Studies in Nonlinear Dynamics and Econometrics

Quarterly Journal
July 2001, Volume 5, Number 2
The MIT Press

Intraday and Interday Basis Dynamics: Evidence from the FTSE 100 Index Futures Market

Ian Garrett
School of Accounting and Finance
University of Manchester
ian.garrett@man.ac.uk

Nicholas Taylor
Business School
Cardiff University

Abstract.   We examine the intraday and interday dynamics of both the level of and changes in the FTSE (Financial Times–Stock Exchange) 100 index futures mispricing. Like numerous previous studies we find significant evidence of mean reversion and hence predictability in mispricing changes measured over high (minute-by-minute) and low (daily) frequencies. Contrary to other studies we show explicitly that for high-frequency data, this predictability is due not to microstructure effects but to arbitrage activity. Using a threshold autoregressive model that is consistent with arbitrage behavior, we show that such models imply first-order autocorrelation in mispricing changes similar in magnitude to that actually observed. For low-frequency data, we show that predictability is driven neither by arbitrage activity nor by microstructure effects. Rather, it is a statistical illusion that is the result of overdifferencing a trend-stationary series.

Keywords.   autocorrelation, arbitrage, microstructure, threshold autoregression, FTSE 100 index basis, mispricing

1 Introduction

Early studies of the arbitrage relationship between stock and stock index futures prices established a puzzling empirical phenomenon: changes in stock index futures mispricing and in the stock index basis exhibit

significant negative first-order autocorrelation and are hence predictable.[1] Further, this finding appears to be robust both across markets and across sampling frequencies. For example, such behavior has been documented by MacKinlay and Ramaswamy (1988) for the S&P (Standard & Poor's) 500 index futures basis in the United States, by Yadav and Pope (1990, 1994) for the FTSE (Financial Times–Stock Exchange) 100 in the United Kingdom, and by Lim (1990) for the Nikkei 225 in Japan. In response to this finding, attention has focused on whether such predictability is a direct result of the underlying market microstructure, in particular, nontrading in the underlying stock index. In this vein, Miller, Muthuswamy, and Whaley (1994) documented that first-order autocorrelation in changes in the S&P 500 index futures basis increases from -0.37 to -0.25 when nontrading effects are removed from the underlying index. This suggests that much of the observed predictability is a statistical illusion in that it is driven not by arbitrage but rather by the fact that the underlying index exhibits autocorrelation because of nontrading in some of the constituent stocks.

Recent studies of arbitrage behavior have placed a great deal of emphasis on the nonlinear nature of mispricing dynamics. Since the presence of profitable arbitrage opportunities is determined by transaction costs, we might reasonably expect the behavior of mispricing to be determined by whether profitable arbitrage opportunities are present. This suggests that the mispricing series exhibits nonlinearity. Yadav, Pope, and Paudyal (1994), Dwyer, Locke, and Yu (1996), and Martens, Kofman, and Vorst (1998) used threshold models to capture this nonlinearity. In the simplest version of the threshold model, mispricing can fall within two regimes: it can be within the transaction cost bounds associated with an arbitrage transaction, in which case there are no profitable arbitrage opportunities, or it can be outside them. The general finding of these studies is that mispricing tends to have a near-unit root when arbitragers are unable to take profitable positions but follows a stationary autoregressive process, and hence is predictable, when profitable arbitrage opportunities are available. Given that in the threshold model, the trigger for arbitrage is typically the previous value of mispricing, it follows that these studies have implicitly documented *conditional* predictability (in the form of mean reversion) in mispricing behavior and hence in mispricing changes.

The interesting question that follows from this is whether such threshold autoregressive behavior implies levels of *unconditional* predictability in mispricing changes, that is, whether threshold autoregressive behavior in the level of mispricing implies negative first-order autocorrelation in mispricing changes *irrespective* of the value of mispricing. To our knowledge, no attempt has been made to investigate this. It could be the case, as Miller, Muthuswamy, and Whaley (1994) argue, that arbitrage activity constitutes a tiny fraction of total trading and, as such, is highly unlikely to lead to unconditional predictability in mispricing changes or basis changes. In this instance, such predictability may be due to microstructure effects. If arbitrage activity is sufficiently large, however, then unconditional predictability in basis and mispricing changes should result.

Using both minute-by-minute and daily data on the FTSE 100 index and index futures contracts for the United Kingdom, we examine whether threshold models are capable of generating levels of unconditional first-order autocorrelation observed in both intraday (minute-by-minute) mispricing changes and daily basis changes. We compare the results from the threshold model with those from assuming that the predictability is the result of microstructure effects. Moreover, we conduct the analysis during different times of the trading day to examine the effect of intraday seasonality in arbitrage activity on the results. To anticipate the results, we find that for the high-frequency (minute-by-minute) mispricing series, predictability disappears once arbitrage effects are removed from mispricing changes. By contrast, accounting for microstructure effects does not remove the predictability. This suggests that arbitrage activity can fully account for unconditional

---

[1]Mispricing is the differential between the theoretical, or fair, futures price at time t for delivery at time T and the actual futures price, whereas the basis is the differential between the actual futures price and the underlying spot price. On an intradaily basis, Miller, Muthuswamy, and Whaley (1994) point out that the basis and mispricing will be essentially the same, since dividends and interest are paid overnight. Therefore, ignoring the first few observations the following day will remove the impact of dividends and interest paid, leaving the basis as a measure of mispricing.

predictability. We also examine the robustness of these results to changes in the specification of the threshold model. In particular, we compare the performance of simple threshold models with that of a smooth-transition threshold model and find that a simple two-regime threshold model is sufficient to explain unconditional predictability in intraday mispricing changes. For the daily basis series, we show that neither arbitrage nor microstructure effects can explain the apparent predictability in basis changes, because at the daily frequency, predictability in basis changes is a statistical illusion induced by overdifferencing a trend-stationary series.

The rest of the article is organized as follows. In the next section we outline the cost-of-carry model of stock index futures pricing. Section 3 describes both the threshold autoregressive model of mispricing and the various methods used to remove microstructure effects from mispricing. Section 4 offers a discussion of the data that we use. The results are presented in Section 5, and Section 6 offers some concluding comments.

2 The Economic Model

The contemporaneous relationship between spot and forward prices can be described by the cost-of-carry model, which is also capable of describing the relationship between spot and futures prices providing that the term structure of interest rates is flat and constant. In the absence of arbitrage opportunities and transaction costs, we have

$$F_{t,T}^* = S_t e^{r(T-t)} - \sum_{i=1}^{n} D_i e^{r(T-\tau_i)} \tag{2.1}$$

where $F_{t,T}^*$ is the theoretical (or fair) stock index futures price observed at time t for delivery at time T, S_t is the price of the index, r is the risk-free continuous interest rate applicable over the contract life, $(T - t)$ is the time to maturity of the futures contract, and D_i is the expected cash dividend paid at time τ_i, where $t < \tau_i \leq T$. If the model holds at all times, then $F_{t,T}^* = F_{t,T}$, where $F_{t,T}$ is the market price of the futures contract.

Provided the contract is held to maturity, then in the presence of proportional transaction costs, c, arbitrage activity will take place when one of the following conditions holds:

$$\frac{F_{t,T}}{F_{t,T}^*} < 1 - c \tag{2.2}$$

$$\frac{F_{t,T}}{F_{t,T}^*} > 1 + c \tag{2.3}$$

where c equals the sum of (1) round-trip spot and futures trading costs, (2) market impact costs from trading in the spot and futures markets, and (3) a "stamp tax" of 0.5% that is charged when investors purchase U.K. equities. Arbitragers, however, can and do unwind their spot and futures positions before maturity (Sofianos 1993; Neal 1996), closing out their positions when it is profitable to do so rather than at maturity of the futures contract.[2] Brennan and Schwartz (1988, 1990) model this behavior as arbitragers' having an option, and this ultimately leads to a lowering in the transaction cost bound. Indeed, Dwyer, Locke, and Yu (1996) argue that c represents approximately one half the total round-trip transaction costs incurred by arbitragers.

