Improved methods for obtaining information from distributed dealer markets

Ajay Shah*
Indira Gandhi Institute for Development Research

1 May 2000

Abstract

Prices and liquidity on distributed dealer markets are known to market participants but not to external observers. Hence, the strategy of polling n respondents, coupled with data reduction using a robust location estimator, has been widely employed, especially in the context of cash–settled futures contracts.

In this paper, we offer a market microstructure interpretation of the information obtained by polling, and propose improvements for many elements of the polling process. The choice of estimator in this context reflects a tradeoff between statistical efficiency and vulnerability to manipulation. We offer empirical evidence about this tradeoff. The results suggest that the adaptive trimmed mean (ATM) has significant advantages over the fixed trimming procedures which are widely used by futures exchanges today.

*The author acknowledges the inputs from members of the NSE *Committee for Development of the Debt Market*, Chitra Ramakrishna, Ajay Sikri and Tushar Waghmare of NSE, Anand Pai of Bank of America, and Susan Thomas of IGIDR in the form of enlightening discussions and access to data. Address: IGIDR, Goregaon (E), Bombay 65, India. Email: `ajayshah@mayin.org` and URL `http://www.mayin.org/~ajayshah`

1 Introduction

The intrinsic non–transparency of distributed dealer markets has motivated special efforts to measure prices and liquidity. Prominent in these efforts are "reference rates" obtained by polling dealers. These are intended to be estimates of the price on a distributed dealer market at a point in time. They have numerous applications in research, exchange–traded and OTC derivatives, performance evaluation of dealers, etc.

On futures markets, the technique of cash settlement has had a major impact upon contract design. While cash settlement was originally motivated by situations such as Eurodollars, where delivery is infeasible, cash settlement can reduce basis risk and eliminate the costs of delivery even when the underlying can potentially be delivered. Even when cash settlement is not used, when the spot market for a product is not centralised at an exchange, reference rates obtained from dealers on the spot market have an important influence upon the futures market, in the life of the contract prior to expiration. In fixed income markets worldwide, which are often distributed dealer markets, widely trusted reference rates are central to interest rate derivatives.

The method that is most commonly in use in computing reference rates is termed the "symmetric trimmed mean" (STM). Prominent examples of the STM are found in the Eurodollar futures contract at the Chicago Mercantile Exchange (CME), and the Muni Bond Index used by the Chicago Board of Trade (CBOT). The CME, for example, samples information from twelve banks and computes a trimmed mean where the two highest and two lowest observations are deleted before computing the sample mean.

The two major concerns in this process are the potential for market manipulation and the extent of sampling noise. Dealers who have positions on the futures market have incentives for misreporting. For hedging and speculative users of a cash–settled futures contract, sampling noise and market manipulation of the reference rate serve to enhance basis risk. Hence futures exchanges which use reference rates in cash settled products have incentives to minimise these problems. From an analytical perspective, the details of the sampling and estimation procedure can have a significant impact upon these concerns. This has motivated a literature on improving the methodology underlying reference rates.

In this paper we examine many elements of the sampling and estimation process. We start with Section 2 which offers a market microstructure interpretation of the information which is obtained, and helps refine the question which respondents are asked in the polling.

We also discuss many ingredients of the polling process which can be designed to minimise the vulnerability to market manipulation. Our normative proposals, for improving the construction of reference rates, are labeled from N1 to N7. We address the tradeoff between vulnerability to manipulation and sampling noise. The sample median is well known as an estimator which is least vulnerable to outliers (Lien, 1989), however it has poor statistical efficiency. Cita and Lien (1997) propose that bootstrapping could be used to do 'adaptive trimming', which may yield greater efficiency as compared with fixed trimming. In Section 4 we offer empirical evidence about this tradeoff. Using observed data from a polling process, we measure the vulnerability of alternative estimators to a simplified form of a manipulative cartel. We also measure the statistical efficiency of alternative estimators.

Our results show that the adaptive trimmed mean (ATM) is the most efficient of the estimators evaluated. The ATM is more vulnerable to market manipulation than the median, but this difference is small for cartels which we might consider plausible in many situations.

2 A market microstructure interpretation

In understanding the process of polling dealers on an OTC market, it is useful to use the the organised financial exchange (which uses electronic order matching) as a benchmark. In the open electronic limit order book (OELOB) market, there is an innate transparency of prices and liquidity:

1. Every trade is matched on the trading computer, so that all trade prices and quantities are publicly revealed.

2. The liquidity of the market is embodied in the limit order book. The best buy price a^* and the best sell price b^* are observed, yielding the market–wide bid–ask spread $a^* - b^*$. In addition, market liquidity at all transaction sizes is publicly revealed through the open limit order book. If a market order of size q executes at a price $p(q)$, and if the true price is estimated by $\bar{p} = (a^* + b^*)/2$, then the *market impact cost* is $\lambda(q) = 100(p(q) - \bar{p})/\bar{p}$. All traders can calculate $\lambda(q)$ for all trade sizes q using the publicly visible limit order book. The function $\lambda(q)$ is termed the *liquidity supply schedule*.

