Looking for Spot in the Presence of Futures

Trinity Economic Paper Series
Technical Paper No. 97/3 (Revised: July 9, 1997)
JEL Classification: D83, G13

Krishna Ramaswamy
The Wharton School of The
University of Pennsylvania
Philadelphia, PA 19104-6367
USA

Krishna@finance.wharton.upenn.edu

Patrick Waldron
Dept. of Economics
Trinity College
Dublin 2
Ireland

Patrick.Waldron@tcd.ie

Abstract

Customers carrying out a costly search among dealers for the best bid or offer are unable to tell whether an unfavorable quote reflects a change in market fundamentals or whether they have met a high margin dealer. The optimal search strategy in the presence of a futures market is shown to have a reservation price property, where the reservation price depends on the current futures price. In equilibrium, dealers randomize their quotes in a way that coincides with searchers' expectations, yielding a self-fulfilling expectations equilibrium. This solution is consistent with optimal dealer behavior.

Acknowledgements

We would like to thank Michael Brennan, Greg Connor, John Fingleton, Maureen O'Hara and seminar participants at the London Business School, the Stockholm School of Economics and University College, Dublin for helpful comments.

Looking for Spot in the Presence of Futures

Technical Paper No. 97/3 (Revised: July 9, 1997)
JEL Classification: D83, G13

Abstract

Customers carrying out a costly search among dealers for the best bid or offer are unable to tell whether an unfavorable quote reflects a change in market fundamentals or whether they have met a high margin dealer. The optimal search strategy in the presence of a futures market is shown to have a reservation price property, where the reservation price depends on the current futures price. In equilibrium, dealers randomize their quotes in a way that coincides with searchers' expectations, yielding a self-fulfilling expectations equilibrium. This solution is consistent with optimal dealer behavior.

Typically, transactions in a spot asset (or commodity) involve searching across dealers whereas the related futures contract is centrally and publicly traded.[1] The search for the best spot quote should follow along the lines of economic models of search (see, for example, the summary by Diamond (1987), or Lippman & McCall (1979)). Diamond (1987, p.279) writes

> Walrasian theory assumes that consumers are perfectly informed about the prices of all commodities in the economy. This assumption is central for the law of one price, that a homogeneous commodity sells at the same price in all transactions in a given market . . . In order to make a rigorous arbitrage argument there must be simultaneous purchase and sale of the same commodity at different prices net of transportation costs. If the purchase and sale are at different times there is likely to be risk for the would-be arbitrageur . . . In search theory with a known distribution of prices, there is a cost to finding any trading partner and possibly a large cost to finding one willing to trade at some particular price.

Even if one dealer's bid price (for immediate delivery of a financial security or commodity) exceeds another's ask price by an amount sufficient to cover costs, a would-be arbitrageur who has just bought at the ask still faces the risk that the bid price will fall before he can find the first dealer. In markets where the price of the underlying is subject to substantial

[1]Searching across dealers may, in practice, involve both searching for a particular quality of a commodity (or a bond with a particular coupon) and searching for a favorable price quote. The futures contract typically allows delivery of various grades of commodity. This paper considers the market for a homogeneous commodity and concentrates on the search for the best price.

volatility, the process of searching for the best spot price takes time, and it is necessary for both buyers and sellers to keep abreast of "market-wide" movements while the search is under way.

The purpose of this paper is to present a model of the search for acceptable spot prices in the presence of a futures market. We analyze the optimal strategy of customers with heterogenous beliefs, each wishing to buy or sell a homogeneous commodity or asset, and show that the typical customer follows a reservation price strategy — she stops searching as soon as she finds a price as good as or better than her reservation price. We also analyze the behavior of dealers who, recognizing that customers who differ according to their costs of search are employing an optimal search strategy, are led to quote prices in equilibrium that are consistent with customers' expectations. The model shows that even when the spot market is fragmented so that there is a cross-dealer distribution of spot quotes, buyers and sellers of the underlying, together with the dealers they face, participate in the futures market in such a way that the most active futures traders are the dealers currently offering the best spot quotes. In this sense, the futures price contains information about the spot.[2] By contrast the traditional approach is to view the (unique) spot price as an input into the computation of values for futures (and other derivative assets).

[2] A clearing-house that collects and disseminates quotes and transaction prices might serve an important informational role, and perhaps reduce the aggregate cost of search in these markets. Such prices will not be viewed by potential users as credible unless there exists a mechanism whereby one can force dealers to honor them. While dealers would be unwilling to commit to a fixed spot price in an environment where there is substantial volatility, the futures price does have all the advantages and none of the disadvantages of prices quoted in such a clearing-house arrangement.

The paper proceeds as follows. Section I describes the typical customer's problem in searching for the best price in the presence of a futures market. Section II describes the environment in which dealers operate, and the equilibrium between the pool of searching customers and the dealers they face. Section III provides a simple numerical example of this equilibrium while Section IV concludes the paper.

Section I: The Customer's Search Problem

In this section we describe the strategy of a typical customer who is intent on buying or selling an indivisible amount of a commodity (say 5,000 bushels of wheat or 100 Treasury Bonds). The spot market is made up of a large number of dealers at various locations; each dealer provides a bid and an o¤er price at which he stands ready to buy and sell the commodity respectively.

In the absence of a futures market, the customer's strategy is to engage in a search for the best price across these dealers: this problem can be posed as an optimal stopping problem. At any given time, there is a "best" pair of bid and o¤er price — the minimum of the ask quotes and the maximum of the bid quotes, both taken across dealers, which we call the innermost quotes — that is not knowable without (or even with) search. Search is costly, and it takes time; the distribution of prices and the innermost quotes may change over time as dealers learn about the state of demand and supply from the arrival of customers, and as dealers' inventory levels ‡uctuate.

We focus on the customer's search in the presence of a futures market. The futures price is the outcome of hedging demand from customers and dealers and speculative demand from investors who are not part of our model. The focus is on the spot market; the futures price is freely observable, and it follows a stochastic process which is taken to be exogenous to this model. We treat the futures market as a competitive market, with no taxes or transactions costs.

A : Customer Strategies

A customer has a choice of four strategies to complete her trade; switching between these strategies during the search process might also be optimal. The customer's possible strategies are:

The forward trading strategy: Simply buy or sell the nearby[3] futures contract, depending on whether she is a buyer or seller of the commodity, and plan to take or make delivery on its maturity date. For convenience, we assume an interest rate of zero: this ensures that the customer buys or sells the commodity at a net price equal to the current futures price, adjusted for the costs of taking or making delivery.

