

Dealer Polling in the Presence of Possibly Noisy Reporting

Jeremy Berkowitz
Federal Reserve Board

February 20, 1998
This Draft: August 12, 1998

Address correspondence to:

Jeremy Berkowitz
Trading Risk Analysis
Federal Reserve Board
20th and C Streets, N.W.
Washington, D.C.
jberkowitz@frb.gov
phone (202) 736-5581
fax (202) 452-5296

Abstract: The values of a vast array of financial assets are functions of rates or prices determined in OTC, interbank or other off-exchange markets. In order to price such derivative assets, underlying price indexes must be surveyed and processed. At present, many standard contracts utilize a technique known as trimmed-means to guard against misreporting, whether unintentional or for market manipulation. This paper points out that the polling problem falls within the statistical framework of robust estimation. Intuitive and economically meaningful criteria for choosing among robust valuation procedures are discussed. In particular, the approach taken is to minimize the worst-case scenario arising from a false report. The finite sample performance of the procedures which qualify, the trimmed-mean and the Huber-estimator, are examined in a set of simulation experiments.

Acknowledgements: I gratefully acknowledge helpful input from Greg Duffee and Matt Pritsker. Any remaining errors and inaccuracies are mine. The opinions expressed do not necessarily represent those of the Federal Reserve Board or its staff.

I. Introduction

It would be difficult to overstate the importance of the role played by interbank rates, such as Libor, in the valuation of global financial assets. For example, the vast interest rate swaps market depends critically on the ability to accurately measure the underlying reference rate.¹ The value of these, and many other assets, are functions of privately negotiated or decentralized rates and prices. As a result, many contracts explicitly call for dealer polling.

The importance of the role of dealer polling has previously been recognized in the context of cash settlement of certain futures contracts. In order to allow cash settlement at expiration, the prevailing cash bond price or reference rate must be estimated. For example, Paul (1985) discusses settlement of the International Monetary Market's Eurodollar futures contract, which is based on private negotiated deposit rates. Cita and Lien (1992) provide a detailed description of how the Chicago Board of Trade (CBOT) samples brokers for Municipal Bond Index futures. These authors note that the hedging value of such contracts depends acutely on the accuracy and reliability of the brokers' reports. Inaccurate reporting, whether intentional or accidental, would distort futures prices and discourage its use by market participants.

However, the use of dealer polling is far more extensive than those authors recognized. Dealer polling plays a key role in much wider contexts. Most notably, leading reference rates such as Libor are themselves estimated by polling a small number of dealers.

1. The total notional outstanding value of interest rate derivatives alone was over 26 trillion dollars in 1995. See, the *Central Bank Survey of Foreign Exchange and Derivatives Market Activity* (1996) from the Bank for International Settlements.

Indeed, the generality of dealer polling has been understated even with respect to futures. Throughout the life of a futures contract, exchanges require daily margining adjustments based on fluctuations in the futures value. Since the futures value itself depends on the value of the underlying security, daily dealer polling affects even physically settled futures contracts.

A primary concern of market participants and exchanges is the potential for misreporting and contamination of reported prices. This is made particularly poignant by the widespread practice of polling very small numbers of dealers (generally, between 5 and 20). In order to minimize the effects of inaccurate observations, exchanges and interested parties make use of “trimmed-means” rather than simple averages. This paper emphasizes that trimmed-means can be couched within the statistical framework known as robust estimation. Robustness is the ability to deliver consistent estimates even in the presence of outliers or contamination. While it is known that trimmed-means ameliorate the effects of misreporting, the extensive theoretical insights formulated in the robust estimation literature have not been brought to bear on the problem.

The three-month U.S. dollar Libor rate provides a leading example of the importance of contamination. Every day, the British Bankers’ Association (BBA) polls 16 large banks, who each report 11:00 o’clock Libor rates. The BBA processes these reports into an aggregate Libor rate by discarding the 4 highest, the 4 lowest reports and averaging the remaining 8. Most interest rate swaps, Eurodollar futures and many other contracts rely on

this rate.² To illustrate the importance of contamination, Figure 1 displays the three-month Libor rate that would result had the BBA taken the simple average of the 16 reports from 1995 to 1997.³ Three days in early 1996 appear to have unusual spikes. Inspection of the 16 underlying reports on these days reveals that one bank reported daily rates of 14.00, 20.00 and 20.00 on the respective days. Undoubtedly, these three rates were unintentionally misreported.

