Optimal Speculation
Against an Efficient Market

by

Jeffrey F. Jaffe

and

Robert L. Winkler

Working Paper No. 19-74

RODNEY L. WHITE CENTER
FOR FINANCIAL RESEARCH

University of Pennsylvania

The Wharton School

Philadelphia, Pa. 19174

1. Introduction

Much of the literature in finance suggests that securities markets are efficient in the sense that they fully reflect all publicly available information. However, some studies have shown that certain groups of individuals with private information reap profits above the returns to the market as a whole.[1] This paper presents a model to help sophisticated traders to make the most of their information. The model also illuminates the process by which special information is incorporated into market prices.

Much of this paper is based on the existence of a market-maker who stands ready to buy claims on assets at a given price (usually called the bid price) and to sell these claims at another given price (the ask price). Thus, the paper has implications for the specialist system. While previous studies have emphasized the costs market-makers incur by the employment of their own time and capital,[2] these previous analyses have ignored an important cost to market-makers: the potential losses from trading with individuals who possess special information. This oversight may result from the belief that market-makers are extremely knowledgeable about their asset and thus would not lose to speculators. The present paper focuses on this potential cost and shows that a market-maker can always expect to lose when trading with a rational individual, even if the market-maker is more knowledgeable and even after the bid-ask spread is included. Under the specialist system, speculators have profitable opportunities which would not arise in a more perfect market.

In Section 2, we introduce a model which allows speculators to maximize their profit for a given amount of information when a market-maker is present. This strategy can be applied to any economic market (e.g., foreign exchange market, stock and bond exchanges, option markets, commodity markets, etc.). In Section 3, some implications for financial theory are discussed. Some practical applications of our model are discussed in Section 4, and the paper is summarized in Section 5.

2. Speculative Trading with a Market-Maker

Individuals usually trade assets indirectly through a broker and a market-maker. The broker, who is responsible for executing transactions in the investor's name (as well as for providing advice), is compensated by a commission. As orders to transact in a particular asset occur irregularly, investors who desire to buy (sell) cannot be certain of locating investors who want to sell (buy) immediately. To insure that individuals can transact promptly, the market-maker stands ready to buy the asset at a specified price (bid price) and to sell the asset at another specified price (ask price). The market-maker is compensated by the difference between the ask price and the bid price, which is commonly called the spread. Below, a model focuses on transactions against a market-maker.

To formalize the notion of trading, consider a simple situation involving one asset a risk-neutral investor called A, and a market-maker called B. As B is the market-maker, he sets a price $P_B + T$ at which A can buy the asset and a price $P_B - T$ at which A can sell the asset. Thus, 2T is the spread. (This quantity can alternatively represent the sum of the spread and the broker's commission). Obviously, T is nonnegative, with T=0 corresponding to the case in which A can buy <u>or</u> sell at P_B, the price set by the market-maker (i.e., the market-maker's estimate of the value of the asset).

Let P_A represent the investor's estimate of the value of the asset. The estimates of the investor and the market-maker can then be written in the form $P_A = P + u_A$ and $P_B = P + u_B$, where P is the "true" value[3] of the asset and u_A and u_B are error terms. Assuming that the two individuals are knowledgeable and experienced enough to avoid systematic errors, suppose that $u = (u_A, u_B)$ has a bivariate normal distribution with $E(u_A) = E(u_B) = 0$, $V(u_A) = \sigma_A^2$,

$V(u_B) = \sigma_B^2$, and $Cov(u_A, u_B) = 0$. [In Appendix I, the more general case with $Cov(u_A, u_B) = \rho\sigma_A\sigma_B$ is presented]. Moreover, it is assumed that this distribution is known to A. For example, the distribution could be based on considerable past data in the form of estimates by A and the corresponding observations of B's prices.

After A determines P_A, but without seeing P_B, A's prior distribution for P is a normal distribution with mean P_A and variance σ_A^2. After learning the value of P_B, A's posterior distribution for P is a normal distribution with mean P_A^* and variance σ_A^{*2}, where

$$P_A^* = \frac{k^2 P_A + P_B}{k^2 + 1}$$

and

$$\sigma_A^{*2} = \frac{k^2 \sigma_A^2}{k^2 + 1}$$

with $k = \sigma_B/\sigma_A$. The derivation of the posterior distribution is given in Appendix I.

The amount of the asset that is traded is not important for our purposes, so it is assumed that just one unit (share) of the asset changes hands.[4] Alternatively, given the assumption that A is risk-neutral, it might be more realistic to assume that he would purchase an infinite amount on buying opportunities and sell short an infinite amount on selling opportunities.

