Improved beta estimation on the Johannesburg Stock Exchange: a simulation study

D.C. Bowie & D.J. Bradfield*

Department of Statistical Sciences, University of Cape Town, Private Bag, Rondebosch 7700, Republic of South Africa

In this article we focus on beta estimation in the thinly-traded environment of the Johannesburg Stock Exchange (JSE). We build on existing literature by evaluating a beta estimation procedure known as the trade-to-trade which has not until now been considered in the context of the JSE. We contrast our results with two known estimation procedures, i.e. the Cohen et al. and the traditional ordinary least squares (OLS). The trade-to-trade methodology, the estimator proposed by Cohen et al. and OLS are objectively assessed for shares typical of the JSE on the basis of unbiasedness and efficiency in the controlled environment of a simulation study. The trade-to-trade technique is found to be superior on both counts and is recommended as the appropriate technique for beta estimation on the JSE.

In hierdie artikel fokus ons op die skatting van beta-koëffisiënte van aandele op die Johannesburgse Effektebeurs (JEB) wat 'dun' verhandel is. Ons ondersoek die sogenaamde transaksie-tot-transaksie-beta-skattingsprosedure wat tot dusver nie op die JEB getoets is nie. Hierdie prosedure word met twee ander bekende skattingsprosedures vergelyk, naamlik die prosedures van Cohen et al. en gewone kleinste kwadrate. Die drie metodes word deur middel van 'n simulasiestudie vergelyk in terme van onsydigheid en doeltreffendheid. Dit blyk dat die transaksie-tot-transaksie-metode beter vertoon op albei kriteria en dié metode word dus aanbeveel vir beta-skatting op die JEB.

*Author to whom correspondence should be addressed.

Introduction

It is well known that infrequent/thin trading biases the estimates of financial risk statistics (particularly the beta of a share) when traditional methods of estimation are used. Several researchers have developed correction techniques to compensate for the biases in the estimates of beta caused by thin trading. The more notable contributions in the area of beta estimation are Dimson (1979), Scholes & Williams (1977) and Cohen et al. (1983). The two correction procedures which have been widely used in the literature are the 'trade-to-trade' (T-T) technique (advocated by Dimson & Marsh, 1983) and the Cohen-type estimators, of which the Dimson aggregated coefficients[1] and the Scholes-Williams methods can be regarded as special cases.

On the Johannesburg Stock Exchange (JSE) Bradfield (1990) provides evidence of the extent of the bias in beta estimates caused by thin trading. In a further paper by Bradfield & Barr (1989) the appropriateness of the beta estimation procedure developed by Cohen et al. (1983), which is aimed at correcting the bias caused by thin trading, was assessed for JSE stocks. Bradfield & Barr (1989) demonstrated that the correction procedure yields, on average, less biased estimates of beta in comparison with the traditional ordinary least squares (OLS) procedure.

In this article we extend the results of Bradfield & Barr (1989) by assessing the efficiency of a further beta correction procedure known as the 'trade-to-trade' approach which was advocated by Dimson & Marsh (1983). We demonstrate, using the controlled environment of a simulation study, that the trade-to-trade beta estimator is superior to the Cohen et al. (1983) procedure in correcting for the effects of thin trading for typical JSE stocks.

The two criteria for an estimator of beta which are important discriminators are firstly unbiasedness and, secondly, efficiency. While both the T-T and Cohen-type estimators have been shown by their developers to be unbiased

and consistent, i.e. the distribution of the estimates converges onto the true beta as the sample size increases, their efficiencies in terms of estimation error have not yet been fully considered (to our knowledge). The aim of this article is to quantify the biases which do exist in the OLS estimates of betas for shares typical of the JSE and to compare the efficiencies of the two competing correction procedures, i.e. the T-T and the Cohen estimators.

