $d(x)$. It is extremely unlikely that there will exist a price for which $d(x) = N$. Instead the price will exist partially between two discrete prices. Let x^* be the price which satisfies $d(x^* - 1) < N < d(x^*)$. All the orders for price level x^* can be satisfied with some remaining shares left unsold. These shares can be sold at the price level $x^* - 1$, however there is an oversupply of buyers at this price. To overcome this two queues, one for buyers and one for sellers, are created and the orders can be filled as prices change. If aggregate demand falls (meaning that $N - d(x^{t+1})$ increases, and more shares will be sold at the lower price) then more shares are allocated to orders in the buyers' queue and those orders are added to the back of the sellers' queue. Conversely if aggregate demand rises then more shares are sold at the higher price and the orders are removed from the head of the sellers' queue and placed at the back of the buyers' queue.

B Other practical matters

This short paper has only scratched the surface of how such a market might work. In practice the devil is in the details. In this appendix we try to cover as many issues as possible, and mention a potential solution for each. Some of the solutions might not be optimal, but will provide a concrete starting point for improvements. We do not believe that there is any single issue which is fatal for the suggestion.

B.1 Material changes in value

A strong criticism of this market structure is that it does not allow a fast response to unforeseen circumstances. If an oil company creates an ecological disaster then fast-moving investors will be able to liquidate their holdings in the oil company and the shares will automatically be bought by other investors. For genuinely unforeseeable events a Tobin tax is punitive and arbitrarily favors investors who have fortunately waited enough time and are able to update their demand function without a penalty. In such circumstances, a halt in the trading of the share might be appropriate, along with the opportunity to update the demand function without penalty.

However it is important to stress that this is only appropriate for genuinely unforeseeable events. If a company surprisingly wins a large and lucrative contract, this does not constitute an unforeseeable event, merely an unlikely event. In practice this distinction is not clearly defined, and might require an independent committee to make such rulings.

B.2 Dividends

Shares have value as they represent both ownership of the assets of a company as well as the right to a portion of the profits. These profits are distributed as dividends. In the conventional stock market there is an increase in the price of a share preceding the dividend distribution and an immediate price drop after the share goes ex-dividend. In order to avoid the price movements associated with a dividend payout, it is recommended that investors accumulate the right to a dividend with each time period that a share is held.

Thus if a share pays a yearly dividend then the dividend will be distributed to all investors who have held the share in past year. This payment will be proportional to the amount of shares they have held and the length of time they have held them.

If a share has been sold short then the payment of dividends must be deducted from the investors account also in a manner proportional to the number and duration of shares held.

B.3 Creating a portfolio

An investor would create a portfolio by creating demand functions for several companies. The investments could be funded from a single pool of money which is automatically invested in government debt. This government debt would return a constant stream of payments every quarter, with the income calculated in a similar manner to dividends.

It is also possible that several shares lower their price and the investor's demand functions suggest buying more shares than there is capital for. In this case the investor is over-committed and needs to reduce their positions. There could be an automatic reduction if the investor is over committed, by dividing each demand curve by the proportion that the investor is over committed.

Using a combination of short-selling (through negative demand functions) and investing it is possible for an investor to leverage up in an unbounded fashion. Thus the market would also be required to limit leverage.

B.4 Voting Rights

The ownership of shares which could be sold at any minute (without the shareholder's active participation) poses an interesting problem for board oversight. A simple first suggestion would be that voting rights are only granted to shares held for any price (the demand function is constant).

Voting rights are further complicated by the suggested notion of allowing negative demand functions for short-selling. This allows many people to own the share, thus there might be 20% of the company's shares sold short, and 120% shares long. Since people typically own many shares, a simple solution

for this example would be to allocate a voting right for each share held with probability $p = \frac{100}{120}$.

B.5 Raising new capital

Under this proposal a company's management team might be able to automatically raise new capital. If there are N shares of a company and a single share of this company is worth \$100, then it is plausible that giving the company another \$100 in exchange for a new share will leave the value of all other shares unchanged. Letting V represent the market capitalization of the company:

$$\frac{V}{N} = \$100 \quad (12)$$

and the new issue results in a per-share valuation of:

$$\frac{V + \$100}{N + 1} = \frac{V}{N} - \frac{V}{N(N + 1)} + \frac{\$100}{N + 1} = \$100. \quad (13)$$

In practice the value of the new share would be strongly influenced by the company's plans for the new capital. This means that the intention to raise new capital should be advertised well in advance so that investors have the ability to change their demand functions without penalty.

Since the demand functions have been fully specified by the market, a management board could raise new capital by simply issuing more shares to the current demand functions. Using the argument above the share price need not change substantially. However the board has a very clear idea of how much money will be raised should new shares be issued, since the aggregate demand function is public knowledge.

References

- Keynes, J. (1936). *The general theory of employment, interest and money*. Atlantic Publishers & Distributors.
- Malkiel, B. and McCue, K. (1985). *A random walk down Wall Street*. Norton New York.
- Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends?
- Tobin, J. (1978). A proposal for international monetary reform. *Eastern Economic Journal*, 4(3/4):153–159.
- Wissner-Gross, A. D. and Freer, C. E. (2010). Relativistic statistical arbitrage. *Phys. Rev. E*, 82(5):056104.