

The Economics of the Dealer Function

A dealer facilitates market liquidity by intermediating between transactors to whom time is important in exchange for charging buyers a higher price than he pays sellers. A value-based investor may also fulfill this function, but at a larger bid-asked spread than that imposed by the dealer. Relative to the value-based investor, the dealer has limited capital, hence limited ability to absorb risk; he will thus limit the position—long or short—he is willing to take.

When the dealer's position reaches a maximum, he will lay off to the only other transactor motivated by price—the value-based investor. The dealer's price is tied to the value-based investor's price at these layoff points. As the value-based investor shifts his prices in response to new information, the dealer's interior prices shift along with his layoff prices.

An investor should realize that, when he trades with the crowd, he is trading at the value-based investor's spread, which may be many times the size of the explicit dealer's spread. More generally, the actions of the crowd—whether it is buying or selling, and in what volume—will determine whether the price of trading quickly is high or low, hence whether the value of his information justifies trading.

A MARKET-MAKER may be defined as someone who accommodates transactors to whom time is important in return for the privilege of charging buyers a higher price than he pays sellers. By this definition, both dealers and value-based investors (VBTs) are market-makers. Yet their roles differ in several important respects—

- in amount of capital, hence ability to absorb losses;
- in length of holding, hence exposure to getting bagged;
- in the spreads (i.e., the difference in bid and asked price imposed on simultaneous purchases and sales).

In particular, the VBT's spread is larger than the dealer's spread; we call them, respectively, the "outside" and "inside" spread. In the absence of dealers, transactors in a hurry would buy and sell at prices that differ by the full outside, or VBT, spread—even if purchase and sale took

place only seconds apart. By intermediating between hurried buyers and hurried sellers, a dealer enables them to benefit from each other's trading, even if the trades aren't simultaneous.

Dealers are thus valuable to transactors in a hurry, because they greatly reduce the spreads encountered by those transactors. By doing so, they also greatly improve the liquidity of the markets in which they deal. Alas for the dealer and for market liquidity, a seller is not always followed by a buyer. Indeed, even if the arrival of buyers and sellers is random, a seller may be followed by a long run of sellers (or a buyer by a long run of buyers), with the result that the dealer builds up a large position.

Compared with the VBT, the dealer has very limited capital with which to absorb an adverse move in the value of the asset. Furthermore, the dealer's spread is too modest to compensate him for getting bagged. The dealer consequently sets limits on the position—long or short—he is willing to take. When his position reaches a limit, he lays off to the only other transactor in the market who is motivated by price—the value-based investor. (Strictly speaking, when his position grows uncomfortably short, he "buys in"; to avoid circumlocution, we shall use the term "lay off" algebraically.) In effect, the

Jack Treynor is President of Treynor Capital Management, Inc., a professor of finance at the University of Southern California and a member of the Editorial Board of this journal.

value-based investor is the market-maker of last resort. Figure A combines these elements in a diagram.

In the problem we address, the value-based investor's bid and asked price and the standard size of orders coming to the dealer for accommodation are givens. We also take as given the maximum position—long or short—the dealer is willing to assume. We ask two questions:

- How will the dealer's mean price—the mean of his bid and asked—vary with his position?
- How big will the dealer's spread be? What determines it?

We shall assume that VBTs get new information as soon as the dealer's customers. Otherwise, of course, accumulations in the dealer's position will not be unaffected by the arrival of new information.

When information reaches the VBT, his bid and asked prices shift to reflect it. Because the dealer's price is tied to the VBT bid and ask at his layoff points, his prices move along with the VBT prices. In general, therefore, dealer prices are responding to two different forces—changes in the VBT's estimate of value and changes in the dealer's position.

Determining the Dealer's Spread

Dealers have salaries, telephone bills and other costs, just like any businessman. Unless these costs have a significant variable component, however, a dealer's dominant variable cost will be the cost of laying off. If, in dealing, price is related to variable cost, then the price the dealer exacts for his services—the dealer's spread—will be related to the cost of laying off—the outside spread.

In the limiting case of perfect competition among dealers, the revenues the dealer receives from his accommodations will equal the costs of laying off. Because the outside spread will typically be many times the dealer's spread, however, revenues will equal costs only if layoffs are far less frequent than accommodations. More precisely, the ratio of the two spreads must equal the inverse ratio of the respective transaction frequencies. To obtain this ratio, we need to know the frequency of transactions.

