

SETS, Arbitrage Activity, and Stock Price Dynamics

Nick Taylor*, Dick van Dijk[†], Philip Hans Franses[‡], and André Lucas[§] **

Abstract

This paper provides an empirical description of the relationship between the trading system operated by a stock exchange and the transaction costs faced by heterogeneous investors who use the exchange. The recent introduction of SETS in the London Stock Exchange provides an excellent opportunity to study the impact of an electronic trading system upon transaction costs and the time taken to carry out a trade. Using the cost-of-carry model of futures prices we estimate (non-linearly) the transaction costs and trade speeds faced by arbitrageurs who take advantage of mispricing of FTSE100 futures contracts relative to the spot prices of the stocks that make up the FTSE100 stock index. We divide the sample period into pre-SETS and post-SETS sample periods and conduct a comparative study of arbitrageur behaviour under different trading systems. The results indicate that there has been a significant reduction in the level of transaction costs faced by arbitrageurs and in the degree of transaction cost heterogeneity since the introduction of SETS. Finally, generalised impulse response functions show that both spot and futures prices adjust more quickly in the post-SETS period.

Key Words: Non-linear error-correction, smooth transition, SETS, transaction costs.

JEL classification: G14; C32.

*Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom, email: nick.taylor@wbs.warwick.ac.uk

[†]Tinbergen Institute, Erasmus University Rotterdam, Burg. Oudlaan 50, NL-3062 PA Rotterdam, The Netherlands, email: djvandijk@few.eur.nl

[‡]Rotterdam Institute for Business Economic Studies and Econometric Institute, Erasmus University Rotterdam, Burg. Oudlaan 50, NL-3062 PA Rotterdam, The Netherlands, email: franses@few.eur.nl

[§]Department of Finance and Tinbergen Institute, Vrije Universiteit, De Boelelaan 1105, NL-1981HV Amsterdam, The Netherlands, email: alucas@econ.vu.nl

**All correspondence should be addressed to Nick Taylor. We thank INQUIRE UK for providing financial support and FTSE International for contributing part of the data used in this paper. André Lucas thanks the Dutch Organization for Scientific Research (N.W.O.) for financial support.

1 Introduction

Exchanges throughout the world have introduced (for example, London and Frankfurt) or are about to introduce (for example, Sydney) electronic trading systems. There remains some uncertainty, however, concerning the benefits (or otherwise) of such systems versus traditional trading systems. This paper provides empirical evidence on the cost and efficiency improvements brought about by electronic trading systems. More specifically, the paper will measure the transaction costs associated with arbitrage activity in spot and futures markets in the UK before and after the introduction of an electronic trading system. These measures will be of particular importance to those who design trading systems with the hope of exploiting inefficiencies of asset markets.

On October 20, 1997, the London Stock Exchange introduced a new electronic trading system (SETS). The system enables traders to place buy or sell orders for any FTSE100 shares in an electronic order book. These orders are then automatically matched with other orders placed. Before the introduction of this system orders were advertised on computer terminals but actual trades were carried out over the telephone. Under this old system market-makers would absorb the impact of large trades by putting their own capital at risk. Such generosity was compensated for by large bid-ask spreads. Gemmill (1998) reports a 39 basis point spread for large companies and a 79 basis point spread for small companies before the introduction of SETS. By contrast, the respective spreads after the introduction of SETS were 32 basis points and 53 basis points.

The reduction in average bid-ask spreads should have an effect on all arbitrage activity. The activity examined in this paper concerns those trades that are conducted in order to lock into risk-less profits that arise because of perturbations in the contemporaneous relationship between FTSE100 spot and futures prices. Arbitrage activity involves simultaneous positions in both the spot and futures index. The length of time these positions are held depends upon whether or not it is profitable to unwind the position before the maturity of the contract. Brennan and Schwartz (1988, 1990) thus consider such a position as both an arbitrage position and an option to unwind the position when positive profits can be obtained. As Neal (1992) and Sofianos (1993) find that most arbitrage positions are not held until maturity it follows that the option to unwind must have some positive value. This additional value presumably lowers the absolute value of the bounds outside which it is profitable to trade. Moreover, as the cost of exercising the option is the difference between the buy and sell prices of the security then any decrease in bid-ask spread lowers the cost of unwinding the position and, thus, the arbitrage bounds.