As it takes time for arbitragers to take appropriate positions in the stock and stock index futures contracts, this arbitrage opportunity is necessarily lagged by d time periods. Therefore, provided c is small, (2.2) and (2.3) can be expressed as[3]

$$|z_{t-d}| > c \tag{2.4}$$

---

[2]Unless of course it is profitable to wait for the contract to mature.
[3]This assumes that the upper and lower transaction cost bounds are the same. Whether this assumption is reasonable in the context of the econometric model is something we test later in the article.

where $z_t = \ln F_{t,T} - \ln F_{t,T}^*$ is mispricing and d is the delay inherent in the arbitrage process. This formulation of the arbitrage condition provides the foundation for the econometric model used in this article to describe arbitrage behavior.

3 Modeling Mispricing

In this section we provide details of an econometric model frequently used to describe arbitrage activity and an econometric model used to describe various microstructure effects.

3.1 The threshold autoregressive model

In the previous section we argued that (2.4) provides the foundation for the econometric model in this article. The key to the econometric model is the behavior of mispricing within the transaction cost bound and outside of it. If mispricing remains within the transaction cost bound, then it will follow a random-walk process, because there will be no arbitrage activity driving prices back into equilibrium. If mispricing goes above or below the transaction cost bound, however, then mispricing will follow a stationary process, because arbitrage activity will force prices back into equilibrium. For instance, if mispricing goes above the upper transaction cost bound, then arbitragers will sell the futures contract and buy the index, hence forcing futures prices down and spot prices up until mispricing no longer violates the transaction cost bound.

The simplest way to model this dynamic price behavior is by means of a threshold autoregressive (TAR) model (see, inter alia, Tong 1983, 1990; Tong and Lim 1980; Tsay 1989; Teräsvirta 1994; and Hansen 1997, 1999, 2000). Formally, z_{t-d} is continuous on $\Re$, so that partitioning the real line defines the number of distinct regimes into which z_{t-d} can fall. The process is in regime h, denoted by r_h, when $r_{h-1} \leq z_{t-d} < r_h$, and when in r_h it follows a pth-order linear autoregressive process:

$$z_t = \phi_0^{\{h\}} + \phi_1^{\{h\}} z_{t-1} + \cdots + \phi_p^{\{h\}} z_{t-p} + \epsilon_t^{\{h\}} \tag{3.1}$$

where the parameters superscripted with $\{h\}$ may vary across regime, $\epsilon_t^h \sim \text{i.i.d.}(0, \sigma^{2\{h\}})$, $h = 1, 2, \ldots, N_r$, and N_r is the number of regimes. This model is quite flexible, in that it may be nonstationary within a regime in the sense that some of the roots of $\phi^{\{h\}}(L) = 1 - \phi_1^{\{h\}} - \cdots - \phi_p^{\{r\}}$ may lie outside the unit circle, but stationary overall because of the alternation of "explosive" and "contractionary" regimes, which generates limit cycle behavior (Chan et al. 1985).

In the absence of arbitrage activity, mispricing follows a random walk, and mispricing changes follow a white-noise process. By contrast, in the presence of arbitrage activity, mispricing changes have nonzero first-order autocorrelation. The use of a TAR model enables us to capture the nature of the arbitrage process and to predict the level of first-order autocorrelation in mispricing changes for given estimated parameter values.

3.2 Microstructure effects

Miller, Muthuswamy, and Whaley (1994) argued that predictability in high-frequency mispricing changes is (partially) due to microstructure effects. In particular, they demonstrated that predictability will occur if stocks in the index trade infrequently and, to a lesser extent, if observed futures prices "bounce" between the bid and ask prices of the contract.

The first of these microstructure effects is often referred to as the "nontrading effect" (Fisher 1966; Dimson 1979; Cohen et al. 1978, 1979; Lo and MacKinlay 1990; Stoll and Whaley 1990). At very high frequencies, say minute by minute, not all stocks in an index/portfolio will trade during each minute. The well-known result of this effect is positive autocorrelation in observed returns. Miller, Muthuswamy, and Whaley (1994) use a modified AR(1) model to capture this effect. They correct for this autocorrelation by fitting an AR(1) model of the form

$$\Delta s_t = \phi_0 + \phi_1 \Delta s_{t-1} + \epsilon_t \tag{3.2}$$

where Δs_t denotes the logarithmic index return and ϵ_t is an i.i.d. error term. The innovations from this regression are defined as

$$\hat{\epsilon}_t^* = \frac{\hat{\epsilon}_t}{(1 - \hat{\phi}_1)} \tag{3.3}$$

and these innovations represent the index return adjusted for nontrading.

The second microstructure effect considered by Miller, Muthuswamy, and Whaley is referred to as "bid-ask bounce." This effect is more likely to show up in futures prices than index levels, as these are individually traded securities. Roll (1984) demonstrated that observed prices randomly bouncing between the bid and the ask prices leads to negative first-order autocorrelation in returns. Miller, Muthuswamy, and Whaley model this effect on futures returns as an MA(1) process:

$$\Delta f_t = \theta_0 + v_t + \theta_1 v_{t-1} \tag{3.4}$$

where Δf_t denotes the observed logarithmic futures return and v_t is an i.i.d. error term. The larger the bid-ask spread, the larger the bid-ask bounce effect as measured by θ_1 and the greater the extent of the negative first-order autocorrelation in futures returns. This effect is removed by fitting an MA(1) model to logarithmic futures returns and taking the residual, $\hat{v}_t$, as the futures return net of the bid-ask bounce effect.

Taking these two effects together, Miller, Muthuswamy, and Whaley demonstrate analytically that negative first-order autocorrelation in mispricing changes is likely to occur under quite general conditions. They provide an illustration by considering mispricing changes for the S&P 500 Index futures contract at a 15-minute frequency. In particular, they show that predictability is reduced (but not eliminated) when nontrading effects are removed from spot returns.

4 Data

We analyze the relationship between spot and futures prices using both intraday and interday data. The intraday data will be described first. The futures price of the nearest FTSE 100 contract was obtained for every transaction carried out between January 5 and April 24, 1998. These data were obtained from the London International Financial Futures and Options Exchange (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.[4] To synchronize the futures and spot prices, the futures price series was converted to a price series with a frequency of one minute. The (spot) level of the FTSE 100 index was obtained from FTSE International. The trading hours of the futures market and the spot market are 8:35 A.M. to 4:10 P.M. and 8:00 A.M. to 4:30 P.M., respectively.[5] Thus one can obtain overlapping futures and spot data covering the period from 8:35 A.M. to 4:10 P.M. For convenience, however, we considered only the period between 9:01 A.M. and 4:00 P.M. A total of 77 trading days were considered, yielding a total of 32,417 (421×77) one-minute-frequency observations.

The validity of the mispricing series thus constructed relies heavily on the use of appropriate ex ante dividends and interest rates. To this end we make use of data supplied by Goldman Sachs. These data are used by arbitragers employed by Goldman Sachs when making judgments about the mispricing (or otherwise) of FTSE 100 futures contracts. Goldman Sachs constructs ex ante dividends by making individual forecasts for each of the dividends paid by companies in the FTSE 100 index and then weight these by market capitalization. The interest rate applicable over the contract life used by Goldman Sachs is the interpolated

---

[4]The volume cross-over method of changing futures contracts results in one change during the period under study, involving a switch from the March 1998 contract to the June 1998 contract on March 11, 1998. On this day the volume of trading in the March contract was 6,312 contracts and the volume of trading in the June contract was 13,355 contracts.

[5]The futures market reopens at 4:32 P.M. under the automated pit trading (APT) system. This additional period of trading is not considered here, however, because of the lack of data between 4:11 P.M. and 4:31 P.M.
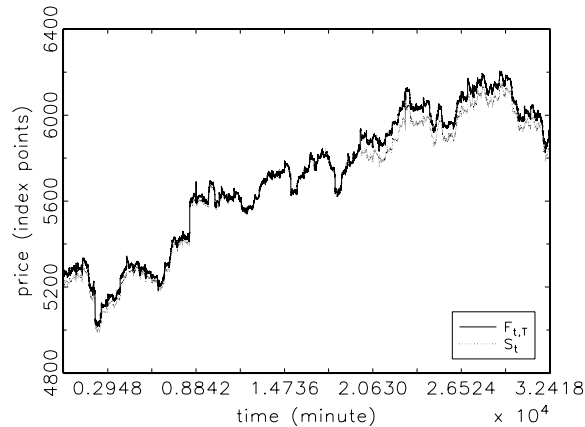
Figure 1
Intraday prices.