The Frequency of Layoffs

Perhaps the simplest way to think about this problem is in terms of accommodation trades of a fixed size. Such trades cause the dealer's position to jump from one inventory position to an adjacent position. The continuum of dealer

Figure A Dealer's Spread and Maximum Position





Figure B Dealer's String of Transactions



positions is thus reduced to a number of discrete positions, like beads spaced evenly along a string. Purchases and sales arrive in random order (but equal frequency), so moves up or down the string occur in random order. This is illustrated in Figure B.

At the ends of the string are beads corresponding to the dealer's maximum tolerable positions. We call them the "layoff positions." (We assume the dealer is willing to adjust his layoff positions so they are separated by a whole number of standard accommodations.) When the dealer's position reaches either extreme, the next transaction may move it back toward a neutral position or forward beyond the dealer's maximum. In the latter case the dealer either buys in (paying an asking price above his own asking price) or lays off (realizing a bid price below his own bid). Thus every share (every unit) laid off represents a loss to the dealer.

If there are no fixed costs of laying off that can be spread over the units laid off, it behooves the dealer to lay off only the units (long or short) acquired in accommodating the current trade. After such a layoff, his position is restored to the layoff position, from which subsequent accommodations will sometimes move him back toward the neutral position at no additional layoff cost.

To define this process algebraically, let X be the dealer's position, X^* the dealer's maximum position, S the standard accommodation and $G(X)$ the frequency with which the dealer finds himself in that position in the steady state. Then, for interior positions, we can write:

$$G(X) = 0.5 G(X - S) + 0.5 G(X + S),$$

reflecting the fact that buy and sell accommodations are equally likely. In other words, the frequency with which the position X occurs depends on the frequency with which the adja-

cent positions occur, times the probability (0.5 in each case) of moves from those positions toward the X position, rather than away from it.

We can rewrite this relation as follows:

$$0.5 G(X) - 0.5 G(X - S) = 0.5 G(X + S) - 0.5 G(X).$$

Now its meaning is clearer: The rate of change of $G(X)$ is everywhere the same. Only a straight-line function of X satisfies this condition. Furthermore, the symmetry between buy and sell orders dictates that this function be symmetric with respect to positive and negative values of X . The only straight-line function that satisfies this condition is a horizontal line: $G(X)$ is a constant; the probability of each position is the same.

If layoffs are the same size as a standard accommodation then, when a dealer reaches his layoff position, his next accommodating transaction is equally likely to (1) move him one position closer to neutrality or (2) force him to lay off, in which case the net effect is to return him to the layoff position. If X^* is the upper layoff position, then in the steady state we have:

$$G(X^*) = 0.5 G(X^*) + 0.5 G(X^* - S),$$

$$G(X^*) = G(X^* - S) = G(X).$$

A similar result holds for the lower layoff position.

If all possible positions, including layoff positions, occur with the same frequency, then layoff positions occur with a frequency equal to the standard accommodation divided by twice the dealer's layoff position, times two, because there are two layoff positions. But layoff positions actually lead to layoffs only half the time. Thus layoffs occur with a frequency equal to the standard accommodation divided by twice the dealer's layoff position:

$$\text{Layoff Frequency} = S / 2X^*.$$

The spread, $p_a - p_b$, that enables the dealer to break even is:

$$p_a - p_b = S/2X^* (P_a - P_b),$$

where $P_a - P_b$ is the outside, or VBT's spread. If the dealer charges more than this, he's covering at least some of his other costs. Not surprisingly, the competitive inside spread is proportional to the outside spread. But it also increases with the size of the standard accommodation and varies inversely with the maximum position the dealer is willing to take.

Determining the Dealer's Mean Price

Now, what about *price*—i.e., the mean of the dealer's bid and asked? The dealer's current price should relate in a rational way to what the price is expected to be in the future. Otherwise, his current price will create profit opportunities across time for those who trade with him.

At positions between his layoff positions, the dealer should set price according to the price he expects to be setting one trade later. The next trade will, of course, move his position up or down with equal probability. If the price the dealer would set in those positions is known, then the price he sets in his current position must be the probability-weighted average of those two prices; otherwise he will create easy profits for those trading against him.

Let the prices for the three positions $X - S$, X and $X + S$ be $p(X - S)$, $p(X)$ and $p(X + S)$, respectively. Because the adjacent positions are equally likely, we have:

$$2p(X) = p(X - S) + p(X + S).$$

This is clearly another straight-line function of position:

$$p(X) - p(X - S) = p(X + S) - p(X).$$

The positions immediately beyond the respective layoff positions have known prices corresponding to the prices at which value-based investors will accommodate the dealer. The straight-line price function must satisfy those prices. If we let S be the standard transaction quantity, then we have:

$$\begin{aligned} p(X^* + S) &= P_{\text{BID}} = P_b, \\ p(-X^* - S) &= P_{\text{ASK}} = P_a \end{aligned}$$

and

$$p(X) = \frac{P_a + P_b}{2} - \left(\frac{P_a - P_b}{X^* + S} \right) \frac{X}{2}.$$

The ratio $(P_a - P_b)/(X^* + S)$ measures the sensitivity of the dealer's mean price to changes in his position. It depends on both (1) the outside spread, reflecting the risk character of the asset, and (2) the dealer's willingness to take a position.