The hedged search strategy: Buy or sell the nearby futures contract and begin searching across dealers in the spot market. After acquiring a quote, the customer must decide whether she should accept the quote and reverse her futures position, whether she should purchase another quote, or whether she should abandon the search and take or

[3]Contracts for at least one maturity are assumed always available, but we will not need to di¤erentiate contracts by maturity date. If the nearby futures matures before a customer who has hedged has completed her search for an acceptable spot price, the futures position is costlessly rolled into the next maturity.

make delivery (essentially, switch to the forward trading strategy).

The unhedged search strategy: Take no position in the futures market, but begin to search across dealers in the spot market, being constantly watchful of the fluctuating futures price. After acquiring a quote, the customer must decide whether to accept it or to continue searching at cost and, in the latter case, whether to buy or sell futures ("lock in" the futures price), essentially switching to the hedged search strategy.

The waiting strategy: Postpone trade, taking the view that prospects of a favorable price movement are so good at the moment that neither searching nor hedging is economically justifiable. There is an implicit paradox in this strategy, which can be overcome by assuming that customers for whom the waiting strategy is initially optimal ultimately either revert to less optimistic beliefs or find their cost of search has increased. This scenario is not considered in the present paper.

The optimal choice of strategy depends on the (relative) values of three parameters:

1. the nominal search cost, $c \geq 0$, incurred in obtaining a price quote from a dealer.

   c may incorporate both the transactions costs incurred in communicating with the dealer and either the waiting costs of doing without the commodity or the storage costs incurred by holding on to it for another period.

2. the view taken by the customer on the stochastic properties of the exogenous futures price.

   To fix matters, let G_t be the futures price observed costlessly by everyone at date t. The customer is assumed to think of this as following a continuous-time stochastic process with independent increments and subjective drift rate δ per unit time, so that

$$E_{t_0}\left[G_{t_1}\right] = G_{t_0} + \delta(t_1 - t_0): \qquad (1)$$

3. the cost, D, of taking or making delivery on a futures contract.

   The forward trading strategy thus allows buyers to purchase at a net price of $G_t + D$ and sellers to receive a net price of $G_t - D$. In practice, the majority of futures positions are reversed before the maturity date, so we think of D as being a large cost relative to the nominal search cost c.

The analysis of the dealers' pricing decision in Section II below requires that there be a pool of heterogeneous customer types searching, in the sense that at least one of c and δ varies across customers. It will become apparent that D can also be allowed to vary across customers, but this is not essential to the model.[4]

[4] We assume that the model is common knowledge, and that the c, δ and D with which a customer enters the market remain fixed for the duration of her search.

B : The Customer's Search: Assumptions

There are N dealers in the spot market, and dealer i posts bid and offer quotes as deviations,[5] M_{it}^{b} and M_{it}^{a}, from the current futures price G_t, so that his bid and offer levels are $G_t - M_{it}^{b}$ and $G_t + M_{it}^{a}$ at time t. For lack of a better term, we label these deviations the quotes. Because everyone observes the futures price, the assumption that the quotations are given as bid and offered deviations from the futures price is innocuous; indeed, in many spot markets where the prices are volatile (including crude oil and many grains) this is the accepted practice.

No dealer wishes to be the source of arbitrage profits, which implies

$$M_{it}^{a} + M_{it}^{b} > 0 \quad \forall i, t. \tag{2}$$

It is assumed that a customer's primary objective is to choose a strategy to minimize the expected outlay on a purchase or maximize the expected revenue from a sale of the commodity; as a secondary objective she chooses the strategy to minimize the variance of the outlay or the revenue. These objectives do not conform to expected utility maximization; the assumed preferences are lexicographic in mean and variance. We chose this way of posing the customer's objective so that we could bring useful results from the theory of search to bear on the problem.[6] If the customer's wealth is large, then maximizing the expected value in the context of a purchase or sale of a small amount of one commodity would be consistent with locally risk-neutral behavior for small risks; and if two or more alternatives had

[5] Note that these deviations can be either positive or negative.
[6] See De Groot (1970) for a discussion of alternative objective functions.

the same expected value, then choosing the one with the least variance serves as a tie-breaking rule.[7]

We assume that one unit of time is taken up in acquiring a spot quote and that a customer acquires the ith dealer's quote pair $\left(M^b_{it}, M^a_{it}\right)$ without revealing which side of the market she is on. We also assume that she views each quote pair acquired as an independent drawing from a fixed distribution, whose marginal c.d.f.s are[8] $F_{\tilde{M}^a}$ and $F_{\tilde{M}^b}$; and that this distribution is independent of the level of the futures price G_t. The customer can keep the futures price continually and costlessly in view and she can trade in the futures market instantaneously with no transactions costs, but if she either makes or takes delivery at the maturity date of the futures contract, she pays the delivery cost D.

C : The Customer's Search: Properties

We now consider the conditions in which each of the potential strategies listed above proves optimal. The assumptions above are sufficient for us to show that the search for the best spot quotes has some special properties when there is a futures market available.

The assumption that the purchased quotes are drawn independently from an identical distribution which is unchanged as G_t fluctuates is at the heart of the results. This assumption enables us to show that the

[7] In practice, we only need this tie-breaker for customers with $\alpha = 0$ (a likely common occurrence); for all other customers, it is sufficient to assume risk-neutrality.

[8] For convenience in exposition, we have treated $\tilde{M}^a$ and $\tilde{M}^b$ as continuous random variables. In what follows it will become clear that with a discrete number of customer types there will correspond a discrete number of quotes: in that case we will interpret the quotes as discrete r.v.s and replace integrals with summations as necessary.

right to recall a previously-sampled quote and to trade with the dealer who had o¤ered it has zero value; and that the customer pursues a strategy of searching until she …nds a price below a reservation price. In dealer markets, the quotes revealed by a dealer are only good for a short duration — it would be foolish for a dealer in a volatile market to stand ready to trade at yesterday's prices — so readers may conclude that the right to recall would simply not be available in any event. However, our model is cast in terms of quotes M_{it} which are deviations from the current futures price, so it may be feasible for a dealer to stand ready to trade at a dollar price which is the sum of the current futures price and a previously displayed quote.

The assumption that the distribution from which quotes are sampled is uncorrelated with the futures price level might be questioned from another standpoint. As the futures price rises or falls, dealers might choose to alter their quotes in anticipation of increased trading interest on one side of the market. However, any individual dealer cannot a¤ect the frequency of arrival or composition of high and low search cost customers at his door by altering his privately-revealed quotes. So he has little incentive to respond to small changes in the futures price.