In contrast, on the distributed dealer market, trades are not centrally reported, so the latest trade price is not observed. Each dealer may offer quotes but there is no reliable way to observe the market–wide best bid b^* and best ask a^*. Hence the true bid–ask spread is also not observed. There is no possibility of observing the entire $\lambda(q)$ function. It is natural to poll dealers as a way of obtaining information about the distributed dealer market.

2.1 Dealer reporting prices for trading as a principal

One possible sampling strategy is to ask n respondents to report prices and quantities at which they are willing to trade as principals. In this case, respondent i would reveal the information (b^i, q_b^i) and (a^i, q_s^i). This information can be interpreted as a pair of limit orders. On the OELOB market, it is common for traders to have placed multiple limit orders at any point in time. Hence respondents could be allowed to report more (or less) than two "limit orders". Information from n respondents could be pooled to create a "limit order book", which would estimate prices and liquidity on the distributed dealer market at the instant of sampling.[1]

This interpretation leads to several difficulties:

1. Since only n respondents are sampled, the "limit order book" thus obtained would always understate the true liquidity of the market. It would only be an unbiased estimate of the true $\lambda(q)$ function as n approaches the complete population.

2. A location estimator such as the sample mean, applied to the buy prices (or the sell prices) on a "limit order book", has little economic interpretation. Summarising the n values of (b^i, a^i) of a limit

order book should be done using the best sell price $\hat{b}^* = \max b^i$ and the best buy price $\hat{a}^* = \min a^i$. These estimators are vulnerable to manipulation. We might sometimes obtain $\hat{b}^* > \hat{a}^*$, a negative spread, if one b^i or one a^i is sufficiently contaminated.

2.2 Dealer reporting prices in agency function

An alternative sampling strategy is based on the idea that the distributed dealer market has a true liquidity supply schedule $\lambda(q)$. Dealers observe $\lambda(q)$ but external observers do not. Respondents are asked to report *their* perception of a^* and b^* at a fixed quantity q_0. The information reported by respondent i would be an estimator of the true state of the market, as observed by him, and need have no link to his own trading intentions. In this case, it is meaningful to apply a location estimator to the polled values. We treat the information from each respondent as a noisy estimate of (a^*, b^*) and use estimators such as the sample mean to obtain a measure of location.

This approach yields estimates for b^* and a^*, which can additionally yield $\hat{\bar{p}}$ and the bid–ask spread. Hence the information revealed by the polling process consists of four numbers: b^*, the best sell price for a quantity q_0; a^*, the best buy price for a quantity q_0; $a^* - b^*$, the bid–ask spread, a measure of the liquidity of the market; and $\bar{p}$, an estimator of the "equilibrium price".

2.3 Normative implications

This reasoning has several normative implications for the way polling is done:

N1 In polling, a respondent should not be asked the question "What is your buy and sell price at a trade size q_0?". He should be asked "In your assessment of the state of the market, what is the best buy and best sell price that could be obtained *from the market* at a trade size q_0?". The former question expects the respondent to act as a principal; the latter implies that the respondent acts as an agent.

This argument is consistent with some existing market practice. For the Municipal Bond Index futures contract, the CBOT polls six brokers for "their assessment" of the bid price on each of 40 bonds that make up the index. Similarly, the British Bankers Association offers the following instructions to respondents:[2]

> An individual BBA LIBOR Contributor Panel Bank will contribute the rate at which it could borrow funds, were it to do so by asking for and then accepting inter–bank offers in reasonable market size just prior to 11:00.

On the other hand, reference rates do exist (e.g. the Reuters "MIOR" on India's inter–bank market) which simply use information available on quote–dissemination systems. These quotes are indicative prices for trades against the dealer as a principal; they are not the dealer's asessment of a^* and b^*. Hence, these reference rates suffer from the important conceptual problems described above.

N2 The respondents who are polled should be homogeneous in their credit risk, so that their information is comparable with other respondents. OTC markets suffer from counter-party credit risk, so that if some participants are seen as being more risky than others, the rates that they will elicit from counter-parties will be higher. For example, Berkowitz (1999) reports that from January 1995 to January 1998, in the environment of the East Asian crisis, Japanese banks reported three–month

LIBOR rates which were 6.2 basis points above those reported by non-Japanese banks. This introduces noise in the process of estimation of reference rates.