The expected return to A from buying the asset is $P_A^* - (P_B + T)$, so it is optimal for A to buy when $P_A^* - (P_B + T) > 0$. This is equivalent to

$$\frac{k^2 P_A + P_B}{k^2 + 1} - P_B > T,$$

which simplifies to

$$P_A - P_B > \beta T$$

where

$$\beta = \frac{k^2 + 1}{k^2} = 1 + \frac{1}{k^2} .$$

The expected return to A from buying the asset is positive when $P_A - P_B$ is greater than T by at least T/k^2. The smaller k is, the larger the price difference that is required before it is advantageeous for A to buy the asset. But $k = \sigma_B/\sigma_A$, so k decreases as the precision of the market-maker's forecasts increases relative to the precision of A's forecasts, where precision is inversely related to error variance.

The situation in which A sells the asset is analogous to the buying situation. The expected return to A from selling is positive when

$$P_B - P_A > \beta T.$$

Thus, combining the buying and selling situations, A has a positive expected return (i.e., A will trade) whenever

$$\left|P_A - P_B\right| > \beta T.$$

What is the probability, calculated before seeing P_A and P_B, that A will trade? This probability is $\Pr(|P_A - P_B| > \beta T)$, which by symmetry is equal to $2\Pr(P_A - P_B > \beta T)$. Since $P_A - P_B$ is normally distributed with mean zero and variance $\sigma_A^2 + \sigma_B^2 = \sigma_A^2(k^2 + 1)$, the probability that A will trade can be written in the form $2\Pr\left[z > (k^2 + 1)^{1/2} T k^{-2} \sigma_A^{-1}\right]$, where z is a standard normal random variable. Letting τ represent the probability that A will trade and letting $\theta = (k^2 + 1)^{1/2} T k^{-2} \sigma_A^{-1}$, we see that $\partial\tau/\partial T = (\partial\tau/\partial\theta)(\partial\theta/\partial T)$. But $\partial\tau/\partial\theta < 0$ and $\partial\theta/\partial T > 0$, so that τ is a decreasing function of T, as would be expected. Moreover, $\partial\theta/\partial k < 0$, implying that τ is an increasing function of k. As A's error variance decreases relative to the market-maker's error variance, A is more likely to trade. The appearance of σ_A^{-1} in the expression for θ suggests that with all other variables held constant, a smaller σ_A makes trading less likely. But "all other variables held constant" means that k, a function of σ_A^2, is held constant, so that σ_B^2 decreases as σ_A^2 decreases. The increased precision of both A and the market-maker means

The increased precision of both A and the market-maker means that the distribution of $P_A - P_B$ is tighter. But T is being held constant, implying that $P_A - P_B$ is less likely to overcome the spread and/or transaction costs. If σ_B/σ_A is substituted for k in the expression for θ, differentiation yields $\partial\theta/\partial\sigma_A > 0$, which implies that $\partial\tau/\partial\sigma_A < 0$. Therefore, the probability that A will trade is a decreasing function of σ_A^2, as would be expected.

The above results imply that any investor with some information, no matter how little can profit from that information by employing a decision rule such as the one arrived at by the model of this section. All investors have a positive probability of trading with the market-maker (i.e., a positive probability of encountering a situation with a positive expected return from buying or selling at the prices offered by the market-maker). However, the probability of trading is a decreasing function of the investor's variance of error of estimation. The more accurately an individual forecasts the price of an asset and the smaller T is, the more likely it is that a trade will occur. Thus, the probability of trading is a function of the degree of information that the investor possesses. For investors with little information, the probability of trading is small (i.e., the number of opportunities for profitable trading is small). On the other hand, for very knowledgeable investors, the number of opportunities for profitable trading may be large. The volume of trading will depend on the forecasting abilities of the investors relative to that of the market-maker.

Because the market-maker can expect to lose to all rational investors, a purely speculative market with a market-maker is unstable. The market-maker will let $T \to \infty$ in order to reduce his trading losses unless he is either subsidized in some manner and/or the size of T is restricted.[5] Liquidity traders, i.e., individuals who trade without information in order to adjust their risk level or total investments level, as well as investors who either misperceive their risk level or total investments level, as well as investors who either misperceive their forecasting ability or do not use rules such as those given in this section, would subsidize the market-maker in the real world.