LIBOR rate. For instance, if a 25-day interest rate is required, then Goldman Sachs interpolates between the two-week and the four-week rates.

The interday data consist of the end-of-day futures prices of every contract traded on LIFFE from the June 1986 contract to the December 1998 contract. The corresponding spot price of the index was obtained from Datastream. We considered the last 80 trading days for each of the futures contracts.

5 Empirical Results

The spot and futures prices for the FTSE 100 index over the period from January 5 to April 24, 1998, are plotted in Figure 1. There appears to be a close relationship between the two prices, with convergence occurring when the contract is rolled over. This convergence in the difference between spot and futures prices can be seen more clearly in Figure 2. The difference declines steadily as the maturity of the contract approaches. By contrast, the difference between the market and theoretical futures prices appears to have a constant zero mean over the period and suggests that on average the FTSE 100 Index futures contract is correctly priced when the Goldman Sachs ex ante dividend and interest rate data are used.

The first-order autocorrelations for futures returns, spot returns, and mispricing changes are provided in Table 1. These correlations are calculated over the whole period (9:01 A.M. to 4:00 P.M.), during the afternoon period (12:01 P.M. to 4:00 P.M.), and during each hour of the trading day. Futures returns exhibit significant negative autocorrelation during the afternoon period only, whereas spot returns tend to be significantly positively autocorrelated throughout the day. From our earlier discussion, it seems that these autocorrelations reflect microstructure effects. To illustrate more clearly, consider again futures returns. Futures returns are most negatively autocorrelated between 12:01 P.M. and 2:00 P.M. This period coincides with the highest effective bid-ask spread (and lowest trading volume) on the futures contract and hence is a time when bid-ask bounce effects are likely to be at their most pronounced.[6]

---

[6]See Abhyankar, Copeland, and Wong 1999 and Tse 1999 for further evidence of intraday patterns in bid-ask spreads and trading volumes in the FTSE 100 Index market.
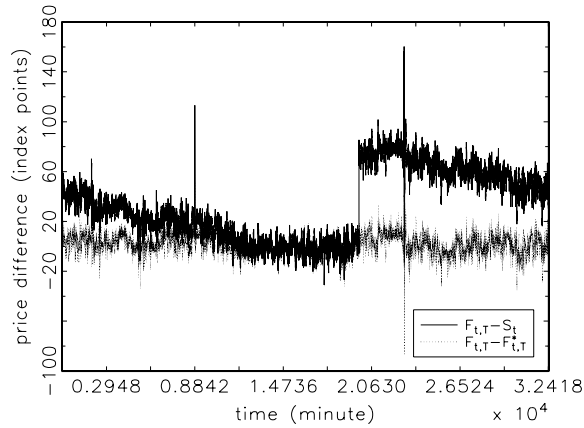
Figure 2
Intraday basis and mispricing.

Table 1
Intraday summary statistics

Interval	T	Unadjusted			Adjusted			Volume	Spread
		$\hat{\rho}_1(\Delta f_t)$	$\hat{\rho}_1(\Delta s_t)$	$\hat{\rho}_1(\Delta z_t)$	$\hat{\rho}_1(\Delta f_t)$	$\hat{\rho}_1(\Delta s_t)$	$\hat{\rho}_1(\Delta z_t)$		
9:01 to 10:00	4,616	0.0023	0.0958***	−0.0707***	−0.0008	−0.0043	−0.1100***	21.5017	3.3085
10:01 to 11:00	4,612	0.0354**	−0.0801***	−0.0930***	−0.0000	−0.0193	−0.0577***	20.6667	3.3736
11:01 to 12:00	4,620	−0.0064	0.0936***	−0.0568***	−0.0001	−0.0053	−0.0987***	18.5983	3.3183
12:01 to 1:00	4,620	−0.0459***	0.0740***	−0.0586***	−0.0015	0.0046	−0.0671***	13.8485	3.4409
1:01 to 2:00	4,618	−0.1587***	0.1275***	−0.1573***	−0.0066	−0.0047	−0.1385***	14.8530	3.4190
2:01 to 3:00	4,620	0.0251*	0.0918***	−0.0553***	0.0005	−0.0045	−0.1092***	23.7545	3.1958
3:01 to 4:00	4,620	0.0451***	0.0916***	−0.0458***	0.0004	−0.0038	−0.1091***	29.9803	3.1997
9:01 to 4:00	32,326	−0.0044	0.0347***	−0.0800***	0.0003	0.0018	−0.0934***	20.4576	3.3223
12:01 to 4:00	18,480	−0.0141*	0.1002***	−0.0734***	−0.0002	−0.0030	−0.1113***	20.6091	3.3138

Note: The number of observations (T), first-order autocorrelation ($\hat{\rho}_1$) in logarithmic futures returns (Δf_t), logarithmic spot returns (Δs_t), mispricing changes (Δz_t), mean volume of futures contracts traded per minute, and mean effective bid-ask spread (in basis points) on these futures contracts are calculated using minute-frequency data. First-order autocorrelations are also calculated using futures returns, spot returns, and mispricing changes adjusted for possible microstructure effects as described in Miller, Muthuswamy, and Whaley 1994. The significance of the first-order autocorrelation is tested using the Box-Pierce Q-statistic and is denoted by *** (1% significance), ** (5% significance), and * (10% significance).

Mispricing changes exhibit significant negative first-order autocorrelation during each hour of the trading day. The autocorrelation is at its most negative during the first few hours of trading and between 1:01 P.M. and 2:00 P.M. This intraday pattern could again be due to microstructure effects. The 1:01 to 2:00 period corresponds to the smallest autocorrelation in futures returns, which, as we have just seen, coincides with the highest effective bid-ask spread on the futures contract and with the highest autocorrelation in spot returns. The significant negative autocorrelation for the first few trading hours is most likely the result of the high bid-ask spreads observed during the beginning of trading in the spot market. For instance, Naik and Yadav (1999) found evidence of high bid-ask spreads on all stocks making up the FTSE 100 index during the first few hours of the trading day. These high spreads are most likely the result of a lack of liquidity during this period. Shah (1999) argued that the lack of liquidity during the early morning period may be due to the "absence of an opening auction facility for market on open orders."

To summarize thus far, then, there is evidence of positive first-order autocorrelation in index returns, negative first-order autocorrelation in futures returns, and negative first-order autocorrelation in mispricing changes. These findings seem to be consistent with a microstructure explanation for predictability in mispricing changes, since the autocorrelation in mispricing changes is at its strongest when spreads are high and liquidity is low in the futures and spot markets. Whether these findings are actually consistent with a microstructure explanation of predictability in mispricing changes is the subject of the next section.

5.1 Removing microstructure effects

Miller, Muthuswamy, and Whaley (1994) argued that first-order autocorrelation in S&P 500 mispricing changes is due to microstructure-induced first-order autocorrelation in spot and futures returns. To examine the validity of this claim using U.K. data, we adjust the spot returns for nontrading and the futures returns for bid-ask bounce effects using the methodology outlined in Section 3.2. The results for the adjusted series are given in Table 1. Whereas autocorrelations in spot and futures returns are eliminated by the appropriate adjustment, mispricing changes still exhibit significant negative first-order autocorrelation. Indeed, the autocorrelation actually becomes more negative during most hours of the trading day and especially during the afternoon period. Therefore, contrary to the S&P 500 market, the FTSE 100 market is not characterized by microstructure-induced autocorrelation in mispricing changes. This leaves open the possibility that the autocorrelation is arbitrage-induced, which is the subject of the next section.

5.2 Removing arbitrage effects

An important aspect of the argument by Miller, Muthuswamy, and Whaley that predictability in mispricing changes is microstructure-induced is based on the fact that if spot and futures prices follow a random walk, then in the absence of arbitrage activity, mispricing should also follow a random walk. Thus, mispricing will persist indefinitely. By contrast, if arbitragers exist in the market, then mispricing will be removed within a very short period of time. As such, mispricing will necessarily follow a mean reverting, stationary process. It follows from this that an indirect, but straightforward, way of examining whether arbitragers are present in the market is to examine whether mispricing contains a unit root since, if it does, mispricing will persist indefinitely, which coincides with the null hypothesis that no (substantive) arbitrage activity takes place. We test this hypothesis using the forward and reverse Dickey-Fuller regression methodology described by Leybourne (1995).[7] The results from this test are given in Table 2. Spot and futures prices appear to have unit roots while returns on these assets are stationary. For mispricing, however, the null hypothesis of a unit root can be rejected at the 1% level of significance. This means that mispricing is a stationary process rather than a random walk, a result that suggests arbitrage activity is of some significance for the FTSE 100 market.