Consider …rst the unhedged search strategy. Suppose a buyer is engaged at time t in searching for the best o¤er price and has the right to recall and purchase at any previously o¤ered quote, so that she can buy at time t a unit of the commodity at any price $G_t + M^a_{is}; s \cdot t$. Let $M^a_{¤t}$ be the smallest (i.e. least) quote she has encountered, since she started searching, up to time t. Then her expected net gain from continuing to search from one more period is

$$[G_t + M^a_{*t}] - E_t\left[\min\left\{\tilde{G}_{t+1} + M^a_{*t}; \tilde{G}_{t+1} + \tilde{M}^a_{i;t+1}\right\}\right] - c \qquad (3)$$

where the first two terms give the expected reduction in price from continuation and the final term is the cost of acquiring another quote. Expression (3) can be rewritten as

$$E_t\left[\max\left\{0; M^a_{*t} - \tilde{M}^a_{i;t+1}\right\}\right] - (c + \delta): \qquad (4)$$

The first term in expression (4) represents the expected reduction in the quote from continuation, and the second term represents the net unhedged cost of search, $c + \delta$, which now includes the customer's forecast of the drift in the futures price.

Now consider the hedged search strategy. The hedged buyer receives a marking to market payment of $\tilde{G}_{t+1} - G_t$ at time $t + 1$, and thus her expected net gain from continuing to search for one more period is given not by expression (4) but by:

$$E_t\left[\max\left\{0; M^a_{*t} - \tilde{M}^a_{i;t+1}\right\}\right] - c: \qquad (5)$$

Thus c itself can be viewed as representing the net hedged cost of search.

Comparing expressions (4) and (5), it is clear that the solution to the customer's problem depends only on the effective cost of search, c', which is given by

$$c' = \begin{cases} c + \delta & \text{for an unhedged buyer} \\ c & \text{for a hedged buyer.} \end{cases}$$

As the search progresses, the best quote to date, M^a_{*t}, and hence the first term in the expressions (4) and (5) for the expected gain to continuing the search, will not increase. For either hedged or unhedged search, there will come a point where the expected net gain to continuation turns negative, provided that the effective cost of search is a non-negative constant — which is the case as long as the subjective drift rate $\delta \geq -c$ for the unhedged buyer. If the net unhedged cost of search, $c + \delta$, is negative then the reservation quote policy spelled out here will not hold, and the waiting strategy will be optimal until such time as either c or δ is revised and $c + \delta$ becomes positive.

Defining

$$W_{\tilde{M}^a}(\theta) \equiv \int_{-\infty}^{\theta} (\theta - m)\, dF_{\tilde{M}^a}(m) = E\left[\max\left\{0; \theta - \tilde{M}^a\right\}\right]; \qquad (6)$$

we can say that the reservation quote for a buyer whose effective cost of search is c' is the unique[9] value $m_R^a(c')$ which satisfies

$$W_{\tilde{M}^a}(m_R^a(c')) - c' = 0. \qquad (7)$$

The optimal strategy for the buyer is to accept the first ask quote that fails to exceed her reservation quote $m_R^a(c')$. From equations (6) and (7), it can be seen that the net search cost c' is analogous to the price of a European put option on the quote with one unit of time remaining and a strike price equal to $m_R^a(c')$. It follows from this analogy that the larger is the variance of the distribution of quotes across dealers, the lower is the reservation quote.[10] The right to recall is a technical device in the

[9] Uniqueness follows from the fact that $W_{\tilde{M}^a}$ is an increasing function (see De Groot (1970, p.246)).

[10] To see this, note that in order to keep the price of the put, c', the same while

exposition and it is never used.

Following a similar argument, the optimal strategy for a seller is to accept the first bid quote which is better than her reservation quote $m_R^b(c')$, which satisfies

$$W_{\tilde{M}^b}(m_R^b(c')) = c'; \tag{8}$$

where the effective cost of search, c', equals $c - \delta$ for an unhedged seller and c for a hedged seller.

Let the random variable τ denote the duration of search for a buyer following the reservation quote strategy, and let

$$\pi^a(c') \equiv \Pr\left[\tilde{M}_t^a \leq m_R^a(c')\right] \tag{9}$$

denote the probability that an ask quote is acceptable to a customer with net search cost c'. Because the quotes are assumed to be i.i.d., $\pi^a(c')$ is time-invariant. The expected duration of search is given by[11]

$$E[\tau] = \frac{1}{\pi^a(c')}. \tag{10}$$

---

increasing the variance of the "asset", we must decrease the strike price.

[11] To see this note that

$$E[\tau] = \pi^a(c') + 2(1 - \pi^a(c'))\pi^a(c') + 3(1 - \pi^a(c'))^2\pi^a(c') + \cdots \equiv \pi^a(c')S$$

where

$$S \equiv \left\{1 + 2(1 - \pi^a(c')) + 3(1 - \pi^a(c'))^2 + \cdots\right\} = \frac{1}{\pi^a(c')} + (1 - \pi^a(c'))S.$$

It follows that

$$S = \left(\frac{1}{\pi^a(c')}\right)^2$$

so that

$$E[\tau] = \frac{1}{\pi^a(c')}.$$

The customer's expected total outlay in acquiring the commodity, assuming she begins to search at time t and follows the reservation quote strategy, is

$$E_t\left[c'\tau + G_{t+\tau} + m_R^a(c')\right] - E_t\left[m_R^a(c') - \tilde{M}^a_{i;t+\tau} \,\middle|\, \tilde{M}^a_{i;t+\tau} \le m_R^a(c')\right]. \qquad (11)$$

The first term in expression (11) is the expected total search cost plus the reservation price; and the second term is the expected savings from hitting a quote strictly less than the reservation quote. Note that the second term will not depend on t or on τ. Because we have delinked the futures price process from the search process, we can rewrite expression (11) as

$$c'E[\tau] + G_t + m_R^a(c') - \frac{\int_{-1}^{m_R^a(c')}(m_R^a(c') - m)dF_{\tilde{M}^a}(m)}{\Pr\left[\tilde{M}^a_t \le m_R^a(c')\right]}. \qquad (12)$$

Using equations (7) and (9) to simplify the numerator and denominator, respectively, of the final term in expression (12), the expected total outlay is

$$c'E[\tau] + G_t + m_R^a(c') - \frac{c'}{\pi^a(c')} = G_t + m_R^a(c'); \qquad (13)$$

substituting for expected duration of search from equation (10). This says that, before she begins to search, the customer's expected total outlay is equal to the current reservation price given by $G_t + m_R^a(c')$.

