Thin-Trading Bias in Beta *v* Estimation Error: a Note*

Piet Sercu†, Martina Vandebroek‡ and Tom Vinaimont§

First draft: July 2004. This draft: February 2006.
Preliminary — comments very welcome

*The authors thank Geert Bekaert and Geert Dhaene for useful suggestions, but accept full responsibility for any remaining errors.

†KU Leuven, Graduate School of Business Studies, Naamsestraat 69, B- 3000 Leuven; Tel: +32 16 326 756; piet.sercu@econ.kuleuven.ac.be.

‡KU Leuven, Graduate School of Business Studies, Naamsestraat 69, B- 3000 Leuven; Tel: +32 16 326 975; martina.vandebroek@econ.kuleuven.ac.be.

§Department of Economics & Finance, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon, Hong Kong; email: eftomv@cityu.edu.hk

Abstract

Two regression coefficients often used in Finance, the Scholes-Williams (1977) quasi-multiperiod "thin-trading" beta and the Hansen-Hodrick (1980) overlapping-periods regression coefficient, can both be written as instrumental-variables estimators. We check the performance of these IV -estimators and the validity of the t-tests in small and medium samples, gauge the robustness of the Scholes-Williams estimator outside its stated assumptions, and report performances relative to the standard OLS and the unit-beta model in a hedge-fund style application.

Keywords: Market Model, thin trading.
JEL-codes: C15, G11.

Thin-Trading Bias in Beta *v* Estimation Error: a Note

Introduction

The "beta" slope coefficient of the market model—the regression of an asset's return onto the market return—is crucial in tests or applications of the CAPM, and the residual of that regression became a popular workhorse in event studies since Fama, Fisher, Jensen and Roll (1973). The standard, in much of the literature, is an OLS beta, which assumes IID idiosyncratic returns. If the beta is to be independent of the observation interval used in the returns, more is needed: market returns must be IID too, and n-period returns must be sums of n one-period returns, a feature that strictly holds only for log-change returns. With IID market returns and errors and additivity over time, of course, also the total stock returns are IID , whatever the holding period. In the rest of the paper we refer to the model with IID and additive returns as the standard model.

This standard model has also been at the basis of beta variants that were proposed to handle measurement errors in returns, caused by thin trading, bid-ask bounce, differential reaction speeds to news, and so on. Its no-autocorrelation assumption echoes the stylized view of the early efficient-markets literature as popularized by *e.g.* Fama (1965). True, observed returns do not fully meet these early efficient-markets assumptions, but the violations are not massive, and many apparent infractions are potentially explained by thin trading.[1] Thus, the standard model is adopted by Scholes and Williams (1977), whose modified beta accounts for non-trading that lasts at most one period. Fowler and Rorke (1983) generalize this to longer no-trading intervals. The Scholes-Williams-Fowler-Rorke (SWFR) solution is an instrumental-variable estimator with, as the instrument, a moving sum of the contemporaneous market return plus L leading and L lagging market returns, L being the longest period of non-trading.

[1]For instance, autocorrelation in market-index returns is exactly what one would expect when not all stocks get quoted every day: the day $t+1$ reported return for stocks that did not trade on day t contains the stock's true day-t return and, therefore, introduces a spurious echo of the day-t market evolution into the reported index return for the next day. For the same reason, actively traded stocks seem to lead the market return by one day, and reported returns on thinly traded stock seem to be related to the once-lagged market. The residual for a thinly traded stock is negatively autocorrelated, when the stock seems to belatedly catch up with the market. And, lastly, the market-model residuals for actively traded stocks are positively autocorrelated, if only because the cross-sectional average autocorrelation is positive.

The SWFR estimator has been used also in different contexts. Apte, Kane and Sercu (1993) consider a setting where PPP holds for traded-goods prices and where CPI inflation sluggishly reacts to traded-goods prices. They show that a SWFR -like estimator is needed to extract the full impact of inflation on the exchange rate. In optimal-hedge problems, the variance-minimizing hedge ratio is likewise found by regressing price changes of the exposed asset on price changes of the hedge instrument (Johnson, 1960, and Stein, 1961). Sercu and Wu (2000) obtain good results from the SWFR estimator when some currencies partly follow others with a lag, as was the case in the Exchange Rate Mechanism.

There are two or three alternative estimators that are not strictly consistent but still reduce the thin-trading bias, and might be useful in reducing other errors-in-data biases too. Dimson (1979) proposes a multiple regression with leading and lagging market returns as additional regressors; his beta is the sum of these multiple coefficients. Another alternative is to measure returns at lower frequencies, using *e.g.* weekly holding-period returns rather than daily ones. Indeed, under the stylized assumptions, a multiperiod true beta is no different from a single-period one, as is easily checked: the covariance in the numerator and the variance in the denominator simply go up by the same multiple, the number of periods. The advantage of working with returns from longer holding periods is that while the noise generated by thin trading is not affected, the true returns become larger, implying a better signal-to-noise ratio (see e.g. Stoll and Whaley, 1990). Similar possible motivators are data problems like bid-ask bounce, reporting lags, or differential adjustment speeds due to differences in trading costs or liquidity: with longer holding periods, the signal-to-noise ratio improves. The cost is that extracting non-overlapping longer holding periods from a data base substantially reduces the number of observation. Hansen and Hodrick (1980), discussing a related problem, point out that the use of overlapping multi-period returns mitigates the problem; they further show that, in such overlapping-return regressions, GLS is not the way to deal with the induced autocorrelation in the errors. OLS , in contrast, remains unbiased, and an asymptotic standard error for the OLS estimator can be used that takes into account the serial autocorrelation.