The introduction of SETS offers an opportunity to study how arbitrage activity and stock price dynamics are affected by a change in transaction costs. In this

paper we consider whether the introduction of SETS has changed the trading bounds outside which arbitrage activity trades place.¹ The paper is organised as follows. The next section outlines an economic model of arbitrageur behaviour based on the cost-of-carry model. Section 3 describes the specific econometric model used. Section 4 provides a description of the data used. The penultimate section contains the empirical results while some concluding remarks are given in the final section.

2 The Economic Model

The (contemporaneous) relationship between spot and forward prices can be described by the cost-of-carry model. This model is also capable of describing the relationship between spot and futures prices providing that the term structure of interest rates is flat and constant. Under the no-arbitrage condition with no transaction costs, the model has the following specification

$$F_t = S_t e^{(r-\delta)(T-t)} \quad (1)$$

where F_t is the futures price, S_t is the spot price, r is the risk-free interest rate, δ is the expected dividend yield on the underlying asset, and $(T - t)$ is the time to maturity of the futures contract. If the contract is held to maturity then in the presence of proportional transaction costs, c , arbitrage activity will take place when the following condition holds

$$1 - c > \frac{F_t}{S_t} e^{(r-\delta)(T-t)} > 1 + c \quad (2)$$

As it takes time for arbitrageurs to take appropriate spot and futures positions, this arbitrage opportunity is necessarily lagged by d time periods. Therefore, providing c is small, the above inequality can be expressed in the following (logarithmic) form

$$|z_{t-d}| > c \quad (3)$$

¹SETS only pertains to spot positions in any FTSE100 shares. It does not effect the mechanism by which futures contracts in the FTSE100 index are traded. The exchange within which these contracts are traded (LIFFE) currently operates open outcry trading. This difference in trading mechanisms should not detract from the fact that costs of arbitrage trading are likely to be reduced under SETS. This is because mispricing will lead to simultaneous trading in both spot and futures markets. As such, a decrease in the cost of trading the FTSE100 shares in the spot market will reduce overall transaction costs.

where $z_t = \ln F_t - \ln S_t - (r - \delta)(T - t)$ and is referred to as the basis or the pricing error, and d is the delay inherent in the arbitrage process. As arbitragers are expected to unwind the positions before the maturity of the contract, c represents approximately one half the total round-trip transaction costs incurred by arbitragers (Dwyer, Locke and Yu, 1996).

3 The Econometric Model

Previous empirical studies have concluded that spot and futures stock indices are each non-stationary while the respective basis is stationary (Dwyer, Locke and Yu, 1996, and Martens, Kofman, and Vorst, 1998). This implies that spot and futures prices are cointegrated with a cointegrating vector equal to $(1, -1)$. If series are cointegrated then they necessarily have an error-correction representation (Engle and Granger, 1987). Such an error-correction representation directly links changes in futures and spot prices to deviations from the arbitrage relation (1), i.e., to pricing errors. Equation (3), however, states that arbitrage activity only occurs if it is profitable. Equivalently, arbitrage positions in spot and futures stock markets are taken only when the pricing error is outside a particular bound. Thus, spot and futures prices only adjust to past disequilibria depending on the state or regime of the world one is in. To model this behavior we use a smooth transition error-correction model (STECM). Ignoring lag dependence in differenced series (for simplicity only), the model can be expressed as

$$\Delta f_t = \alpha_f z_{t-d} F(z_{t-d}) + \epsilon_{1t} \quad (4)$$

$$\Delta s_t = \alpha_s z_{t-d} F(z_{t-d}) + \epsilon_{2t} \quad (5)$$

where Δf_t is the differenced logarithmic futures price series, Δs_t is the differenced logarithmic spot price series, ϵ_{1t} and ϵ_{2t} are (possibly) cross correlated iid series, $d = \{1, 2, \dots\}$, and $F(z_{t-d})$ is a continuous transition function bounded between 0 and 1. This model allows for smooth transition between low z_{t-d} dependence ($F(z_{t-d}) = 0$) and high z_{t-d} dependence ($F(z_{t-d}) = 1$). More specifically, the strength of the relationship between Δf_t and z_{t-d} will range from zero to α_f as $F(z_{t-d})$ changes (in a smooth fashion) from 0 to 1. This model is capable of allowing for regime dependent arbitrage as given by equation (3). It is expected that for z_{t-d} around zero the value of the transition function will take values close to zero. When z_{t-d} takes relatively large positive and negative values the transition function should take values close to unity. In allowing for smooth transition in $F(z_{t-d})$ the model can allow for heterogeneity in investors exposure to transaction costs. See Anderson (1997) for an application using US Treasury bill data.