Thus, the higher the customer's effective cost of search c', the higher her reservation quote[12] and, hence, the higher her expected total outlay and the shorter her expected search duration. Given our assumptions about customer preferences, her primary objective leads her to min-

[12]Since $W_{\tilde{M}^a}$ and hence its inverse m_R^a are increasing functions.

imize her expected total outlay by minimizing her effective cost of search. If $\delta < 0$, then a buyer will not hedge, ensuring an effective cost of search of $c + \delta < c$; and if $\delta > 0$, then a buyer will hedge, ensuring an effective cost of search of $c < c + \delta$. Hence we can define the buyer's true cost of search as

$$c^* \equiv \begin{cases} c & \text{if } \delta \geq 0 \\ c + \delta & \text{if } \delta \leq 0 \end{cases} \tag{14}$$

$$= c + \min\{\delta, 0\}. \tag{15}$$

This says that a buyer who expects prices to fall will not hedge but a buyer who expects prices to rise will. Similarly, the seller's true cost of search is

$$c^* \equiv \begin{cases} c & \text{if } \delta \leq 0 \\ c - \delta & \text{if } \delta \geq 0 \end{cases} \tag{16}$$

$$= c - \max\{\delta, 0\}. \tag{17}$$

If a buyer's beliefs, say, are neutral ($\delta = 0$), then her expected total outlay from hedging or not hedging is the same. We need our tie-break rule: she will choose the strategy for which the random total outlay has the lower variance. This total outlay, if she arrives and begins to search at time t without hedging, is:

$$c\tau + G_{t+\tau} + \tilde{M}^a_{i;t+\tau} = c\tau + G_t + \tilde{\varepsilon}_{t;\tau} + \tilde{M}^a_{i;t+\tau}; \tag{18}$$

where the term $\tilde{\varepsilon}_{t;\tau} \equiv G_{t+\tau} - G_t$ is the change in the futures price from t to $t + \tau$. On the other hand, hedging fully at t by buying a futures contract, rolling over at intervening maturity dates, and reversing the hedge at $t + \tau$ yields an income of $-\tilde{\varepsilon}_{t;\tau}$ from the marking-to-market flows of the futures position over the lifetime of the search. This changes the random total outlay of the hedged searcher to

$$c\tau + G_t + \bar{M}^{*}_{i,t+\tau}. \tag{19}$$

The variance of the total outlay for the unhedged strategy is

$$\begin{aligned} &\mathrm{Var}_t\left[c\tau + G_t + \tilde{\varepsilon}_{t,\tau} + \bar{M}^{*}_{i,t+\tau}\right] \\ &= \mathrm{Var}_t\left[c\tau + \tilde{\varepsilon}_{t,\tau} + \bar{M}^{*}_{i,t+\tau}\right] \\ &= \mathrm{Var}_t\left[c\tau + \bar{M}^{*}_{i,t+\tau}\right] + 2c\mathrm{Cov}_t\left[\tau, \tilde{\varepsilon}_{t,\tau}\right] + 2\mathrm{Cov}_t\left[\bar{M}^{*}_{i,t+\tau}, \tilde{\varepsilon}_{t,\tau}\right] + \mathrm{Var}_t\left[\tilde{\varepsilon}_{t,\tau}\right]. \end{aligned} \tag{20}$$

The first term on the right hand side of equation (20) above is the variance of the hedged outlay. A sufficient condition for the hedged search strategy to have lower variance is

$$\mathrm{Cov}\left[\tau, \tilde{\varepsilon}_{t,\tau}\right] = \mathrm{Cov}\left[\bar{M}^{*}_{i,t+\tau}, \tilde{\varepsilon}_{t,\tau}\right] = 0, \tag{21}$$

which says that the innovations in the futures price are uncorrelated with the search duration and the randomly sampled quotes. We can then conclude that a buyer pursues a hedged search strategy when $\delta = 0$, for the outlay using the hedged search strategy has the same mean as, and a lower variance than, using the unhedged search strategy. The corresponding analysis for a seller is straightforward and is omitted.

We conclude our analysis of the properties of the customer's search by considering the forward trading strategy. Suppose that at time t a customer whose true cost of search is K^{*} is indifferent between searching optimally and buying a futures contract and accepting delivery at its maturity date, T, paying the delivery cost D. Her outlay on the forward trading strategy is the sum of the marking-to-market flows plus the final settlement price G_T plus the delivery charge:

$$-(G_T - G_t) + G_T + D = G_t + D;$$

while her expected outlay under the optimal search strategy is

$$G_t + m^a(K^a):$$

Thus, the buyer's critical search cost K^a is the solution to

$$D = m^a_R(K^a):$$

The seller's critical search cost K^b can be defined analogously.

The preceding analysis leads to the following propositions (see Fig. 1):

Proposition 1 Customers for whom search is optimal[13] adopt the hedged search strategy if they have pessimistic or neutral beliefs ($\delta \leq 0$ for sellers, $\delta \geq 0$ for buyers) and the unhedged search strategy if they have optimistic beliefs.

$\square$

Search costs combine a nominal cost and the subjective drift rate of the customer. The following two propositions say that low costs and a favorable drift rate will lead to unhedged search or, at the extreme, a waiting strategy; high costs and an unfavorable drift rate will lead to

[13] In practice, these should be the vast majority; we explain why when considering the dealers' pricing decision below.

hedged search or, at the extreme, a forward trading strategy.

Proposition 2 The most optimistic customers, in particular those with the lowest nominal costs of search relative to their beliefs ($c < \delta$ for sellers, $c < -\delta$ for buyers), adopt a waiting strategy until such time as they revert to less optimistic beliefs.

$\square$

Proposition 3 The highest nominal search cost customers, in particular those with pessimistic beliefs relative to their costs of search ($c + \max\{\delta; 0\} > K^b$ for sellers, $c + \max\{-\delta; 0\} > K^a$ for buyers), adopt a forward trading strategy and do not search on the spot market.

$\square$

Proposition 4 Every customer's optimal strategy is determined by the values of c, δ and D and hence is known when she enters the market (in other words, switching between strategies never occurs).

$\square$

Fig. 1 summarizes the above results; it highlights the fact that the optimal choice of strategy separates the set of possible search types and beliefs into four simply delineated regions. It is important to note that the indicated linear boundaries between these regions apply only under our specific preference assumptions.