Most of the literature thus far has focused on bias, but the issue of standard error can be important too. For instance, Brown and Warner (1980) show, in a Monte-carlo experiment about event studies, that using OLS -estimated betas introduces more noise than the errors caused if all betas are postulated to be equal to unity. Thus, more in general, it is possible that a biased estimator is still more powerful than an unbiased one. Accordingly, one innovation in this paper is our discussion of standard error (se) next to bias as an evaluation criterion. A related point is our check whether true standard errors differ dramatically from the asymptotic error margins produced by OLS or IV software. A third innovation is the discussion of overlapping-observation regression in IV form, next to the HH OLS form.

Summing up, in this note we report Monte-Carlo results on the comparative performance of (*i*) one-day OLS , (*ii*) Scholes-Williams-Fowler-Rorke (SWFR), (*iii*) Dimson multiple OLS , (*iv*) overlapping-return regressions, and (*v*) an instrumental-variable variant of the overlapping-return regression. We judge the contending estimators on the basis of bias in the coefficient, estimator standard deviation, and reliability of the reported large-sample standard deviation. We also set up a real-world performance race. Specifically, we consider the problem of a portfolio manager who has a target beta—zero, in our application, like for a hedge fund—and uses empirical betas to try and keep her portfolio market neutral. In that particular application we also let the simple unit-beta model enter the race: it's surely biased, but has zero estimation error. We also test whether the standard deviations of the estimates provide any information useful to the hedge-fund manager.

Our recommendation from the Monte-Carlo experiments is to use the overlapping-returns regression with OLS or IV : its bias is minimal, and its precision beats everything else, even for very active stocks. OLS on non-overlapping daily returns is biased (a known result) but rather precise. The SWFR estimator does best with respect to bias, but its standard error is quite large. Dimson's beta is almost as good as SWFR when judged by bias, but has the advantage of being less imprecise. In all cases, the standard error produced by the standard routines is a reliable guide to the actual one, on average. OLS is quite consistent in this, in the sense that the standard deviation, across samples, of the standard error is smallest. SWFR does worst in this respect, and the Dimson and overlapping-OLS models' performance is in-between.

All this bears on Monte-Carlo sampling. The real-world application is still ongoing; preliminary results support the unit-beta model, but more work is needed,

The contending estimators are lined up in Section 1. Section 2 reports the simulation results. We conclude in section 3.

1 The contenders

The familiar market-model regression is

$$R_{j,t} = \alpha + \beta_j R_{m,t} + \epsilon_{j,t}, \tag{1.1}$$

where $R_{j,t}$ is the simple percentage change, *cum* dividend, in the j-th stock price over period t and $R_{m,t}$ is the return on the market portfolio. The stylized assumptions are that the only correlation among these variables is the contemporaneous one between $R_{j,t}$ and $R_{m,t}$, without any auto- or cross-correlation.

1.1 Scholes-Williams-Fowler-Rorke (SWFR)

Thin trading induces bias in the OLS beta, as illustrated in *e.g.* Dimson (1979), arising from (related) errors in both the regressand and in the regressor.[2] The interactions of these errors mean that for active stocks the beta is biased upward, while for thinly traded stocks the bias is negative, as Scholes and Williams show. Assuming that the duration of inaction never exceeds one period, they then derive a consistent instrumental-variable estimator for the market-model beta, where the instrument is a moving sum of the market returns for days $t-1$, t, and $t+1$. (Since the mis-timing of the day-t return does not affect the sum, the instrument is uncorrelated with the error in the reported market return.) Fowler and Rorke (1983) generalize this solution to longer no-trading intervals by extending to moving-sum window.

Scholes and Williams (1977) and Fowler and Rorke (1983) show that, when prices may really date from up to L periods ago, the downward bias is avoided if beta is estimated as

$$\beta^{SW}_{j,H} = \frac{\text{cov}(R_j, Z^{SW}_H)}{\text{cov}(R_m, Z^{SW}_H)} \text{ with } Z^{SW}_{H,t} = \sum_{l=-L}^{L} R_{m,t+l}. \tag{1.2}$$

The intuition why IV works here is that (i) any auto- and cross-correlations present in actually observed returns are, by assumption, induced by thin trading, and (ii) the $2L+1$-period moving sum of market returns picks up all correlations induced by thin trading. The measured return on stock j has to be linked to the past L true market returns because j's last measured return may contain its bottled-up true returns over the past L periods, which are logically linked to the past L true market returns. But also the leading market returns must be included, because the true market return for a period s can be smeared out over all periods $s, ..., s+L$ if the