Given the finding that mispricing is a stationary process, the natural next step is to model mispricing in the presence of transaction costs using the TAR model outlined in Section 3.1. We estimate the TAR over the whole period, for each hour of the trading day, and for the afternoon period, following the methodology described in Hansen 1997. From our earlier discussion, we consider a two-regime model; the regimes will be referred to as the "inside regime" and the "outside regime." When lagged mispricing is within the transaction cost bound, the process is in the inside regime. When lagged mispricing is greater (in absolute terms) than the transaction cost bound, then the process is in the outside regime. In both cases the first lag of absolute mispricing is taken as the threshold variable. Longer lags were used but were found to produce an inferior fit. Further, since we are interested in explaining first-order autocorrelation, the number of lags of mispricing used in the TAR model is set equal to unity.

The results from estimating the TAR model are given in Table 3. As the errors from the model are heteroskedastic, all inference is carried out using heteroskedasticity-consistent standard errors. The Lagrange

---

[7]Leybourne (1995) showed that this testing methodology is more powerful than the standard Dickey-Fuller methodology.

Table 2
Testing for nonstationarity

Interval	d = 0			d = 1		
	$\Delta^d f_t$	$\Delta^d s_t$	$\Delta^d z_t$	$\Delta^d f_t$	$\Delta^d s_t$	$\Delta^d z_t$
9:01 to 10:00	−0.8983	−0.7467	−16.8399***	−67.6873***	−61.6831***	−72.9229***
10:01 to 11:00	−0.9079	−0.8744	−17.0691***	−65.5850***	−73.5099***	−74.4642***
11:01 to 12:00	−0.9150	−0.8922	−13.1765***	−68.3814***	−61.8619***	−71.9197***
12:01 to 1:00	−0.9091	−0.8873	−11.0174***	−71.1344***	−63.0969***	−72.0441***
1:01 to 2:00	−0.9394	−0.8774	−13.6719***	−79.7413***	−59.7800***	−79.6285***
2:01 to 3:00	−0.7978	−0.7156	−15.5093***	−66.2594***	−61.9782***	−71.8138***
3:01 to 4:00	−0.7334	−0.6638	−17.3553***	−64.9462***	−61.9892***	−71.1410***
9:01 to 4:00	−0.7863	−0.7031	−39.0341***	−180.5319***	−173.6754***	−194.8339***
12:01 to 4:00	−0.7419	−0.7575	−27.9556***	−137.8559***	−122.9339***	−146.2966***

Note: Nonstationarity tests were performed on various series using the forward and reverse Dickey-Fuller regression methodology described by Leybourne (1995). The series considered are logarithmic futures prices (f_t), logarithmic spot prices (s_t), and mispricing (z_t). The test statistics and an indication of whether the null hypothesis of a unit root was rejected are given. Significance is denoted by *** (1% significance), ** (5% significance), and * (10% significance).

Table 3
Estimated parameters of the TAR model

Interval	Inside Regime			Outside Regime			$\hat{c}$	LM test
	$\hat{\phi}_0^{r_1}$	$\hat{\phi}_1^{r_1}$	R^2	$\hat{\phi}_0^{r_2}$	$\hat{\phi}_1^{r_2}$	R^2		
9:01 to 10:00	0.0000	0.9214***	0.8389	−0.0016	0.5372*	0.4595	56.86***	9.75*
10:01 to 11:00	0.0000	0.9397***	0.8208	−0.0008	0.4400***	0.5124	105.75***	13.92***
11:01 to 12:00	0.0000	0.9504***	0.8830	0.0000	0.7759***	0.9006	30.19***	23.67***
12:01 to 1:00	0.0000	0.9714***	0.9174	0.0001	0.8361***	0.9203	26.23***	29.58***
1:01 to 2:00	0.0000	0.9644***	0.8937	0.0005	0.6967***	0.6478	26.85***	24.16***
2:01 to 3:00	0.0000	0.9217***	0.8360	0.0003	0.5452***	0.6142	36.90***	20.43***
2:01 to 4:00	0.0000	0.9254***	0.6941	0.0002	0.8254***	0.8547	18.32***	28.13***
9:01 to 4:00	0.0000	0.9243***	0.8439	−0.0005	0.4752***	0.4541	105.75***	55.70***
12:01 to 4:00	0.0000	0.9442***	0.8477	0.0001	0.7989***	0.8239	26.23***	75.53***

Note: The TAR model

$$z_t = \phi_0^{\{b\}} + \phi_1^{\{b\}} z_{t-1} + \epsilon_t^{\{b\}}$$

where $b = r_1$ if $|z_{t-1}| \leq c$ and $b = r_2$ if $|z_{t-1}| > c$, z_t is mispricing, and $\epsilon_t^{\{b\}}$ is an error term in regime b, is estimated. The estimated threshold $\hat{c}$ is multiplied by 10,000 and is thus measured in terms of basis points. Heteroskedastic-consistent standard errors associated with the TAR coefficients are calculated. The null hypothesis of no threshold autoregression is tested using a Lagrange multiplier test, and the associated p-value is calculated using the bootstrap technique. Significance is denoted by *** (1% significance), ** (5% significance), and * (10% significance).

multiplier (LM) test of the null hypothesis that a linear autoregression is appropriate against the TAR alternative clearly rejects the linear model. This shows that the threshold model is the appropriate model to use, as is clearly demonstrated by the results from estimating the threshold model. There is strong support for the notion that mispricing follows a different process depending upon whether arbitrage opportunities are present. Taking the behavior of mispricing in the inside regime first, the coefficient on lagged mispricing is close to unity, suggesting that mispricing behaves in a fashion similar to a random walk. Second, mispricing follows an autoregressive process in the outside regime, with a lower coefficient on lagged mispricing. These two findings are consistent with additional adjustment in mispricing when arbitrage profits are greater than the transaction costs incurred. Third, transaction costs, as defined by the estimated threshold value, $\hat{c}$, have an intraday pattern that is consistent with the behavior of bid-ask spreads in the spot market documented by

Naik and Yadav (1999).[8] As these spreads are considerably greater than futures spreads, we would expect spreads in the spot market to exert a greater influence on transaction costs incurred by arbitragers than the spreads in the futures market.

Arbitrage transaction costs incurred during the first few hours of the trading day are considerably greater than those incurred during other periods of the day. Indeed, early-morning transaction costs appear to be prohibitively high, and it would seem that very little arbitrage is carried out during the first few hours of the day. From a statistical point of view this is somewhat problematic, as the outside regime contains a very limited number of observations (typically less than 10) during these periods. Moreover, when the model is estimated over the whole trading day, the transaction costs appear to be extremely high because of the influence of the morning data. To avoid this effect the model is estimated during the afternoon period only. The results indicate that there is a considerable decrease in the transaction costs incurred in the afternoon period in comparison to the morning period. As such, in the subsequent analysis the focus is placed on the afternoon period because of the lack of degrees of freedom present in the outside regime when other sample periods are considered.

Having seen that the TAR model provides a good description of the behavior of mispricing in the presence of transaction costs, we turn our attention toward investigating the effects of arbitrage on autocorrelation in mispricing changes. To investigate whether the observed predictability in mispricing changes vanishes once the effects of arbitrage have been accounted for, we need to remove the effects of arbitrage activity from mispricing changes. We do this by estimating the TAR model for the afternoon period, with mispricing changes replacing the level of mispricing.[9] The coefficient on lagged mispricing changes is then an estimate of the first-order autocorrelation in mispricing changes in the different regimes. Estimating the TAR model using mispricing changes over this period gives

$$\Delta z_t = 0.9867 - 0.0005\Delta z_{t-1} + \hat{\epsilon}_t^{r_1}, \quad \text{if} \quad |z_{t-1}| \leq 22.83 \tag{5.1}$$
$$(0.3302) \quad (0.0105)$$

$$\Delta z_t = -9.5276 - 0.4303\Delta z_{t-1} + \hat{\epsilon}_t^{r_2}, \quad \text{if} \quad |z_{t-1}| > 22.83 \tag{5.2}$$
$$(3.2650) \quad (0.1037)$$

where z_t is measured in basis points and heteroskedastic-consistent standard errors are given in parentheses. The threshold value ($\hat{c}$) of 22.83 basis points is significantly different from zero at the 1% level, and the LM test statistic for testing the null hypothesis that there are no thresholds is 64.59 with a p-value less than 0.01. This again provides strong evidence in support of the threshold model. The number of observations in the inside and outside regimes is 17,421 and 1,059, respectively. The number of observed arbitrage violations suggests that approximately 6% of trades are carried out by traders arbitraging perturbations in the FTSE 100 spot-futures relationship.