N3 Some reference rates, such as LIBOR, only estimate the offer price. It appears that there is a very small cost in sampling both the bid and the offer, and reporting reference rates for both the bid and the offer. Derivative contracts would then have a choice of b^* or a^* or $\bar{p}$ for use as underlyings. In addition, sampling both the bid and the offer yields one additional time–series, that of the bid–ask spread.

Imprecision in reported rates is a source of risk for users who design financial contracts based on a reference rates. Hence, it is desirable to report measures of imprecision of these estimates, such as the σ of the sampling distribution of $\hat{a}^*$ and $\hat{b}^*$.

Derivative products are sometimes designed with special provisions for extreme events (Ensminger, 1999). For example, in India, the inter–bank market has experienced periods in which market liquidity dropped to near–zero levels. In these extreme situations, where the "normal spot market" can be thought to have broken down, we may see extremely high standard deviations of the reference rates, and extremely high bid–offer spreads. Hence, the dissemination of these time–series would help economic agents in designing more complete financial contracts.

Sometimes a respondent has only paid the search cost of knowing the $\lambda(q)$ function for positive q (or negative q).[3] Such a respondent would only report a^i or b^i, but not both. This information enters into estimation without any special procedures.

3 Choice of estimator

Once information is obtained from a set of respondents, the next question in obtaining a reference rate is the choice of estimator. Methods derived from the problems of robust estimation in statistics (Berkowitz, 1999) have proved to be useful here.

While the field of robust estimation offers a wide variety of estimation strategies, trimmed means dominate the reference rates known today.[4] This class contains many possible estimators, ranging from the simple mean (the trimmed mean with no trimming) to the median (maximal trimming). For simplicity, we only consider symmetric trimming. Suppose X is a random sample of size $n : X = (X_1, \ldots, X_n)$. Let $X_{(1)}, X_{(2)}, \ldots, X_{(n)}$ denote the order statistics. The α-trimmed symmetric mean $T(\alpha, \alpha; X)$, $\alpha \in \Lambda = \{i/n : i = 0, \ldots, (n-1)/2\}$, is based on trimming $2\alpha n$ observations, leaving a censored sample of $n(1-2\alpha)$ observations:

$$T(\alpha, \alpha; X) = \frac{X_{(1+\alpha n)} + \ldots + X_{(n-\alpha n)}}{n(1-2\alpha)}$$

[Table 1 about here.]

[Table 2 about here.]

Table I shows how several important estimators can be viewed as special cases of $T(\alpha, \alpha; X)$. Table II shows estimators used in major real–world applications, where we see the prominent role played by the STM.

The choice of α is a tradeoff between statistical efficiency and vulnerability to market manipulation. At one extreme, with $\alpha = 1/2$, the median is well–known to be robust to "outliers" and hence relatively robust to manipulation, but it has the poorest statistical efficiency. Estimation procedures which yield much improved statistical efficiency are more vulnerable to manipulation.

3.1 Manipulation

The basic problem with any polling process lies in verifying the accuracy of information from respondents. A respondent who has a long (short) position on the futures market has an incentive to inject a positive (negative) bias in the prices reported by him. Even small distortions in the reference rate can imply economically significant payoffs on a futures position (Cita and Lien, 1992).

Lien (1989) analyses one specific setting: where the median is used as the location estimator, and the incentive to falsify reports derives purely from open positions on the corresponding futures market. He derives an approximation for the bias b_i introduced by respondent i:

$$b_i \approx \left(\frac{S_i}{c_i}\right) \left(\sum_{j=1}^{k} \frac{(k!)^2}{(k-j)!(k+j)!} 2^{k-j}\right)^{-1}$$

where S_i is the open position on the futures market of respondent i, c_i is the cost to the respondent of being untruthful, and $2k + 1$ respondents are polled by the exchange. This expression leads to three implications for the design of the polling mechanism: (a) respondents with smaller positions S_i are more truthful, (b) if the cost of misreporting c_i were higher then more truthful reports are obtained,[5] and (c) polling more respondents reduces the bias.

3.1.1 Cartel formation

In an environment where a *fixed* set of respondents are sampled, manipulative cartels can come about. There is a game–theoretic aspect in understanding the information revelation of respondents, especially in the context of the restricted participation in many distributed dealer markets which leads to a small community of dealers dominating the market.

These concerns have been present from the earliest days of cash settlement using reference rates. Cornell (1997) studies an episode where respondents may have colluded in providing false information close to the expiration date of a futures contract.