Earlier studies of the relationship between futures prices and spot prices use the

threshold error-correction model (TECM), see for example, Dwyer, Locke and Yu (1996) and Martens, Kofman and Vorst (1998). This model allows a very limited number of different regimes and hence transaction costs. The STECM allows for an infinite number of different regimes. As such, the STECM represents a more realistic representation of the heterogeneity of investors that each face different transaction costs.

The following exponential function is used to obtain a parametric specification of the transition function

$$F(z_{t-d}; \gamma) = 1 - \exp[-\gamma z_{t-d}^2] \quad (6)$$

where $\gamma > 0$. The parameter γ measures the speed of transition from no adjustment ($F(z_{t-d}) = 0$) to full adjustment ($F(z_{t-d}) = 1$). Equivalently, γ measures the degree of heterogeneity in transaction costs. Low γ values imply a wide range of transaction costs faced by investors. By contrast, high γ values imply a more uniform transaction cost structure.

The introduction of SETS should lower the transaction costs faced by all investors. Moreover, small (private) investors are expected to face similar transaction costs to those faced by large (institutional) investors. Such transaction cost homogeneity is conveniently measured by a large γ . It follows that γ should be larger after the introduction of SETS. Moreover, if γ is larger in the post-SETS period then transaction costs must be lower in this period. This is because the transition function equals zero when there is no pricing error ($z_{t-d} = 0$). As such, a large γ value means that the transition function is necessarily above the small γ value transition function. The null hypothesis tested in this paper is that γ takes the same value in the pre-SETS and post-SETS periods.

4 Data

The futures price of the nearest FTSE100 contract is obtained for every transaction carried out. These data were obtained from LIFFE. The contract is changed when the volume of trading in the next nearest contract is greater than the volume of trading in the nearest contract.² To synchronise the futures and spot prices, the futures price series is converted to a price series with a frequency of one minute. As one does not know whether the price is a bid or ask price, the

²The volume cross-over method of changing futures contracts results in one change in the pre-SETS period and no changes in the post-SETS period. The change involves a switch from the September 1997 contract to the December 1997 contract on September 19, 1997. On this date 1,422 September 1997 contracts are traded and 6,132 December 1997 contracts are traded. The post-SETS period futures prices make exclusive use of the March 1998 futures contract.

average of the last two prices is taken as the futures price. The (spot) level of the FTSE100 index was obtained from FTSE International. The trading hours of the futures market and the spot market are, 8.30am to 5.30pm and 8.00am to 4.30pm, respectively. Thus one can obtain overlapping futures and spot data covering the period, 8.30am to 4.30pm. However, since the introduction of SETS it has been noted that spreads are unusually high during the first hour of trading. For this reason only prices observed between 9.00am and 4.30pm are used in the analysis. This results in 451 observations per day. The pre-SETS sample period covers the period September 8, 1997, to October 17, 1997. To allow traders to adapt to the new system, the post-SETS sample period will start on January 5, 1998, and end on February 13, 1998. These sample periods correspond to six weeks of data both before and after the introduction of SETS. Following Dwyer, Locke and Yu (1996) and Martens, Kofman and Vorst (1998), we remove overnight returns. This gives a total of 13,500 ($450 \times 6 \times 5$) one minute frequency returns in each of the sample periods. The analysis is also conducted using five minute frequency returns over the same sample periods.

The pricing error is constructed using the daily demeaned futures and spot prices. This methodology follows Dwyer, Locke and Yu (1996). Subtracting the daily mean from the futures prices ensures that any constant in the logarithmic price due to expected dividends or interest rates is removed. The pricing error is set equal to the difference between the demeaned futures price and the demeaned spot price. Henceforth, the demeaned logarithmic futures and spot prices will be denoted by f_t and s_t , respectively, while the pricing error will be denoted by z_t .

5 Empirical Results

Time series plots of logarithmic futures and spot prices are presented in Figure 1. Sharp changes in these prices occur when the trading day changes. Problems associated with these price discontinuities are avoided in this paper because only intraday returns are considered.