Nevertheless, it is far from clear that the trimmed-mean is the optimal solution to the contamination problem. In this paper, both the theoretical and finite-sample performance of alternative robust estimators are compared to that of trimmed-means. The finite-sample performance is gauged through Monte Carlo experiments designed to mimic the settlement price problem.

The remainder of the paper is organized as follows. Section II presents a formal description of truthful and false price reports. Two leading robust procedures are discussed. In section III, the sensitivity of robust estimators to multiple false reports is examined. Section IV presents the results of the Monte Carlo experiments. Section V concludes.

II. Theoretical Framework

To make the notion of robustness concrete, it is useful to begin with a formal definition of contamination. Assume that the true distribution of the cash price is F_0 . The exchange wants to guard against inaccurate reports of unknown form. In other words,

2. The Chicago Mercantile Exchange's Eurodollar futures contracts began using BBA poll results in January 1997 (*CME Consolidated Rulebook* (1997)).

3. The *Open Bloomberg* database provides daily observations on the 16 underlying Libor rates.

contamination is drawn from some other distribution, \mathbf{F} . Denote the probability of false reporting as ϵ , so that $\epsilon=0$ corresponds to no contamination. In this situation, the exchange only knows that observations are drawn from a family of compound distributions,

$$\mathbf{G}_\epsilon = \{\mathbf{g} : (1-\epsilon)\mathbf{F}_0 + \epsilon\mathbf{F}\}, \quad 0 < \epsilon < 1. \quad (1)$$

\mathbf{G}_ϵ is the set of distributions formed by $1-\epsilon$ accurate reports combined with a fraction ϵ of faulty reports.

This sort of definition leads to a copious theory of optimal robust estimators. If the exchange wanted to minimize the variance or bias (or both) assuming perfect reporting, then the simple average is generally optimal.⁴ Unfortunately, the simple average is highly sensitive to false reports or outliers. Robust procedures protect against contamination or market manipulation (at the cost of higher variance). Two particularly successful notions of optimality which have emerged in the robust estimation literature are minimax bias and minimax variance. To possess minimax bias, an estimator must minimize the worst (maximum) possible bias for any type of contamination.

Minimizing the worst possible bias entails finding the estimator $\hat{\mu}$ which solves the following problem,

$$\hat{\mu} = \operatorname{argmin} \max_{\mathbf{G}_\epsilon} \{\hat{\mu} - \mu(\mathbf{F}_0)\}, \quad (2)$$

4. In fact, if the data are normally distributed, the simple mean is the most efficient unbiased estimator. However, this is not the case if the underlying distribution is fat-tailed as, for example, with the logistic distribution.

where $\mu(F_0)$ is the true mean of F_0 . The maximum here ranges over all possible types of misreporting.

Huber (1981) shows that in very general settings, the median is the solution to equation (2).⁵ This result is perhaps not surprising. After all, the median is obviously insensitive to misreporting of any size as long as the ordering of reports is preserved. Moreover, the median is certainly the easiest robust estimator to compute. It does not require the user to specify any further design parameters as, for example, does the truncated-means (which requires specification of the amount of truncation).

It is not clear, however, that minimax bias is what an exchange or market participants will care about. An equally plausible objective would be to minimize the worst possible variance of an estimator.⁶ Estimators with this property are called minimax for variance and solve the problem,

$$\hat{\mu} = \operatorname{argmin}_{\hat{\mu}} \max_{G_0} \{\operatorname{var}(\hat{\mu}, F_0)\}. \quad (3)$$

In general, the variance of the estimator depends on the true distribution as well as the type of contamination.

If minimax variance is the criterion of interest, it is possible to show that there are

5. Specifically, the median is minimax for bias if the true uncontaminated distribution is unimodal and symmetric.

6. Still other criteria have been suggested. For example, Hampel (1968, 1971) proposes minimization of (asymptotic) variance subject to some tolerance level for sensitivity to outliers. Unfortunately, this approach yields estimators which are excessively variable in realistic settings.

three estimators which qualify (see Jureckova and Sen (1996)).⁷ The next section considers the first of these, the trimmed-mean.