[2]On the regressand side, a stock that is not traded on day t reports a return of zero, apparently without relation with the market's movement. Actually, the unobserved true day-t return shows up on day $t+1$ (as part of the return reported for that day), but again seemingly without relation with the true day-t market return. Thus, at least two returns are mis-measured. On the regressor side, thin trading likewise induces two types of error, each induced by the errors in the regressands. First, as yesterday's non- traded stocks catch up today, an echo of yesterday's market return is added into today's reported market return; this error is assumed to be uncorrelated with today's true return. Second, due to non-trading today, part of today's true market return is missing, an error that is negatively correlated with the true market return. It will also be obvious that the thin-trading error in the regressand is positively correlated with the one in the regressor.

market index contains thinly-traded stocks.

The consistency of the SWFR estimator is not the issue, so our Monte-Carlo tests are intended as checks for small-sample unbiasedness and especially for the validity of the standard error, whose consistent estimator is:[3]

$$\text{asymptotic stdev}(\beta^{SW}_{j,H}) = \sqrt{\frac{\sigma^2_\epsilon(1 + 2\sum^{L}_{l=1,L}\rho_{l,j}\rho_{l,Z})}{\sum^{N-L}_{1+L}(R_{m,t} - \overline{R}_m)^2 R^2_{R_M,Z}}} \tag{1.3}$$

where σ^2_ϵ is the residual variance, $\rho_{l,X}$ the lth-order autocorrelation of the returns from asset $X = \{j, m\}$ and $R^2_{R_M,Z}$ the squared correlation between the market return and the instrument.

1.2 The Hansen-Hodrick overlapping-observation regression

Under standard assumptions, in the absence of thin trading the true one- and multi-period betas are the same. For reasons outlined in the introduction, an overlapping-return regression may be preferred,

$$\left(\sum_{l=0}^{H-1} R_{j,t+l}\right) = \alpha + \beta^{HH}_{j,H}\left(\sum_{l=0}^{H-1} R_{m,t+l}\right) + \epsilon_{j,t,H}. \tag{1.4}$$

The overlapping returns obviously generate autocorrelation in the error terms. Hansen and Hodrick (1980) reject GLS as biased, in the case of overlapping returns, and propose OLS with an autocorrelation-consistent standard deviation which is consistently estimated as

$$\Theta = (X'X)^{-1}X'\Omega^{-1}X(X'X)^{-1} \tag{1.5}$$

where X is the $N \times 2$ matrix of observations $\{1, R_{m,t}\}$ and Ω is a band matrix with the estimated variances and autocovariances of the regression errors u:

$$\Omega = \sigma^2_\epsilon \begin{bmatrix} 1 & \rho_1 & ... & \rho_{H-1} & 0 & ... & ... & ... & 0 \\ \rho_1 & 1 & \rho_1 & ... & \rho_{H-1} & 0 & ... & ... & 0 \\ ... & \rho_1 & 1 & \rho_1 & ... & \rho_{H-1} & 0 & ... & 0 \\ \rho_{H-1} & ... & \rho_1 & 1 & \rho_1 & ... & \rho_{H-1} & 0 & 0 \\ 0 & & & & & & & & 0 \\ 0 & 0 & \rho_{H-1} & ... & \rho_1 & 1 & \rho_1 & ... & \rho_{H-1} \\ 0 & ... & 0 & \rho_{H-1} & ... & \rho_1 & 1 & \rho_1 & ... \\ 0 & ... & ... & 0 & \rho_{H-1} & ... & \rho_1 & 1 & \rho_1 \\ 0 & ... & ... & ... & 0 & \rho_{H-1} & ... & \rho_1 & 1 \end{bmatrix}. \tag{1.6}$$

[3] The autocorrelations are new relative to the textbook IV case. See Scholes and Williams for the derivations for $L = 1$. The generalization to $L \geq 2$ follows easily.

1.3 An Instrumental-Variable Overlapping-Return Estimator

We easily derive a consistent IV estimator of the overlapping-return model. We start from the explicit objective that we want to estimate the beta over a horizon of H periods, and accordingly write the H-period price change as the sum of H one-period changes. Next, we use equalities like $\text{cov}(X(+1), Y(+3)) = \text{cov}(X, Y(+2))$; and lastly we regroup and interpret the estimator as an IV one:

$$\begin{aligned}
\beta_{j,H}^{IV} &= \frac{\text{cov}(R_j + R_j(+1) + ... + R_j(+H-1), R_m + R_m(+1) + ...R_m(+H-1))}{\text{var}(R_m + R_m(+1) + ...R_m(+H-1))} \\
&= \frac{\sum_{k=0}^{H-1}\sum_{l=0}^{H-1}\text{cov}(R_j(+k), R_m(+l))}{\sum_{k=0}^{H-1}\sum_{l=0}^{H-1}\text{cov}(R_m(+k), R_m(+l))} \\
&= \frac{\sum_{k=-H+1}^{H-1}(H- \mid k \mid)\text{cov}(R_j, R_m(+k))}{\sum_{k=-H+1}^{H-1}(H- \mid k \mid)\text{cov}(R_j, R_m(+k))} \\
&= \frac{\text{cov}(R_j, Z_H^{IV})}{\text{cov}(R_m, Z_H^{IV}))},
\end{aligned} \tag{1.7}$$

where

$$Z_H^{IV} = \sum_{k=-H+1}^{H-1}(H- \mid k \mid)R_m(+k). \tag{1.8}$$

The asymptotic standard error is the same as for SWFR .