The coefficient on lagged mispricing changes is insignificantly different from zero in the inside regime. This suggests that there is no first-order autocorrelation when arbitrage opportunities are not present. By contrast, mispricing changes are characterized by significant negative first-order autocorrelation in the outside regime, which is where arbitrage activity occurs. This finding lends support to the notion that negative first-order autocorrelation in mispricing changes is an arbitrage phenomenon. Although this confirms the presence of *conditional* prediction in mispricing changes, however, it could be argued that given the small number of observations in the outside regime, arbitrage activity is likely to be infrequent and of insufficient magnitude to cause unconditional negative first-order autocorrelation in mispricing changes. This proposition is examined in more detail by considering the level of *unconditional* first-order autocorrelation in mispricing changes implied by various nonlinear models of arbitrage.

---

[8]The reported estimated threshold is multiplied by 10,000 and is therefore reported in terms of basis points.
[9]The threshold variable is still defined as the lagged level of mispricing, because this remains as the trigger for arbitrage activity.
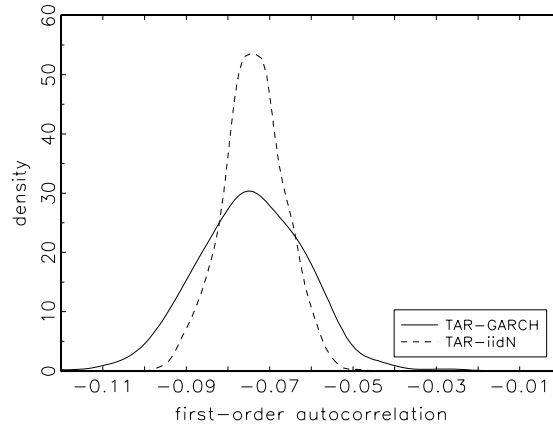
Figure 3
TAR-implied autocorrelation.

5.3 Implied unconditional autocorrelations

To examine the ability of the TAR model to generate the levels of unconditional first-order autocorrelation that we observe in mispricing changes, we use the parameters of the TAR model estimated using the level of mispricing observed between 12:01 P.M. and 4:00 P.M. on each trading day to calculate the unconditional distribution of first-order autocorrelation in mispricing. This distribution will give some indication of the validity of the model in terms of explaining the negative autocorrelation in mispricing changes. If the likely implied first-order autocorrelation is close to the empirical first-order autocorrelation, then the model provides a satisfactory explanation for the negative autocorrelation. Otherwise, the arbitrage model, like the microstructure model, is inadequate to explain the negative correlation in mispricing changes.

Given the nonlinear nature of the TAR model, we calculate the unconditional distribution of the first-order autocorrelation by Monte Carlo simulation. The estimated parameters and residual variances obtained from the actual data are used in the simulation. However, as theory suggests that the process should follow a random walk when no profitable arbitrage opportunities are available, we impose the restriction that the process follows a random walk in the inside regime, that is, we set ϕ_1^n equal to 1. We generate 1,000 TAR processes, each with 18,480 observations. The first-differences of these series are then taken and the first-order autocorrelation and the Box-Pierce Q-statistic testing for first-order autocorrelation are calculated. A kernel estimate of the distribution of the first-order autocorrelation is then calculated using a normal window, with bandwidth selected using the method of likelihood cross-validation.[10]

The estimated unconditional distribution of the first-order autocorrelation and the Box-Pierce Q-statistic are given in Figures 3 and 4. We make two assumptions regarding the error term in the TAR model. The first is that the errors are i.i.d. normal, and the second is that the errors follow a GARCH(1,1) process, with parameters used in the simulation matching the parameters obtained from fitting a GARCH(1,1) model to the residuals from an AR(1) model of mispricing. This second assumption captures the observed heteroskedasticity in the residuals of the TAR model.[11] Both of the implied distributions have modal points close to the empirical first-order autocorrelation given in Table 1. For example, the TAR-GARCH model has a

---

[10]See Silverman 1986 for further details on density estimation using the kernel method.
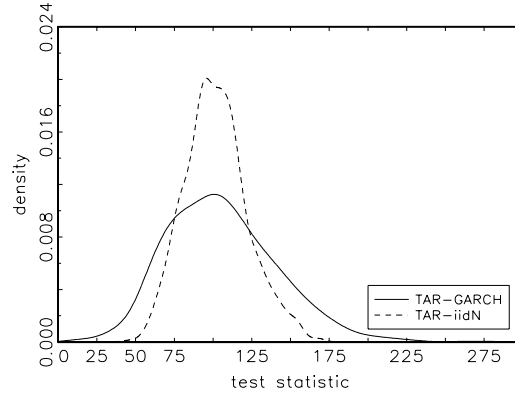[11]These results are available upon request.

Figure 4
TAR-implied test statistic.

median implied first-order autocorrelation coefficient of -0.0745 (see Figure 3) compared to an empirical first-order autocorrelation of -0.0734. Similar results are obtained when the implied distribution of the Box-Pierce Q-statistic is considered. This evidence lends further support not only to the arbitrage explanation of negative first-order autocorrelation in mispricing changes but also to the ability of a simple two-regime TAR model to capture this behavior. In the next two sections, we turn our attention to an examination of whether more sophisticated empirical models of arbitrage behavior offer an improvement over the simple two-regime TAR model.

5.4 Short-selling restrictions

It could be argued that the above two-regime TAR model is inappropriate, as it imposes the restriction that the upper and lower transaction bounds are of the same magnitude. Indeed, if short selling of the index is restricted, then one would expect the lower transaction bound to be greater (in absolute value) than the upper transaction bound. In the presence of short-selling restrictions, per unit negative mispricing (which involves short selling the index) is less profitable than per unit positive mispricing. This results in asymmetric transaction bounds.

To allow for asymmetric transaction bounds we fit a three-regime TAR model to the level of mispricing observed between 12:01 P.M. and 4:00 P.M. on each trading day. Having fitted this model we then test the null hypothesis that mispricing follows a two-regime TAR process against the alternative that mispricing follows a three-regime TAR process. Testing inference is carried out using the (general) heteroskedastic-consistent bootstrap methodology of Hansen (1999). We use 1,000 bootstrap replications to generate the distribution of the test statistic under the null hypothesis.

The fitted three-regime TAR model is given by the following set of equations:

$$z_t = -56.1222 + 0.6183 z_{t-1} + \hat{\epsilon}_t^{r_1}, \quad \text{if} \quad z_{t-1} \leq -17.84 \tag{5.3}$$
$$(23.0995) \ (0.1083)$$

$$z_t = 1.6403 + 0.9565 z_{t-1} + \hat{\epsilon}_t^{r_2}, \quad \text{if} \quad -17.84 < z_{t-1} \leq 18.75 \tag{5.4}$$
$$(0.3464) \ (0.0039)$$

$$z_t = 69.8995 + 0.6301 z_{t-1} + \hat{\epsilon}_t^{r_3}, \quad \text{if} \quad z_{t-1} > 18.75 \tag{5.5}$$
$$(11.1702) \ (0.0487)$$

where z_t is measured in basis points, heteroskedastic-consistent standard errors are given in parentheses, and the number of observations in the three regimes is 711, 16,037, and 1,731, respectively. The test statistic associated with the null hypothesis that the two-regime TAR model provides an adequate fit equals 77.11 and has an associated p-value of 0.62. Therefore, we can conclude that short-selling restrictions have an insignificant effect on the symmetry of the transaction bounds. The fact that short-selling restrictions are unimportant can also be seen by noting the similarity of the threshold values in the above equation. The failure to reject the null hypothesis means that these threshold values are insignificantly different from each other and, therefore, that a two-regime TAR model provides an adequate description of mispricing.

