

# Search Costs: The Neglected Spread Component

Mark D. Flood, Ronald Huisman, Kees G. Koedijk, and Richard K. Lyons

November 2, 1998

## Abstract

Dealers need to search for quotes in many of the world's largest markets (such as spot foreign exchange, US government bonds, and the London Stock Exchange). This search affects trading cost. We estimate the share of total trading cost attributable to search. Our experiments show that the share is large—roughly one-third of the effective spread. Past work on estimating spread components typically omits the search component. Our estimates suggest this omission is important.

## Correspondence

Professor Richard Lyons  
Haas School of Business  
U.C. Berkeley  
Berkeley, CA 94720-1900  
Tel: 510-642-1059, Fax: 510-643-1420

Flood is at Concordia University and the University of North Carolina at Charlotte (mdflood@ email.uncc.edu), Huisman is at Erasmus University Rotterdam (R.Huisman@fac.fbk.eur.nl), Koedijk is at the Limburg Institute of Financial Economics (LIFE) at Maastricht University (c.koedijk@berfin.unimaas.nl), and Lyons is at UC Berkeley and NBER (lyons@haas.berkeley.edu). We thank CREED at the University of Amsterdam for use of their laboratory, securities dealers for their participation, and Roger Otten for assistance with the experiments. We also thank the following for valuable comments: Peter Bossaerts, Ian Domowitz, Joel Hasbrouck, Chuck Schnitzlein, Ingrid Werner, and seminar participants at Maastricht University, the University of Liege, and HFDF-II. Flood gratefully acknowledges financial assistance from the Social Sciences and Humanities Research Council of Canada. Lyons gratefully acknowledges financial assistance from the National Science Foundation. All remaining errors belong to the authors.

## 1. Introduction

The bid-ask spread and its components are of considerable significance. In Europe, for example, attempts to lower spreads to prepare for monetary union are driving major exchange redesign. In the US, both the NYSE and NASDAQ have recently implemented trading-rule reforms to lower trading costs (such as smaller minimum tick sizes). How reforms lower the spread, however, depends on how reforms affect particular spread components. These links are not yet well understood. Spread components are significant at the macro level as well. They provide insights into trading environments that are difficult to generate in other ways. Consider the finding of an adverse-selection component in FX spreads. Most people believe this market is free of private information—the finding of an adverse-selection component suggests otherwise.

This paper measures a component of the spread that is too-often overlooked—search costs. Early work on spreads does not include a search component, emphasizing three other components instead: order-processing costs, inventory-holding costs, and adverse-selection costs.<sup>1</sup> To be fair, early work had reason to ignore search costs: it addressed the NYSE, a consolidated market in which all participants trade at the best bid/offer, so search is a non-issue. But most of the world's largest markets are not consolidated, including spot foreign exchange, US government bonds, and the London Stock Exchange. In non-consolidated markets dealers have to search for the best counterparties, and this imposes costs. Accordingly, to estimate spread components in these markets one needs to account for these costs. (Even the NYSE is not as consolidated as once thought: off-floor block trades in these stocks also involve search costs; see

---

<sup>1</sup> The literature on spread components is extensive. See Campbell, Lo, and MacKinlay (1997, ch. 3) and Huang and Stoll (1997) for recent surveys. Glosten and Harris (1988) and Hasbrouck (1988) were first to disentangle these three components empirically.

Keim and Madhavan 1996.)

It is surprising that search theory is applied so narrowly in microstructure given that so many markets are not consolidated.<sup>2</sup> Application to date focuses on block trading. In an early theoretical paper Burdett and O'Hara (1987) use a search model to address the forming of a syndicate of buyers. As typical in search models, the block seller's strategy is an optimal stopping rule, in this case one that determines when the seller stops contacting new potential syndicate members. Search costs in their model take the form of price impact from revealing the trade to more traders. Keim and Madhavan (1996) model the search process differently. In their block-trading model search is done by an intermediary who charges commissions to cover costs. Departing from formal modeling, Campbell et al. (1991) examine off-market trading using experiments that incorporate search. In particular, they introduce an off-floor alternative to determine conditions under which off-floor trading might occur. These papers clearly deepen our understanding of block trading. In our judgment, though, they do not go far enough—they do not recognize how pervasive search costs are. This is a central message of our paper.