We have implicitly assumed in our discussion of the model presented here that once the bid and ask prices are set, the market-maker must meet the excess demand or supply by adjusting his inventory. However, a market-maker could alternatively shift the price of a security so that the rate at which investors sell to him is equal to the rate at which they buy from him. Here the profits of a shrewd investor are made at the expense of other investors, not at the market-maker's expense. In practice, a specialist does not take a position when he crosses limit orders on his book with incoming market orders.[6]

In addition, the specialist on the NYSE is often allowed to cross transactions for his personal account with limit orders on his book.[7] To trade optimally in this last case, the specialist himself can use the model presented in this section.

3. Implications for Financial Theory

In order to determine optimal speculation rules for investors, we treated the specialist's estimate of the value of a security as a constant in Section 2. However, this estimate may change when the market-maker receives an order from the floor. Our paper can shed light on this process whereby specialized information is incorporated into market prices.

For example, consider a hypothetical world where the trader must reveal his price estimate, P_A, after trading with the specialist. If the specialist knows his own error variance, σ_B^2, and the error variance of the trader, σ_A^2, this specialist can meld his own prior price estimate, P_B, with P_A to determine a posterior estimate, P_B^*. Since the trader and the specialist agree on the parameters of the distribution of u, each will weight P_B and P_A in the same way in determining his own posterior distribution. Thus, P_B^* will equal P_A^*, as given in Section 2. This implies that even when the information itself is not exchanged, the full value of the information is discounted in the posterior estimates. Thus, in this hypothetical world, our model shows the precise manner in which information is incorporated into market prices.

Though a real world specialist would know a great deal less about individual traders, he would often learn enough to change his quote after transactions are made. Upon receiving an order from the floor, the specialist learns such information as the type of trade (market order, limit order, etc.), its urgency (day order, good till cancelled order, fill or kill, etc.), the brokerage house handling the transaction, whether the order was public or private, and so on. The name of the trader may be revealed in a negotiated transaction, such as a specialist block transaction. In a secondary offering, some information on the type of vendor is revealed approximately a week after the distribution.[8] Thus, a specialist can use records of price movements following certain types of transactions to update his quotes. Our model could be used here, though information such as σ_A^2 would not refer to the errors of a single individual but to a class of individuals and/or transactions. As the specialist accumulates information in this fashion, the specialist's error variance, σ_B^2, should decrease, implying that k should decrease and an investor with a fixed error variance using the model of Section 2 is less likely to trade.

Though the model developed herein has been placed in the setting of an organized exchange, it has implications for all areas of economic activity where assets or claims are traded. For instance, an analogy can easily be drawn between investing in securities and gambling on sporting events. The model of Section 2 can be illustrated in terms of gambling on sporting events by considering betting on football games with a bookie. The bookie is analogous to the market-maker, for he sets the point spreads, which can be thought of as prices, and bettors can choose to bet on either side of a given point spread. The bookie is compensated by a percentage taken from winning bets and/or by winning ties. (A tie occurs when the actual point spread is identical to the bookie's point spread). This compensation is comparable to the market-maker's spread and/or commission. Also as in the case of the market-maker, the bookie's compensation is related in part to his ability to forecast, although in this case the outcomes of football games rather than the values of assets are the objects of the forecasts. Often a bookie would like to set a point spread so that equal amounts are wagered on either side of the point spread. In this way, the bookie is not gambling. No matter what the actual outcome of the game is, half of the bettors are, in essence, betting against the other half, and the bookie's "return" is a percentage of the gains of the winning bettors and/or the possibility of keeping all money bet on either side if a tie occurs.

If the bookie sets a point spread and a majority of bettors prefer one particular side of the point spread, a common procedure is to adjust the point spread in order to "even off" the bets, much as a market-maker might adjust his ask price and bid price in response to investors' actions. A bookie who winds up with more bets on one side of the point spread than on the other side becomes a gambler, winning if the majority of bettors turns out to be wrong and losing if they turn out to be right.