The article is organized as follows: In the second section we outline details of the simulation study used to compare the beta estimators. In the third section the results of the simulation study are presented and discussed. The fouth section contains some evidence concerning the sensitivity of the simulation analysis to the choice of parameter values used in generating the data. We offer concluding remarks in the last section.

Outline of methodology

The major problem in comparing estimators of the beta coefficient is that the parameter that is being estimated, i.e. the 'true' beta, is unknown. One does not, therefore, have a yardstick with which to compare estimates of beta, when using empirical data.

Further, a simple comparison of empirical estimates of betas provided by the various estimation methods will not give a clear indication of whether or not the particular method is giving unbiased estimates and, more importantly, which method is the more efficient.

A simulation technique is therefore necessary to address the above-mentioned problem. A simulation study allows data with known betas to be constructed so that the various estimators can be assessed objectively. In this article a simulation methodology similar to that employed by several researchers (e.g. Dimson, 1979) in the area of the evaluation of beta estimators is used. Briefly the approach involves simulating returns for a large number of shares at various

levels of thin trading assuming known beta coefficients. Thereafter the three methods of beta estimation (T-T, Cohen and OLS) are contrasted by comparing the estimated parameters with the known parameters based on which the data was simulated.

Details of the simulation methodology are outlined below.

Description of simulation

The simulation of stock market returns involved two stages: firstly, returns for the simulated 'shares' were generated using the market model and secondly, thin trading was introduced into the returns so that we could measure the relative efficiencies of the estimators of beta.

In the first stage we chose the beta for all simulated share returns to be a unity. On this basis we simulated 60 calendar months of daily returns for a market index using the normal[2] distribution. The parameters for the simulation were derived from the monthly returns on the *JSE-Actuaries Overall Index* from its inception in September 1978 to December 1990. The relevant parameters are shown in Table 1. To assess the ability of the estimators to deal with various levels of non-systematic variation we ran the simulations at a range of levels of residual standard deviation as shown in Table 1.

From the simulated market returns we generated daily returns for a large number of shares under the assumption of the market model with a beta of unity and a zero alpha, i.e.:

$$r_t = r_t^m + e_t$$

where:

r_t = the return on the share for day t,

r_t^m = the market return for day t,

e_t = a random normal deviate with standard deviation σ_e.

A beta of unity and alpha of zero meant that a regression of the simulated share returns on the market returns would result in a beta estimate that could easily be interpreted as a proportion of the true beta.

In the second stage of the simulation procedure we introduced the effects of thin trading into the simulated share returns. This was done by allocating a measure of thin trading to each of the simulated shares. The measures of thin trading were defined to be the proportion of days on which no 'trades' occurred for each share. The allocation of the measures of thin trading to the simulated shares was based

on the distribution of thin trading measures among the shares actually listed on the JSE on 30 November 1990 which is shown in Table 2. Briefly, equal numbers of simulated shares were assigned to the ten deciles and for each 'share' a particular level of thin trading was allocated randomly from the range of levels within the assigned decile. Details of the procedure used to introduce the effects of thin trading are outlined in the Appendix.

The simulated daily returns were then aggregated to form a series of simulated monthly returns. The set of monthly returns were then regressed against the monthly market returns to obtain estimates of alpha and beta.

Results of simulations

In order to consider the impact of various levels of unique risk[3] on the estimators of beta the simulation was conducted under the different parameter scenarios listed in Table 1. By repeating the simulation study at various levels of σ_e (the standard deviation of the security's residuals), we were able to compare how the estimators performed across various classes of unique risk, that is from the high unique risk (high σ_e) shares to the lower unique risk (low σ_e) shares.

In the first part of this section we present the results of the simulation at the level of unique risk $\sigma_e = 0.02$. This level of unique risk results in estimates of beta which have efficiencies of the same order of magnitude as the empirical estimates obtained for typical shares on the JSE.