We view search costs as a distinct, fourth component of the spread. Our view rests on three arguments. First, search-theoretic models are fundamentally different than those used to derive the other three spread components. This is *prima facie* evidence that as a cost category search is distinct. Second, search is neither necessary nor sufficient for any of the other three costs to arise; indeed, that search could be neglected so long within microstructure is evidence of this. Finally, search costs cannot be neatly subsumed in one of the other categories (order-processing, inventory-holding, and adverse-selection). For example, order-processing cost, either fixed or

---

<sup>2</sup> In economics, search theory is most extensively applied to labor markets (see Pissarides 1985 and Mortenson 1986, among many others) and product markets (see Diamond 1982 and Rubinstein and Wolinsky 1987).

variable, can be affected by the cost of search for particular counterparties, due to netting or other settlement economies. Inventory-holding cost too is affected by search costs that arise while managing inventory. Adverse-selection costs are affected as well. Consider an informative order received by one dealer that becomes known by all other dealers only after a sequence of interdealer trades; within this sequence, the adverse selection faced by any one dealer depends on the duration of his exposure, and search costs affect this duration. In the end, treating search as just another dimension of adverse selection—as some have suggested—misses much of its richness.

Measuring the search component of the spread is not an easy task. The approach we adopt here is experimental. A valuable feature of this approach is that it provides us as researchers more information than is available to participants themselves. This additional data allows us to isolate the search cost component in a way not possible with conventional empirical methods. Specifically, for every transaction we can measure both the actual price and the best price in the market at that time. The gap between actual price and best price is our measure of search costs. When we use this "Gap Method" to estimate search cost's share of the effective spread we find its magnitude is striking—accounting for about one-third.

Understanding why transactions are not executed at the best price requires some perspective on our experiment. In our market, to obtain quotes dealers must call one-another on a bilateral basis.<sup>3</sup> It is thus not possible to observe the best bid and ask in the market directly. Not only is shopping for quotes time-consuming, it also makes older quotes increasingly uncertain. (Uncertainty increases because quotes expire if the quoting dealer revises his quote in the interim,

---

<sup>3</sup> That dealers shop for bilateral quotes is shown by Lyons (1998) in the foreign exchange market. In his sample more than 70 percent of quote requests do not generate a trade.

and revisions are not automatically communicated.) This induces dealers to trade at inferior prices, implying an effective spread wider than the best bid/ask.

In addition to our Gap Method for measuring search costs, to check robustness we also estimate search costs using what we call the "Transparency Method." The Transparency Method relies on a change in transparency regime to measure search costs—from low pre-trade transparency, which imposes search costs, to full transparency, which eliminates them. More specifically, the experimental design described above is one of low pre-trade transparency—dealers cannot see others' quotes, except bilaterally. In our full-transparency design, dealers see all dealer quotes, and strict price priority is enforced. With no need to search for quotes, and no scope for negotiation, the benefits to search drop to zero. This is an important robustness check: implicit in our Gap Method is the premise that without search costs the sequence of best prices would remain unchanged. This is not necessarily true since best price is endogenous, and the sequence may change due to changes in search costs/benefits. In the end, results from the Transparency Method confirm those from the Gap Method—the reduction in effective interdealer spreads from eliminating search via full transparency is roughly one-third.

We design our experimental approach to minimize exposure to common methodological concerns. This is especially true in two dimensions. First, our subjects are professional traders, so they do not lack professional expertise. Second, our software is designed to mimic trading systems that actual traders employ, to the extent possible. Our results are thus unlikely to be contaminated by extensive learning about the experimental setting.

The remainder of the paper is in four sections. Section 2 outlines two leading models for estimating effective spreads. Section 3 describes our experimental design. Section 4 presents our results. And Section 5 concludes.

## 2. Effective Spreads: Two Leading Models

Estimating the share of search costs in spreads requires both an estimate of search costs and an estimate of the total spread. How we estimate search costs is sketched above and further clarified in section 4. This section describes our two estimators of the total spread. Because there is not yet consensus on methods of spread estimation, we use two different estimators for robustness. The first, Roll's (1984) estimator, is perhaps the best known, but it is also designed for a special trading environment. The second, Huang and Stoll's (1997) estimator, is both recent and relatively general in its applicability.

Consider first Roll's (1984) model.<sup>4</sup> Roll shows that first-order auto-correlation in returns can be used to measure the effective spread. Specifically, if  $S$  is the effective spread and  $P_t$  is the transaction price at time  $t$ , then Roll's estimator for the effective spread,  $S_{\text{roll}}$ , is:

$$(1) \quad S_{\text{roll}} = 2\sqrt{-\text{cov}(\Delta P_t, \Delta P_{t-1})}$$

A crucial assumption in Roll's model is that the conditional probability of an order being a buy is always 0.5, regardless of past orders.<sup>5</sup> The measure is thus designed to capture pure bid-ask bounce, arising for example from order-processing costs (the order-processing component of the spread). This is a strong assumption, however, particularly where adverse-selection and inventory-holding components are likely. In our experiment, for example, an adverse-selection component *is* likely since customer trades, generated by computer, provide information about fundamental value.