In the next section we compare the performance of these contenders as far as bias is concerned, true standard error, and reliability of the standard error produced by standard software. This last issue is especially relevant when the method is applied to data where, unlike in the standard model, the market factor exhibits genuine autocorrelation.

2 Simulation results

2.1 Monte Carlo results for the standard model

Each simulation each consist of 2000 experiments. One such experiment consists of the following steps. First, an IID market factor f is generated with, zero mean and constant volatility equal to 0.01 *p.d.*, i.e. about 0.15 *p.a.*. This market factor has a fat-tailed Student's distribution with seven df. From this, returns for 2000 assets are generated, all with a unit beta and idiosyncratic noise with 1.5 times the market standard deviation, generating a *p.a* volatility of about 0.30 for an individual stock. There are three classes of trading thinness: 1000 of the stocks trade every day, 500 trade with a probability of 0.9, and 500 trade with probability 0.75.

If a stock is not traded, the recorded return is zero; simultaneously, its true return is added to a buffer until the stock does trade, at which time the cumulative return is recorded. Starting from 2000 initial value weights taken from NYSE in 1992, we then update the value weights. The initial market caps are ranked in descending order, so that large stocks tend to be active ones, and the smallest firms most prone to missing prices. The market return is computed as a value-weighted average of asset returns. For the market-model regressions we generate 250 such "daily" observations for all 2000 stocks, and pick three stock files out of these, one per thinness class, to run the regressions. We then start a new experiment, until we have gone through 2000 of them. For SWFR , L is set at 1, 2, 3, and 4 each side,[4] while for the overlapping-observation regressions H is likewise set at 2, 3, 4, 5 (one week) and 20 (monthly holding periods). For each regression we compute the slopes and theoretical (asymptotic) standard error. For each set of 2000 such computations (per type of regression and thinness class) we then produce the mean and standard deviation.

Predictably, OLS does worst in terms of bias, with beta approximating unity minus the probability of no trade. There is no evidence of any upward bias for the active stocks.[5] Equally predictably, SWFR does a good job in eliminating even severe thin-trading biases like $p = .25$ with modest values like $L = 3$. But Dimson's beta does as well in terms of bias. For both estimators, the reduction in bias upon increasing the window size $2L + 1$ comes quite rapid: with no-trading probabilities not exceeding 0.25, there seems to be no point in going beyond $L = 4$. The overlapping-return regressions, whether OLS or IV , are good at handling moderate thin-trading like $p = .10$; but for the $p = .25$ stocks, convergence toward unity seems to be slow: even monthly returns ($H = 20$) achieve only .97, on average.

Equally interesting are the standard deviations of the betas. There is a good connection between the theoretical standard errors (SE) and the actual one computed from the cross-section of betas. Further experiments, not shown here, reveal that this relationship breaks down in small samples—especially for IV estimators—or when the noise is irrealistically small. OLS is quite precise, despite the slight autocorrelation and heteroscedasticity induced by thin

[4]The chance that a stock with no-trade probability 0.25 would not trade for 3 (4) periods is $0.75^2 0.25^3 = 0.0088$ ($0.75^2 0.25^4 = 0.0022$), so L=3 looks reasonable. For the other models we also used $H = 20$ (monthly returns); but the SWFR equivalent, 19 leads and lags produced gigantic se's.

[5]Scholes and Williams demonstrate the theoretical possibility, and we actually see it in simulations with very small noise (not shown here). The fact that the average beta is always below unity, in our table, is compatible with a value-weighted average beta equal to unity because of the positive covariance between weights and estimated beta: $1 = \sum_j w_j \beta_j = (\sum_j w_j)\overline{\beta} + \sum_j (w_j - \overline{w})(\beta_j - \overline{\beta})$.