While the efficiency of an estimator depends on the level of unique risk, the average estimated beta resulting from that estimator does not.[4] Hence, in the second part of this section, we combine the estimates of beta from all four simulations, i.e. the average betas across all four levels of unique risk as shown in Table 1. This enables us to calculate a more accurate average estimated beta and thus to obtain a more accurate estimate of the difference between the average estimate and the known beta of unity, i.e. the bias of the estimator.

Thereafter, in the fourth section, we report on the differences between the estimators by interpreting their efficiency at the different levels of thin trading.

Table 1 Parameter combinations used in simulations

Simulation:	1	2	3	4
$E[r_t]$ [a]	.000709	.000709	.000709	.000709
$Std[r_t]$ [b]	.015272	.015272	.015272	.015272
σ_e^c	.02	.01	.005	.0005

a The mean return on the market.

b The standard deviations of the market returns through time.

c The standard deviations of the residuals in the market model.

Table 2 Distribution of q (the percentage of days on which a share does not trade) on the JSE

The figures given represent the average, minimum and maximum values of q for each decile of the shares (ranked by q).

Decile	Average value of q in decile	Minimum value of q in decile	Maximum value of q in decile
1.	1.96	0.00	5.74
2.	11.25	6.28	18.03
3.	24.29	18.31	30.60
4.	36.90	30.60	41.62
5.	47.21	41.90	51.64
6.	57.35	51.64	62.84
7.	67.13	62.84	70.77
8.	75.51	70.77	80.06
9.	83.54	80.06	87.98
10.	94.22	88.25	99.81

The average of the beta estimates as well as the cross-sectional standard deviations (i.e. the standard deviation of the estimated betas) were computed for each decile of thin trading and for each level of unique risk. Since all the data were simulated using a beta of unity, the variance of the estimated betas within each decile provides an estimate of the efficiency of the estimator.

Table 3 Average betas (cross-sectional standard deviations) of simulations with residual standard deviation of 0.02

Decile	OLS	T-T	Cohen
1.	1.00 (0.18)	1.00 (0.18)	1.00 (0.23)
2.	0.98 (0.17)	1.00 (0.17)	1.00 (0.25)
3.	0.98 (0.19)	1.00 (0.19)	1.00 (0.26)
4.	0.99 (0.17)	1.00 (0.17)	1.00 (0.23)
5.	0.95 (0.19)	1.00 (0.19)	0.97 (0.26)
6.	0.94 (0.19)	1.00 (0.18)	1.02 (0.26)
7.	0.89 (0.20)	0.99 (0.20)	0.97 (0.26)
8.	0.79 (0.18)	1.00 (0.18)	0.95 (0.26)
9.	0.77 (0.17)	1.00 (0.17)	0.98 (0.22)
10.	0.49 (0.27)	1.00 (0.26)	0.89 (0.34)
Average	0.88 (0.19)	1.00 (0.19)	0.98 (0.26)

The average beta is computed to be the arithmetic mean of all the estimated betas in the decile.

The cross-sectional standard deviation is the standard deviation of the estimated betas in each decile.

Shown in Table 3 are the results of the simulation of 5 000 shares in each decile with σ_e = 0.02. The cross-sectional standard deviations of beta in decile 1 are equivalent to estimates of the standard errors of the beta estimates[5] and are typical for shares on the JSE (for evidence, see Bradfield & Bowie, 1993).

To assess the information in Table 3 we use a well-known measure of the relative performance of estimators, the mean squared error (MSE). The MSE is a measure of the combined effects of bias (defined to be the difference between the average estimated value and the true beta) and inefficiency, and can be calculated using the relationship (see Appendix):

$$MSE = bias^2 + variance$$

The smaller the value of the MSE for a partipular estimator relative to the MSEs of other estimators the better that estimator is in comparison with the others. The MSE of the OLS estimator in decile 6 is, for example:

$$MSE = (1 - 0.94)^2 + 0.19^2$$
$$= 0.03970$$

The estimated MSEs corresponding to Table 3 are presented in Table 4. It is clear that the trade-to-trade estimator outperforms the other estimators in terms of MSE. It has no