---

<sup>4</sup> Like the rest of the literature, we focus on transaction prices and effective spreads. Effective spreads can differ from quoted spreads for many reasons. One reason emphasized in the literature is that dealers negotiate prices within their quoted spreads (see, for example, Fialkowski and Petersen 1994). The reason emphasized in this paper—search costs—is quite different.

The second spread model we use is that of Huang and Stoll (1997). Theirs is a quite general model that accommodates not only order-processing but also adverse-selection and inventory-holding components. To derive it, they define a trade-direction indicator  $Q_t$  that equals +1 when a buyer trades at the ask and  $-1$  when a seller trades at the bid.<sup>6</sup> This indicator variable plays a central role in the regression they use to estimate the spread,  $S_{HS}$ :

$$(2) \quad \Delta P_t = S_{HS}(\Delta Q_t/2) + S_{HS}I(Q_{t-1}/2) + e_t$$

The parameter  $\lambda$  captures here both adverse-selection costs and inventory-holding costs. Note that both  $S_{HS}$  and  $\lambda$  have no time subscript, reflecting Huang and Stoll's assumption of a constant total spread as well as constant asymmetric-information and inventory-holding components. The adjustment of quotes to trades is captured by the second term on the right side (the first term on the right side does not involve  $\lambda$  because it picks up only the bounce between the bid and ask of the constant spread).

An attractive feature of this model is that it nests many different modeling approaches, including the Roll model, in which case the parameter  $\lambda$  equals zero.<sup>7</sup> It is therefore more robust than the Roll measure, particularly in environments like the one we analyze. We estimate Eq. (2), using ordinary least squares, to generate our second measure of the effective spread  $S_{HS}$ .

---

<sup>5</sup> Many papers analyze the robustness of Roll's estimator and propose improvements. See for example Choi, Salandro, and Shastri (1988), Harris (1990), and Jong, Nijman, and Roell (1995).

<sup>6</sup> In our estimation of the Huang-Stoll model we can assign values to  $Q_t$  directly (i.e., without reference to the quote midpoint) because we observe the direction of every trade.

<sup>7</sup> Huang and Stoll show this nesting on page 1001. To see it, set  $\lambda=0$  and calculate the first-order autocovariance of both sides of Eq. (2), using the fact that  $\text{cov}(\Delta Q_t, \Delta Q_{t-1}) = -1$  when buys and sells are equally likely.

### 3. Experimental Design

In a nutshell, our market has multiple live dealers and computer-generated customer trades. The basic design is similar to that used in Flood et al. (1998), though here we introduce some new parameters. This section outlines the structure of the experiments. (For further details on the software see Flood et al. 1998.)

We conducted three experiments on three different dates in 1996 and 1997 using the computer laboratory at the University of Amsterdam's CREED (Center for Research in Experimental Economics and Political Decision Making). The need for three separate experiments arises from our desire to check robustness in various directions. Though details appear below, here is an overview:

**Experiment 1:** This experiment is designed to measure search costs using the Gap Method.

Transparency of quotes is low, and dealers must search for good quotes. The design includes asymmetric information, arising from informative customer orders.

**Experiment 2:** This experiment is designed to measure search costs using the Transparency

Method. Unlike experiment 1, quote transparency is full, which eliminates quote search. Like experiment 1, the design includes asymmetric information, arising from informative customer orders.

**Experiment 3:** This experiment is designed to measure search costs using the Gap Method. The

design is the same as experiment 1 except that it adds exogenous price variation. This introduces inventory risk that is not tied directly to revelation of asymmetric information.

Adding experiment 2 allows us to check the robustness of the Gap Method by providing the data necessary for the Transparency Method. (The Gap Method uses data from a single experiment to

measure search cost's share of the spread, whereas the Transparency Method compares data across two distinct experiments.) Adding experiment 3 allows us to check the robustness of the experiment-1 design. Specifically, the design of experiment 3 eliminates the feature that inventory risk shrinks to zero in experiment 1 as customer information is discovered.