Table 1: **Simulation results: base case**

	Mean $\hat{\beta}$			mean SE($\hat{\beta}$)s			stdev across $\hat{\beta}$			stdev of SE($\hat{\beta}$)s		
prob no trade	.00	.10	.25	.00	.10	.25	.00	.10	.25	.00	.10	.25
	One-period regression											
OLS	.99	.89	.74	.10	.10	.10	.09	.11	.12	.008	.008	.009
SWFR ±1	.99	.98	.92	.17	.18	.18	.17	.17	.18	.020	.021	.023
SWFR ±2	.99	.98	.97	.22	.23	.25	.22	.23	.23	.036	.037	.040
SWFR ±3	.99	.99	.98	.27	.28	.30	.27	.28	.28	.055	.058	.062
SWFR ±4	.99	.98	.98	.32	.33	.35	.32	.32	.32	.078	.082	.087
Dimson±1	.99	.98	.92	.14	.14	.15	.14	.14	.15	.012	.012	.013
Dimson±2	.99	.98	.97	.17	.17	.18	.17	.17	.17	.016	.016	.017
Dimson±3	.99	.99	.98	.19	.20	.21	.19	.20	.20	.019	.020	.021
Dimson±4	.98	.99	.98	.22	.23	.24	.22	.22	.22	.023	.024	.026
Dimson±19	.99	.99	.99	.51	.53	.55	.52	.51	.52	.100	.105	.110
	H-period overlapping-returns regression											
IV 2	.99	.93	.83	.12	.12	.13	.12	.12	.13	.010	.010	.011
IV 3	.99	.95	.87	.14	.15	.15	.14	.15	.15	.012	.013	.014
IV 4	.99	.96	.90	.16	.17	.18	.16	.17	.17	.015	.016	.017
IV 5	.99	.96	.92	.18	.19	.20	.18	.18	.18	.018	.018	.020
IV 20	.99	.98	.96	.41	.42	.44	.39	.40	.41	.074	.077	.084
OLS Adj2	.99	.95	.88	.13	.13	.13	.14	.14	.15	.014	.014	.014
OLS Adj3	.99	.96	.90	.15	.15	.16	.16	.16	.16	.018	.018	.019
OLS Adj4	.99	.96	.92	.17	.17	.17	.18	.18	.18	.023	.023	.023
OLS Adj5	.99	.97	.93	.19	.19	.19	.20	.20	.20	.027	.027	.027
OLS Adj20	.99	.98	.96	.33	.33	.33	.38	.39	.39	.089	.091	.089

Key. 250 "daily" returns for 2000 stocks are generated by a unit-beta market model with a thick-tailed IID market factor with stdev 0.01/day plus idiosyncratic noise (R^2=0.3). 1000 of the 2000 stocks trade daily, 500 nine days out of ten, and 500 three days out of four. On no-trade days a zero return is reported, and the invisible price change is cumulated until the first trade. Value-weighted market returns are computed from these 2000 returns, and three individual-stock return series are stored, one for each thinness class. This experiment is repeated 10,000 times. OLS uses daily data. SWFR uses as the instrument a moving-window sum of market returns with $\pm L$ leads/lags. Dimson uses a multiple regression with $\pm L$ leading/lagging market returns as additional regressors; the beta is the sum of these $2L + 1$ coefficients. The overlapping-return IV regression uses $H - 1$ proximity-weighted leading and lagging market returns as the instrument. OLS Adj runs OLS with autocorrelation-adjusted SEs. The IV SEs also account for thin-trading-induced autocorrelation in regressand and regressor.

Table 2: **Simulation results: autocorrelation in R_m**

	Mean $\hat{\beta}$			mean SE($\hat{\beta}$)s			stdev across $\hat{\beta}$			stdev of SE($\hat{\beta}$)s		
prob no trade	.00	.10	.25	.00	.10	.25	.00	.10	.25	.00	.10	.25
	One-period regression											
OLS	.99	.90	.75	.10	.10	.10	.10	.11	.12	.008	.008	.009
SWFR ±1	.99	.98	.93	.16	.17	.18	.16	.16	.17	.020	.021	.022
SWFR ±2	.99	.99	.98	.20	.21	.22	.20	.20	.21	.032	.033	.035
SWFR ±3	.99	.99	.99	.23	.24	.26	.23	.23	.24	.043	.045	.049
SWFR ±4	.99	.99	.99	.26	.27	.28	.26	.26	.26	.055	.058	.061
Dimson±1	.99	.98	.93	.13	.14	.14	.13	.13	.14	.012	.013	.013
Dimson±2	.99	.99	.97	.16	.17	.17	.16	.16	.17	.016	.017	.018
Dimson±3	.99	.99	.98	.18	.19	.20	.18	.18	.19	.021	.021	.022
Dimson±4	.99	.99	.98	.20	.21	.21	.20	.20	.21	.025	.026	.027
Dimson±19	1.00	.99	.99	.35	.36	.38	.35	.35	.35	.084	.087	.090
	H-period overlapping-returns regression											
IV 2	.99	.94	.85	.12	.12	.13	.12	.12	.13	.010	.011	.011
IV 3	.99	.96	.89	.14	.14	.15	.14	.14	.14	.013	.013	.014
IV 4	.99	.96	.92	.16	.16	.17	.15	.16	.16	.016	.016	.017
IV 5	.99	.97	.93	.17	.18	.19	.17	.17	.18	.019	.019	.021
IV 20	1.00	.99	.98	.32	.34	.35	.32	.32	.32	.072	.075	.078
OLS Adj2	.99	.96	.89	.12	.13	.13	.13	.14	.14	.014	.014	.014
OLS Adj3	.99	.96	.92	.14	.14	.15	.15	.15	.16	.019	.019	.019
OLS Adj4	.99	.97	.93	.16	.16	.16	.17	.17	.17	.023	.023	.023
OLS Adj5	.99	.97	.94	.17	.17	.17	.18	.18	.19	.027	.027	.027
OLS Adj20	1.00	.99	.98	.25	.25	.25	.30	.30	.30	.080	.080	.080