Of course, bettors who are able to forecast the outcomes of football games accurately can, by following decision rules such as those generated by the model of Section 2, obtain a positive expected return. The bettor who knows the relevant distribution of forecast errors can make a forecast on the basis of his judgments concerning a game and then revise the forecast after seeing the bookie's point spread. If the revised point spread is far enough from the bookie's point spread to overcome the effect of the bookie's percentage of winnings and/or policy of winning ties, the bettor should bet on the game.[9]

In the case of investing in securities, the market-maker can offset losses to information traders by profiting from liquidity traders. As in the pari-mutuel situation mentioned in Footnote 9, nonlinear utility and a positive utility for gambling per se may cause bettors to behave differently than the model of Section 2 would indicate. Thus, there may be a class of bettors who would continually lose to bookies. (Similarly, there may be a class of investors who would continually lose to market-makers). This makes the entire situation more attractive to the bookie, overriding the possibility of losing to a few very knowledgeable bettors.[10]

4. Practical Applications

Although the model presented in Section 2 was conceived primarily for theoretical purposes, it should be useful to speculators in practice. In order to use the model, it is necessary for an investor to determine the required inputs to model, and the key input is a probability distribution. The spread and/or transaction costs, represented by T, should be known in any given investment situation, and the remaining factors affecting the results of the model are the parameters of the joint distribution of the forecast errors of the investor and the market-maker.

The determination of probability distributions for future asset prices is discussed in detail in two papers by Winkler,[11] so to conserve space we will not repeat those discussions here. In general, an investor should utilize any information and means for processing information that are available, implying that a probability distribution for future asset prices should typically involve the use of past data, subjective judgments of the investor and any expert who may be consulted, and output from any statistical forecasting models that may be available. In the particular situation posited in this paper, a great deal of relevant past data is available or will become available quickly as the trading process goes on, so an investor's probability distribution should depend heavily on such data. By systematically looking at the past history of his own and the market's errors of estimation, an investor should be able to investigate whether a normal distribution provides a reasonable approximation of the joint distribution of forecast errors. Moreover, the past data can be used to calculate estimates of the parameters of the joint distribution. In practice, an investor might make forecasts at regular intervals, such as intervals of one month. It would then be necessary to investigate these forecasts as well as month-to month shifts in the asset price. It should be mentioned that models concerning equilibrium pricing of risky capital assets, such as those discussed by Sharpe, Lintner, and Black,[12] may be useful in investigating the performance of investors and market-makers. Empirical work implementing these models, at least least in stock exchanges, has focused on the "residual" (i.e., the movement in the stock price not accounted for by either the stock's level of risk or the movement in the price level of the entire stock market).[13] The residual, which is the return above or below the equilibrium return for all assets with a given risk, appears to be a reasonable estimate of the market-maker's error of estimation and can easily be estimated for any security where past data are available.

Of course, to the extent that all investors possess the same amount of information, determining a distribution of forecast errors and utilizing the model of this section would seem to be a sterile exercise as far as the practical applicability of the analysis is concerned. At least with respect to the stock markets in the United States, there is evidence that few individuals or models possess enough special information to predict asset prices with any degree of accuracy.[14] However, some investors and models have been particularly successful in this regard, so it should be useful to briefly review a few such cases.

First, there is ample evidence that corporate insiders possess and act upon special information.[15] However, each individual insider may possess information on only one or a few securities. In addition, these insiders may receive information at irregular intervals, and it may be of varying quality. Thus, the process may not be stationary, making it more difficult to determine a probability distribution of forecast errors.

Though many investment counselors, mutual funds, and brokers do not appear to possess significant information,[16] evidence suggests that the Value-Line Advisory Service[17] may possess such information. As the performance of many other advisory services has not been examined, it is possible that other services also possess special information. These services investigate many different securities at regular intervals, so that past data could be used to arrive at a distribution of future forecasting errors. At present, many of the services only indicate whether an investor should buy or sell a particular security; a method of forecasting the price of the security would be needed in order to apply our model.

As a final example of successful prediction, a few models of security valuation presented in the financial literature appear to be successful. For example Jones and Litzenberger[18] have shown that stock prices do not adjust immediately when firms have sizeable increases in accounting earnings. Though Jones and Litzenberger only decide which securities should be bought, a model yielding forecasts of the prices of securities could be developed. Black and Scholes[10] have shown that their theoretical model for pricing options has information content. As their model yields a numerical value for options, an empirical implementation could follow easily.

Even if investors do possess special information, they may not use it to full advantage. A model of the type presented here is of value in formulating optimal trading rules so that investors are able to utilize their information in an attempt to increase expected return. For example, the results of Black and Scholes indicate that although proper use of their option pricing equation can yield positive returns with no transaction costs, the equation cannot be used to earn profits net of transaction costs. However, in the calculation of profits after transaction costs, all options were either bought or sold. Using our model, Black and Scholes could have bought only those options with values estimated to be above the market price by at least βT. By taking a position in fewer securities, their option model might have been profitable even with positive transaction costs.