### **3.A Features Common to All Three Experiments**

The subjects in each experiment (7-9) are professional traders from various firms in Amsterdam (ABN Amro, de Generale Bank, Optiver, and Oudhoff Effecten). Each subject had his/her own trading screen and keyboard (a sample screen appears in Figure 1). The subjects trade as dealers in 12 five-minute rounds, with a single security in each round.<sup>8</sup> The rounds are independent in the sense that information about the security's value in one round is unrelated to values in other rounds.

#### Customer Trades

In addition to the seven human dealers there are customers in the market whose trades are computer generated every 3.5 seconds on average. Customers do not provide price quotes or trade with each other. An important feature of our design, across all three experiments, is that when a customer trades with a dealer it is always at the best price in the market.

Any given customer trade is either informed or uninformed. This is determined at random just prior to each trade. The *ex-ante* probability that a particular customer trade is informed is determined by a parameter,  $\alpha$ , which we hold constant throughout each five-minute trading

---

<sup>8</sup> Experiment 3 has a slightly different design in this respect. It has instead 16 four-minute rounds. Though we have no evidence that this change had substantive effect, it is not strictly true that Experiment 3 is a perfectly controlled shift of the inventory-risk parameter described below.

round. The (human) dealers know the value of  $\alpha$  at the start of each round. If a customer trade is uninformed then the computer determines at random—with even odds—whether to buy or sell, and then buys (sells) one share at the lowest ask (highest bid) among the dealer quotes. If a customer trade is informed then the computer compares the terminal value to the best bid and ask prices, buying (selling) one share if above (below) the lowest ask (highest bid).

### Trade Transparency

Each dealer has a transaction history window that contains all his trades, both interdealer and customer. The transactions of other dealers are not public information. The window includes for each transaction the identity of the counterparty, the quantity traded (one share), and the transaction price. In the case of a customer trade, the counterparty designation provides no information about whether the trade was informed, only that the counterparty is a customer. There is no delay in transaction disclosure.

## **3.B Features That Distinguish the Three Experiments**

### Dealer Quotes and Interdealer Trades

Within each 5-minute round, trading occurs at prices quoted in esquires. Each dealer must maintain a current bid and ask quote in a central queue. Prices can be revised—but not withdrawn—at any time.

*In experiments 1 and 3 the queue is not public information; in experiment 2, the full-transparency design, the queue is public information.*

For inter-dealer trading, a dealer must request prices bilaterally by pressing a key and then entering the identifying letter of another dealer. The quote then appears on the caller's screen for 7 seconds and is available for trading. (The quoting dealer is not informed of the call—the quoted

price is simply taken from the price queue.) To buy the dealer presses the “B” key; to sell the “S” key. As in Glosten and Milgrom (1985), we normalize the trade size to exactly 1 share per transaction; however, traders can initiate several one-share trades during the 7 seconds that the quote is active.

### Information Structure

At the end of each trading round each dealer’s position is converted into a fictitious laboratory currency, esquires, at the security's terminal value.

*In experiments 1 and 2 the terminal value has only a single component, and that value alone governs the trading of informed customers as described above. In experiment 3 the terminal value has two components, one which governs the trading of informed customers, and another reflecting exogenous price risk.*

In all three experiments, the component that governs informed customer trading is distributed uniformly between 0 and 200 (esquires per share). We refer to this value as learnable. In experiment 3 the terminal value also has an unlearnable component. The unlearnable component is distributed uniformly over three possible values,  $\{-50, 0, \text{ and } +50\}$ . Dealers are informed *ex-ante* about the distributions of the two components, but are not told the realized terminal values until their inventories are converted at the end of trading.

The learnable component is learnable in the sense that dealers can filter this information from observing customer trades. For the unlearnable component, in contrast, there is no conditioning information revealed that is useful for prediction. Rather, the unlearnable component is an end-of-round price shock that prevents discovery of the learnable component from eliminating all inventory risk.

Payoffs for each of the 16 rounds are converted from esquires into Dutch guilders

according to the following scheme. The dealer with the highest profits in each round receives 7 guilders, and the dealer with the lowest receives 1 guilder. Guilder payoffs for the remaining 5 dealers are interpolated linearly between these two extremes.<sup>9</sup> Thus, the conversion into guilders is normalized to ensure that all participants receive a positive guilder payoff for their participation (and to constrain the absolute guilder cost to the experimenters).

## 4. Results

### 4.A Summary Statistics

Table 1 contains summary statistics from experiment 3.<sup>10</sup> Row 1 shows that across rounds the parameter  $\alpha$  alternates between 0.3 and 0.7 (save round 9). Interestingly, interdealer volume as a share of the total ranges from 74 to 92 percent, roughly matching the spot FX market's empirical share of about 80 percent (the remainder is of course from customers). Note that our continuous-trading set-up has the notable benefit of yielding large amounts of data: our sample of transactions includes 10,922 observations.