We can speculate how these boundaries might vary for risk-averse, expected utility maximizing customers. For example, a risk-averse customer might prefer a hedged search strategy to an unhedged search strategy even when δ is favorable. At another extreme, she might opt to take delivery with a futures contract rather than search even when

$D < m_R^a(c^*)$: Given that delivery is relatively infrequent in the real world, this may not be a fruitful path to pursue.

Section II: The Dealer's Pricing Problem

We proceed to tackle the dealer's problem. Each dealer must set a pair of quotes, taking as given a distribution of customer types among (potential) searchers. Dealers passively respond to the order flow for the commodity rather than seek out customers, although they might trade futures contracts. In order to keep the problem simple, we allow the customers to differ only along their true search costs, c^*. From the results of the previous section, a customer with a higher true search cost sets a higher reservation quote for her search; this means that dealers could equally well view the customer pool as belonging to different reservation quote types.

In order to describe the dealer's problem and the ensuing equilibrium it is best to describe the process we have in mind. Customer arrival is random. In the aggregate, we assume that customers are generated by a Poisson process with parameter λ. Once a customer is generated, his or her type is determined by the outcome of two further random experiments, independently of each other and of the underlying Poisson process.

The first of these experiments determines the new customer's true search cost, c^*. There are H possible customer types with (true) search costs $0 = c_1 < c_2 < \dots < c_H$. Search cost is determined by the outcome of a multinomial experiment, and equals c_h with probability ρ_h. It will become clear in the next subsection why we need the possibility of searchers who have zero true search cost, or who are almost optimistic enough to adopt the waiting strategy. A second, Bernoulli, experiment determines which side of the market the new customer is on. With probability $1/2$ she is a buyer and with probability $1/2$ again she is a seller.[14]

We say that a customer is of type hr ($h = 1, \ldots, H$; $r = a, b$) if her true search cost is c_h and she is searching for a type r quotation ($r = b$ for a seller seeking a bid quotation or $r = a$ for a buyer seeking an ask quotation). We label the reservation quote at which type hr customer stops search as m_h^r.

A : Clientele and Arbitrage Arguments

A dealer posts a bid and an ask price at which he is prepared to deal a single unit of the commodity. We now show that the quotes posted by dealers are always chosen from the set of reservation quotes of the different possible customer types. A dealer knows that (for $h > 1$) m_h^r is acceptable to any customer who would accept a quote between m_{h-1}^r and m_h^r, and so setting any quote between these two values leads to a loss of revenue without leading to any reduction in the probability that the quote will be accepted, and therefore is sub-optimal. Similarly, no dealer will post a quote lower than m_1^r, the reservation quote of the lowest cost customer type, since m_1^r is accepted with probability one (i.e. by all customer types).

The term innermost quotes was used to describe the highest bid quote and the lowest ask quote which are (potentially) available in the

[14] The rationale behind this assumption is discussed in Subsection II.C below.

spot market, m_1^b and m_1^a respectively. Similarly, the lowest acceptable bid and highest acceptable ask quotes (m_H^b and m_H^a) are called the outermost quotes.[15]

Any equilibrium in which the smallest possible true search cost is a known positive constant, say $c_1 > 0$, and in which customers follow reservation price strategies is unstable. To see this, consider a dealer whose optimal strategy is to quote the reservation price, m_1^r, of the minimum cost customer. This dealer knows that any price below $m_1^r + c_1$ is equally acceptable (for the customer is better off paying this amount now than paying a search cost of c_1 and a price of at least m_1^r next period): an unending cycle of raising price quotes ensues and the equilibrium "unravels." If $c_1 = 0$, or equivalently if there are customers whose optimism leaves them indifferent between searching and waiting, then the equilibrium does not unravel.

We can also show that, if there exists a customer with zero nominal search cost, then the innermost ask exceeds the innermost bid at every date, even though these quotes are given by different dealers. Otherwise, such a customer holding one unit of the commodity could earn riskless arbitrage profits: by waiting long enough she will eventually encounter and be able to sell at the innermost bid, at no cost. If, simultaneously with selling spot, the customer buys a futures contract before embarking on a new search as a buyer, she avoids the risk of an adverse movement in the futures price process and by waiting long enough will encounter the

[15] A dealer who is not prepared to deal on one side of the market or the other can find an extreme price at which he knows he will do no business, namely any price outside the outermost quote on that side of the market.

innermost ask, again at no cost. By buying at this price and closing out the futures position, the customer is guaranteed a profit of the difference between the innermost quotes. Formally, if the customer sells at t_1, and buys back at t_2, she is assured of an income of

$$\left(G_{t_1} - m_1^b\right) + \left(G_{t_2} - G_{t_1}\right) - \left(G_{t_2} + m_1^a\right) = -\left(m_1^a + m_1^b\right). \qquad (22)$$

The first term here represents the sale price, the second term represents marking-to-market receipts, and the third term represents the price at which the commodity is repurchased. If arbitrage is not permitted, the right hand side of equation (22) must be non-positive, or equivalently

$$-m_1^b \le m_1^a; \qquad (23)$$

or

$$G_t - m_1^b \le G_t + m_1^a \quad \forall t; \qquad (24)$$

which says the innermost bid cannot exceed the innermost ask. A similar argument could be used for a customer with zero nominal search cost, but without an initial endowment of the commodity, who seeks to buy first and then to resell, provided that she can store one unit free of storage costs. Inequality (24) is stronger than inequality (2) which applies to the quotes of an individual dealer.

B : A Dealer's Operating Environment

The dealer's pricing decision is complex and potentially influenced by many variables. These include principally the number of units of inventory, but also the cost of storage, cash reserves, and perhaps such variables as bankruptcy costs.[16] To simplify the problem here, we assume that dealers can borrow, using inventory as collateral, but that lenders require dealers to hedge fully to avoid any risk of a fall in prices. Note that a fully hedged dealer can never become financially embarrassed as a direct result of a sudden change in the futures price. This is because an increase of ΔG_t in the futures price raises the credit limit of a dealer with q units of inventory by $q \times \Delta G_t$. However, the latter quantity is exactly the marking to market payment which the dealer (if fully hedged) is required to make on his short futures position as a result of the price rise. Similarly, a futures price decline earns a perfectly hedged dealer a marking to market payment which exactly offsets the reduction in his credit limit.[17]

Dealing is risky and a dealer who hedges perfectly is still subject to bankruptcy risk since he may spend a prolonged period without consummating a sale, all the while paying storage costs for inventory. Lenders eventually force such a dealer to liquidate inventory by delivering on futures contracts. We avoid the complicating dependency of the value of inventory and the optimal pricing decision on the time to the next futures maturity date by assuming that dealers can deliver against futures

[16] See, for example, Amihud & Mendelson (1980), Ho & Stoll (1981) and Madhavan & Smidt (1993).