A. Trimmed-Means

The best known procedure for processing dealer reports, and the one currently used to calculate Libor, is that of trimmed-means. The trimmed-mean is formed by discarding a given number of the highest and lowest observations and then taking the average of the remaining data. Specifically, the symmetric trimmed mean can be written,

$$\mu_T(\alpha, \alpha; \mathbf{X}) = \frac{1}{n(1-2\alpha)} (x_{1+\alpha n} + \dots + x_{n-\alpha n}) \quad (4)$$

where $\alpha \in \{0, 1, \dots, (n-2)/2\}$ and x_i is the i -th largest observation. For example,

$\mu_T(.1, .1, \mathbf{X})$ eliminates the lowest 10% of the values and the highest 10% of the values. In other words, this trimmed mean is the average of the values between the 10th quantile and 90th quantile. The median is, in fact, a trimmed-mean formed by taking α to be .5 (assuming an odd number of observations for simplicity).

A second closely related estimator, which has been explored by Cita and Lien (1997) in the context of cash futures settlement, is the Winsorized-mean. The Winsorized-mean can be written,

7. This paper only considers two of these. The third approach, known as R-estimation, is not intuitively appealing and tends to be computationally burdensome.

$$\mu_w(\alpha, \alpha; \mathbf{X}) = \frac{1}{n} (r\mathbf{x}_{r+1} + \mathbf{x}_{r+1} + \dots + \mathbf{x}_{n-r} + r\mathbf{x}_{n-r}) \quad (5)$$

where $r = \alpha n$. Rather than discarding the observations outside the α -quantiles, the Winsorized mean repeats the α -quantiles for each outlier.

Unfortunately, both the trimmed-mean and Winsorized-mean tend to be quite sensitive to faulty data in small samples. In fact, it is possible to show that a 10% trimmed-mean based on fewer than 20 observations can break down in the presence of even 2 outliers.⁸ This means there is no bound on the bias or variance of the estimate, given sufficiently bad outliers. Intuitively, if the trimmed mean discards only the single highest and lowest values, 2 overstated (or 2 understated) trades contaminate the estimate.

Although the trimmed-mean is the only approach currently in use, a variety of other robust estimators have been discussed in the statistical literature.⁹ The next subsection discusses an alternative set of robust procedures.

B. Maximum Likelihood

Maximum likelihood estimators include many familiar statistics, including the simple mean, and generally possess good finite-sample properties. The maximum likelihood estimator which minimizes the worst possible variance is known as the Huber-estimate. Again, this is an objective function likely to be of interest to both exchanges and other market

8. The CBOT Municipal Bond Index futures contracts are settled based on 6 observations, while Libor is calculated from 16 underlying rates.

9. Jureckova and Sen (1996) provide a thorough overview of robust estimation.

participants. It is defined as the value, μ , which satisfies the following equation,

$$\sum_{i=1}^N f(x_i - \mu) = 0 \quad (6)$$

where the function $f(\cdot)$ is defined by, $f(x - \mu) = \max(-\beta, \min(\beta, x - \mu))$. β is a cutoff point analogous to the α -level of an L-estimator.¹⁰

To understand the Huber-estimate it is helpful to view it as a close relative of the Winsorized-mean. Given a specified tolerance level, β , the Huber-estimate is an average that first replaces outliers (values greater than β or less than $-\beta$) with the values β and $-\beta$, respectively. The Winsorized-mean is similar but less flexible because it requires β be equal to an observation, x_i . The Huber-estimate allows β to take on *any value* that minimizes the maximum variance.

For the Huber-estimator, the optimal cutoff is a function of the degree of contamination. In the case of normally distributed data, an analytic formula is available for β . It is defined implicitly by the function,

$$\frac{2\phi(\beta)}{\beta} - 2\Phi(-\beta) = \frac{\epsilon}{1-\epsilon}, \quad (7)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard Normal density and distribution functions, respectively.

10. Other M-estimators are defined by replacing the function $f(\cdot)$ in equation (6) with other functions. For example, the simple average is the M-estimator which results from defining $f(\cdot) = 1$.

Of course, in practice the degree of contamination, ϵ , is generally unknown. Nevertheless, equation (7) can be used for rough guidance. For illustration, Table 1 reports optimal cutoff values for some plausible contamination rates. Table 1 also displays the optimal truncated-mean for comparison.¹¹ It is interesting to compare these values to the actual degree of trimming commonly in use. The British Bankers' Association calculates Libor by setting $\alpha = .25$. Table 1 indicates that the BBA's choice is optimal if they believe that there is approximately a 20% chance of misreporting on any given day.¹²

At first blush, a 20% chance of misreporting may seem unrealistically high. However, high trimming rates may reflect a fear of more general notions of "contamination" than simple misreporting. In particular, recent crises in Asia have caused Japanese banks to pay somewhat higher Libor rates. Of the 16 dealers polled by the BBA, three are Japanese banks. From January 1995 to June 1998, Japanese banks have reported three-month Libor rates 6.2 basis points higher than non-Japanese banks. Whether unusual risk premia are properly treated as contamination is not addressed further in this paper. I merely note that dealer polling is significantly more complicated when the underlying market is segmented or differentiated.