5. Summary and Discussion

The model presented in this paper concerns trading with a market-maker, who is forced to set prices at which any investor can buy and sell an asset. When investors trade with the market-maker, speculative trading can be profitable, even if the market-maker is a better forecaster.[20] The volume of such trading, however, will depend on the ability of investors (vis-a-vis the market-maker) to forecast the price of an asset. The model provides optimal trading rules for an investor and can be used to determine the probability than an investor will be able to trade profitably in a given situation.

The model discussed here is purposely simple so that the main points are not obscured by unnecessary details. Potential extensions include the simultaneous consideration of several securities, the inclusion of nonlinear utility functions for money on the part of investors, and the development of multiperiod models that take into account the potential effect of future decisions on a current decision.[21] This list is by no means exhaustive, and it is clear that such extensions would complicate matters and would make the model more difficult to solve. However, such extensions would not be expected to change materially the essential nature of the results obtained in this paper.

Appendix I - Derviation of A's Posterior Distribution

Consider the model of Section 2 with the generalization that $Cov(u_A, u_B) = \rho\sigma_A\sigma_B$. Before P_B is known, A's prior distribution for P is a normal distribution with mean P_A and variance σ_A^2,

$$f(P|P_A) \propto \exp\left\{-(P - P_A)^2/2\sigma_A^2\right\}.$$

The likelihood function is proportional to $f(P_B|P, P_A)$, which is a normal distribution with mean $P + \rho k(P_A - P)$ and variance $k^2\sigma_A^2(1 - \rho^2)$, where $k = \sigma_B/\sigma_A$:

$$f(P_B|P, P_A) \propto \exp\{-[P_B - \{P + \rho k(P_A - P)\}]^2/2k^2\sigma_A^2(1 - \rho^2)\}.$$

Thus, the posterior distribution is

$$f(P|P_A, P_B) \propto f(P|P_A)f(P_B|P, P_A)$$

$$\propto \exp\left[-\frac{k^2(1 - \rho^2)(P - P_A)^2 + [P_B - \rho kP_A - (1 - \rho k)P]^2}{2k^2\sigma_A^2(1 - \rho^2)}\right]$$

$$\propto \exp\left[-\frac{P^2[k^2(1 - \rho^2) + (1 - \rho k)^2] - 2P[k^2(1 - \rho^2)P_A + (1 - \rho k)(P_B - \rho kP_A)]}{2k^2\sigma_A^2(1 - \rho^2)}\right].$$

Completing the square and simplifying,

$$f(P|P_A, P_B) \propto \exp\{-(P - P_A^*)^2/2\sigma_A^{*2}\},$$

where

$$P_A^* = \frac{k(k - \rho)P_A + (1 - \rho k)P_B}{k^2 - 2\rho k + 1}$$

and

$$\sigma_A^{*2} = \frac{k^2(1 - \rho^2)\sigma_A^2}{k^2 - 2\rho k + 1}.$$

Hence, A's posterior distribution for P is a normal distribution with mean P_A^* and variance σ_A^{*2}. To obtain the values of P_A^* and σ_A^{*2} given in Section 2, simply set $\rho = 0$.

When $\rho \neq 0$, the optimal trading strategy in Section 2 changes slightly It is optimal for A to buy when

$$\frac{k(k - \rho)P_A + (1 - \rho k)P_B}{k^2 - 2\rho k + 1} - P_B > T,$$

which simplifies to

$$P_A - P_B > \beta T \quad \text{if } k - \rho > 0$$

and

$$P_A - P_B < \beta T \quad \text{if } k - \rho < 0,$$

where

$$\beta = \frac{k^2 - 2\rho k + 1}{k(k - \rho)} \quad .$$

If $k - \rho > 0$, the optimal strategy for A is identical in form to that given in Section 2; the only difference is a slightly more complicated expression for β. The relationship among β, ρ, and k is as follows when $k - \rho > 0$:

$$\beta \begin{Bmatrix} > \\ = \\ < \end{Bmatrix} 1 \qquad \text{iff} \qquad \rho \begin{Bmatrix} < \\ = \\ > \end{Bmatrix} \frac{1}{k} \quad .$$

This can be seen by writing β in the form $1 + [(1 - \rho k)/(k^2 - \rho k)]$. The denominator of the second term is strictly positive (since $k > \rho$), so the sign of the second term is simply the sign of $1 - \rho k$. But it might be expected that the market-maker is a better forecaster (in the sense of having a smaller error variance) than most investors, in which case $\sigma_B < \sigma_A$, or $k < 1$. If $k < 1$, then $1/k > 1$, so $\rho < 1/k$ by definition, implying that $\beta > 1$. Even if A has a smaller error variance than the market-maker, β will still be greater than one unless the correlation is large enough so that $\rho \geq \sigma_A/\sigma_B$. This implies that when A has a smaller variance than the market-maker <u>and</u> the correlation is larger than σ_A/σ_B, A can take advantage of his superior forecasting ability and obtain a positive expected return even in some cases where $P_A - P_B < T$. In most situations, however, we would expect to see $\beta > 1$. Moreover, if $k \leq 1$, $d\beta/dk > 0$, implying that β increases as the market-maker becomes a better forecaster (in terms of the ratio of error variances) relative to the investor.