Table 4 MSEs for the estimators based on the estimates of bias and variance in Table 3

Decile	OLS	T-T	Cohen
1.	.0324	.0324	.0529
2.	.0293	.0289	.0625
3.	.0365	.0361	.0676
4.	.0290	.0289	.0529
5.	.0386	.0361	.0685
6.	.0397	.0324	.0680
7.	.0521	.0401	.0685
8.	.0765	.0324	.0701
9.	.0818	.0289	.0488
10.	.3330	.0676	.1277
Average	.0505	.0361	.0680

bias, i.e. the average value of the T-T estimates is equal to the parameter being estimated, and the standard errors of the estimate are no larger than those of either of the other methods (see Table 4).

It is also evident from Table 3 that the OLS estimates have approximately the same cross-sectional standard deviations as the trade-to-trade estimates, but suffer markedly from bias. On the other hand the Cohen estimator has far more cross-sectional variance, but appears to be unbiased for the first nine deciles of thin trading. The betas in the tenth decile are biased downwards and indicate that further lagged coefficients are necessary at very high levels of thin trading.[6]

Estimation of the bias component of the different estimators

The bias component in the estimates of beta is not related to the residual standard deviation since the bias is dependent only on the average length of time between the last trade in an estimation interval and the end of the interval; therefore we computed average betas for each type of estimator across

Table 5 Average simulated betas in each decile across all levels of unique risk

Decile	OLS	T-T	Cohen
1.	0.999	1.000	1.001
2.	0.988	1.000	0.999
3.	0.985	1.000	1.001
4.	0.980	1.000	1.005
5.	0.970	1.000	1.021
6.	0.916	1.000	0.980
7.	0.898	1.000	1.005
8.	0.865	1.000	0.992
9.	0.748	1.000	0.971
10.	0.492	1.000	0.866
Average	0.884	1.000	0.984

all four levels of residual standard deviation. The average betas are shown in Table 5.

From Table 5 we observe that in terms of unbiased estimation, the trade-to-trade estimator is superior. The Cohen estimator shows no clear bias from the expected unity except in decile 10, where we would anticipate having to use more than the one lagged coefficient used in these simulations. The increasing bias with increasing levels of thin trading is evident in the OLS estimates.

Assessing the Impact of different levels of unique risk on the efficiency of the estimators

Given in Table 6 is a comparison of the cross-sectional standard deviations of the estimators at the two extreme levels of unique risk shown in Table 1, i.e. 0.0005 and 0.01 (see Table 6).

From Table 6 it is evident that at the level of unique risk corresponding to the residual standard deviation of 0.01, the trade-to-trade estimator is only marginally more efficient than the OLS estimator, except for the most thinly traded shares. On the other hand the Cohen estimator has substantially higher standard error than the other two estimators.

At the lower residual standard deviation (0.0005), i.e. at low levels of unique risk, the Cohen estimator appears to be marginally more efficient than the OLS estimator (on average), but both of these have far larger standard errors than the trade-to-trade estimates which are more efficient at all levels of thin trading.

Clearly the trade-to-trade estimator is superior both in terms of efficiency and unbiasedness at all levels of thin trading and unique risk.

The trade-to-trade estimator requires information regarding the times of the trades in a share and the value of the market proxy at that time. In some cases the information about the timing is unavailable. It is therefore useful to be

able to decide between the Cohen and OLS methods for estimating a beta without requiring precise information regarding the times of the trades. By comparing the MSEs of the estimators it is possible to ascertain at which levels of thin trading the OLS estimator is superior to the Cohen estimator and vice versa.