11. For the truncated-mean, the optimal degree of trimming, α^* , can be written as a function of the fraction of inaccurate reports. Specifically, again assuming the true distribution is normal, $\alpha^* = (1 - \epsilon)\phi(-\beta) + \epsilon/2$.

12. The CBOT's procedure, setting $\alpha = .16$, corresponds to an expected 10% chance of misreporting.

C. Bootstrapping the Cutoff Point

In practice, it is straightforward to estimate β by bootstrapping, as Cita and Lien (1997) do to select the truncation point for the trimmed-mean. The bootstrap algorithm is designed to provide estimates of the finite-sample variance of the Huber-estimate for alternative cutoff points. To do so, it is necessary to estimate the finite-sample variance of the trimmed-mean and select the truncation point associated with the lowest variance.¹³ Once this is done, the optimal Huber-estimator is easily calculated. The bootstrap algorithm is as follows,

- 1) Resample with replacement, n draws, $\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_n^*$ from the actual observations, $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$.
- 2) Calculate the truncated-mean of the bootstrap data, $\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_n^*$, for a range of cutoff points, $\alpha_1, \alpha_2, \dots, \alpha_k$. The values given in Table 1 may be used for guidance in constructing the range, (α_1, α_k) .
- 3) Repeat steps 1 and 2 a large number of times, R.
- 4) Pick the cutoff point β^* which corresponds to the optimal α^* . This is implemented via Huber's (1981) formula, $\beta^* = -\hat{F}^{-1}(\alpha^*)$, where \hat{F}^{-1} is the inverse of the "least informative" distribution function.¹⁴

To illustrate its use in practice, Figure 2 displays a sample of daily observations on the three-month Libor rate as currently estimated by the BBA, together with the bootstrapped

13. The cutoff point of the Huber-estimator cannot be directly bootstrapped. To see why, note that the finite-sample variance is strictly decreasing in the cutoff point. As a result, the bootstrap will select the smallest possible cutoff point which is, in general, not minimax.

14. The least informative distribution is described in the appendix.

Huber-estimate. The graphs seems to indicate that the two approaches generate similar estimates on most days. However, the BBA's method displays higher peaks in, for example, mid-1995, and lower troughs than the Huber-estimator. In addition, the BBA's trimmed-mean occasionally takes discrete jumps rather than varying smoothly, as occurred in May and June of 1995.

Several variations and refinements to this bootstrap algorithms are also possible. For example, rather than resampling with replacement from the observed data, it is also possible to resample from a normal distribution fit to the data.¹⁵

II. Breakdown Points

In the foregoing discussion, the probability of observing an inaccurate report, ϵ , was treated as a small fixed number. It is of obvious interest, however, to imagine what would happen with a large fraction of inaccurate reports. Can some robust estimators handle more misreporting than others? To answer this and other questions, Hampel (1971) introduced the notion of a breakdown point. The breakdown point defines the largest number of inaccurate reports, *no matter how extreme*, that can be handled. If false reporting exceeds the breakdown point, the bias of the estimate can be infinitely large.

Formally, it is the largest number of false reports, m , for which a limit can be put on the bias,

$$\max_m \{m: \max |\mu - \mu(F_o)| < \infty\}. \quad (8)$$

15. For a detailed discussion of refinements to the standard bootstrap, see Berkowitz and Kilian (1998).

For example, the breakdown point of the truncated mean is $\frac{1}{2}$, half the number of discarded observations (see, Huber (1981)). For the CBOT's Municipal Bond contract, the exchange omits the single largest and smallest values. As a result, the breakdown point for the both contracts is $\frac{1}{6} = 1$ false report. Regardless of how bad is one faulty report, the worst possible bias or variance cannot be very large. However, with *two* dealers misreporting, there is no bound on the bias or variance of the estimate whatsoever.