[17] An over-hedged dealer may run into difficulties after a sudden increase in the futures price, when he may be required to make a marking to market payment greater than the increase in his credit limit. Likewise, a partially hedged dealer may have credit problems after a sudden fall in the futures price, when the marking to market payment received may be insufficient to match the reduction in his credit limit.

contracts at any time at an appropriate delivery cost.[18]

C : Equilibrium: Customer Flows and Stocks

We show now that there are sensible steady-state outcomes for the composition of customers in the arriving ‡ow and in the searching stock or pool; and that dealers facing the pool of searching customer types adopt a pricing strategy that conforms to the expectations of the arriving customers. The pool of customer types currently searching displays di¤erent (expected) proportions in general from the in‡ows and out‡ows.

As noted above, equilibrium requires that half of all customers in the in‡ow are buyers and half are sellers. This arises because in a steady state there must be zero drift in the level of dealer inventories and the rate of spot purchases must exactly match the rate of spot sales. In a general equilibrium, any imbalances between the arrival rates of spot buyers and sellers would also a¤ect the futures price: the determination of both the arrival rates and the futures price are exogenous to our model. The symmetry between buyers and sellers in the customer in‡ow need not carry through to the steady-state pool of searchers.[19] In order to characterize the dealers' pricing decision and the equilibrium, we need to calculate the proportions of buyers and sellers in the steady-state pool.

---

[18] Cootner (1967) has argued that even if the cost of delivering into the futures contract is lower for dealers than it is for searchers — which is likely to be the case — a dealer will be better o¤ waiting for an importunate customer than simply delivering into the futures contract.

[19] In certain circumstances — for example if buyers and sellers face di¤erent delivery costs or if dealers' storage costs depend on inventory levels — buyers may face longer searches than sellers, ceteris paribus, and so it may be the case that they account for more than half of the pool of searchers.

To derive them, we first need parameters arising from the dealer's choice of optimal pricing policy. For each customer type hr, let π_h^r denote the probability that, from the point of view of a customer of type h seeking a type r quotation and following a reservation quote strategy, a particular price quotation is acceptable:

$$\pi_h^r \equiv \Pr\left[\tilde{M}_t^r \leq m_R^r(c_h)\right]. \tag{25}$$

The expected search duration for the h-type customer is the reciprocal of the corresponding[20] probability, $\frac{1}{\pi_h^r}$; r = a, b.

We now collect some results that relate to the stocks and flows of customer types, relegating all proofs but one to the Appendix.

Proposition 5 The expected value of the number, $\tilde{N}_h^r$, of type hr customers in the current pool of searchers is

$$E\left[\tilde{N}_h^r\right] = \frac{\lambda\rho_h}{2\pi_h^r}, \quad r = a, b. \tag{26}$$

This says that, ceteris paribus, high cost, high reservation price customers are a smaller proportion of the pool: they find acceptable quotes sooner. Adding the relevant expectations gives corresponding results for the expectations of the total numbers of buyers and sellers ($\tilde{N}^a$ and $\tilde{N}^b$) and total number of customers of type h ($\tilde{N}_h$) in the current pool of searchers:

$$E\left[\tilde{N}^a\right] = \frac{\lambda}{2}\sum_{h=1}^{H}\left[\frac{\rho_h}{\pi_h^a}\right] \tag{27}$$

$$E\left[\tilde{N}^b\right] = \frac{\lambda}{2}\sum_{h=1}^{H}\left[\frac{\rho_h}{\pi_h^b}\right] \tag{28}$$

[20] As shown in footnote 11.

$$E\left[\tilde{N}_h\right] = \frac{\lambda\rho_h}{2}\left[\frac{1}{\pi_h^a} + \frac{1}{\pi_h^b}\right] \tag{29}$$

Proposition 6 The flow of customers of type hr at a particular dealer's door is Poisson with parameter

$$\frac{\lambda\rho_h}{2n\pi_h^r},$$

where n denotes the total number of dealers from whom the customer may choose.

□

To calculate the parameter of the Poisson process generating the aggregate flow of customer arrivals at a particular dealer's door, the parameters in Proposition 6 cannot simply be summed over customer types, but must be weighted by the conditional proportions of the various customer types among the pool of searchers. These proportions can be easily deduced from Proposition 5. The following proposition is needed in the construction of a sensible equilibrium.

Proposition 7 If $\tilde{\rho}_h^r$ is the proportion of type hr customers in the current pool of searchers, then $\frac{\rho_h}{\tilde{\rho}_h^r}$ is increasing in h.

□

We have defined ρ_h as the probability that an h-type customer is generated in the incoming flow. However, a customer selected at random from among the current pool of searchers is relatively more likely to have a low search cost, since the expected duration of search, $\frac{1}{\pi_h^r}$, is, by equation (25), shorter for higher cost customers.[21]

[21] For customer of type H, the search duration is 1.

Proof (Proposition 7) Of every N buyers or sellers generated, $\rho_h N$ are expected to be of type hr. These customers are expected to seek $\frac{1}{\pi^r_h}$ price quotes each, giving a total expected number of $\sum_{h=1}^{H} \frac{\rho_h}{\pi^r_h} N$ customer-dealer encounters for every N customers generated. Of these encounters, the fraction

$$\frac{\frac{\rho_{h'}}{\pi^r_{h'}} N}{\sum_{h=1}^{H} \frac{\rho_h}{\pi^r_h} N} = \frac{\frac{\rho_{h'}}{\pi^r_{h'}}}{\sum_{h=1}^{H} \frac{\rho_h}{\pi^r_h}} \tag{30}$$

is accounted for by customers of type h'. However, if we think of this cross-section of searches as those taking place at a particular point in time, rather than those conducted by a particular cohort of entrants, it can be seen that the proportion given in equation (30) is just $\rho^r_{h'}$. It follows that

$$\frac{\rho_{h'}}{\rho^r_{h'}} = \pi^r_{h'} \sum_{h=1}^{H} \frac{\rho_h}{\pi^r_h} \tag{31}$$

and since π^r_h is increasing in h, $\frac{\rho_h}{\rho^r_h}$ is increasing in h, from which it follows that low cost searchers constitute a larger proportion of the stock than of the flow.

Q.E.D.