If $k - \rho < 0$, then $\beta < 0$, and the optimal buying rule is to buy if $P_A - P_B < \beta T < 0$. This seems to be a strange result; A should <u>buy</u> the security when $P_A - P_B$ is <u>negative</u> enough! Noting that $k < \rho \leq 1$, we see that the market-maker has a smaller error variance than A, and ρ is high enough that it is highly likely that P_A is on the same side of P as is P_B but that P_A is further from P. For instance, if $k < 1$ and $\rho = 1$, A <u>knows</u> that P_A is on the wrong side of P_B, so A utilizes this information to <u>buy</u> when $P_A < P_B$. In this situation, $P_A^* = (P_B - kP_A)/(1 - k)$, and $P_A^* - P_B$ and $P_A - P_B$ are of opposite sign.

Appendix II - Trading Without a Market-Maker

To contrast the model of Section 2 with mechanisms where all participants are able to trade on an equal basis, we consider the situation in which the trading mechanism is a sealed-bid procedure whereby the two investors, A and B, write down estimates, P_A and P_B, of the "true" price of the asset and trading occurs at a price mdiway between the two estimates. If $P_A > P_B$, B sells the asset to A at a price of $(P_A + P_B)/2$, and if $P_B > P_A$, A sells to B at that price. The two investors are both aware of the bivariate distribution of errors.

From the posterior distribution presented in Appendix I and the assumption that trading of one unit of the asset occurs at a price of $(P_A + P_B)/2$, A's expected return from the trade, as determined after P_A and P_B are known, is simply $P^*_A - [(P_A + P_B)/2]$ if $P_A > P_B$ (in which case B sells the asset to A) and $[(P_A + P_B)/2] - P^*_A$ if $P_A < P_B$ (in which case A sells the asset to B). The expected return for A simplifies to

$$\pi_A = \begin{cases} \dfrac{(k-1)(k+1)(P_A-P_B)}{2(k^2-2\rho k+1)} & \text{if } P_A > P_B, \\[2ex] \dfrac{(k-1)(k+1)(P_B-P_A)}{2(k^2-2\rho k+1)} & \text{if } P_A < P_B. \end{cases}$$

Obviously, $(k+1) > 0$ (since $k = \sigma_B/\sigma_A > 0$), $(P_A-P_B) > 0$ if $P_A > P_B$, and $(P_B-P_A) > 0$ if $P_A < P_B$. The term in the denominator, $k^2 - 2\rho k + 1$, is nonnegative if $\rho \leq (K^2+1)/2k$. But $(k^2+1)/2k \geq 1$, with equality holding only when $k = 1$. Thus, $k^2 - 2\rho k + 1 > 0$ except when $\rho = k = 1$, the uninteresting case in which there is no trading since $P_A \equiv P_B$. As a result, all of the factors of π_A are strictly positive except for k-1, implying that the sign of k-1 is the sign of π_A:

$$\pi_A \begin{Bmatrix} > \\ = \\ < \end{Bmatrix} 0 \qquad \text{iff} \qquad k-1 \begin{Bmatrix} > \\ = \\ < \end{Bmatrix} 0.$$

Recalling that $k = \sigma_B/\sigma_A$, we see that A's expected return from the trade is positive if $\sigma_A < \sigma_B$ and negative if $\sigma_A > \sigma_B$. Moreover, since A and B start with the same joint distribution for (u_A, u_B) and since the trading procedure is a zero-sum game, $\pi_B = -\pi_A$, where π_B represents B's expected return from the trade. Therefore, the investor with the smaller standard deviation of estimation error has a positive expected return, whereas the other investor has a negative expected return.