To establish the level of thin trading at which the Cohen estimator becomes superior to the OLS estimator we find the OLS estimate of beta at which the MSEs of the two estimators are equal. To do this we denote the OLS estimate of beta at which the MSE of OLS is the same as the MSE of the Cohen estimator by b. From Table 5 we can reasonably assume that the Cohen estimator is unbiased and so the MSE for the Cohen estimator is simply found by squaring the standard error of the estimate. We take the standard error for the Cohen estimator to be 0.26 and the standard error for the OLS estimator to be 0.19 (which correspond to the average cross-sectional standard deviations at a level of unique risk equal to 0.02). The expressions for the respective MSEs are:

$$\text{MSE (OLS)} = (1 - b)^2 + 0.192^2, \text{ and}$$
$$\text{MSE (Cohen)} = 0.26^2$$

Equating the two expressions and solving for b yields 0.82 as the value of the OLS estimate of beta for which the two MSEs are equal. Thus we have established that for about the first seven or eight deciles when the OLS estimates of beta are greater than 0.82 (see Table 5) the OLS technique yields 'better' estimates than the Cohen estimator in terms of a smaller MSE.

Conclusion

The results of this simulation study show conclusively that the trade-to-trade approach to beta estimation is both the more efficient and consistent of the well-known thin-trading correction procedures. This method of estimation does however require extensive data (at least data on a daily basis) and a well-traded and continuously updated market index. However, we did find that the Cohen estimator does correct for the biases induced by thin trading, but at the expense of efficiency. At most levels of thin trading, the OLS estimator outperforms the Cohen estimator in terms of its mean squared error of the estimate.

The Cohen estimator performs relatively more efficiently when there is less non-systematic randomness, i.e. at low levels of unique risk. However, a disadvantage of the Cohen estimator is that the determination of the appropriate number of lags (and leads) is problematic and hence it is not a technique which can be easily automated for all shares on the JSE.

The simulation study provides insights into the effectiveness of the various ways of estimating a security beta. It is evident that thin trading severely biases ordinary least squares estimates of beta and a correction technique should therefore be used. On the basis of the two statistical discriminators, i.e. unbiasedness and efficiency, the trade-to-trade approach to estimating beta coefficients was found to be superior to the Cohen-type estimator. We therefore recommend its use in estimating risk statistics on the JSE.

Table 6 Cross-sectional standard deviations of betas at different levels of unique risk

Method	OLS		T-T		Cohen	
Unique risk	0.01	0.0005	0.01	0.0005	0.01	0.0005
Decile						
1.	0.0910	0.0074	0.0907	0.0038	0.1385	0.0089
2.	0.0894	0.0189	0.0884	0.0040	0.1344	0.0213
3.	0.0916	0.0261	0.0915	0.0053	0.1315	0.0280
4.	0.0880	0.0280	0.0847	0.0042	0.1253	0.0265
5.	0.1042	0.0341	0.0936	0.0044	0.1331	0.0378
6.	0.1009	0.0437	0.0927	0.0045	0.1414	0.0427
7.	0.1103	0.0536	0.0966	0.0043	0.1481	0.0441
8.	0.1103	0.0623	0.0932	0.0047	0.1279	0.0560
9.	0.1345	0.0910	0.1023	0.0050	0.1690	0.0789
10.	0.2030	0.1732	0.1243	0.0061	0.2223	0.1765
Average	0.1171	0.0707	0.0964	0.0047	0.1497	0.0691

Acknowledgement

Notes

1. Cohen et al. (1983) point out a small inconsistency in Dimson's (1979) derivation of the aggregated coefficients method and present a revised estimator.
2. The use of the normal distribution was based on convenience, comparability with other studies (e.g. Dimson) and justified by results reported in the Appendix.
3. Bradfield & Barr (1989) ascertained that the efficiency of the Cohen-type estimators was affected by the level of non-systematic (unique) risk even though, on average, the bias in the estimate of beta caused by the thin trading had been removed.
4. See Dimson & Marsh (1983).
5. In the deciles which represent the shares which are the most thinly traded the biases of the estimators of beta are themselves subject to an additional non-negligible variation. The variation arises because the bias is very sensitive to differences in the level of thin trading at the levels represented by those deciles. In the Appendix we give some indication of the magnitude of the increased variation. The inflation of the variance will not, however, affect comparisons between the estimators. For the other deciles the inflation of the variance is negligible because the bias is not as sensitive to changes in the level of thin trading, i.e. the bias is very nearly the same for all levels of thin trading within the range of thin trading represented by the decile.
6. For the Cohen estimator no leading coefficient was used because the market proxy was constructed to be completely efficient, and only one lag was used.
7. We generated 20 000 uniform deviates between the bounds of each decile using the SYSTAT package, and computed q^k and μ_1 for each generated q and then calculated the variance of $(1 - \mu_1)(1 - q^k)$.