D : Dealer Profit Conditions

The distribution of quotes can be viewed in two ways. In the steady state, the unconditional probabilities of observing each possible quote can be interpreted as the proportion of a dealer's lifetime for which he posts the relevant quotes. Similarly, from the customer's perspective, those same probabilities describe the cross-sectional distribution of price quotes across dealers at any point in time at which she seeks a price quote, or the expected proportion of dealers in each state.

Our strategy for defining a set of profit-maximizing dealer conditions together with customer search-optimality conditions is as follows. We assume that the fully hedged dealers face a cost structure that is independent of the scale of their inventory, so that dealers set prices so as to maximize expected revenue. Because they know that customers belong to one of H reservation quote types, their posted prices correspond to those reservation prices, as discussed earlier. A given dealer will switch his quote from the reservation quote of the h-th customer type to the j-th type if that increases his expected revenue. In equilibrium, the expected revenue at all quotes will equalize. The choice of quotes by dealers is random, but from a discrete distribution which assigns to each reservation price m_h^r the same cumulative probability, π_h^r, as that on which searchers base their optimal stopping condition (7).

To fix matters, let z_j^r be the fraction of dealers quoting the reservation quote m_j^r. Because smaller quotes are acceptable to a type hr customer, her probability of finding an acceptable quote is

$$\pi_h^r = \sum_{j=1}^{h} z_j^r. \tag{32}$$

Now we restate her optimality condition or stopping rule (equation (7)), in terms of the discrete probabilities applicable to H customer types:[22]

$$\sum_{j=1}^{h} \left[m_h^r - m_j^r \right] z_j^r = c_h; \quad h = 2, 3, \ldots, H. \tag{33}$$

[22] Note that the condition to prevent unravelling, $c_1 = 0$, can now be seen to be just the optimality condition for $h = 1$.

Defining

$$\Delta m_h^r \equiv m_{h+1}^r - m_h^r \quad \text{and} \quad \Delta c_h \equiv c_{h+1} - c_h$$

we can write a set of difference equations[23] which corresponds to a relationship between the quotes $m_h^r; h = 1; \ldots; H$ and the probabilities, given the customers' cost parameters:

$$c_1 = 0 \tag{34}$$

$$\Delta m_h^r \times \pi_h^r = \Delta c_h; \quad h = 1; \ldots; H - 1. \tag{35}$$

The expected proportion of customers of type hr that each dealer expects to see at his door is just ρ_h^r. If a dealer quotes m_j^r, the reservation quote of customer type jr, then the proportion of customers that accept his quote is

$$\mu_j^r = \sum_{h=j}^{H} \rho_h^r; \tag{36}$$

because all customer types with higher search costs are also happy to trade with him. Hence the expected revenue to a dealer quoting m_j^r is

$$\mu_j^r \times m_j^r$$

and if the dealer is to find no gains to switching his quote, it must be the case that his expected revenue is a constant

$$\mu_j^r \times m_j^r = \gamma^r; \qquad \text{independently of } j. \tag{37}$$

[23] By differencing the system of equations (33), or the equivalent system

$$m_h^r \pi_h^r - \sum_{j=1}^{h} m_j^r z_j^r = c_h.$$

This relation is an optimality condition for dealers. Note that μ_j^r depends on the probabilities z_j^r; so the relation connects these probabilities and the dealers' quotes and is in effect a supply relationship, while the customer's optimality condition (33), which also connects z_j^r and m_j^r, is in effect a demand relationship. The values of γ^a and γ^b (which need not be the same) must be sufficient to cover the dealer's cost, and we discuss this further below.

The values of z_j^r and m_j^r are to be found simultaneously, leading to a self-fulfilling expectations equilibrium; computationally this requires that the customers' assessment of the probabilities of finding acceptable quotes from dealers corresponds to the fraction generated by revenue-maximizing dealers.

The values of γ^a and γ^b, when multiplied by the appropriate expected rates of customer arrival at a dealer's door, determine the rate at which the dealer earns revenue. However, it is clear from the model that increasing γ, ceteris paribus, just leads to a proportional increase in all the equilibrium quotes. Economic arguments must provide both upper and lower bounds on γ.

E : Properties of Equilibrium

A customer who buys from a dealer at the quoted ask must not be able to earn arbitrage profits by simultaneously selling and delivering into the futures contract, so in equilibrium $G_t + m_h^a > G_t - D \; \forall h$. Similarly, a customer taking delivery from an expiring futures cannot profit by selling to a dealer at his quoted bid, or $G_t - m_h^b < G_t + D \; \forall h$. Tighter bounds on spot quotes than these can be found by considering the customer who is indifferent between the searching and forward trading strategies. If the maximum quote, m^r_H, is raised above the futures market delivery cost, D, then high-cost customers are driven away from the hedged search strategy to the forward trading strategy and the equilibrium collapses. A new equilibrium will result for the remaining, lower cost, customers, but there is no guarantee that the new m^r_H will not also exceed D, and so on.[24] Thus, in equilibrium, $m^r_h < D \; \forall h, r$.

While on the one hand equilibrium is not sustainable if equilibrium margins are too high relative to delivery costs to make search viable, on the other hand it is also unsustainable if equilibrium margins are too low to allow dealers to cover their costs. In our simple model, we can assume that dealers face fixed costs only, say at a rate k per unit time. If the expected equilibrium revenue rate is less than k, then dealers eventually go bankrupt. On the other hand, if the expected equilibrium revenue rate is greater than k, then other (risk-neutral) dealers will enter the market, and either bid up k or bid down γ^a or γ^b until the expected profit rate equals zero. Any such increase in the number of dealers will not affect the rest of the model, given the assumption of sequential search by customers. However, an increase in the customer arrival rate, given fixed dealing costs, will raise the profitability of dealing and lead to entry.

Finally, consider the interaction between the spot and futures

[24] While we do not model this process formally, we can see this occurring in the example in the next section. Even if raising γ^r allowed another stable equilibrium to emerge, the fall in the overall arrival rate, due to high cost customers not entering the search process, might cause a fall in the revenue rate.

markets. Our assumption that bank collateral requirements force dealers to hedge and the result that pessimistic customers hedge to reduce risk are sufficient to guarantee that at least one side of every spot deal is immediately undone on the futures market. The largest share of spot market deals is done by the dealers posting the innermost quotes, m_1^r. This follows from the fact that dealers' market shares are proportional to μ_i^r, the proportion of searchers in the search pool prepared to accept m_i^r. In other words, the innermost dealers do the most spot deals and the most futures deals. Indirectly, they communicate the fundamental information contained in the customer generation process, which they collectively observe, through to the futures price.