### 4.B Estimating Search Costs: The Gap Method

Our experiment provides data that isolate the search cost component in a way not possible with conventional empirical methods. For every transaction we can measure both the actual price

---

<sup>9</sup> Specifically, if dealer  $i$ 's esquire profits are denoted  $E_i$ , and the best and worst trader's earnings are  $E_{\max}$  and  $E_{\min}$ , then  $i$ 's guilder payoff is:  $G_i = 1 + 6(E_i - E_{\min}) / (E_{\max} - E_{\min})$ . Thus,  $G_i$  is an affine transformation of  $E_i$ , albeit a different transformation in each round.

<sup>10</sup> We present only one set of summary statistics to conserve space, and we select this experiment because we view the continuing-inventory-risk and low-transparency features as the most realistic. Summary statistics for the other experiments are of course available on request. In experiments 1 and 2 the parameter  $\alpha$  is also allowed to take an intermediate value of 0.5.

and the best price in the market at that time. The gap between them provides a measure of search costs—our Gap Method. Recall that in all three experiments, customers trade against the best dealer price by design. Dealers, on the other hand, have to search the available prices on a bilateral basis. The effective spread for customer trades thus provides a benchmark for evaluating the effective interdealer spread. At the margin, a dealer ceases searching and executes a trade when the cost of continuing search just equals the benefit. The gap between the two effective spreads measures the benefit of eliminating the search costs. This is our rationale for assigning the term “search costs” to this gap.

Table 2 presents four Gap-Method estimates of search cost's share in effective spreads. Clearly, search costs are a significant component of the interdealer spread. The estimates come from experiments 1 and 3, and use Roll and Huang-Stoll spread measures. Recall that experiments 1 and 3 differ only in their treatment of inventory risk: in experiment 1 inventory risk shrinks to zero as price is discovered, whereas in experiment 3 there is always residual inventory (price) risk. In all four cases, search costs account for a strikingly large share of the interdealer spread. Because the Roll measure is not designed to be robust in our trading environment, per the introduction, we consider the HS estimates more compelling.<sup>11</sup> Of the two experiments, we consider experiment 3 the more realistic in terms of combining price discovery with inventory

---

<sup>11</sup> To gauge the Roll model's reliability in our environment, we test what is perhaps its key implication, namely that equally likely buy and sell orders should produce autocorrelation in transaction-price changes of -0.5. The model passes the test: the average first-order autocorrelation across the sixteen rounds is -0.50, and we cannot reject -0.5 at the five-percent level (all transactions included). Note that this occurs despite the occurrence of price discovery, which involves gradual drift in prices toward fundamental value, which should increase autocorrelation. One possibility is that this price-discovery effect on autocorrelation is swamped by the preponderance of interdealer trades, which are so frequent that the effect of price drift is negligible. Consistent with this is our finding that average autocorrelation of price changes between uninformed customer trades is just -0.28; since customers trade only every 3.5 seconds, more time elapses between these trades, and the added price discovery increases first-order autocorrelation.

risk. (Note that the smaller share in experiment 3 is consistent with a larger role for inventory-holding costs.) In the end, though, the message is the same: search costs are substantial.

#### **4.C Estimating Search Costs: The Transparency Method**

In addition to measuring search costs using the Gap Method, we also examine search costs by comparing spreads across experiments that differ only in quote transparency. Experiment 1 has low quote transparency—dealers cannot see others’ quotes, except bilaterally. In experiment 2 quote transparency is full, so dealers see all dealer quotes, and strict price priority is enforced. Consequently, search disappears in this full-transparency regime. (Note that quotes are not negotiable. If they were, then even under full transparency there might still be an incentive to search.) Under this approach, then, we measure search costs as the reduction in spreads from full transparency, which eliminates search costs.

Table 3 presents our results using the Transparency Method. This method provides confirmation of our results from the Gap Method. The reduction in interdealer spreads from eliminating search via full transparency is substantial—about one-third. We consider this an important robustness check: implicit in our gap-based measure is the premise that without search costs one could trade at an unchanged sequence of best prices. This is not necessarily true since best price is endogenous, so the sequence of best prices is likely to change with changes in search costs/benefits. Of course, the Transparency Method too has its shortcomings. For example, though search falls to zero under full transparency, other spread components might also be affected. It is comforting nevertheless that both methods attach a substantial share of the spread to search costs. The broadly similar results are generated from very different methods.