In contrast, the breakdown point of robust M-estimators is $\frac{n}{2}-1$. For the CBOT Municipal bond contract, with a sample of 6 cash prices, the Huber-estimate can handle two inaccurate reports. With a sample of 16 (as with Libor), the Huber-estimate can handle as many as 7 false reports. In fact, Jureckova and Sen (1996) prove that robust M-estimators have the largest breakdown point possible. This suggests a strong a priori justification for considering Huber-estimates of settlement prices where multiple false reports are plausible.

It is worth emphasizing that this definition of breakdown points is formulated for finite-samples. Huber (1981) also describes an asymptotic formula for the breakdown point. In that case, the breakdown point is the limiting fraction of bad reports that yield a finite (bounded) asymptotic bias. For the purposes of daily dealer polling, however, the sample size is always quite small, suggesting an important role for finite-sample definitions.

III. Simulation Experiment

In this section, a Monte Carlo study is conducted to assess the relative ability of these estimators to deliver accurate underlying rate/price estimates in the presence of outliers. In each simulation, the simple mean, median, trimmed-mean and Huber-estimator are calculated

from a sample size, n , of 16.¹⁶ The number of Monte Carlo trials is also set to 2000.

Both the BBA's trimmed-mean and the Cita and Lien (1997) bootstrap trimmed-mean are considered. For the Huber-estimate, the cutoff point, β , is estimated by the bootstrap described in section 2. The number of bootstrap replications is set to 1000. In order to consider alternative price/rate processes, the Monte Carlo is conducted assuming actual true prices follow Normal or Student-t distributions with mean zero.

The form of the contamination is designed to mimic realistic misreporting. Contaminated prices are drawn from a Normal or Student-t with a random underlying mean, ranging from zero to four. For example, with four contaminated observations, the family of compound distributions is,

$$\mathbf{G}_e = \{g:(12/16)\mathbf{F}_0 + (4/16)\mathbf{F}\}, \quad \mathbf{F} \sim N(\mu, \sigma), \quad 0 < \mu < 4. \quad (10)$$

where σ is the standard deviation of the true distribution (as well as the contaminating distribution). This design reflects the realistic assumption that, on any given day, the severity of contamination is unknown.

Table 2 displays the results when no contamination is present. This is useful both as a starting point and of independent interest. After all, on many days there will be no misreporting. The table indicates the bias, standard deviation, mean-squared-error of each estimator across Monte Carlo trials. In addition, the line labeled "Worst 1%" shows the average bias of the 20 worst estimates (out of 2000) for each estimator. This is a measure of

16. Because the sample size is even, the modified median, the average of the two middle observations, must be used. Setting n to 17 would dispose of this problem but would be unrealistic since the exchanges currently use even-valued sample sizes.

how well the estimator handles worst-case scenarios. The column marked “BBA trimmed-mean” is the trimmed-mean with $\alpha=4/16$, currently used to estimate Libor rates. For the Normal distribution, the simple mean displays the lowest mean-squared error across simulations. This is exactly as the statistical theory predicts. On the other hand, if rates follow a Student-t distribution, the simple mean is no longer the minimum variance estimator. In addition, even with no contamination, the simple mean need not always have the lowest worst-1% of misses. In the Student-t case, the Huber-estimator is best along this dimension.

The results for four false reports are presented in Table 3. In this context, the BBA trimmed mean has lower MSE than the simple mean for both distributions. However, the simpler median performs equally as well. This pattern hold for both underlying distributions. As far as the worst-1%, the median and the status-quo BBA estimator are best followed by the Huber-estimator. Again, this is in fact what the statistical theory indicates.

We expect the median and Huber-estimator to have low worst-case misses, since they are minimax estimators. At the same time, the status-quo procedure is expected to perform well as long as there are no more than 4 false reports. The BBA procedure is to always drop the four highest and four lowest prices, thereby, in all likelihood, eliminating the four false reports. This approach should not be accurate if the breakdown point of 4 is exceeded. On the other hand, the bootstrapped truncated-mean, the median and the bootstrapped Huber-estimator sacrifice some efficiency to guard against larger numbers of false reports.

Table 4 displays the Monte Carlo results based on eight inaccurate reports. Here the pattern of results clearly favors the Huber-estimator. The Huber-estimator has the lowest mean-squared error regardless of the underlying distribution. More importantly, it guards well

against worst-case scenarios. That is, the Huber-estimator has the lowest worst-1%. Not surprisingly, in this context the status-quo approach performs poorly by comparison.