Key. 250 "daily" returns for 2000 stocks are generated by a unit-beta market model with a thick-tailed market factor with stdev 0.01/day plus idiosyncratic noise (R^2=0.3). The partial autocorrelation schedule in the market factor starts at $\rho_1 = 0.04$, and linearly falls for higher lags, crossing the zero line at lag 15. The resulting autocorrelation of monthly market factors is 0.16. Otherwise, the procedure is like in Table 1.

trading. But for thinly-traded assets, autocorrelation and heteroscedasticity also make the OLS SEs underestimate the true error margins. SWFR betas are 1.5 to 3 times noisier than OLS ones, a feature easily explained using the asymptotic formula (1.3). Specifically, OLS is identical to IV with the regressor as instrument, so the main difference between the OLS and IV SE's comes from the factor R^2 in the numerator of (1.3), which drops from 1 (OLS) to about 1/3 (IV with $L = 1$) or 1/7 (IV with $L = 3$). These R^2s, in turn, imply that the IV asymptotic errors rise by a factor $\sqrt{3} = 1.7$ for $L = 1$ and $\sqrt{7} = 2.6$ for $L = 3$. The result is that the SWFR estimator is actually the noisiest of the lot, two or three times worse than Dimson's beta. The reason why we show no results for SWFR $L = 19$ (i.e. one month either way) is that the variance simply is absurdly high (five decimals before the dot). Dimson offers much more precision for the same unbiasedness.

In terms of SE, the IV version of overlapping-observation regression does better than SWFR 's IV estimator, for the same reason: putting more weight on the contemporaneous market return, its instrument has higher correlation with the regressor than has the SWFR instrument with the same window size L. The asymptotic SEs for the IV - and OLS versions of overlapping-return regressions are quite similar, which is as expected because the estimators are asymptotically equivalent. In terms of actual estimation error, in contrast, the IV version does better than the OLS one unless the window becomes quite large ($H = 20$, one month). Lastly, the pecking order on the basis of actual standard error is also echoed by the standard deviations of the SEs. That is, when the estimator is more imprecise, also the estimated theoretical SE becomes more sensitive to sample coincidences, with OLS and Dimson doing best, followed by overlapping returns, and finally SWFR .

2.2 Field experiment: setting up market-neutral portfolios

In this section we verify whether any of the differences observed above are noticeable in practice, and (in one experiment) whether the standard deviations of the beta estimates are useful. We consider the problem of a hedge-fund manager who has selected N underpriced stocks and now wants to add positions in N more hedge stocks so as to make the entire portfolio market-neutral.

2.2.1 Two-asset market-neutral portfolios

In an exploratory round we set N equal to unity and use a target beta of zero. This corresponds to a hedge-fund manager who has identified two stocks that seem to be mispriced relative to each other, and wants to go long one an short the other in such a way that the combination is market-neutral. We take all US stocks of a given industry, and rank them by size (market cap). We pick assets ranked 1, 4, 5, 8, 9, 12, ... for the long positions, matching each of them with a size-wise close stock of the same industry for the short positions, notably those ranked 2, 3, 6, 7, 10, 11, etc. of a given industry, and match it with the next one.

For the portfolio to have a zero beta, the weights of the two risky assets have to satisfy

$$\frac{w_l}{w_s} = -\frac{\beta_s}{\beta_l}. \tag{2.9}$$

In a first round we set $w_l = 1$, implying $w_s = -\beta_l/\beta_s$ and a risk-free position $w_0 = 1 - w_l - w_s = \beta_l/\beta_s$. This turned out to be naive: some estimated beta pairs are so egregious, say 2 and 0.01, that hedge ratios of 200 would be implied, generating absurd variances for the "hedged" portfolios. No real-world fund manager would have believed these estimates, though. So we truncate the estimated betas at 0.25 and 4. After this first correction, the result can still depend very much on which stock happens to be stock 1 or stock 2. If, for instance, the betas are 2 and 0.5, the weights would be (1, 4), producing a much higher portfolio variance than if the betas had been the other way around and the weights, accordingly, (1, 0.25). As a compromise we rescaled the first-pass weights by their geometric average. Elementary algebra show that this gives us (with β^{tr} denoting a truncated beta):

$$w_l = \sqrt{\frac{\beta_s^{tr}}{\beta_l^{tr}}};\ \ w_s = -\sqrt{\frac{\beta_l^{tr}}{\beta_s^{tr}}};\ \ w_0 = 1 - w_l - w_s, \tag{2.10}$$

so that the dependence of the portfolio in which stock happened to be the high-beta one disappears. Betas are estimated using two years of daily date. Given the portfolio weights we then hold the two stocks for one month, and note the portfolio returns. We then re-estimate the betas, re-rank the stocks by value, and form new portfolios, etc. This produces 100 non-overlapping out-of-sample tests per pair; and since we have about 2000 US stocks with a 20-year history, we have about 1000 of these two-asset portfolios. We compute, for each beta estimator and portfolio, the time-series variance of the portfolio returns. For these, we report the mean, the median, and the number of times the particular estimator produced the lowest-variance portfolio (averaged across the portfolios). The results for the naive and the adjusted procedures are shown in Table 3.