What these results show is that it is expensive to buy when everyone else is in a hurry to buy and expensive to sell when everyone else is in a hurry to sell.

We have begged the question of how big a position the dealer should tolerate. The answer probably has something to do with whether value-based investors, who help determine the dealer's mean price, get new information as quickly as information-based investors. It probably also has something to do with the risk character of the dealer's other assets, and with the size of his capital. Rich people make the best dealers.

Pricing Large Blocks


In the real world, of course, the individual accommodation trades brought to the dealer will vary in size. This raises the question: How should the dealer price trades larger than the standard trade? In particular, should he price the trade on the basis of the average position incurred in accommodating a trade, or the final position? If the average, then the cost of large trades is the same as the cost of small trades. If the final, the effective cost is much higher.

If large trades are frequent, then they will affect the probabilities we have assumed for transitions from one dealer position to another. If such trades are sufficiently infrequent that we can safely ignore their effects on the probabilities, then we can treat the occasional large trade as a string of standard trades.

Consider, for example, the cost of a large round trip as depicted in Figure C. If the dealer's purchase and sale prices are based on his average position, then the cost of the round trip is the shaded area in the figure. This implies a cost per share equal to the area divided by the number of shares traded or the inside spread. If, however, the dealer's prices are based on his final positions (i.e., after the customer's purchase is completed and after the customer's sale is completed), then the round-trip cost of trading is depicted by the shaded area in Figure D.

Unfortunately for the customer, the key to rational pricing behavior on the part of the dealer in this situation is our earlier comment

Figure C Cost of Large Round Trip,
Based on Dealer's Average Position



that "the dealer should set price according to the price he expects to be setting one trade later." This price is determined by the dealer's final position when the trade is completed—not his average position during the trade. Although the *next* trade is equally likely to be in the same direction or the opposite direction, he knows that the *current* trade is equivalent to an unbroken run of standard-sized trades in the same direction.

A dealer will price a large block on the basis of his final position (rather than an average of his intermediate positions) because, in contrast to the assumptions underlying the standard accommodation model, he knows that his position will not fluctuate randomly around the intermediate positions. Instead, it will fluctuate randomly around the final position. The prices corresponding to that position should thus apply to the whole block. This, of course, implies that the size effects in prices are not reversible: The customer doesn't get back when he sells the block what he paid when he bought it.

What about the cost of trading a series of smaller blocks—i.e., of trading so the dealer doesn't know how big the whole series is until the last trade? In this case, the customer can get intermediate prices for intermediate trades, paying the final price only for the final trade. Of

course, if the smaller blocks are big enough to push the dealer to his maximum, then nothing has been gained, because the outside spread price would have governed if the entire block had been handled as a single trade. In the meantime, too, the customer runs the risk that the information motivating his trade will get impounded in the price (i.e., the mean of the VBT's bid and ask) before he completes his trade.

In sum, (1) a large block will move the dealer's position, hence his price, in a direction that will increase the price of the trade, unless (2) the dealer's position is already at the maximum limit to which the block would otherwise move him, in which case (3) the size of the block has no effect on the dealer's price.

Valuation Errors in VBTs' Estimates

So far we have assumed that VBTs estimate the value of the asset in question correctly. This implies that they agree, in which case the cumulative probability distribution of their assessments is the Z-shaped distribution given in Figure E. Actual bid and asked prices, set one-half the outside spread below and above the assessment, will have their own cumulative probability distributions, which will also be Z-shaped, echoing the shape of the assessment distribution.

Figure D Cost of Large Round Trip,
Based on Dealer's Final Positions





Figure E Cumulative Distribution of Value-Based Estimates of Value,
When Estimates are Correct




If the assessments of value-based traders are in error, however, they will be dispersed around a central assessment, and their cumulative probability distribution will no longer be Z-shaped. It will instead be S-shaped, as in Figure F, with gradually rounded corners and long,

tapering tails. As before, the distributions of value-based traders' bid and asked prices will echo the assessment distribution. They too will now be S-shaped rather than Z-shaped, with rounded corners and long, tapering tails.