## 5. Conclusions

We examine trading costs in markets where dealers search for quotes. Earlier work shows spreads are affected by three costs—order-processing, adverse-selection, and inventory-holding. Our results indicate a fourth cost component is present in bilateral dealer markets: search costs.

In particular, our objective is to measure the share of total trading costs attributable to search. We do this in two ways. Both produce a striking message: search costs account for a substantial share of the effective interdealer spread. Our first measure, the Gap Method, is based on the gap between transaction price and best price. Dealers in our experiment trade at prices that differ from the best price because they do not observe all quotes and shopping for quotes bilaterally takes time. At the margin a dealer stops searching and trades when the cost of continuing search equals the benefit. This stopping rule provides a basis for interpreting the gap between transaction price and best price as a measure of search costs. Of course, it is not necessarily true that the sequence of best prices would remain unchanged if search costs were zero (an implicit assumption of the Gap Method). For this reason we devise a second method for measuring search costs, the Transparency Method. Under our regime of full quote transparency, the benefits of search drop to zero, so no search takes place. The resulting no-search spreads can then be compared to the with-search spreads from low transparency. The results from this second measure confirm those from the first.

A central message of this paper is that the effects of search are much larger and more pervasive than is currently recognized. Work to date on search within microstructure focuses narrowly on block trading. There is ample room to expand application. Indeed, with microstructure's extension beyond the NYSE into some of the world's largest markets, like spot FX, ignoring search cost is increasingly untenable.

The presence of a search component in spreads opens some fascinating topics for future research. For example, a fully articulated theory might distinguish (at least) two types of search. The distinguishing feature is their effect on subsequent beliefs. For the type present in standard search-theoretic models, observing an additional price does not alter one's beliefs about value—one is simply sampling from a distribution of prices in an attempt to find the best. The second type, in contrast, involves searching for valuation signals in a common value, asymmetric information environment. Presumably, both of these types are relevant in most non-consolidated markets.

**Table 1****Experiment 3: Parameters and Summary Statistics***This table summarizes parameters and key statistics for each of the experiment's 16 rounds*

	1	2	3	4	5	6	7	8	9	10	11	12	13
<b>Parameters</b>													
$\alpha$ : prob. informed	0.3	0.7	0.3	0.7	0.3	0.7	0.3	0.7	0.7	0.3	0.7	0.3	0.7
Terminal Value	200	142	82	102	138	121	81	102	171	55	2	114	205
- learnable part	150	142	82	52	188	121	31	152	171	55	52	64	155
- unlearnable part	50	0	0	50	-50	0	50	-50	0	0	-50	50	50
<b>Data</b>													
# quotes set	59	40	55	61	54	58	60	64	63	60	59	47	52
# transactions	438	1038	480	505	720	566	370	690	434	626	781	1080	976
- % interdealer	78	91	81	81	87	86	74	86	78	85	88	91	90

The parameter  $\alpha$  is the probability that a given customer trade is informed. Terminal value is expressed in esquires, the experimenter has two parts, learnable and unlearnable, where the learnable part is correlated with customer trades per the experiment description. The total number of transactions is 10,922.

**Table 2**

**Search Cost Share of the Spread: Gap Method**

---

	<u>Roll</u>	<u>HS</u>
<u>Search Cost Share</u>		
<b>Experiment 1</b>	56%	64%
<b>Experiment 3</b>	46%	28%

---

The two effective spread measures are from Roll (1984) and Huang and Stoll (1997). Search cost share is the share of the interdealer spread accounted for by search costs. Under the Gap Method this is calculated as GAP divided by the average interdealer spread, where GAP is the difference between the average interdealer spread and the average customer spread. The average customer spread is strictly lower than the average interdealer spread in experiments 1 and 3 since by design customers always get the best price in the market. Dealers, on the other hand, must search for best price in the low quote-transparency regimes of experiments 1 and 3. (Note that across rounds the average probability of a customer being informed is 0.5). Experiment 3 differs from experiment 1 in that it adds exogenous price risk, the unlearnable component of terminal value, which serves to maintain inventory risk throughout price discovery.