5.1 Testing for Non-stationarity

Augmented Dickey Fuller (ADF) tests are performed on various one minute and five minute frequency series. In each case a constant is included and the lag lengths are selected on the basis of the Schwarz Information Criterion (SIC). The one minute frequency results given in Table I show that futures and spot prices are non-stationary.³ These prices are not the same prices as those plotted

³Similar results hold when five minute frequency data are used and are available upon request.

in Figure 1. The non-stationarity tests are applied to intraday prices. These prices are calculated as follows. First, logarithmic returns are calculated. Second, overnight returns are removed. Third, intraday prices are calculated by numerically integrating the intraday returns. Possible cointegration between these prices is investigated by testing for non-stationarity in the pricing error using the ADF test. Van Dijk and Franses (1998) show that unit root tests such as the ADF test perform well in the presence of non-linearity in the adjustment process. The results indicate that the null hypothesis of non-stationarity can be rejected with a high level of confidence. Therefore, the cointegrating vector $(1, -1)$ provides a combination of non-stationary futures and spot prices that is stationary. As such, these prices have an error-correction representation.

5.2 Testing for Non-linearity

It is possible to test the null hypothesis that returns follow a linear error-correction process against the alternative that returns follow a smooth transition error-correction process. The testing procedure used in this paper is based on Luukkonen, Saikkonen and Teräsvirta (1988), Swanson (1999) and van Dijk and Franses (1998). Observation of the transition function given by (6) shows that the null hypothesis of linearity is equivalent to testing the null hypothesis that γ equals zero. Moreover, if the transition function is replaced by a third-order Taylor approximation then the STECM for futures returns can be expressed as⁴

$$\Delta f_t = \phi' \mathbf{w}_t + \phi'_1 \tilde{\mathbf{w}}_t z_{t-d} + \phi'_2 \tilde{\mathbf{w}}_t z_{t-d}^2 + \phi'_3 \tilde{\mathbf{w}}_t z_{t-d}^3 + \epsilon_t \quad (7)$$

where $\mathbf{w}_t = (1, \tilde{\mathbf{w}}_t)'$, $\tilde{\mathbf{w}}_t = (z_{t-d}, \Delta f_{t-1}, \dots, \Delta f_{t-p}, \Delta s_{t-1}, \dots, \Delta s_{t-p})'$, ϕ is a $(m+1 \times 1)$ vector of coefficients, ϕ_1 , ϕ_2 , and ϕ_3 are $(m \times 1)$ vectors of coefficients, and $m = 2p+1$. The original null hypothesis of linearity, $H_0 : \gamma = 0$, is equivalent to the null hypothesis that all coefficients of the auxiliary regressors, $\tilde{\mathbf{w}}_t z_{t-d}^j$, $j = \{1, 2, 3\}$, equal zero, that is, $H_0 : \phi_1 = \phi_2 = \phi_3 = 0$. An LM-type test is used to test this hypothesis. The one minute frequency results are presented in Table II for $p = 1$ and $d = \{1, 2, 3, 4, 5\}$.⁵ The results indicate that the null hypothesis can be rejected with a high degree of confidence. There is little to choose between each of the d values. In each case there is a clear rejection of the null hypothesis except when $d = 4$ in the post-SETS period.

⁴The same approach is followed for the spot return series.

⁵When the five minute frequency returns are used we restrict the value of d to one. Longer delays are unrealistic, for instance, if $d = 2$ then this implies a delay of ten minutes in the arbitrage process. The results indicate a rejection of the null hypothesis and are available upon request.

5.3 Estimating the Error-Correction Models

Exponential transition function error-correction models are estimated in both the pre-SETS and post-SETS periods. In both cases $p = 1$ and $d = \{1, 2, 3, 4, 5\}$.⁶ The exponential transition function error-correction model is specified as follows⁷

$$\Delta f_t = \boldsymbol{\pi}'_1 \mathbf{w}_t^* + \boldsymbol{\pi}'_2 \tilde{\mathbf{w}}_t^* F(z_{t-d}; \gamma) + \alpha_f z_{t-d} F(z_{t-d}; \gamma) + \eta_t \quad (8)$$