We will analyse cartels in one highly simplified setting, where members of a single manipulative cartel intentionally introduce a constant bias in the data that they report. Such a cartel is characterised by two numbers: S, the size of the cartel, and $b_i = \Delta$, the fixed bias introduced by each cartel member. Larger cartels and larger distortions are expected to have a bigger impact upon the reported reference rate. If X is uncontaminated data that might have been obtained from dealers, let $X'(S, \Delta)$ denote the contaminated data. The bias associated with a given estimator is $B(\alpha; S, \Delta) = T(\alpha, \alpha; X'(S, \Delta)) - T(\alpha, \alpha; X)$.

With the sample mean, it is easy to show that $B(0; S, \Delta) = S\Delta/N$. Obtaining analytical expressions for the vulnerability of other estimators requires making assumptions about the distribution from which responses are obtained. Hence, Cita and Lien (1992) explore simulations based on three distributions for cash prices: the standard normal, the uniform and the double exponential. Similarly, Berkowitz (1999) shows simulation results based on Normal or t distributions.

3.1.2 Implications for the polling mechanism

In the light of these ideas, several aspects of the polling mechanism can be designed to minimise the extent of manipulation and cartel formation:

N4 If the identities of respondents and the values reported by them are made publicly available, then it reduces the costs of enforcing a cartel. It would be harder for a cartel member to fink if other cartel members could observe the information reported by him. This is related to the standard argument that cartel formation is less important in markets which trade under conditions of anonymity. Hence it would be advisable for the exchange to not reveal the responses. CME and CBOT protect the secrecy of the sampled information; the British Bankers Association does not.[6]

N5 Cita and Lien (1992) observe that polling *brokers*, who act on an agency basis for a variety of users, is intrinsically less vulnerable to abuse than polling *dealers*, who may be expected to have principal positions on the market. This is related to the above expression for the bias b_i; a broker with no principal positions has $S_i = 0$.

N6 Cartel formation is harder to the extent that the pool of respondents that is polled is enlarged, and when the polling randomly selects n respondents on each day from a large pool. This procedure is used at the CME, for instance, where 12 banks are randomly chosen from a pool of 20.

N7 The use of robust location estimators, such as the STM, reduces the impact that outliers have on the reported reference rate, and reduces the incentive for a respondents to misreport rates.

3.2 Bootstrap methods

The sample size used in a typical reference rate is fairly small (see Table II). In these situations, inference procedures based on asymptotic statistics are unsatisfactory. As with Cita and Lien (1997), we use the bootstrap for obtaining non–parametric inference about $T(\alpha, \alpha; X)$.[7]

In our present context, given the small datasets, simple estimators, and modern computers, the availability of computer power is not a constraint in using bootstrapping. Cita and Lien (1997) observe that with 6 observations, the entire bootstrap distribution has just 462 atoms. We find that using 20 data points, 5000 bootstrap replications are completed in 1 second using generic PC hardware.[8]

Cita and Lien (1997) explore the adaptive trimmed mean (ATM) where α is chosen so as to minimise the bootstrapped sample variance $\sigma^2(\alpha)$ of $T(\alpha, \alpha; X)$. A trimmed mean which uses the optimal proportion of trimming, as determined by the bootstrap method, is called the bootstrap adaptive trimmed mean, and the optimal α is denoted as α^*.

4 Empirical evidence

In this paper, we work with the MIBID/MIBOR reference rates, which are calculated by India's National Stock Exchange (NSE). The methodology used by NSE is summarised in Section 4.1. The data consists of the underlying responses for 22 days in June and July 1998. Using this data, we offer evidence about the gains in statistical efficiency owing to the ATM in Section 4.2 and the vulnerability to manipulation in Section 4.3.

4.1 NSE's methodology for MIBID/MIBOR

The methodology used by NSE is consistent with the normative proposals of this paper.

1. All respondents are asked to quote interest rates which reflect prevailing market conditions, for a fixed transaction size of Rs.100 million (see N1).

2. A random subset of 26 respondents is polled (see N6). Convenience sampling determines the number polled on a given date. Restrictions upon borrowing from the call money market in India imply that some dealers report the lending rate but not the borrowing rate; these dealers are not excluded from the polling (see N3). On average, on any given day, 18 dealers report the bid and 21 dealers report the ask. Information about the values supplied by respondents is not publicly disclosed (see N4).

3. Sampling variances associated with $\alpha \in \{0, 1/N, 2/N, 3/N, 4/N\}$ are calculated using bootstrapping in order to choose α^* which is used in the ATM (see N7).

4. The bid, offer, and their standard deviations are publicly reported (see N3).

4.2 Statistical efficiency

[Table 3 about here.]