The bid and asked distributions will still be

Figure F Cumulative Distribution of Value-Based Estimates of Value,
When Estimates are in Error



Shaded areas represent negative layoff price (and buy-in cost) for the dealer.

set one-half the outside spread to the left and right, respectively, of the distribution of value assessments. But something curious has happened: The distance between the upper tail of the bid distribution and the lower tail of the asked distribution has narrowed. Error in value-based traders' assessments has reduced the dealer's cost of laying off.

Strictly speaking, if the error distributions have infinitely long tails, the dealer's cost of laying off has been eliminated. What really matters to the dealer are the prices that elicit the necessary volume of bids and offers from the tails. It should be clear that, as the average (e.g., standard) error in value-based traders' assessments increases, the cost of laying off in the required volume declines.

Many investment professionals assert that their work in assessing security values not only serves their clients, but also makes security markets more efficient. It is certainly true that, without value-based investors, dealers would have no one to lay off to. But now we see that, if competition among dealers is sufficiently brisk, any reduction of outside spread resulting from value-based investors' errors will be passed along to the inside spread, thereby improving market efficiency. In this case, of course, dealers will be indifferent between less error on average in value-based investors' assessments and more; it is their customers who are the ultimate beneficiaries of larger assessment errors.

If, on the other hand, dealers are not subject to the pressures of competition, then the savings they realize from an increase in the average size of value-based traders' assessment errors will not be passed on to their customers in the form of smaller inside spreads. Markets will not be more efficient. And dealers will no longer be indifferent between less error and more—which is to say, between higher standards of investment analysis and lower ones.

The Economics of Investing

What are the lessons for the reader who wants to be a successful investor, rather than a successful dealer?

First, in reckoning the cost of any trade, there are two spreads to consider—the inside spread and the outside spread. The latter, which is what the dealer pays to trade at a time of his choosing, is also what the investor pays to trade with the crowd. Normally invisible to the inves-

tor, it is often an order of magnitude or two bigger than the more readily visible inside spread.

Second, when an investor comes to the market with insights not yet impounded in the price, what he pays for speed depends on what the crowd, often motivated by different information, is paying for speed. In particular, it depends on whether the information-motivated crowd is eager to buy or eager to sell. If he is buying when the crowd is selling, for example, he is in effect market-making to the crowd. He is receiving, rather than paying, some or all of the outside spread. And if it turns out that the information motivating the crowd was not yet in the price, he will get bagged along with those other investors who make it their business to accommodate information-motivated investors—namely, value-based investors.

Third, and more generally, what the information-motivated crowd is doing—whether it is buying or selling, and in what volume—determines the current price of trading fast. The investor needs to know this price in order to judge whether the time value of his own insight is high enough to make the price worth paying. If it is, he trades. If it isn't, he doesn't trade.

Fourth, because orders motivated by liquidity tend to arrive as a random mixture of buys and sells, whereas orders motivated by information don't, the latter are much more likely to push the dealer to the extreme of laying off to or buying in from the value-based investor. These considerations are important for the value-based investor attempting to set his spread so that gains from liquidity trading will be large enough to offset losses on information trading.

Fifth, because liquidity-motivated trading is by definition uncorrelated with information-motivated trading, hence with trading by the crowd, its expected cost is the inside spread. But the cost of "pseudo" information-motivated trading is the outside spread. The volume of "pseudo" trading is critical to the viability of the value-based investor.

Finally, the functions of a trading desk should be to

- (a) estimate the inside spread on all securities of interest to investors;
- (b) estimate the outside spread;
- (c) maintain running estimates of outside bid and ask on all securities of current trading interest: the price of trading quickly is the

- difference between the price of the trade and the mean of the outside bid and ask;
- (d) obtain from research estimates of the time value of current recommendations and
 - (e) match (c) and (d). ■

Appendix

Our argument for each position occurring with the same frequency is highly heuristic, to say the least. When the dealer's possible positions are reduced to a limited number of discrete states, the basic structure of the problem is that of a Markov process. It is well known that the steady-state probabilities of such a process are related to the transition probabilities by the requirement that the product of the vector of steady-state probabilities and the matrix of tran-

sition probabilities be the same vector of steady-state probabilities.

We can thus test our heuristic conclusion that the dealer's position frequencies are all equal by testing the truth of the matrix equation given in Table A1. Inspection confirms that the equation is satisfied.

Table A1

$$\begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \begin{pmatrix} 0.5 & 0.5 \\ 0.5 & 0 & 0.5 \\ & 0.5 & 0 & 0.5 \\ & & 0.5 & 0 & 0.5 \\ & & & 0.5 & 0.5 \end{pmatrix} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}$$