5.5 Alternative models of arbitrage behavior

Recent literature on arbitrage activity has made use of smooth transition autoregressive (STAR) models (see Anderson 1997 and Taylor et al. 2000) to model mispricing dynamics. These models allow for heterogeneous transaction cost exposure by imposing a smooth parametric function on the regime space. The TAR framework essentially uses a discontinuous transition function in which the process makes abrupt switches from one regime to another. Smooth transition models, on the other hand, allow for gradual changes of regime.

A commonly used smooth transition function is the exponential STAR model (see Teräsvirta 1994 for further details). The version of the model that results in the best fit here is

$$z_t = (\phi_0 + \phi_1 z_{t-1})(1 - g(\gamma; z_{t-1})) + (\phi_2 + \phi_3 z_{t-1})g(\gamma; z_{t-1}) + \epsilon_t \qquad (5.6)$$

where the transition function $g(\gamma; z_{t-1}) = 1 - \exp(-\gamma z_{t-1}^2)$. It is clear from this transition function that the higher the level of mispricing in the previous period, the greater the value of the transition function. Indeed, the transition function is bounded from zero (when $z_{t-1} = 0$) to unity (when $|z_{t-1}|$ is large). Therefore, when there is no arbitrage, time dependency is measured by the coefficient ϕ_1, and when there is full arbitrage, time dependency is measured by ϕ_3. At all other times the coefficients reflect a mixture of arbitrage activity.

The specification given in (5.6) is estimated using the level of mispricing observed between 12:01 P.M. and 4:00 P.M. on each trading day. The methodology is the same as that used by Teräsvirta (1994). Ignoring the (insignificant) constant terms, the estimated model is (heteroskedasticity-consistent standard errors in parentheses)

$$z_t = 0.9472 z_{t-1}(1 - g(\gamma; z_{t-1})) + 0.5512 z_{t-1}g(\gamma; z_{t-1}) + \hat{\epsilon}_t \qquad (5.7)$$
$$\;\;\;\;(0.0048) \qquad\qquad\qquad (0.1760)$$

where

$$g(\gamma; z_{t-1}) = 1 - \exp(-0.0672 z_{t-1}^2) \qquad (5.8)$$
$$(0.6669)$$

The first point to note is that γ is insignificantly different from zero, suggesting that the STAR model may not be an inappropriate way of modeling mispricing. Notwithstanding this, the others parameters take on reasonable values with the no-arbitrage coefficient (ϕ_1) being close to unity, whereas the full-arbitrage coefficient (ϕ_3) is somewhat less than unity but is still positive. In both cases the coefficients are significantly different from zero.

We use the estimated parameters from (5.7) and (5.8) to generate the unconditional distribution of the first-order autocorrelation using a 1,000-repetition Monte Carlo simulation. As with the previous simulation, we assume that mispricing follows a random walk in the absence of arbitrage. Thus, we set $\phi_1 = 1$. The number of observations used in each repetition is 18,480. The errors in the STAR model are drawn from a GARCH(1,1) process with parameter values matching the empirical estimates, as before. The kernel estimate of the unconditional distribution of first-order autocorrelation is plotted in Figure 5. In comparison to the distribution implied by the TAR-GARCH model, the STAR-GARCH model implies a negative autocorrelation
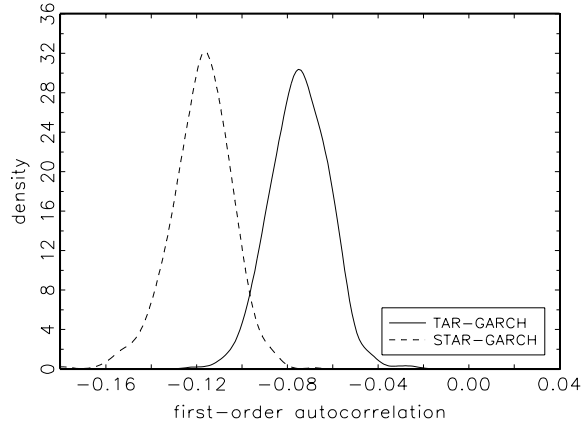
Figure 5
STAR-implied autocorrelation.

that is far less than the empirical autocorrelation. In particular, the median STAR-GARCH-implied first-order autocorrelation is -0.1171 compared to the empirical value of -0.0745. Moreover, the density function implied by the STAR model suggests that there is little chance of observing the empirical first-order autocorrelation. Similar results are obtained when alternative specifications of the transition function are considered. For instance, when a logistic STAR model[12] is used, there is an excessive amount of negative first-order autocorrelation in mispricing changes. These results suggest that the STAR model is not capable of generating the appropriate level of first-order autocorrelation observed empirically, and hence it is an inadequate and inappropriate model of arbitrage in this case.

5.6 Explaining predictability in interday basis changes
The dividends paid to owners of FTSE 100 shares tend not to be seasonal (Shah 1999). Therefore, it is reasonable to assume that the sum of the dividends paid out during the life of the futures contract is approximately equal to some constant (annualized) dividend amount, D, multiplied by the time to maturity of the contract. Under this assumption, (2.1) can be rewritten as

$$F_{t,T}^* = S_t e^{r(T-t)} - D(T-t)e^{r(T-t)} \tag{5.9}$$

Expressing the dividend level as a proportion of the current price gives

$$F_{t,T}^* = S_t e^{r(T-t)} - S_t \delta(T-t)e^{r(T-t)} = S_t e^{r(T-t)}\big(1 - \delta(T-t)\big) \tag{5.10}$$

where $\delta = D/S_t$ and denotes the (annualized) dividend yield on the index. If $\delta(T-t)$ is small, then $(1 - \delta(T-t)) \approx e^{-\delta(T-t)}$, and (5.10) becomes

$$F_{t,T}^* = S_t e^{(r-\delta)(T-t)} \tag{5.11}$$

When considering daily mispricing, differences between the logarithmic market price ($f_{t,T}$) and the logarithmic theoretical price ($f_{T,t}^*$) are unlikely to follow the same process as that of minute-by-minute

---

[12]See Teräsvirta 1994 for more details on STAR models with a logistic transition function.

mispricing. As arbitrage profits are unlikely to persist over a day, mispricing is unlikely to be autocorrelated. It is therefore reasonable to assume that mispricing is a mean zero **i.i.d.** process which we can denote by ϵ_t. Using this assumption, taking logarithms of both sides of (5.11) and rearranging, we obtain the following expression for the basis:

$$b_t \equiv f_{t,T} - s_t = \mu_t(T - t) + \epsilon_t \tag{5.12}$$

where $\mu_t = r - \delta$ and represents a time-varying trend coefficient. If we assume that μ_t follows a process with random perturbations away from the mean, that is, $\mu_t = \mu + \nu_t$, where ν_t is an error term that is identically distributed over time but can be autocorrelated, then (5.12) becomes

$$b_t = (\mu + \nu_t)(T - t) + \epsilon_t = \mu(T - t) + \xi_t \tag{5.13}$$

where $\xi_t = \nu(T - t) + \epsilon_t$ and represents a mixture of an **i.i.d.** error term (ϵ_t) and an error term with time-varying first and second moments ($\nu_t(T - t)$). Taking first differences of (5.13) gives an expression for changes in the basis:

$$\Delta b_t = -\mu + (\xi_t - \xi_{t-1}) \tag{5.14}$$

This equation shows that the change in the basis is an MA process with a unit root. As the basis is stationary around a trend, taking differences of this process amounts to overdifferencing the series. It is this overdifferencing that induces the autocorrelation in the process. Ignoring the time variation in the first two moments of ξ_t, it is a trivial exercise to show that first-order autocorrelation in basis changes equals $-\frac{1}{2}$.[13]

To summarize, there are three testable predictions that arise from this theory. First, the net-of-trend basis will be heteroskedastic if $\nu_t(T - t)$ dominates ξ_t. We test this proposition using the Goldfeld-Quandt test statistic. Second, basis changes will have an MA unit root. We test this proposition using the Kwiatkowski, Phillips, Schmidt, and Shin (KPSS) test.[14] This test is reasonably robust to time variation in the first and second moments. Third, the first-order autocorrelation coefficient for basis changes should be $-\frac{1}{2}$. Results from tests of these propositions are provided in Table 4.