The results of the model of Section 2 differ from those of the model of this Appendix, where individuals with large variances of estimation error can expect to lose in trading with individuals with small variances of estimation error. The key difference in the two models, of course, involves the trading procedure. In Section 2, the market-maker is forced to set prices at which any investor can buy and sell the asset. An investor can observe these prices and revise his judgments concerning P before deciding whether to trade. Furthermore, he can determine a decision rule such as the following: buy only if $P_A - P_B$ is larger than a certain breakeven point and sell only if $P_B - P_A$ is larger than another breakeven point. In the sealed-bid trading procedure, the investor is not able to see P_B before deciding whether or not to trade. A's decision must be made before P_B is known, and the two estimates, P_A and P_B, jointly determine the price and direction of the trade.

Footnotes

*The authors are grateful for the helpful advice of Fischer Black and Roger Ibbotson. This paper is part of a manuscript presented at the Wells Fargo Conference, 1973. The authors wish to thank the participants for their many useful comments. We also gratefully acknowledge the helpful comments of an anonymous referee.

[1] See Black [3], Glass [10], Hausman [11], Jaffe [13], [14], Lorie and Neiferhoffer [20], Jones and Litzenberger [16], Rogoff [23], and Shelton [26].

[2] See Demsetz [6] and Tinic [27].

[3] To interpret the meaning of "true" price, one can imagine an instantaneous world where immediately after any trade between A and B, the asset can be sold to the government for cash, the exact amount of which is unknown to A and B at the instant of the trade. Because A is risk-neutral, this value can serve as the "true" value of the asset. If the asset can be sold for a fixed, but as yet unknown, value of cash in some future period, this value of cash, when discounted by an appropriate rate, can serve as the "true" value now.

[4] At least on the New York Stock Exchange, the specialist must stand ready to buy (sell) from any customer up to 100 shares at the bid (ask) price.

[5] Here T can be thought of as a payment to the market-maker for having to reveal P_B to traders before they make their decisions. For a model in which P_B and P_A are revealed simultaneously, see Appendix II.

[6] An excellent explanation of how the specialist performs this brokerage function can be found in Neiderhoffer and Osborne [22].

[7] In general, the specialist can only trade with the limit orders on his book when there are no incoming market orders. This and more specific rules of the NYSE concerning the above point can be found in Leffler and Farwell [18].

[8] See Scholes [24] for a discussion of this point.

[9] By comparison, some forms of gambling on sporting events are similar to the sealed-bid trading procedure of the model presented in Appendix II. For example, in pari-mutuel betting at a racetrack, the final odds, which can be thought of as a price, are determined for each race by the actions of the bettors. Furthermore, the odds are computed <u>after</u> all bets have been placed, so they are not known for certain in advance. Therefore, the results in Appendix II indicate that the bettors who are better forecasters of the outcomes of races than other bettors can expect to earn a positive return at the expense of other bettors. (It is more difficult to earn a positive return in this context than in the context of the model of Appendix II, however, for the track takes its own cut, or commission, before the total amount wagered is divided among the successful bettors). Under the assumptions of Appendix II, unsuccessful bettors would stop betting due to losses, and eventually no individuals would be willing to place bets. In the racetrack situation, however, nonlinear utility functions for money and a positive increment of utility due simply to the recreational value of being at the racetrack and to the "joy" of gambling explain in part why racetracks do not all go out of business.

[10] Empirical evidence (See Winkler [28] suggests that market point spreads determined by bookies are quite accurate relative to point spreads determined by sportswriters and naive bettors. Such bettors probably bet more often than the model of Section 2 implies that they should for speculative purposes.

[11] Winkler [29], [30].

[12] Sharpe [25], Lintner [19], and Black [2].

[13] See Fama [8] for a review of empirical work using residual analysis. Specific studies using this concept are Fama, Fisher, Jensen and Roll [9], Ball and Brown [1], Scholes [24], Mandelker [21], Jaffe [13], [14], and Ibbotson [12].

[14] See Fama [8].

[15] See Glass [10], Rogoff [23], Lorie and Neiderhoffer [20], Jaffe [13], [14], and Scholes [24].

[16] See Cragg and Malkiel [5], Elton and Gruber [7], and Jensen [9].

[17] See Shelton [26], Hausman [11], and Black [3].

[18] Jones and Litzenberger [16].

[19] Black and Scholes [4].

[20] In contrast, a poor forecaster can expect to lose to a better forecaster in the sealed-bid trading procedure described in Appendix II.

[21] See Winkler and Barry [31].

References

1. Ray Ball and Phillip Brown. "An Empirical Evaluation of Accounting Income Numbers," Journal of Accounting Research, 6 (1969), 159-178.