where $F(z_{t-d}; \gamma) = 1 - \exp[-\gamma z_{t-d}^2 / \sigma_{z_{t-d}}^2]$, $\mathbf{w}_t^* = (1, \tilde{\mathbf{w}}_t^*)'$, $\tilde{\mathbf{w}}_t^* = (\Delta f_{t-1}, \Delta s_{t-1})$, $\boldsymbol{\pi}'_1$ is a (3×1) vector of coefficients, $\boldsymbol{\pi}'_2$ is a (2×1) vector of coefficients, α_f is the adjustment coefficient, and $\sigma_{z_{t-d}}^2$ is the variance of the pricing error. Inclusion of $\sigma_{z_{t-d}}^2$ follows Teräsvirta (1994) and enables interpretation of γ in a scale-free environment. As such, comparisons of γ 's over various sample periods is allowed. Both of these models are estimated using NLS. The estimated adjustment and γ coefficients are presented in Table III. The results of ARCH tests performed on the residuals from the estimated models indicated that there is significant heteroscedasticity present. Therefore, heteroscedastic-consistent standard errors are presented in Table III.

When using one minute frequency returns information criteria are minimised when the delay equals five minutes in the pre-SETS period and two minutes in the post-SETS period. This suggests that the trade speed has been reduced since the introduction of SETS. The adjustment coefficients have the expected signs. That is, $\alpha_f < 0$ and $\alpha_s > 0$. Adjustment in the spot equation is considerably larger, in absolute terms, than adjustment in the futures equation during the post-SETS period. This suggests that the introduction of SETS has lowered the relative costs of trading the spot index. Similar results are obtained when five minute frequency returns are used.

5.4 Comparing Transaction Cost Profiles

The results given in Table III indicate that when the exponential transition function is used the degree of transaction cost heterogeneity is greater in the pre-SETS period than in the post-SETS period. That is, $\hat{\gamma}$ is smaller in the former period. As the transition function must take a value of zero when there is no mispricing then this result implies that the transactions costs faced by arbitrageurs in the post-SETS period are smaller than those faced in the pre-SETS period.

The profiles presented in Figure 2 (one minute frequency data) plot the estimated transition function against the pricing error. This figure shows that there

⁶The delay is set equal to one when five minute frequency returns are used.

⁷Spot returns are similarly defined.

is full adjustment outside a narrow range of mispricing in the post-SETS period. By contrast, this range is considerably larger in the pre-SETS period. Moreover, there is a sharper change from no adjustment ($F(z_{t-d}) = 0$) to full adjustment ($F(z_{t-d}) = 1$) in the post-SETS period than in the pre-SETS period.

To formally test this equality of γ values over the two sample periods a simple t-test based on heteroscedastic-consistent standard errors is performed for various delay values. In each case the same delay values are assumed in each period. In addition, the optimal delays, as given by the SIC, are used in each period and the t-statistic is calculated. The one minute frequency results are presented in Table IV. If delays of two, three or five minutes are assumed then transaction costs are significantly lower in the post-SETS period. However, this conclusion cannot be drawn if delays of one or four minutes are assumed or if optimal delays are assumed. When five minute frequency returns are used the t-statistic equals 2.70 thus indicating that transaction costs are significantly lower in the post-SETS period.

5.5 Generalised Impulse Responses

Generalised impulse response functions are calculated using the smooth transition error-correction models estimated in the pre-SETS and post-SETS periods. In both cases, one minute frequency data are used and the delay (d) is selected by the SIC. Shocks equal to -0.4 , -0.35 , -0.3 , \dots , 0.35 , and 0.4 are assumed to affect both spot and futures markets. The effects that these shocks have on subsequent spot and futures returns are measured at various points within the pre-SETS and post-SETS sample periods.⁸ The distribution of these innovations is estimated using a quartic kernel function at various time periods after the shock hits the system.⁹ A uniform distribution taking values between -0.4 and 0.4 (inclusive) is observed when the shock occurs. Subsequent distributions are less uniform and have a smaller range as the effects of the shock gradually disappear. This rate of decay gives an indication of the speed of adjustment in the respective markets. The estimated distributions at two and five minutes after the initial shock occurs are presented in Figures 3 and 4.¹⁰

The introduction of SETS causes more rapid adjustment in both spot and

⁸This estimation process is based on a sub-sample of the pre-SETS and post-SETS periods. The ‘histories’ used in the current context equal the 1st, 101st, 201st, \dots , 13,301st, and 13,401st observations. These histories are selected from a sample consisting of 13,500 observations. Selection of these histories is based on a need to reduce the computation time.

⁹For further details of kernel functions and the optimal bandwidth used see equation 3.31 of Silverman (1986).