IV. Conclusion

Valuation of securities and obligations that depends on negotiated rates or OTC traded assets relies on dealer polling. In order to cope with unintentional misreporting and to guard against market-manipulation, dealer reports are routinely processed through robust estimation. At present, the truncated-mean is the only such approach in use. However, both theoretical underpinnings and simulation studies suggest that serious consideration should be given to the median and the Huber-estimator. While the Huber-estimator requires specification of a truncation point, the bootstrap provides a straightforward and reliable procedure for its selection.

The reliance on dealer polling is widespread in financial markets. Perhaps most notably, quoted Libor rates are formulated from a small sample of bank reports. Robust estimation may provide market participants, exchanges and regulators interested in minimizing the effects of worst-case scenarios a useful approach to processing dealer polling.

References

- Berkowitz, J. and Kilian, L. (1998). "Recent Developments in Bootstrapping Time Series," Board of Governors, *Finance and Economics Discussion Series*, 1996-45.
- Cita, J. and Lien, D. (1992). "Constructing Accurate Cash Settlement Indices. The Role of Index," *The Journal of Futures Markets*, 12, 339-360.
- Cita, J. and Lien, D. (1997). "Estimating Cash Settlement Price. The Bootstrap and Other Estimators," *The Journal of Futures Markets*, 17, 617-632.
- Hampel, F. R. (1968). *Contributions to the Theory of Robust Estimators*. Dissertation, University of California at Berkeley.
- Hampel, F. R. (1971). "A General Qualitative Definition of Robustness," *Annals of Mathematics and Statistics*, 42, 1887-1896.
- Huber, P. J. (1964). "Robust Estimation of a Location Parameter," *Annals of Mathematical Statistics*, 35, 73-101.
- Huber, P. J. (1981). *Robust Statistics*. New York. Wiley.
- Jureckova, J. and Sen, P. K. (1996). *Robust Statistical Procedures, Asymptotics and Interrelations*. New York. Wiley.
- Paul, A. B. (1985). "The Role of Cash Settlement in Futures Contract Specification," in *Futures Markets. Regulatory Issues*, Peck, A. E. (ed.). Washington, DC. American Enterprise Institute.

Appendix

Least Informative Distributions

Assume that the true underlying rates or prices of interest are drawn from a distribution, \mathbf{F}_0 . They may be falsely reported with probability ϵ , in which case they are drawn from some other distribution, \mathbf{F} . Polling results are therefore drawn from the set of compound distributions, \mathbf{G}_ϵ shown in equation (1). The objective is to formulate a minimax estimator in this context.

Huber (1964) shows that, under general conditions, there is some unique distribution $\hat{\mathbf{F}}$ in \mathbf{G}_ϵ which attains the lowest Fisher information, $\mathbf{I}(\cdot)$, defined by

$$\mathbf{I}(\hat{\mathbf{F}}) = \int \left(\frac{\hat{\mathbf{F}}(\mathbf{x})'}{\hat{\mathbf{F}}(\mathbf{x})} \right)^2 \hat{\mathbf{F}}(\mathbf{x}) d\mathbf{x}.$$

Given a distribution, Fisher information provides an efficiency bound for all possible estimators -- no location estimator can have asymptotic variance lower than $1/\mathbf{I}(\mathbf{f})$ if the data is distributed \mathbf{f} .

Since $\hat{\mathbf{F}}$ has the lowest Fisher information among the possible true but unknown price distributions, it is the least favorable for the purpose of estimation. Of all distributions in the family \mathbf{G}_ϵ , $\hat{\mathbf{F}}$ gives rise to the worst possible *efficient* estimator. This estimator satisfies the minimax variance criterion: it has the lowest variance under the worst-case scenario.

The distribution $\hat{\mathbf{F}}$ is known as the “least informative” distribution. It can be explicitly solved for by taking the Gateaux derivative of Fisher information and setting it equal to zero, $\left(\frac{d}{d\theta} \mathbf{I}(\mathbf{F}_\theta) \right)_{\theta=0} = 0$. In the Normal case, Huber (1981) shows that the least informative distribution is,

$$\frac{1-\epsilon}{\sqrt{2\pi}} e^{-x^2/2}, \quad \text{for } |x| \leq \beta$$

$$\frac{1-\epsilon}{\sqrt{2\pi}} e^{\beta^2/2}, \quad \text{for } |x| > \beta.$$

Of course, for small ϵ (and hence large β), the least informative density is approximately itself a Normal density. In practice, this approximation is reasonably accurate for ϵ less than a few percent.