References

Bradfield, D.J. & Barr, G.D.I. 1989. 'Risk estimation in the thinly traded JSE environment', *S.A.J. Bus.Manage.*, Vol. 20: 169–173.

Bradfield, D.J. & Bowie, D.C. 1993. May editing of *Financial Risk Service*. Cape Town: University of Cape Town.

Bradfield, D.J. 1990. 'A note on the estimation problems caused by thin trading on the JSE', *De Ratione*, Vol. 3: 22–25.

Cohen, K.J., Hawawini, G.A., Maier, S.F., Schwartz, R.A. & Whitcomb, D.K. 1983. 'Friction in the trading process and the estimation of systematic risk', *Journal of Financial Economics*, Vol. 12: 263–278.

Dimson, E. 1979. 'Risk measurement when shares are subject to infrequent trading', *Journal of Financial Economics*, Vol. 7: 197–226.

Dimson, E. & Marsh, P.R. 1983. 'The stability of UK risk measures and the problem of thin trading', *Journal of Finance*, Vol. 38: 753–782.

Scholes, M. & Williams, L. 1971. 'Estimating beta from non-synchronous data', *Journal of Financial Economics*, Vol. 5: 309–327.

Appendix

A. Procedure for introducing thin trading into daily returns:

The simulated returns are generated using the market model which is usually written in the form:

$$r_t = \alpha + \beta\, r_t^m + e_t$$

The parameters necessary for generating the returns are the alpha and beta of the market model, as well as the return on market index (denoted r_t^m). For ease of comparison of the resultant estimators we chose alpha and beta to be fixed values, namely, zero and unity, respectively.

For each share where q is the thin-trading parameter representing the probability that the share does not trade in a day and r_t is the true, underlying daily return on the share on day t:

a. set $z = 0$

b. set $t = 0$

c. $t = t + 1$

d. $z = z + r_t$

e. generate u, a uniform random number between 0 and 1

f. if $u > q$ then share is deemed to have traded
    $r_t^o = z$; allocate a value to the observed return
    $z = 0$
    else share has not traded and therefore
    $r_t^o = 0$; zero observed return

g. goto c.

B. Evidence for use of normal distribution

In the simulation the normal distribution was used to simulate the daily returns which were then aggregated to form simulated observed monthly returns. In order to ascertain the validity of this assumption, we calculated a few summary statistics for both daily and monthly returns of the *JSE Actuaries Overall Index*. The skewness, kurtosis and the Wilk-Shapiro statistic (a test for normality) were computed (see Table 7).

The results indicate that the assumption of normality is a reasonable one for the monthly returns (in the index), but not for the daily data where even when outlying returns are excluded, the measures of peakedness (kurtosis) reveal that the tail areas are heavier than expected. The distribution of the residual component is not tested here.