Table III summarises the reduction in the standard deviation of the estimator obtained with the ATM as opposed to alternative estimators. For example, if the sample mean were used on all datasets, the bootstrapped standard deviation of the sample mean averages to 0.0321 percentage points.The sample mean is 15.9% noisier than the ATM, fixed trimming is inferior with values like 17.0% ($\alpha = 2/N$) or 12.2% ($\alpha = 4/N$). As is well–known, the sample median has the poorest statistical efficiency, with a standard deviation which is 26.4% above the ATM. The gains in efficiency from using the ATM seem to be economically significant.

4.3 Vulnerability to market manipulation

In our data, there is a unique opportunity to measure the vulnerability to manipulation of alternative estimators. This is derived from the fact that in the early weeks in which the NSE MIBID/MIBOR was calculated, there were no financial products which derived their value from the reference rates. Hence we may assume that over this period, dealers had little incentive to falsify their reports. In this case, the data from this period reflects "normal" sampling noise, without attempts by dealers to manipulate the reference rate.

This enables a Monte Carlo simulation to measure the bias introduced by various kinds of manipulative cartels. For bootstrap datasets drawn from the distribution of a dataset X, random cartels of size S are chosen, and the true reports of cartel members are distorted by a factor of Δ. All estimation strategies, including selection of α^* for the ATM, are applied to the distorted datasets, yielding estimates of $B(\alpha)$ for each estimator.

Details of this simulation are found in Appendix A. These results are summarised in Table IV. For example, the last line in this table tells us that when a cartel of 10 dealers colludes to falsify their reports by one percentage point, the sample mean is distorted by 0.523 percentage points while the median is influenced by 0.519 percentage points. The last column shows the improvement that the median obtains over $B(\alpha^*)$.

[Table 4 about here.]

In all cases, the median is the most robust estimator. This evidence is consistent with the analysis of Lien (1989) who focuses upon the median as the most robust estimator.

Our implementation of the ATM involves choosing α^* over the set $\{0, 1/N, 2/N, 3/N, 4/N\}$. Hence, the ATM is expected to be quite vulnerable to cartels of size 5. For example, in our results, when $\Delta = 0.10$, at $S = 4$, the bias with the ATM is 2 basis points. This jumps to 2.6 basis points for cartels of size 5. In the notation of Hampel (1971), the cartel of size 5 is beyond the "breakdown point" of the ATM as defined by us.[9]

[Figure 1 about here.]

A typical set of results is shown in Figure 1, which pertains to cartels of size three. This shows the influence function, as Δ varies from 0 to 1 for the seven alternative estimators, using a fixed cartel size of three. We see that the simple mean (SIM) is the most vulnerable to manipulation, and the median (MED) is the least vulnerable. The STM which drops four observations is robust to cartels of size three, and fares better than the ATM, which is superior to the other estimators.

The last column in Table IV shows the gains obtained by the median as compared with the ATM. These gains are economically significant in situations with moderate–sized cartels who introduce large distortions. For example, at $S = 4$ and $\Delta = 1$, where four dealers collude to falsify their reports by one percentage point, $B(\frac{1}{2})$ is 0.1377 percentage points smaller than $B(\alpha^*)$.

Earlier, in Table III, we noted that the ATM has the best statistical efficiency among alternative estimators. In the simulation reported of Table IV, we see empirical evidence for the extent to which the median has the least vulnerability to manipulative cartels. This evidence illustrates the tradeoffs in choice of reference rate estimator. If an economy has poor law enforcement, and fairly large cartels are likely to attempt fairly significant manipulation, then the sample median is the most robust estimator and is desirable. In an economy where the enforcement against fraud is effective enough to ensure that cartels are small, and cartels only employ low levels of Δ, the ATM has significant appeal. In economies where cartels such as $S = 4$ and $\Delta = 0.5$ are plausible, futures exchanges establishing reference rates should use the median.

5 Conclusion

The transparency of prices and liquidity that is innate in an OELOB market is unattainable on a distributed dealer market. The full $\lambda(q)$ function is unobservable using information from a sample of respondents. The best information that can be obtained using a small sample of respondents is the best–bid and best–ask at a stated transaction size q_0.

In this paper, we have clarified the market microstructure interpretation that underlies this sampling process, where we find that it is meaningful to interpret the respondent as an agent who observes the $\lambda(q)$ function on a distributed dealer market, and reports his assessment of the true b^* and a^* at a stated q_0.

We discussed statistical procedures which could convert the information obtained from respondents into a relatively efficient location estimate, which reduces the vulnerability to market manipulation. We offered empirical evidence which suggests that the ATM yields the best statistical efficiency, and the median the least vulnerability to manipulation. The economic significance of the superiority of the median, with regard to manipulation, is pronounced with medium–sized cartels introducing large distortions into prices. To the extent that this may be considered uncommon, the ATM could be considered the estimator of choice.