¹⁰Similar distributions are estimated at one, three, and four minutes after the initial shock occurs. These figures are available upon request.

futures markets. This can be observed by comparing Panels A and C with B and D, respectively, in Figures 3 and 4. In each case the range of values taken by the innovations is smaller in the post-SETS period than in the pre-SETS period. Moreover, in the pre-SETS period the adjustment in the futures market is faster than adjustment in the spot market. That is, after two and five minutes the range of the innovations in the futures equation is smaller than the range of innovations in the spot equation. By contrast, adjustment is similar in both markets in the post-SETS period.

6 Conclusion

The transaction costs faced by arbitragers trading the FTSE100 spot and futures markets have been significantly reduced since the introduction of SETS. Analysis of generalised impulse response functions leads to two additional findings. First, shocks to the futures and spot markets have less effect in the post-SETS period. Indeed, the effects of such shocks almost disappear after five minutes. Second, the futures market is less affected by shocks than the spot market in the pre-SETS period. However, both markets appear to be equally affected in the post-SETS period. These two findings are consistent with the objectives of SETS. That is, lower cost trading in the spot market.

The results are also encouraging with respect to LIFFE's planned introduction of the Connect trading system covering the FTSE100 futures market in May 1999. Analysis of the impact of this system upon transaction costs will be the subject of future research. In particular, one would expect an increase in the level of adjustment in the futures market equal to that seen in the spot market since the introduction of SETS.

References

- Anderson, H., 1997, Transaction Costs and Non-linear Adjustment towards Equilibrium in the US Treasury Bill Market, *Oxford Bulletin of Economics and Statistics* 59, 465–484.
- Brennan, M. J., and E. Schwartz, 1988, Optimal Arbitrage Strategies Under Basis Variability, *Studies in Banking and Finance* 5, 167–180.
- Brennan, M. J., and E. Schwartz, 1990, Arbitrage in Stock Index Futures, *Journal of Business* 63, S7–S31.
- van Dijk, D., and P. H. Franses, 1998, Non-linear Error-Correction Models for Interest Rates in the Netherlands, in *Non-linear Econometric Modelling* (S. Hylleberg ed.), Cambridge University Press (Cambridge).
- Dwyer, G., Locke, P., and W. Yu, 1996, Index Arbitrage and Non-linear Dynamics between the S&P500 Futures and Cash, *Review of Financial Studies* 9, 301–332.
- Engle, R., and C. Granger, 1987, Cointegration and Error-Correction: Representation, Estimation, and Testing, *Econometrica* 55, 251–276.
- Gemmill, G., 1998, Touched by Success, *Financial News Briefing Notes*, October 5.
- Luukkonen, R., Saikkonen, P., and T. Teräsvirta, 1988, Testing Linearity against Smooth Transition Autoregressive Models, *Biometrika* 75, 491–499.
- Martens, M., Kofman, P., and T. Vorst, 1998, A Threshold Error-Correction Model for Intraday Futures and Index Returns, *Journal of Applied Econometrics* 13, 245–263.
- Neal, R., 1992, Direct Tests of Index Arbitrage Models, (Working Paper) University of Washington.
- Silverman, B., 1986, *Density Estimation for Statistics and Data Analysis*, Chapman and Hall (London).
- Sofianos, G., 1993, Index Arbitrage Profitability, *Journal of Derivatives* 1, 6–20.
- Swanson, N., 1999, Finite Sample Properties of a Simple LM Test for Neglected Non-linearity in Error-Correcting Regression Equations, *Statistica Neerlandica*, forthcoming.

Teräsvirta, T., 1994, Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models, *Journal of the American Statistical Association* 89, 208–218.

Table I: Non-stationarity ADF Tests

Period	Series				
	f_t	Δf_t	s_t	Δs_t	z_t
Pre-SETS	-1.15	-93.57	-0.70	-27.82	-5.58
Post-SETS	-1.14	-98.13	-1.09	-61.31	-4.42

NOTE: The lag length used in this test is selected on the basis of the Schwarz Information Criterion. The 5% critical value equals -2.87 .

Table II: Linearity Tests

Period	Eq.	d				
		1	2	3	4	5
Pre-SETS	Δf_t	44.18 (0.00)	46.81 (0.00)	29.05 (0.00)	23.23 (0.01)	35.95 (0.00)
	Δs_t	434.67 (0.00)	361.93 (0.00)	281.99 (0.00)	273.64 (0.00)	240.22 (0.00)
Post-SETS	Δf_t	19.40 (0.02)	15.40 (0.08)	26.15 (0.00)	12.30 (0.20)	21.17 (0.01)
	Δs_t	129.44 (0.00)	122.98 (0.00)	112.61 (0.00)	40.25 (0.00)	31.88 (0.00)

NOTE: The LM-type test statistics are reported. The numbers in parentheses are the p -values associated with this test.