2. Fischer Black. "Capital Market Equilibrium with Restricted Borrowing," Journal of Business, 45 (July 1972), 444-455.

3. Fischer Black. "Yes, Virginia, There is Hope: Tests of the Value Line Ranking System," unpublished manuscript, University of Chicago, 1971.

4. Fischer Black and Myron Scholes. "The Valuation of Option Contracts and a Test of Market Efficiency," Journal of Finance, 27 (May 1972), 399-417.

5. John Cragg and Burton Malkiel. "The Consensus and Accuracy of Some Predictions of the Growth of Corporate Earnings," Journal of Finance, 23 (March 1968), 67-84.

6. Harold Demsetz. "The Cost of Transacting," Quarterly Journal of Economics, 82 (February 1968), 33-53.

7. Edwin J. Elton and Martin J. Gruber. "Earnings Estimates and the Accuracy of Expectational Data," Management Science, 18 (April 1972), 409-424.

8. Eugene F. Fama. "Efficient Capital Markets: A Review of Theory and Empirical Work," Journal of Finance, 25, (May 1970), 383-417.

9. Eugene F. Fama, Lawrence Fisher, Michael Jensen, and Richard Roll. "The Adjustment of Stock Prices to New Information," International Economic Review, 10 (February 1969), 1-21.

10. Gary A. Glass. "Extensive Insider Accumulation as an Indicator of Near Term Stock Price Performance," unpublished doctoral dissertation, Ohio State University, 1963.

11. Warren H. Hausman. A Note on 'The Value Line Contest: A Test of the Predictability of Stock Price Changes,'" Journal of Business, 42 (July 1969), 317-320.

12. Roger Ibbotson. "Price Performance of Common Stock New Issues," unpublished doctoral dissertation, University of Chicago, May 1974.

13. Jeffrey F. Jaffe. "Special Information and Insider Trading," Journal of Business, in press.

14. Jeffrey F. Jaffe. "The Effect of Regulation Changes on Insider Trading," The Bell Journal of Economics and Management Science, Vol. 5, No. 1 (Spring 1974), 93-122.

15. Jeffrey F. Jaffe. "The Value of Information for Trading Purposes," Journal of Financial Economics, in press.

16. Charles Jones and Robert Litzenberger. "Quarterly Earnings Reports and Intermediate Stock Price Trends," Journal of Finance, 25 (March 1970), 143-149.

17. Michael Jensen. "The Performance of Mutual Funds in the Period 1945-61," Journal of Finance, 23 (May 1968), 389-416.

18. George L. Leffler and Loring C. Farwell. The Stock Market. New York: Ronald Press, 1963.

19. John Lintner. "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," Review of Economics and Statistics, 47 (February 1965), 13-37.

20. James Lorie and Victor Neiderhoffer. "Predictive and Statistical Properties of Insider Trading," Journal of Law and Economics, 11 (1969), 35-51.

21. Gershon Mandelker. "The Effect of Mergers on Stock Prices," unpublished manuscript, University of Chicago, May 1973.

22. Victor Neiderhoffer and M.F.M. Osborne. "Market Making and Reversal on the Stock Exchange," Journal of the American Statistical Association, 61 (December 1966), 897-916.

23. Donald L. Rogoff. "The Forecasting Properties of Insiders' Transactions," unpublished doctoral dissertation, Michigan State University, 1964.

24. Myron S. Scholes. "A Test of the Competitive Hypothesis: The Market for New Issues and Secondary Offerings," unpublished doctoral dissertation, University of Chicago, 1969.

25. William F. Sharpe. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk," Journal of Finance, 19 (September 1964), 425-442.

26. John P. Shelton. "The Value Line Contest: A Test of the Predictability of Stock Price Changes," Journal of Business, 40 (July 1967), 251-269.

27. Seha Tinic. "The Economics of Liquidity Services," Quarterly Journal of Economics, 86 (February 1972), 79-93.

28. Robert L. Winkler. "Probabilistic Prediction: Some Experimental Results," Journal of the American Statistical Association, 66 (December 1971), 675-685.

29. Robert L. Winkler. "The Assessment of Probability Distributions for Future Security Prices," in J.L. Bicksler, ed., Methodology in Finance-Investments (lexington, Massachusetts: Lexington Books, D.C. Heath, 1972), 129-148.

30. Robert L. Winkler. "Bayesian Models for Forecasting Future Security Prices," Journal of Financial and Quantitative Analysis, 8 (1973), 387-405.

31. Robert L. Winkler and Christopher B. Barry. "A Bayesian Model for Portfolio Selection and Revision," Journal of Finance, in press.