We also discussed numerous other aspects of the polling process which could be designed in a way that reduces the vulnerability to market manipulation, reduces sampling noise, and reduces the incentives and ability of respondents to form cartels and manipulate prices. We offered empirical evidence about the tradeoffs between the ATM and the median in respect of statistical efficiency and vulnerability to manipulation.

The normative ideas offered in this paper about how reference rates may be efficiently and robustly computed often differ from existing methods used by well–established reference rates. This work may hence be useful in a variety of situations when economic agents, including futures exchanges, face the need to obtain reference rates from dealer markets.

A Appendix: Simulating the impact of cartels

Let X be a vector of information polled from n dealers. We seek to estimate $B(\alpha; S, \Delta) = T(\alpha, \alpha; X'(S, \Delta)) - T(\alpha, \alpha; X)$, for a given cartel (S, Δ), for a given robust estimator parametrised by α. The simulation proceeds in the following steps:

1. Draw independent samples $Y_1, Y_2, \ldots, Y_T$ of size n by sampling with replacement from X.
2. For a sample Y_t:
   (a) Evaluate $\hat{\mu} = T(\alpha, \alpha; Y_t)$.
   (b) Add Δ to the responses of a randomly chosen subset in Y_t of S members.[10] Let Y_t' be the modified vector obtained.
   (c) Evaluate $\hat{\mu}' = T(\alpha, \alpha; Y_t')$.
   (d) This gives us one observation for $b_t = \mu' - \mu$.
3. This gives T draws from the distribution of $B(\alpha; S, \Delta)$ for the dataset X.

When analysing α^*, the ATM, there is a bootstrap inside a bootstrap in this simulation. We create bootstrap datasets $Y_1, \ldots, Y_T$, and then use the bootstrap in computing the ATM for the datasets Y_t and Y_t'.

The choice of T is based on the statistical efficiency desired. Using $T = 100$ for each of the 44 datasets yields 4400 draws from the distribution of $B()$. This gives high statistical efficiency, with standard errors on the estimated bias ranging from 10^{-5} to 10^{-4}.

Notes

[1]One alternative objective of information gathering could be the last traded price (LTP). On an exchange, exact LTP values are trivially observed. On distributed dealer markets, observing the LTP is not easy, but it does carry the advantage that it is relatively objective and hence (in principle) verifiable.

The LTP suffers from "stale prices": each respondent would typically report a LTP which pertains to a different time point. The resulting estimate would not quite reflect conditions prevalent at the time of sampling. In addition, some LTPs would be near the bid and others would be near the ask; this generates noise in the estimation process. Hence, we focus exclusively upon the objective of measuring the market–wide best buy a^* and sell b^* prices. Here, stale prices are not a concern.

[2]This is from `http://www.bba.org.uk` on the World Wide Web.

[3]For example, mutual funds in India are only permitted to lend into the call money market. A dealer working at a mutual fund would have no incentive to expend search costs to know the full $\lambda(q)$ function.

[4]In the literature, there are some studies of estimators which are not trimmed means. Cita and Lien (1998) explore using non–uniform weights in the trimmed mean, and Cita and Lien (1997) evaluate the role of the Winsorized mean.

[5]A direct way of increasing the cost of misreporting is to have some mechanism of cross–checking information supplied by respondents. One possibility consists of asking respondents for "good quotes" which they would be willing to trade at. It may be possible for the exchange to design a mechanism to test the veracity of these reports, e.g. by randomly arranging to trade against some of them. However, this approach breaks down owing to the conceptual problems that follow from observing "limit orders", described in Section 2.1.

Hence, there is no alternative to falling back upon the second strategy, proposed in Section 2.2, where respondents are asked to report their subjective evaluation of the state of the market. These statements are essentially un-verifiable.

[6]The "Open Bloomberg" database provides daily observations on the 16 underlying rates that are used in computing the BBA's LIBOR.

[7]An appendix in Cita and Lien (1997) describes Efron's bootstrap algorithm (Efron and Gong, 1983; Efron, 1987). The distribution of the trimmed mean with such small samples is pronouncedly non–normal, which reduces the usefulness of the standard deviation as a measure of imprecision. A measure of dispersion such as the interquartile range should be more useful at expressing the sampling noise in $\hat{b}^*$ and $\hat{a}^*$ instead of the standard deviation of the sampling distribution.

[8]This refers to a C program, compiled with `gcc -O2`, running on a Pentium at 300 MHz, running Linux. The source code is available from the author.

[9]The breakdown point is the largest number of inaccurate reports, no matter how extreme, than a robust estimator can handle. Beyond the breakdown point, the bias can be unbounded.