The results indicate that heteroskedasticity is present in roughly one third of the contracts considered. The nature of the heteroskedasticity can also be seen in the box plot given in Figure 6. The basis appears to fall in a linear fashion up to maturity, with the distribution of the basis around the mean becoming less dispersed as the time to maturity of the contract decreases. As far as the second proposition is concerned, the results in Table 4 suggest that the null hypothesis of an MA unit root cannot be rejected in roughly half of all cases. Moreover, when contracts that have traded since 1990 are considered, the median KPSS test statistic indicates an inability to reject the null. Thus, in general, basis changes have an MA unit root.

Finally, the first-order autocorrelation in basis changes is calculated. The results indicate that this autocorrelation is significantly less than zero in almost all cases. Moreover, the autocorrelation is close to its predicted value of $-\frac{1}{2}$. For instance, the median value of this autocorrelation equals -0.4 when one considers contracts that have traded since 1990. An explanation as to why this correlation is greater than $-\frac{1}{2}$ is provided in Figure 7, which gives the distribution of basis changes based on a 10,000-repetition Monte Carlo simulation. Eighty values of the basis are generated using (5.12). The error term, ϵ_t, is drawn from an **i.i.d.** normal distribution, and μ_t is drawn from a normal distribution with first-order autocorrelation values of $\rho_\mu = \{0, 0.2, 0.4, 0.6\}$. The first two moments of these series are based on their empirical counterparts.[15]

---

[13]In general, an MA(1) process with coefficient θ has first-order autocorrelation equal to $\theta/(1 + \theta^2)$. Equation (5.14) shows that $\theta = -1$ and hence first-order autocorrelation equals $-\frac{1}{2}$.

[14]See Kwiatkowski et al. 1992 for more details of this test.

[15]The series used were obtained from LIFFE, Goldman Sachs, and Datastream and cover the period from January 1, 1990, to December 31,

Table 4
Interday summary statistics

Contract	Volume	F_H	KPSS	$\hat{\rho}_1$
Jun 1986	606.2750	3.3842***	0.1068	−0.2679**
Sep 1986	357.7875	0.4714	0.1442*	−0.2060*
Dec 1986	411.4750	1.5130	0.1450**	−0.1815*
Mar 1987	684.4500	0.7442	0.0631	−0.4415***
Jun 1987	1091.5625	3.4802***	0.1219*	−0.3149***
Sep 1987	1541.8875	0.8210	0.1617**	−0.3066***
Dec 1987	1773.0250	2.7568***	0.1426*	−0.3762***
Mar 1988	1039.6625	1.4733	0.1902**	−0.5537***
Jun 1988	999.1500	3.0976***	0.1975**	−0.1697
Sep 1988	1234.3125	2.3194***	0.1608**	−0.4990***
Dec 1988	1817.0250	1.1456	0.0864	−0.3566***
Mar 1989	1947.8375	1.0834	0.1874**	−0.3424***
Jun 1989	2170.9125	1.9035**	0.1024	−0.2663**
Sep 1989	1902.2750	7.6173***	0.1241*	−0.2318**
Dec 1989	2510.6125	1.9720**	0.1687**	−0.5318***
Mar 1990	2410.9750	0.0298	0.1559**	−0.5002***
Jun 1990	2297.4375	1.6446*	0.1170	−0.3975***
Sep 1990	2567.4375	0.4052	0.0743	−0.3856***
Dec 1990	3193.5750	2.6993***	0.0569	−0.2505**
Mar 1991	2268.2750	2.3854***	0.1768**	−0.2534**
Jun 1991	2302.2375	1.5327	0.0683	−0.3610***
Sep 1991	3050.0625	1.4504	0.0785	−0.3721***
Dec 1991	4964.7000	0.6786	0.1130	−0.3786***
Mar 1992	3148.2625	2.2938***	0.1319*	−0.3295***
Jun 1992	3184.8750	1.3014	0.1374*	−0.4020***
Sep 1992	5396.6625	0.7200	0.2104**	−0.4713***
Dec 1992	4695.4000	3.4799***	0.0680	−0.4027***
Mar 1993	5253.7875	1.3722	0.1330*	−0.3849***
Jun 1993	4018.5500	2.4891***	0.0669	−0.4119***
Sep 1993	5074.7000	1.7888**	0.1413*	−0.1998*
Dec 1993	7683.8750	0.6811	0.0825	−0.2884***
Mar 1994	7683.7625	0.5199	0.1013	−0.4624***
Jun 1994	6745.8625	0.9111	0.1197	−0.4136***
Sep 1994	5178.3875	0.9566	0.1785**	−0.5630***
Dec 1994	5225.2375	1.8901**	0.0984	−0.4399***
Mar 1995	7110.8750	1.2728	0.1032	−0.3550***
Jun 1995	7374.8375	3.3138***	0.1845**	−0.5203***
Sep 1995	4278.5875	2.4720***	0.0654	−0.4215***
Dec 1995	7665.9125	1.9503**	0.0654	−0.4051***
Mar 1996	5286.6250	1.5515	0.1223*	−0.4816***
Jun 1996	5928.7125	2.5575***	0.1400*	−0.3699***
Sep 1996	5667.3250	1.4341	0.1456*	−0.4162***
Dec 1996	8244.0250	0.3804	0.1214*	−0.4647***
Mar 1997	5870.8750	2.3940***	0.1365*	−0.3411***
Jun 1997	6553.6625	1.1525	0.2040**	−0.3728***
Sep 1997	6920.0500	0.8113	0.0747	−0.3041***
Dec 1997	6434.3375	1.1359	0.0918	−0.4606***
Mar 1998	4461.9250	0.8004	0.2155**	−0.4666***
Jun 1998	17035.0500	1.1564	0.0565	−0.4041***
Sep 1998	20012.2000	0.6585	0.0934	−0.5773***
Dec 1998	21300.2375	1.1150	0.0749	−0.5652***
1986–1989	1234.3125	1.9035***	0.1442*	−0.3149***
1990–1998	5270.2063	1.3368	0.1150	−0.4034***

Note: Average daily futures volume, the Goldfeld-Quandt test statistic (F_H), the KPSS test statistic, and estimated first-order autocorrelation ($\hat{\rho}_1$) applied to daily basis are calculated. Median statistics are given in the last two rows. Significance is denoted by *** (1% significance), ** (5% significance), and * (10% significance).
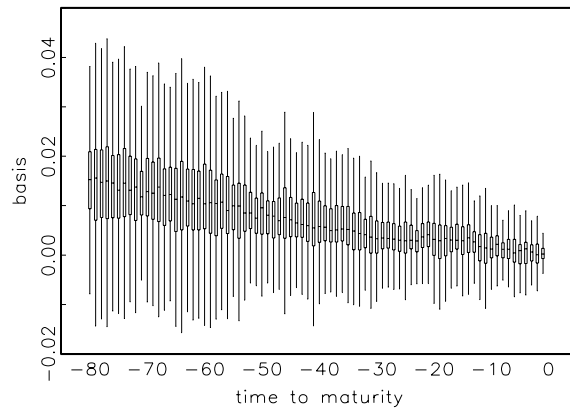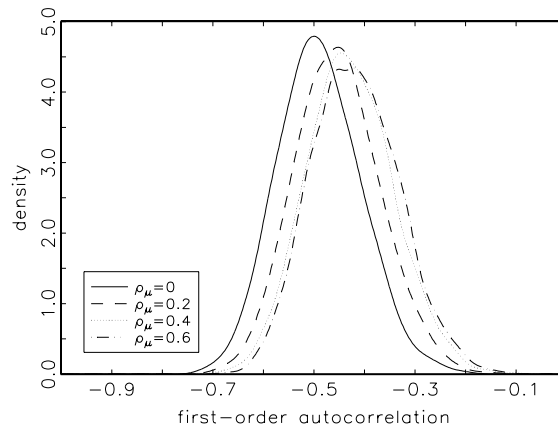
Figure 6
Interday basis convergence.

Figure 7
Interday basis change autocorrelation.