**Table 3**

**Search Cost Share of the Spread: Transparency Method**

---

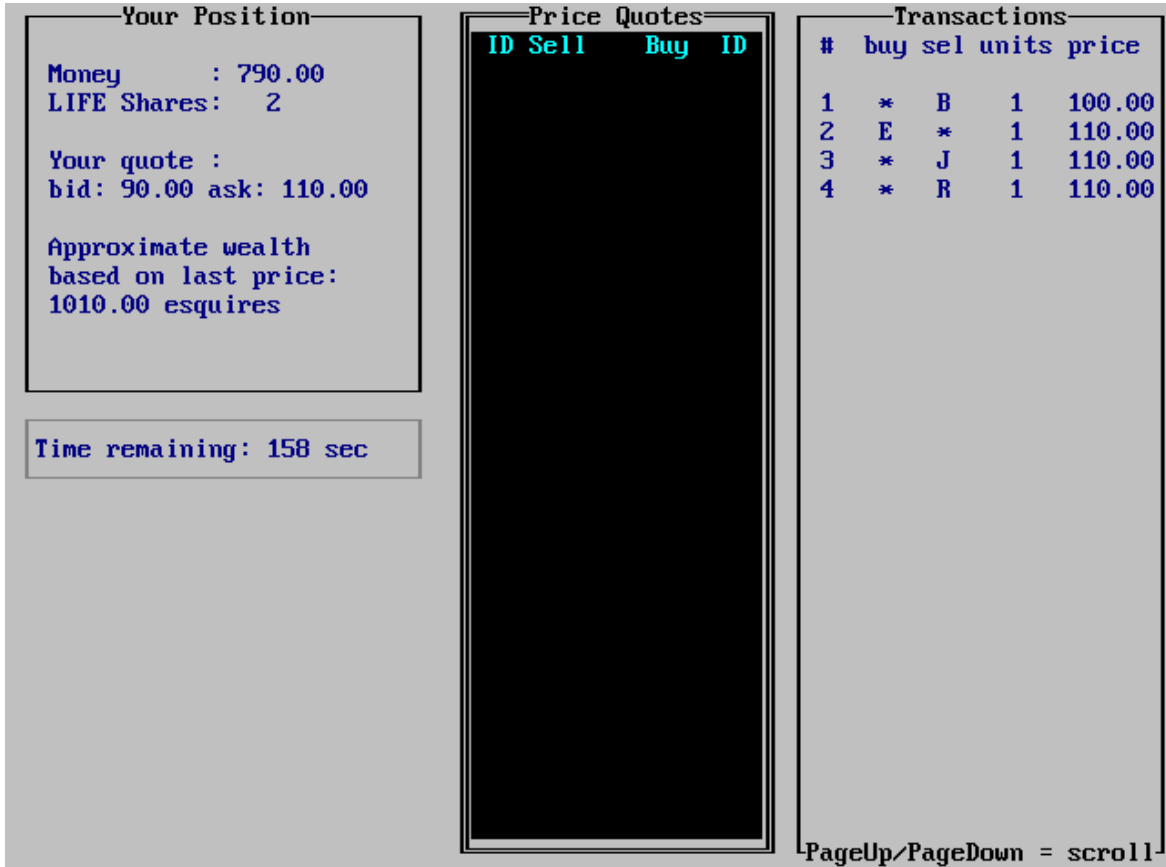
	<u>Roll</u>	<u>HS</u>
<b>Search Cost Share</b>	33%	34%

---

The two effective spread measures are from Roll (1984) and Huang and Stoll (1997). Search cost share is the share of the interdealer spread accounted for by search costs. Under the Transparency Method this is calculated as the percent reduction in effective interdealer spreads from switching to a full-transparency regime, where full transparency means each dealer can observe all other dealer quotes, eliminating the need to search. The data for this calculation come from experiments 1 (low transparency) and 2 (full transparency).

Figure 1

Trading screen



Each dealer has his/her own trading screen. The window on the upper left presents this dealer's cash balance (790 esquires), inventory (2 shares long), current outstanding quote (90 - 110), and approximate profit based on the price of the last transaction in which he/she was involved (1010 esquires). The middle window on the left side shows the time remaining in this round. The black window in the center of the screen is where a quote appears when this dealer calls another dealer. Under the heading « ID Sell Buy ID », ID denotes the identity of the dealer presenting the quote and Sell denotes the quoted bid (at which this dealer can sell); the Buy column contains the quoted ask (at which this dealer can buy the share). Information on past transactions appears at the right of the trading screen. By default, it displays the details of the last 20 transactions. The dealer can scroll through the list with the PageUp and PageDown keys. For all transactions, the identities of the buying (under heading buy) and the selling (sel) dealer is displayed, along with the number of shares involved in the trade and the price at which the trade cleared. For example, the first row indicates that this dealer (his identity is shown as an asterisk \*) bought one share for 100 esquires from dealer B. The seven dealers' identities are denoted by letter ranging from "A" through "G." A customer is denoted with the letter "R." The fourth transaction is thus an example of a trade in which a customer sold to this dealer one share for 110 esquires.