For practical purposes then, we can set $\hat{F}(\cdot) = \Phi(\cdot)$. Thus, it is straightforward to convert the optimal truncated-mean into an optimal Huber-estimator, according to the formula $\beta^* = -\Phi^{-1}(\alpha^*)$. The Monte Carlo results suggest that the approximation works adequately even for fairly large ϵ .

Table 1
Optimal Huber-Estimates for Given Contamination Rates

probability of misreporting	optimal α	optimal β
0.1%	0.005	2.630
0.2%	0.008	2.435
0.5%	0.018	2.160
1.0%	0.031	1.945
2.0%	0.052	1.717
5.0%	0.102	1.399
10%	0.164	1.140
20%	0.256	0.862

Notes: Optimal α refers to the degree of truncation for the truncated-mean. Optimal β refers to cutoff points for Huber-estimates. These results are exact for observations drawn from contaminated standard Normal distributions and were tabulated in Huber (1981).

Table 2
Alternative Estimates of Location
No Contamination

Estimator:	Mean	Median	BBA Trimmed-mean	Bootstrap Trimmed-mean	Bootstrap Huber-Estimate
Normal					
Bias	-0.0025	0.0044	-0.0002	-0.0012	-0.0512
Std. Dev.	0.2448	0.2965	0.2653	0.2546	0.2601
MSE	0.0599	0.0879	0.0704	0.0648	0.0703
Worst 1%	0.6691	0.8034	0.7247	0.7070	0.7449
Student-t					
Bias	-0.0059	-0.0001	-0.0007	-0.0028	-0.1287
Std. Dev.	0.3588	0.3315	0.3064	0.3153	0.3241
MSE	0.1288	0.1099	0.0939	0.0994	0.1216
Worst 1%	1.2634	1.0844	0.9781	1.0086	0.9637

Notes: The Table compares the Monte Carlo performance of alternative estimators of the average price over 2000 simulations. In each simulation, the sample size is 16 and the true distribution is Normal, or Student-t with 4 degrees of freedom. "Std. Dev." and MSE denote the Monte Carlo standard deviation and mean squared error. Worst 1% is the average bias of the estimator's 20 worst misses.

Table 3
Alternative Estimates of Location
Four Contaminated Reports

Estimator:	Mean	Median	BBA Trimmed-mean	Bootstrap Trimmed-mean	Bootstrap Huber-Estimate
Normal					
Bias	0.4959	0.3092	0.3419	0.4427	0.3522
Std. Dev.	0.3745	0.3465	0.3186	0.3696	0.3819
MSE	0.3861	0.2156	0.2184	0.3326	0.2699
Worst 1%	1.4061	1.2179	1.1596	1.3610	1.2823
Student-t					
Bias	0.5130	0.3355	0.3810	0.4543	0.3305
Std. Dev.	0.4552	0.3909	0.3761	0.4300	0.4840
MSE	0.4704	0.2654	0.2865	0.3913	0.3435
Worst 1%	1.7000	1.4618	1.4528	1.6294	1.6181


Notes: Monte Carlo performance of estimators in the presence of four misreported prices.
See note to Table 1.

Table 4
Alternative Estimates of Location
Eight Contaminated Reports

Estimator:	Mean	Median	BBA Trimmed-mean	Bootstrap Trimmed-mean	Bootstrap Huber-Estimate
Normal					
Bias	1.0026	0.9917	1.0031	1.0037	0.9622
Std. Dev.	0.6231	0.6758	0.6410	0.6256	0.6386
MSE	1.3934	1.4402	1.4172	1.3987	1.3336
Worst 1%	2.3278	2.6350	2.4059	2.3278	2.2987
Student-t					
Bias	1.0198	1.0250	1.0209	1.0175	0.9308
Std. Dev.	0.6583	0.7300	0.6645	0.6496	0.6967
MSE	1.4734	1.5836	1.4838	1.4572	1.3518
Worst 1%	2.6314	3.0825	2.6595	2.5622	2.5521