C. Inflation of cross-sectional standard deviation for OLS estimator

We can write the average beta obtained in a decile as

$$\hat{\beta} = plim\, \hat{\beta}_q + e^\beta$$

and the variance of this average will depend on not only the variance of e^β but also the variance of $plim\,\beta_q$ with respect to q which is distributed uniformly across the decile. (The variance of the error term is also dependent on q and hence will also have its estimated variance inflated by the variability of the q in the decile. The effect of the inflation of the standard error is not very large in comparison with large

Table 7 Measures of non-normality for the *JSE-Actuaries Overall Index*

Frequency	No. obs	Skewness[1]	Kurtosis[2]	W-Stat[3]
Daily (all)	1054	-1.95	17.64	0.866
		(-26.574)	(123.010)	(0.0000)
Daily[4] (red)	1047	-0.04	4.54	0.9538
		(0.58)	(30.02)	(0.0000)
Monthly	158	-0.60	1.29	0.9794
		(-3.08)	(3.33)	(0.3566)

1. Normal distribution has skewness coefficient of 0. (t-statistic given in parenthesis.)

2. Normal distribution has kurtosis coefficient of 0. (t-statistic given in parenthesis.)

3. Normal distribution has W-Stat of 1. (Significance level given in parenthesis; a low significance level indicates departure from normality.)

4. The reduced daily set comprises all the daily returns except for 7 outlying returns which were less than -0.06.

standard errors caused by the residual variance at typical levels to be found on the JSE.) It can be shown (see, for example, Dimson & Marsh, 1983) that

$$E[\hat{\beta}] = (1 - q^k)(1 - \mu_t)\beta$$

where k is the number of days in the estimation interval (20 in this simulation case) and μ_t is average proportion of the interval between the last trade in the interval and the end of the interval (in those intervals in which trades occur). On the basis of the simulation (i.e. that the trades occur at random and with constant probability on each day) this parameter can be expressed in terms of q:

$$\mu_t = \left[\frac{q(1-q^{k-1})}{1-q \ (1-q^k)}\right] - (k-1)\left[\frac{q^k}{(1-q^k)}\right] \div k$$

The variance of the first term $(1-q^k)$ in the expression for the probability limit of the estimated beta is quite easily obtainable as follows:

If $Z = q^k$ then $Var[Z] = Var[1 - q^k]$.

Since q is distributed $U(a,b)$, i.e. uniformly between the bounds in the decile a and b, Z will have density function

Table 8 Expected cross-sectional variance for OLS betas with residual standard deviation of 0.01

Decile	Standard error (based on T-T errors)	Variance in bias	Theoretical cross-sectional std dev. of OLS estimates
6.	0.0927	0.0001	0.0932
7.	0.0966	0.0001	0.0971
8.	0.0932	0.0005	0.0958
9.	0.1023	0.0016	0.1098
10	0.1243	0.0185	0.1843

$$f_t(z) = \frac{1}{b-a} \ \frac{1}{k} \ z^{\frac{1}{k}-1} \quad \text{for } a^k \leq z \leq b^k$$

and the expectation and variance can be derived as:

$$E[Z] = \frac{1}{(b-a)(k+1)}[b^{k+1} - a^{k+1}]$$

$$Var[Z] = \frac{1}{(b-a)(2k+1)}[b^{2k+1} - a^{2k+1}] - E^2[Z]$$

yielding, for example when $a = 0.8798$ and $b = 0.98$ (decile 10's bounds), $Var[Z] = 0.02737$.

The variance of μ_t is not easy to write in a closed form and the distribution of the product $(1 - q^k)(1 - \mu_t)$ will be even more difficult to ascertain. We therefore resorted to Monte Carlo techniques[7] to estimate the variance of the bias term for each decile. The results of the Monte Carlo runs are reported in Table 8.

D. A well-known decomposition of the MSE

If X is an estimator of θ and we denote the expectation function by $E[\cdot]$ then

$$
\begin{aligned}
MSE &= E[(X - \theta)^2] \\
&= E[((X - E[X]) + (E[X] - \theta))^2] \\
&= E[(X - E[X])^2] + (E[X] - \theta)^2 \\
&\quad + 2 \ (E[X] - \theta) \ E[(X - E[X])] \\
&= Var[X] + bias^2
\end{aligned}
$$

where $Var[\cdot]$ is the variance function and *bias* represents the difference between the expected value of the estimator and θ, the value of the parameter being estimated.