## References

- Burdett, K., and M. O'Hara, 1987, Building Blocks: An Introduction to Block Trading, *Journal of Banking and Finance*, 13, 397-419.
- Campbell, J., S. LaMaster, V. Smith, and M. Van Boening, 1991, Off-Floor Trading, Disintegration, and the Bid-Ask Spread in Experimental Markets, *Journal of Business*, 64, 495-522.
- Campbell, J., A. Lo and, A. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press.
- Choi, J., D. Salandro, and K. Shastri, 1988, On the Estimation of Bid-Ask Spreads: Theory and Evidence, *Journal of Financial and Quantitative Analysis*, 23, 219-230.
- Cohen, K., R. Conroy, and S. Maier, 1985, Order Flow and the Quality of the Market, in Y. Amihud, T. Ho, and R. Schwartz (eds.), *Market Making and the Changing Structure of the Securities Industry*, Lexington: Lexington, Massachusetts.
- Diamond, P., 1982, Aggregate Demand Management in Search Equilibrium, *Journal of Political Economy*, 90, 881-894.
- Fialkowski, D., M. Petersen, 1994, Posted Versus Effective Spreads: Good Prices or Bad Quotes? *Journal of Financial Economics*, 35, 269-292.
- Flood, M., R. Huisman, K. Koedijk, and R. Mahieu, 1998, Quote Disclosure and Price Discovery in Multiple Dealer Financial Markets, *Review of Financial Studies*, forthcoming.
- George, T., G. Kaul, and M. Nimalendran, 1991, Estimation of the Bid-Ask Spread and its Components: A New Approach, *Review of Financial Studies*, 4, 623-656.
- Glosten, L., and L. Harris, 1988, Estimating the Components of the Bid-Ask Spread, *Journal of Financial Economics*, 21, 123-142.
- Glosten, L., and P. Milgrom, 1985, Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders, *Journal of Financial Economics*, 14, 71-100.
- Hansch, O., N. Naik, and S. Viswanathan, 1998, Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange, *Journal of Finance*, forthcoming.
- Harris, L., 1990, Statistical Properties of the Roll Serial Covariance Bid/Ask Spread Estimator, *Journal of Finance*, 45, 579-590.
- Hasbrouck, J., 1988, Trades, Quotes, Inventories, and Information, *Journal of Financial Economics*, 22, 229-252.

- Howitt, P., 1988, Business Cycles With Costly Search and Matching, *Quarterly Journal of Economics*, 103, 147-165.
- Huang, R., and H. Stoll, 1997, The Components of the Bid-Ask Spread: A General Approach, *Review of Financial Studies*, 10, 995-1034.
- Jong, F. de, T. Nijman, and A. Röell, 1995, A Comparison of the Cost of Trading French Shares on the Paris Bourse and on SEAQ International, *European Economic Review*, 39, 1277-1301.
- Keim, D., and A. Madhavan, 1996, The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects, *Review of Financial Studies*, 9, 1-36.
- Lamoureux, C., and C. Schnitzlein, 1997, When It's Not The Only Game in Town: The Effect of Bilateral Search on the Quality of a Dealer Market, *Journal of Finance*, 52, 683-712.
- Lin, J.-C., G. Sanger, and G. Booth, 1995, Trade Size and Components of the Bid-Ask Spread, *Review of Financial Studies*, 8, 1153-1183.
- Lyons, R., 1998, Profits and Position Control: A Week of FX Dealing, *Journal of International Money and Finance*, February, 97-115.
- Mendelson, H., 1987, Consolidation, Fragmentation, and Market Performance, *Journal of Financial and Quantitative Analysis*, 22, 189-207.
- Pissarides, C., 1985, Short-Run Dynamics of Unemployment, Vacancies, and Real Wages, *American Economic Review*, 75, 676-690.
- Reiss, P., and I. Werner, 1997, Interdealer Trading: Evidence from London, Stanford Graduate School of Business Research Paper No. 1430, February.
- Roll, R., 1984, A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market, *Journal of Finance*, 39, 1127-1139.
- Rubinstein, A., and A. Wolinsky, 1987, Middlemen, *Quarterly Journal of Economics*, 102, 581-593.
- Stoll, H., 1989, Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests, *Journal of Finance*, 44, 115-134.
- Wolinsky, A., 1990, Information Revelation in a Market With Pairwise Meetings, *Econometrica*, 58, 1-23.