Table III: Estimated STECM Parameters

Period	d	$\hat{\alpha}_f$	$\hat{\alpha}_s$	$\hat{\gamma}$	AIC	SIC
Panel A: One Minute Frequency						
Pre-SETS	1	-0.0089 (0.0048)	0.0137 (0.0022)	0.3258 (0.0850)	-362.1515	-362.0314
Post-SETS	1	-0.0045 (0.0087)	0.1299 (0.0108)	0.3594 (0.0995)	-248.4590	-248.3388
Pre-SETS	2	-0.0133 (0.0077)	0.0218 (0.0058)	0.1624 (0.0626)	-300.4555	-300.3353
Post-SETS	2	-0.0060 (0.0051)	0.1128 (0.0111)	0.8034 (0.2127)	-258.1992	-258.0790
Pre-SETS	3	-0.0087 (0.0041)	0.0118 (0.0017)	0.4700 (0.1323)	-357.6144	-357.4943
Post-SETS	3	-0.0014 (0.0046)	0.0900 (0.0103)	0.9180 (0.2993)	-243.5659	-243.4457
Pre-SETS	4	-0.0059 (0.0041)	0.0108 (0.0017)	0.4504 (0.1459)	-357.4856	-357.3654
Post-SETS	4	-0.0027 (0.0048)	0.0722 (0.0107)	0.6903 (0.2694)	-233.0849	-232.9648
Pre-SETS	5	-0.0051 (0.0032)	0.0082 (0.0010)	0.7766 (0.2730)	-364.3381	-364.2179
Post-SETS	5	-0.0011 (0.0041)	0.0488 (0.0055)	2.0688 (0.8235)	-233.1394	-233.0192
Panel B: Five Minute Frequency						
Pre-SETS	1	-0.0174 (0.0588)	0.0974 (0.0462)	0.1113 (0.0515)	-512.4671	-511.5249
Post-SETS	1	0.0372 (0.0534)	0.4054 (0.0781)	0.4981 (0.1409)	-523.1024	-522.1601

NOTE: The numbers in parentheses are heteroscedastic-consistent standard errors. The last two columns contain transformed values of the Akaike Information Criterion (AIC) and the Schwarz Information Criterion (SIC).

Table IV: Transaction Cost Difference Tests

<i>d</i>					
1	2	3	4	5	SIC
0.26	2.89	1.37	0.78	1.49	0.08
(0.40)	(0.00)	(0.09)	(0.22)	(0.07)	(0.47)

NOTE: The null hypothesis that the γ coefficient in the pre-SETS period equals the γ coefficient in the post-SETS period is tested against the alternative that the pre-SETS γ is less than the post-SETS γ . The t-statistics associated with the difference between the pre-SETS γ and the post-SETS γ are reported. The standard error of this difference is calculated using the heteroscedastic-consistent standard error. The numbers in parentheses are the p -values associated with this test.