[10]Berkowitz (1999) uses independent random numbers in his model of the process of contamination. We focus on the behaviour of a cartel which explicitly seeks to manipulate prices; so cartel members can be expected to synchronise their mis-reporting, and independence is violated.

References

Berkowitz, J. (1999): ‘Dealer polling with noisy reporting of interest rates’, *Journal of Fixed Income* **9**:47–54.

Cita, J. and Lien, D. (1992): ‘Constructing accurate cash settlement indices: The role of index specifications’, *Journal of Futures Markets* **12**(3):339–360.

Cita, J. and Lien, D. (1997): ‘Estimating cash settlement price: The bootstrap and other estimators’, *Journal of Futures Markets* **17**(6):617–632.

Cita, J. and Lien, D. (1998): Improving the weighting schemes of cash settlement indices, Technical report, Kansas Corporation Commissions and University of Kansas.

Cornell, B. (1997): ‘Cash settlement when the underlying securities are thinly traded: A case study’, *Journal of Futures Markets* **17**(8):855–871.

Efron, B. (1987): *The Bootstrap, the Jacknife and Other Resampling Plans*, SIAM Publications.

Efron, B. and Gong, G. (1983): ‘A leisurely look at the bootstrap, the jacknife, and cross-validation’, *The American Statistician* **37**(1):36–48.

Ensminger, J. (1999): ‘FX and currency option definitions describing disruption events and fall-backs’, *Journal of Financial Engineering* **8**(1):51–71.

Hampel, F. R. (1971): ‘A general qualitative definition of robustness’, *Annals of Mathematics and Statistics* **42**:1887–1896.

Lien, D.-H. D. (1989): ‘Sampled data as a basis of cash settlement price’, *Journal of Futures Markets* **9**(6):583–588.

Figure 1 Influence function for cartels of size 3

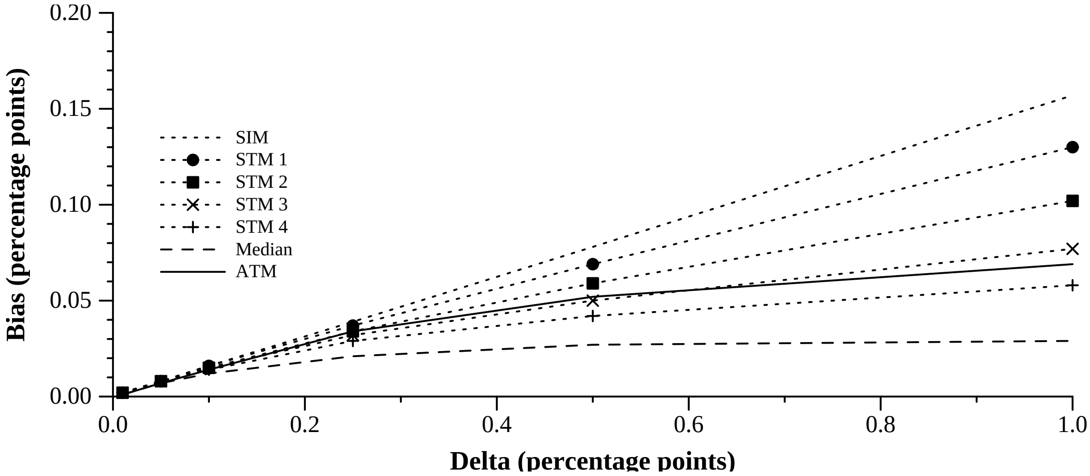

Table I Important special cases of $T(\alpha, \alpha; X)$

Estimator	Interpretation
$T(0, 0; X)$	No trimming, i.e. sample mean (SIM).
$T(\frac{k}{n}, \frac{k}{n}; X)$	Drop the k highest and lowest observations with k fixed, called "symmetric trimmed mean" (STM).
$T(\frac{1}{2}, \frac{1}{2}; X)$	Sample median (MED).
$T(\alpha^*, \alpha^*; X)$	Trimmed mean based on optimal choice of α, the "adaptive trimmed mean" (ATM).

Table II Choice of estimator: examples from real–world markets

Market	Methodology
CME Eurodollar futures contract	STM, $\alpha = \frac{1}{6}$, 12 banks polled randomly from a panel of 20.
CBOT Municipal Bond futures contract	STM, $\alpha = \frac{1}{6}$, 6 brokers polled.
British Bankers Association's LIBOR	STM, $\alpha = \frac{1}{4}$, 8 or 16 banks sampled. Respondent information is public.