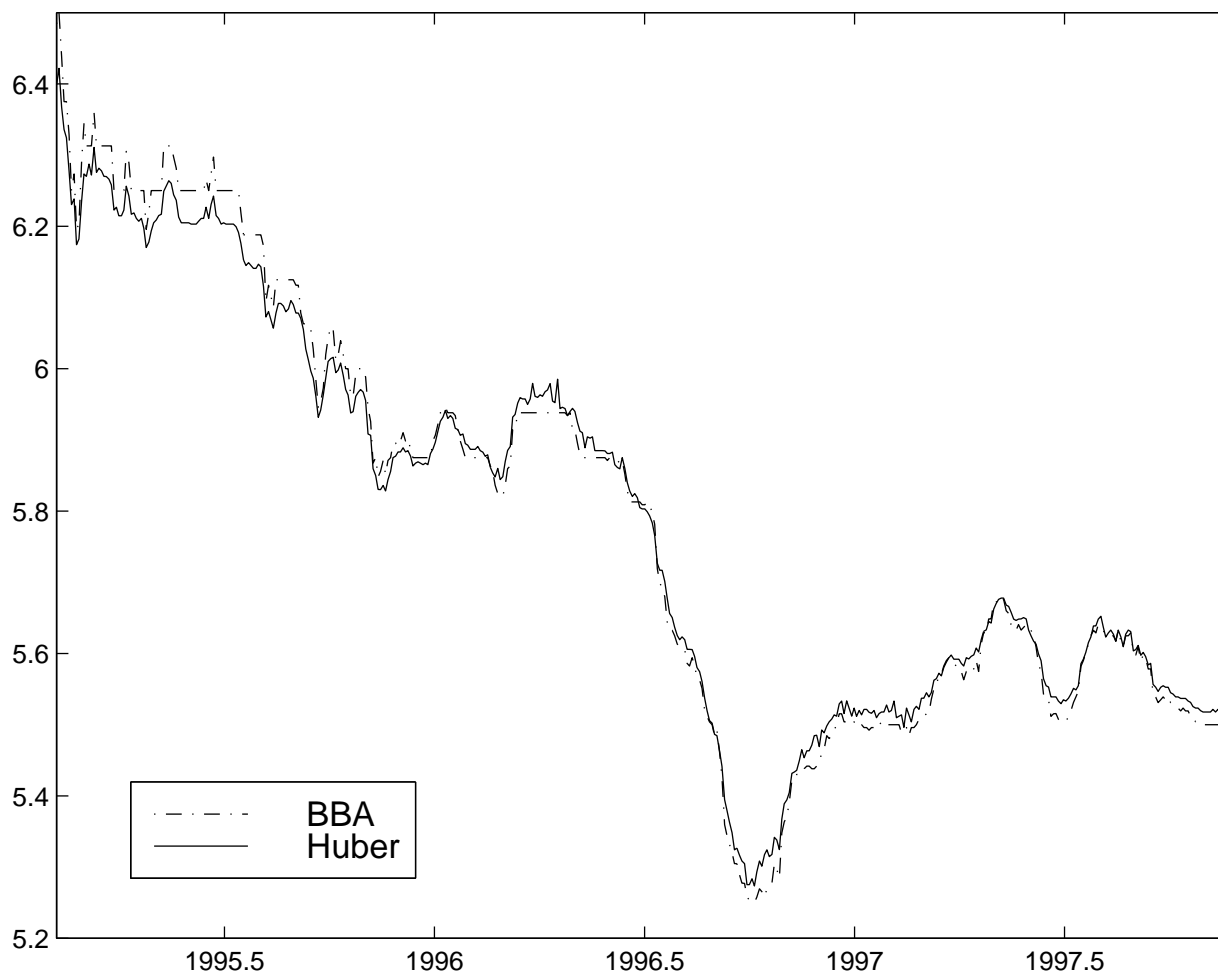
Notes: Monte Carlo performance of estimators in the presence of eight misreported prices.
See note to Table 1.

Figure 1
Average Three-Month Libor Reported
by 16 Large Banks



Notes to Figure: The average rate over 16 large dealers reporting daily three-month Libor to the *British Bankers' Association*. Data are taken from *Open Bloomberg* and extend from January 4, 1995 to June 4, 1998.

Figure 2
Alternative Estimates of Three-Month Libor



Notes to Figure: Three-month Libor fixings calculated two ways. BBA refers to the method currently in use by the British Bankers' Association. Huber denotes the bootstrapped Huber-estimator. See text for details. Data are taken from *Open Bloomberg* and extend from January 4, 1995 to June 4, 1998.