Figure 1: FTSE100 Futures and Spot Index Levels

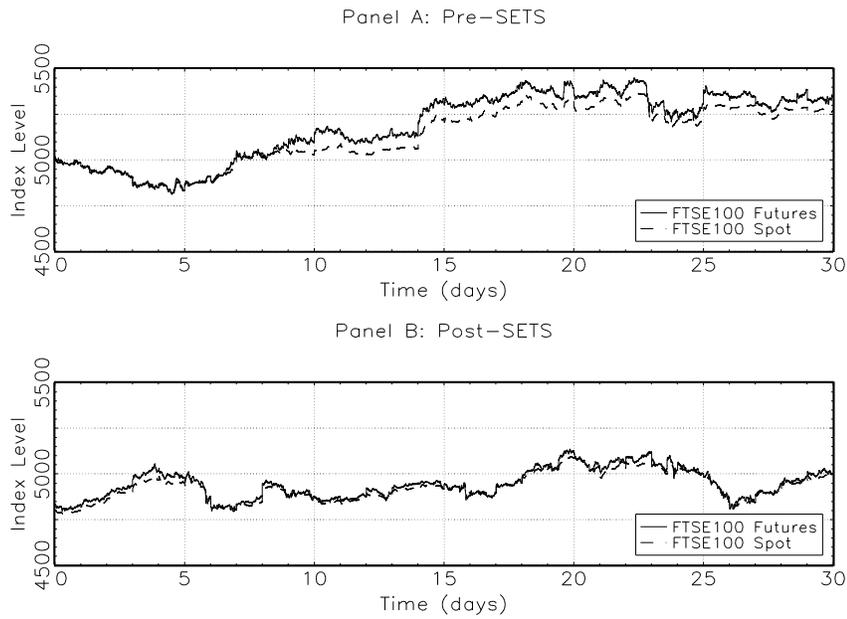


Figure 2: Estimated Transition Functions

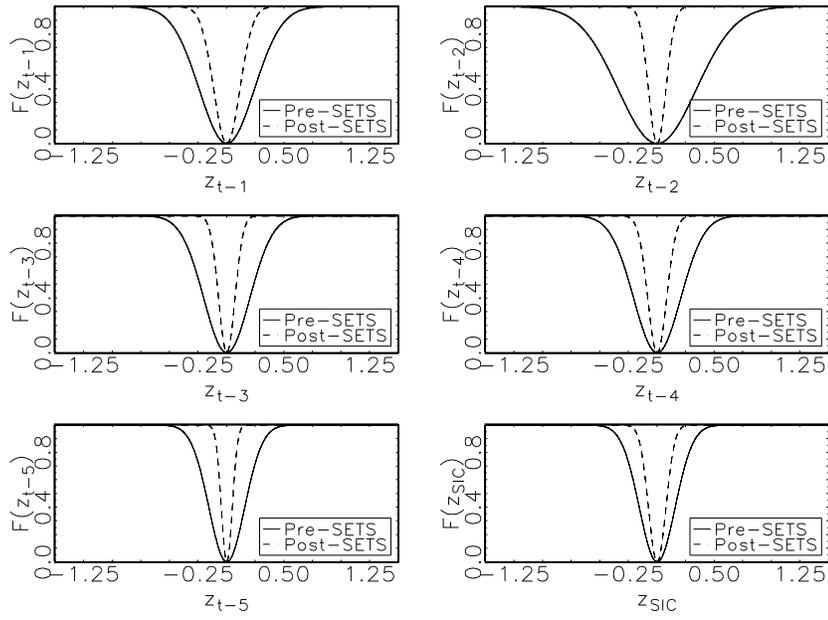


Figure 3: Generalised Impulse Response Distributions (Two Minutes After Shock)

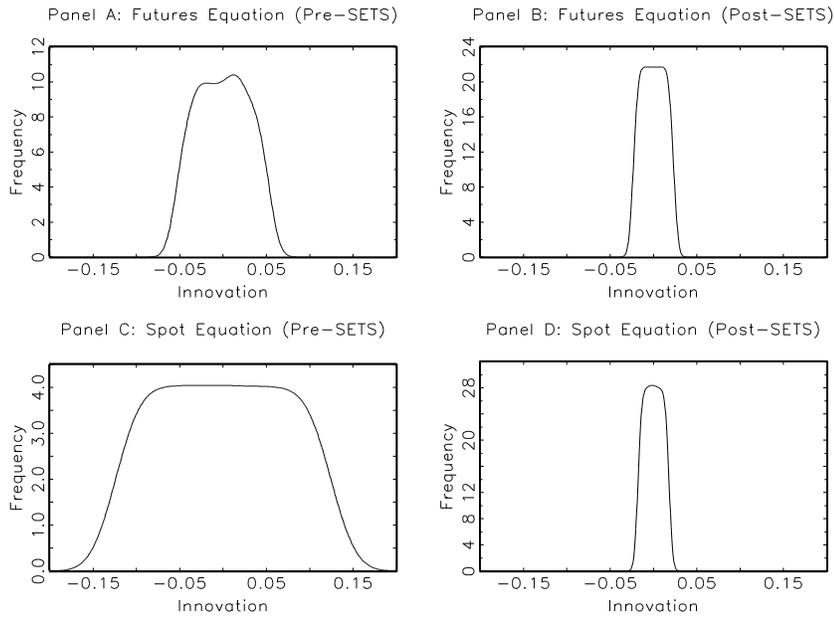


Figure 4: Generalised Impulse Response Distributions (Five Minutes After Shock)