Table III Statistical efficiency associated with alternative estimators

Estimator	α	Standard deviation
SIM	0	0.0321
STM		
	1/N	0.0328
	2/N	0.0324
	3/N	0.0318
	4/N	0.0311
MED	1/2	0.0350
ATM	α^*	0.0277

Table IV Vulnerability of estimators to manipulation

	α							Gain over
Cartel	0	1/n	2/n	3/n	4/n	1/2	α^*	ATM
$S = 1$								
$\Delta = 0.01$	0.001	0.000	0.000	0.000	0.000	0.000	0.000	(+0.0001)
$\Delta = 0.05$	0.003	0.002	0.002	0.002	0.002	0.002	0.002	(+0.0004)
$\Delta = 0.10$	0.005	0.005	0.005	0.004	0.004	0.004	0.005	(+0.0011)
$\Delta = 0.25$	0.013	0.010	0.009	0.008	0.007	0.006	0.009	(+0.0035)
$\Delta = 0.50$	0.026	0.018	0.013	0.011	0.010	0.008	0.013	(+0.0054)
$\Delta = 1.00$	0.052	0.029	0.017	0.012	0.010	0.007	0.012	(+0.0047)
$S = 2$								
$\Delta = 0.01$	0.001	0.001	0.001	0.001	0.001	0.001	0.001	(+0.0001)
$\Delta = 0.05$	0.005	0.005	0.005	0.005	0.004	0.004	0.005	(+0.0009)
$\Delta = 0.10$	0.010	0.010	0.009	0.009	0.008	0.007	0.010	(+0.0022)
$\Delta = 0.25$	0.026	0.023	0.020	0.018	0.017	0.014	0.020	(+0.0065)
$\Delta = 0.50$	0.052	0.041	0.032	0.026	0.022	0.016	0.028	(+0.0118)
$\Delta = 1.00$	0.105	0.075	0.050	0.035	0.026	0.016	0.029	(+0.0130)
$S = 3$								
$\Delta = 0.01$	0.002	0.002	0.001	0.001	0.001	0.001	0.001	(+0.0002)
$\Delta = 0.05$	0.008	0.008	0.007	0.007	0.007	0.006	0.007	(+0.0011)
$\Delta = 0.10$	0.016	0.015	0.014	0.014	0.013	0.012	0.015	(+0.0030)
$\Delta = 0.25$	0.039	0.036	0.032	0.029	0.027	0.021	0.033	(+0.0123)
$\Delta = 0.50$	0.078	0.067	0.056	0.046	0.039	0.026	0.049	(+0.0231)
$\Delta = 1.00$	0.157	0.127	0.097	0.072	0.054	0.028	0.065	(+0.0369)
$S = 4$								
$\Delta = 0.01$	0.002	0.002	0.002	0.002	0.002	0.002	0.002	(+0.0002)
$\Delta = 0.05$	0.010	0.010	0.010	0.010	0.010	0.009	0.010	(+0.0016)
$\Delta = 0.10$	0.021	0.020	0.020	0.019	0.018	0.016	0.020	(+0.0038)
$\Delta = 0.25$	0.052	0.049	0.046	0.042	0.039	0.030	0.049	(+0.0182)
$\Delta = 0.50$	0.105	0.094	0.083	0.072	0.062	0.039	0.093	(+0.0535)
$\Delta = 1.00$	0.209	0.183	0.153	0.123	0.097	0.044	0.182	(+0.1377)
$S = 5$								
$\Delta = 0.01$	0.003	0.003	0.003	0.003	0.002	0.002	0.003	(+0.0003)
$\Delta = 0.05$	0.013	0.013	0.013	0.013	0.012	0.011	0.013	(+0.0012)
$\Delta = 0.10$	0.026	0.026	0.025	0.025	0.024	0.022	0.026	(+0.0036)
$\Delta = 0.25$	0.065	0.062	0.059	0.056	0.053	0.042	0.064	(+0.0217)
$\Delta = 0.50$	0.131	0.122	0.112	0.101	0.090	0.058	0.129	(+0.0717)
$\Delta = 1.00$	0.261	0.239	0.213	0.184	0.154	0.070	0.261	(+0.1908)
$S = 10$								
$\Delta = 0.01$	0.005	0.005	0.005	0.005	0.005	0.005	0.005	(+0.0000)
$\Delta = 0.05$	0.026	0.026	0.026	0.026	0.026	0.026	0.026	(+0.0004)
$\Delta = 0.10$	0.053	0.052	0.053	0.053	0.053	0.052	0.052	(+0.0002)
$\Delta = 0.25$	0.131	0.131	0.132	0.132	0.132	0.130	0.131	(+0.0018)
$\Delta = 0.50$	0.261	0.263	0.264	0.266	0.267	0.260	0.263	(+0.0029)
$\Delta = 1.00$	0.523	0.526	0.529	0.534	0.538	0.519	0.524	(+0.0051)