Section III: An Example

Here we construct a numerical example of the equilibrium described above. We consider only customers who seek to buy, so we drop the superscripts on the various variables; a parallel situation applies to sellers.

The theoretical model is described by the probability distribution of search costs in the incoming flow of customers, in other words by the search cost levels $c_1; \ldots; c_H$ and the associated probabilities $\rho_1; \ldots; \rho_H$. To fix ideas, suppose there are ten types of customers ($H = 10$), each represented by a row of Table 1. For computational simplicity, we work backwards in this example to find the search costs which would result in a given quote distribution. In the table, the objective is to find the search costs which result in one-tenth of the dealers quoting each of the ten reservation quotes ($z_h = 0.1$ for $h = 1, \ldots, 10$, the second column), given that one-tenth of arriving customers are of each type ($\rho_h = 0.1$ for $h = 1, \ldots, 10$). This objective can also be expressed in terms of the probability that the dealer making the hth quote will satisfy the next customer to arrive:

$$\pi_h = \sum_{k=1}^{h} z_k,$$

(in the third column of Table 1), or equivalently, in terms of the expected search durations for each customer type (Column 4):

$$\frac{1}{\pi_h}.$$

From the proportions of each customer type in the inflow (in the fifth column) and the expected search durations, we can (using equation (30)). calculate the proportions of each customer type in the pool of searchers,

$$\hat{\rho}_1, \ldots, \hat{\rho}_{10},$$

which are given in the seventh column of the table. The probability that the quote m_h is accepted is just

$$\mu_h = \sum_{k=h}^{H} \hat{\rho}_k,$$

as shown in the eighth column.

We now arbitrarily set $\gamma = 1$ (essentially just choosing the numeraire in our model in terms of which storage and delivery costs can be expressed) and the equilibrium quotes are

$$m_h \equiv \frac{\gamma}{\theta_h};$$

as shown in the ninth column of the table. Now we are in a position, using the system of difference equations (35), to find (in the final column of the table) what exogenous search costs,

$$c_1 = 0; c_2; \ldots; c_{10};$$

would have resulted in the chosen distribution of quotes across dealers, $z_1; \ldots; z_{10}$: In order that the highest cost customer searches rather than engage in a forward strategy, we must assume that $D > 29{:}29 = m_H^a$: If the delivery cost falls then the equilibrium must be recomputed and its stability is not assured, as discussed previously.

Section IV: Conclusion

In this paper, we have considered the relationship between dealer spot markets and open outcry futures markets, in particular the effects of costly search in the spot market in the presence of an evolving and observable futures market. We have shown that those most likely to trade on the futures market are the dealers who are quoting the most competitive spot prices and the customers who expect the most adverse price movements or who are the most risk averse. The optimal use of the futures market reduces the effective cost of searching for the best price.

The model also provides insights into the nature of competition between different channels of trade. On the one hand, wide spreads on the spot market will make not only trading, but even delivering on, futures contracts attractive. Narrow spot market spreads, on the other hand, may be insufficient to cover a dealer's costs of financing. Equilibrium is reached when there is an appropriate balance between these two opposing forces.

In the traditional literature on price search, guaranteeing the uniqueness and stability of equilibria has always been problematic. In particular, a new equilibrium can usually be found in these models by assuming that dealers double all margins and that searchers revise their beliefs accordingly. In our model, the availability of an alternative trading channel, namely the futures market, rules out this possibility and applies closure, even though in equilibrium, as in practice, delivery on futures contracts is the exception rather than the rule.

Our model is of a partial equilibrium nature and treats as exogenous the many factors that influence the futures price. We leave to future work the study of information effects, and of the processes by which buyers and sellers learn about the cross-dealer distribution of quotes and by which dealers learn about the pool of customers and their willingness to trade at various prices.

Appendix

Proof (Proposition 5) Let T denote the current time.

$$
\begin{aligned}
E\left[\tilde{N}_h^r\right] &= \sum_{k=1}^{\infty}\int_{T-k}^{T-k+1} \Pr[\text{type } hr \text{ customer arrived in } [t, t+dt)] \\
&\quad \times \Pr[\text{customer still searching} \mid \text{arrived then}] \\
&= \sum_{k=1}^{\infty}\int_{T-k}^{T-k+1} \frac{\lambda\rho_h}{2} dt\,(1-\pi_h^r)^{k-1} \\
&= \frac{\lambda\rho_h}{2}\sum_{k=1}^{\infty}(1-\pi_h^r)^{k-1} \\
&= \frac{\lambda\rho_h}{2\pi_h^r}. \qquad (38)
\end{aligned}
$$

Q.E.D.

Proof (Proposition 6)

$$
\begin{aligned}
&\Pr[\text{customer arrives in } [t, t+\Delta t)] \\
&= \sum_{k=0}^{\infty} \Pr[\text{customer generated in } [t-k, t-k+\Delta t)] \\
&\quad \times \Pr[\text{still unconsummated} \mid \text{generated } k \text{ periods ago}] \\
&\quad \times \Pr[\text{chooses this dealer} \mid \text{unconsummated and generated } k \text{ periods ago}] \\
&= \sum_{k=0}^{\infty}\left(\frac{\lambda\rho_h}{2}\Delta t + o(\Delta t)\right)(1-\pi_h^r)^k \frac{1}{n} \\
&= \frac{\lambda\rho_h}{2\pi_h^r n}\Delta t + o(\Delta t). \qquad (39)
\end{aligned}
$$

Q.E.D.

References

Amihud, Y. & Mendelson, H. (1980), 'Dealership market: Market-making with inventory', Journal of Financial Economics 8, pp. 31–53.

Cootner, P. H. (1967), 'Speculation and hedging', Stanford Food Research Institute Studies: Supplement 7(supplement), pp. 65–105.

De Groot, M. H. (1970), Optimal Statistical Decisions, McGraw-Hill.

Diamond, P. (1987), Search theory, in J. Eatwell, M. Milgate & P. Newman, eds, 'The New Palgrave: A Dictionary of Economics', Macmillan, pp. pp. 273–9.

Ho, T. & Stoll, H. R. (1981), 'Optimal dealer pricing under transactions and return uncertainty', Journal of Financial Economics 9, pp. 47–73.

Lippman, S. A. & McCall, J. J., eds (1979), Studies in the Economics of Search, North-Holland.

Madhavan, A. & Smidt, S. (1993), 'An analysis of changes in specialist inventories and quotations', Journal of Finance 48, 1595–1628.