As predicted by the theory, the distribution of the first-order autocorrelation of basis changes is centered around $-\frac{1}{2}$ when $\rho_\mu = 0$. When the degree of autocorrelation in μ_t increases, however, this series starts to resemble a nonstationary series. For instance, if ρ_μ is set equal to unity, then the resulting basis becomes a mix of a stationary series (ϵ_t) and a nonstationary series (μ_t). If μ_t dominates, then any differencing yields a stationary series, and the resulting autocorrelation will approach zero. The weights assigned to this mixture of

1998. The price of the nearest futures contract is taken as the futures price and was obtained from LIFFE. The theoretical price of the futures contract is obtained using (2.1) and is based on the Goldman Sachs data set, and annual dividend yields and interest rates were obtained from Datastream. When the futures contract matures, the next nearest contract is used.

errors are given by the respective variances of the two error terms. The empirical evidence suggests that the stationary error dominates and, as such, the first-order autocorrelation in basis changes is only slightly less than $-\frac{1}{2}$. For instance, when $\rho_\mu = 0.6$, the median value of the autocorrelation is -0.42. This compares to an empirical median value of -0.4 when all contracts traded between 1990 and 1998 are considered. Moreover, this ρ_μ value closely resembles the empirical estimate of the autocorrelation in μ_t. Therefore, using realistic parameter values and the assumption of **i.i.d.** mispricing, the above theory seems to be able to predict the observed first-order autocorrelation in basis changes with a considerable degree of accuracy.

6 Concluding Remarks

In this article we have examined whether observed intra- and interday first-order autocorrelation in mispricing changes and basis changes, respectively, for the FTSE 100 Index futures market is a result of arbitrage behavior or a manifestation of market microstructure effects such as nontrading in the underlying stock index. The microstructure explanation of predictability (see Miller, Muthuswamy, and Whaley 1994) revolves around the proposition that in the absence of arbitrage opportunities, mispricing should follow a random walk, and hence mispricing changes should exhibit zero first-order autocorrelation. Miller, Muthuswamy, and Whaley (1994) argue that the level of arbitrage activity observed is insufficient to generate predictability in mispricing changes and hence such predictability must be a reflection of the autocorrelation nontrading induces in the underlying stock index.

More recent studies of arbitrage behavior have focused on the implications of transaction costs for the behavior of mispricing. Since the presence of arbitrage opportunities is determined by transaction costs, recent studies of the behavior of index futures mispricing have made use of TAR models, since they allow mispricing to behave differently according to whether there are profitable arbitrage opportunities present or not. The general finding of these studies is that when arbitrage opportunities are not present, mispricing follows a near random walk, whereas it follows a stationary autoregressive process when arbitrage opportunities are present. The first question we address is whether this *conditional* predictability in mispricing changes that the TAR model implies translates into *unconditional* predictability in mispricing changes on an intradaily basis.

We find that microstructure effects cannot explain the observed first-order negative autocorrelation in minute-by-minute mispricing changes. Indeed, even though adjusting for microstructure effects purges autocorrelation from observed spot and futures returns, the negative autocorrelation in mispricing changes becomes more pronounced after microstructure effects are adjusted for. When we consider the arbitrage explanation via the use of the TAR model, however, we find that the TAR model is capable of generating unconditional negative first-order autocorrelation in mispricing changes that is very similar to that which is observed empirically. Indeed, once arbitrage effects are accounted for, we find that unpredictable mispricing changes result.

We show that predictability in daily basis changes is the result of neither microstructure effects nor arbitrage activity. Rather, it is a statistical illusion that arises because of overdifferencing a series that is stationary around a trend.

References

Abhyankar, A., L. Copeland, and W. Wong (1999). "LIFFE cycles: Intraday evidence from the FTSE 100 stock index futures market." *European Journal of Finance*, 5: 123–139.

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.

Chan, K., J. Petrucelli, H. Tong, and S. Woolford (1985). "A multiple threshold AR(1) model." *Journal of Applied Probability*, 22: 267–279.

Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb (1978). "The returns generation process, returns variance, and the effect of thinness in securities markets." *Journal of Finance*, 33: 149–167.

Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb (1979). "On the existence of serial correlation in an efficient securities market." *TIMS Studies in the Management Sciences*, 11: 151–168.

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

Dwyer, G., P. Locke, and W. Yu (1996). "Index arbitrage and nonlinear dynamics between the S&P500 futures and cash." *Review of Financial Studies*, 9: 301–332.

Fisher, L. (1966). "Some new stock market indexes." *Journal of Business*, 39: 191–225.

Hansen, B. (1997). "Inference in TAR models." *Studies in Nonlinear Dynamics and Econometrics*, 2: 1–14.

Hansen, B. (1999). "Testing for linearity." *Journal of Economic Surveys*, 13: 551–576.

Hansen, B. (2000). "Sample splitting and threshold estimation." *Econometrica*, 68: 575–603.

Kwiatkowski, D., P. Phillips, P. Schmidt, and Y. Shin (1992). "Testing the null hypothesis of stationary against the alternative of a unit root." *Journal of Econometrics*, 54: 159–178.

Leybourne, S. (1995). "Testing for unit roots using forward and reverse Dickey-Fuller regressions." *Oxford Bulletin of Economic and Statistics*, 57: 559–571.

Lim, K. (1990). "Arbitrage and price behaviour of Nikkei stock index futures." *Journal of Futures Markets*, 12: 151–161.

Lo, A., and C. MacKinlay (1990). "An econometric analysis of non-synchronous trading." *Journal of Econometrics*, 45: 181–212.

MacKinlay, A., and K. Ramaswamy (1988). "Index-futures arbitrage and the behaviour of stock index futures prices." *Review of Financial Studies*, 1: 137–158.

Martens, M., P. Kofman, and T. Vorst (1998). "A threshold error-correction model for intraday futures and index returns." *Journal of Applied Econometrics*, 13: 245–263.

Miller, M., J. Muthuswamy, and R. Whaley (1994). "Mean-reversion of Standard & Poor's 500 index basis changes: Arbitrage-induced or statistical illusion." *Journal of Finance*, 49: 479–513.

Naik, N., and P. Yadav (1999). "The effects of market reform on trading costs of public investors: Evidence from the London Stock Exchange." Working paper. London: London Business School.

Neal, R. (1996). "Direct tests of index arbitrage models." *Journal of Financial Quantitative Analysis*, 31: 541–562.

Roll, R. (1984). "A simple implicit measure of the effective bid-ask spread." *Journal of Finance*, 39: 1127–1139.

Shah, N. (1999). "The European equity trading day: When should we trade?" Working paper. London: Equity Derivatives Research, Goldman Sachs International.

Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall.

Sofianos, G. (1993). "Index arbitrage profitability." *Journal of Derivatives*, 1: 6–20.

Stoll, H., and R. Whaley (1990). "The dynamics of stock index and stock index futures returns." *Journal of Financial and Quantitative Analysis*, 25: 441–468.

Taylor, N., D. van Dijk, P. Franses, and A. Lucas (2000). "SETS, arbitrage activity, and stock price dynamics." *Journal of Banking and Finance*, 24: 1289–1306.

Teräsvirta, T. (1994). "Specification, estimation, and evaluation of smooth transition autoregressive models." *Journal of the American Statistical Association*, 89: 208–218.

Tong, H. (1983). *Threshold Models in Nonlinear Time Series Analysis*. New York: Springer-Verlag.

Tong, H. (1990). *Nonlinear Time Series: A Dynamical Systems Approach*. Oxford: Oxford University Press.

Tong, H., and K. S. Lim (1980). "Threshold autoregression, limit cycles and cyclical data." *Journal of the Royal Statistical Society B*, 42: 245–292.

Tsay, R. (1989). "Testing and modeling threshold autoregressive processes." *Journal of the American Statistical Association*, 84: 231–240.

Tse, Y. (1999). "Market microstructure of FTSE 100 index futures: An intraday empirical analysis." *Journal of Futures Markets*, 19: 31–58.

Yadav, P., and P. Pope (1990). "Stock index futures arbitrage: International evidence." *Journal of Futures Markets*, 10: 573–603.

Yadav, P., and P. Pope (1994). "Stock index futures mispricing: Profit opportunities or risk premia?" *Journal of Banking and Finance*, 18: 921–953.

Yadav, P., P. Pope, and K. Paudyal (1994). "Threshold autoregressive modeling in finance: The price differences of equivalent assets." *Mathematical Finance*, 4: 205–221.

Editorial Policy

The *SNDE* is formed in recognition that advances in statistics and dynamical systems theory may increase our understanding of economic and financial markets. The journal will seek both theoretical and applied papers that characterize and motivate nonlinear phenomena. Researchers will be encouraged to assist replication of empirical results by providing copies of data and programs online. Algorithms and rapid communications will also be published.