The unit-beta model does best on all counts. Dimson with 19 lags comes up best quite often too, despite its high median portfolio variances. OLS does a good job too. The overlapping-return models provide low variances but do not come up best very often.

Table 3: Average and Median variances of market-neutral portfolios (single share exposure hedged by single share matching industry and size)

	estimated betas, unit w_l		
estimator	mean var	median var	#times 1st
OLS	0.01043	0.00544	10
Dimson±1	0.01070	0.00548	5
Dimson±2	0.01175	0.00564	4
Dimson±3	0.01156	0.00576	4
Dimson±4	0.01209	0.00589	6
Dimson±19	0.01686	0.00741	11
OLS Adj2	0.01081	0.00547	1
OLS Adj3	0.01102	0.00553	1
OLS Adj4	0.01114	0.00563	1
OLS Adj5	0.01125	0.00568	2
OLS Adj20	0.01274	0.00623	6
IV 2	0.01048	0.00542	3
IV 3	0.01078	0.00548	1
IV 4	0.01100	0.00552	1
IV 5	0.01113	0.00563	1
IV 20	0.01282	0.00627	5
SWFR ±1	0.01113	0.00564	5
SWFR ±2	0.01276	0.00599	5
SWFR ±3	0.01299	0.00628	5
SWFR ±4	0.01280	0.00627	6
$\beta = 1$	0.00928	0.00536	16

2.2.2 Multi-asset portfolios

In a more ambitious version of the problem the manager wants to set up a market-neutral portfolio with minimal SE for the portfolio beta, meaning that she is reluctant to put much money into stocks whose betas are estimated with a wide error margin. Denote V = the (exogenously given) vector of weights for the underpriced stocks, β_v = the vector of their betas, W = the vector of weights to be chosen for the hedge stocks, β_w = the vector of their betas, and Ω = the covariance matrix of the 2N beta estimates. The general problem is to find, among all portfolios that seem to be market-neutral—see the constraint in the Lagrangian, below—the one with the smallest estimation error for the market exposure:

$$\mathrm{Min}_W\ [V', W'] \left[\begin{array}{cc} \Omega_{vv} & \Omega'_{wv} \\ \Omega_{wv} & \Omega_{ww} \end{array}\right] \left[\begin{array}{c} V \\ W \end{array}\right] - \lambda \cdot [\beta'_w W - \beta'_v V], \tag{2.11}$$

but since V is given exogenously, this simplifies to

$$\mathrm{Min}_W\ [2W'\Omega_{wv}V + W'\Omega_{ww}W] - \lambda \cdot [\beta'_w W - \beta'_v V]. \tag{2.12}$$

The general solution is

$$W = \Omega_{ww}^{-1}[\frac{\lambda}{2}\beta_w - \Omega_{wv}V] \tag{2.13}$$

$$\frac{\lambda}{2} = \frac{\beta'_v V + \beta'_w \Omega_{ww}^{-1}\Omega_{wv}V}{\beta'_w \Omega_{ww}^{-1}\beta_w}. \tag{2.14}$$

We consider three special cases, corresponding to increasing degrees of scepticism about the value of the estimates. In the first variant, the manager thinks the information about covariances among estimation errors is no good, since the market factor has already picked up the most important source of covariance among returns and, therefore, among beta estimation errors. In the second one, the managers extends this scepticism to the SEs themselves: all betas are viewed as equally prone to error, whatever the computer output may say about that. In the last variant even the very beta estimates are distrusted and replaced by a unit value, which produces an equally-weighted hedge portfolio:

$$\Omega_{ww} \text{ diagonal and } \Omega_{wv} = 0 \quad : \quad W^* = -\frac{\Omega_{ww}^{-1}\beta_w}{\beta'_w \Omega_{ww}^{-1}\beta_w}\beta'_v V \tag{2.15}$$

$$\Omega_{ww} = \sigma^2 I \text{ and } \Omega_{wv} = 0 \quad : \quad W^{**} = -\frac{\beta_w}{\beta'_w \beta_w}\beta'_v V \tag{2.16}$$

$$\Omega_{ww} = \sigma^2 I\ ,\ \Omega_{wv} = 0\ ,\ \text{all}\beta_j = 1 \quad : \quad W^{***} = -\frac{1}{N}V'e, \tag{2.17}$$

where e is the unit vector, [1,1, ... 1]. To get a standardized situation, we scale the exogenous long positions such that $\beta'_v V = 1$.

Table 4: Average and Median variances of market-neutral portfolios (both exposure and hedge portfolios contain ten shares of same industry and match size)

	estimated betas, unit w_l		
estimator	mean var	median var	#times 1st
OLS	0.00115	0.00080	8.7
Dimson±1	0.00117	0.00084	5.3
Dimson±2	0.00116	0.00084	3.8
Dimson±3	0.00117	0.00083	3.7
Dimson±4	0.00118	0.00082	5.1
Dimson±19	0.00141	0.00098	14.5
OLS Adj2	0.00115	0.00082	0.8
OLS Adj3	0.00115	0.00082	0.6
OLS Adj4	0.00116	0.00082	0.9
OLS Adj5	0.00117	0.00082	1.4
OLS Adj20	0.00126	0.00092	5.9
IV 2	0.00115	0.00081	2.0
IV 3	0.00115	0.00082	0.9
IV 4	0.00115	0.00082	0.7
IV 5	0.00116	0.00082	0.8
IV 20	0.00127	0.00091	5.1
SWFR ±1	0.00116	0.00084	4.9
SWFR ±2	0.00119	0.00083	5.1
SWFR ±3	0.00121	0.00088	5.9
SWFR ±4	0.00124	0.00088	7.2
$\beta = 1$	0.00111	0.00080	16.8

Again, for each beta estimator and portfolio, we compute the time-series variance of the hedged portfolio returns and report the mean, the median, and the number of times the particular estimator produced the lowest-variance portfolio. The results are shown in Table 4.

Compared to the two-asset market hedged portfolios, the difference in average and median variances across models are much smaller. Clearly all methods benefit from the larger portfolio effect. Based on the average and mean Dimson with 2 lags seems to be doing just marginally better than its one lag variant, the unit-beta model and the IV and OLS regressions with two- to four-day overlapping returns. When combining the average and mean with the number of runs a model scores best, the unit-beta model becomes a clear winner. Dimson with 19 lags comes up best quite often too, despite its high median portfolio variances. As in the two-asset tests, while the overlapping-return models provide low variances they do not come up best very often.

3 Conclusion

We evaluate, in the presence of thin trading, the performance and the validity of the t-tests in small and medium samples of two commonly used regression coefficients: the Scholes-Williams (1977) quasi-multiperiod "thin-trading" beta and the Hansen-Hodrick (1980) overlapping-periods regression coefficient. These two instrumental variable estimators are compared with the estimates from ordinary least squares and Dimson's multiple regressions with leads and lags as additional variables.

When we introduce thin-trading in a Monte Carlo simulation, OLS does worst in terms of bias. SWFR 's estimator and Dimson's beta do a good job and manage to eliminate even severe thin-trading biases. The overlapping-return regressions, whether OLS or IV , are good at handling moderate thin-trading, however convergence toward unity appears to be slow for more serious thin trading.

When looking at standard deviations, we find that OLS and Dimson estimators are quite precise. In contrast, SWFR betas are up to 3 times as noisy. In terms of SE, the IV version of overlapping-observation regression does better than SWFR 's IV estimator.

Finally, in a field experiment we verify whether any of the differences observed above are noticeable in practice. Therefore, we consider the problem of a hedge-fund manager who has selected N underpriced stocks and wants to add positions in N more hedge stocks to make the entire portfolio market-neutral. We find that Dimson with 19 lags and the naive unit-beta model come first most often. In terms of low variances, the overlapping return models perform fine, but not much better than the unit-beta model.

References

[1] APTE, P., M. KANE, and P. SERCU, 1994: Evidence of PPP in the medium run, *Journal of International Money and Finance*, 601-622

[2] BROWN, S. and J. WARNER, 1980: Measuring security price performance, *Journal of Financial Economics* 8, 205-258

[3] DIMSON, E., 1979, Risk measurement when shares are subject to infrequent trading, *Journal of Financial Economics* 7, 197-226

[4] FAMA, E.F., 1965, Tomorrow in the New York Stock Exchange, *Journal of Business Studies* 38, 191-225

[5] FOWLER, D.J., and H.C. RORKE, 1983, Risk Measurement when shares are subject to infrequent trading: Comment, *Journal of Financial Economics* 12, 297-283

[6] HANSEN, L. P. and R. J. HODRICK, 1980, Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis, *Journal of Political Economy* **88**(5), 829-853

[7] JOHNSON, L.L., 1960: The theory of hedging and speculation in commodity futures, *Review of Economic Studies* 27, 139-151

[8] SCHOLES, M. AND J. WILLIAMS, 1977: Estimating betas from non-synchronous data, *Journal of Financial Economics* 5, 308-328

[9] STEIN, J.L. 1961, The simultaneous determination of spot and futures prices, *American Economic Review* 51, 1012-1025

[10] STOLL, H. and R. WHALEY, 1990: The dynamics of stock index and stock index futures returns, *Journal of Financial and Quantitative Analysis* 25, 441-368