Migration of Price Discovery with Constrained Futures Markets?

by

Anthony D. HALL

Paul KOFMAN
School of Finance and Economics
University of Technology, Sydney
PO Box 123
Broadway, NSW 2007
Australia

and

Steve MANASTER
College of Business Administration
Graduate School of Business Administration
University of Colorado at Boulder
Campus Box 419
Boulder, CO 80309-0419
USA

17 October, 2001

Abstract:

This paper investigates the information content of futures option prices when the futures price is regulated while the futures option price itself is not. The New York Board Of Trade provides the empirical setting for this type of dichotomy in regulation. Most commodity derivatives markets regulate prices of all derivatives on a particular commodity simultaneously. NYBOT has taken a fairly unique position by imposing daily price limits on their futures contracts while leaving the options prices on these futures contracts unconstrained. We are particularly interested in the volatility and futures prices of the options-implied risk neutral density when the underlying futures contract is locked limit.

Keywords: option implied density, price limits

JEL-code: G13

? The authors gratefully acknowledge data support provided by the New York Board of Trade. The views expressed in this article are, however, those of the authors and do not necessarily reflect those of the New York Board of Trade. Special thanks are due to Joshua Rodgers and Bill Delaney from NYBOT for data support and comments and Thuy Duong To for excellent research assistance. Hall and Kofman acknowledge financial support from the ARC large grant scheme.

1 Introduction

Two of the three major economic functions of a derivatives market are price discovery and price dissemination. In contrast to financial derivatives markets, for agricultural derivatives markets the underlying cash commodity markets tend to be heterogeneous, hampering a unique price discovery and dissemination. At any one time there may be a multitude of prices depending on producer, quality, location, etc. A typical example is the spot price of cocoa, as reported by the New York Board of Trade on the 20th of June, 2001. On that day, the U.S. Cocoa Merchants' reported spot price varied from a high of $1,166 per metric ton (for main crop Ghana, Grade 1) to a low of $976 per metric ton (for Superior Season Arriba), both prices ex-dock, eastern seaboard. These prices get reported on a daily basis and often involve some kind of averaging. It is not clear whether they are customer or even transaction specific and may therefore not be relevant to the 'representative' customer.

Standardization of the traded commodity and concentration of trade provided by the derivatives markets lead to a convergence of market opinions regarding the 'representative' commodity's equilibrium price. This 'perfect' commodity then serves as a benchmark for market participants against which the cash product which is often of inferior quality (compared to the derivative standard), available at some distant location, can be accurately priced. Given this fundamental role, it seems surprising to observe that many commodity derivatives markets – unlike their financial counterparts – use price limits which effectively censor the range of price discovery. Hall and Kofman (2001) survey the major derivatives exchanges worldwide and find that two thirds use price limit regulation (mostly for agricultural derivatives). Price limits, according to Telser (1981), Lee, Ready and Seguin (1994) and many others, obstruct (by postponing) price discovery. Others, e.g., Ma, Rao and Sears (1989), argue that price limits may actually enhance price discovery by avoiding overreaction to fundamental price signals. Brennan (1986), who agrees with the critics, suggests that price limits may nevertheless play an important role in cost minimizing contract design. In fact, they can be a natural outcome of a standardized futures contract. Through the introduction of temporary uncertainty regarding true losses, traders will be more inclined to meet their margin call when marking-to-market. Of course, this beneficial function of price limits disappears as soon as traders receive a signal regarding the true, but unobserved loss. The stronger the signal the less effective price limits can hide the true loss. These signals could come from the cash market (though these markets tend to be opaque for many commodities as argued above), from nearest-delivery futures contracts for which the limits are lifted, or from related derivatives markets that operate without limits. To prevent these signals originating from related assets, most derivatives exchanges (e.g., the CBOT) simultaneously restrict the futures and futures options to trade within price limits. Of course, Brennan's argument is only valid for contracts where contract default, by losing traders not meeting the margin call, leads to substantial open positions for the clearing house. For futures, this implies that the cost of trade interruption caused by price limits is offset by the reduction in contract default. For futures options, this risk of contract default is nonexistent since the losing party will just let the option expire and, having already paid an upfront premium, not create any exposure. Price limits then impose a real cost on option trading. This would suggest that the futures options traders 'subsidize' the futures traders. Due to the nonlinear relationship between futures options prices and underlying futures prices, it is also tricky to set matching price limits in the futures options market at different exercise prices. Most exchanges do not tailor these price limits correctly, but rather use a fixed price range for all series. An easy way out of this dilemma is to just suspend trading in futures options whenever the futures price limits are invoked.

If the futures option prices are not limited, it is possible to use traded futures options prices to derive the implied futures price as a signal for the limit locked futures price. Due to its sheer volume, contract standardization, and liquidity, the futures market usually serves as the predominant source of price discovery for market participants in related (cash and options) markets. It seems only logical that the futures options market will claim this role whenever the futures market price discovery function gets obstructed by price limits. In practical terms this implies a directly observable migration of volume from the constrained to the unconstrained derivatives market. At the same time, the options implied futures price should provide a market signal for the unobservable, constrained, futures price.

Migration of price discovery due to price limits, has so far attracted little attention in the literature with two exceptions. The first paper by Evans and Mahoney (1996) estimates the implied cotton futures price from cotton futures options using the put-call-parity relation whenever the futures price is locked limit.This type of indirect inference based on observed market prices has become rather popular. Model based volatility structures are nowadays routinely compared to options implied volatilities. More recently researchers also derived implied higher order moments, like implied skewness and implied kurtosis, e.g., Martin, Forbes and Martin (2001). In fact, there now is a large literature that infers the full probability density function from traded options prices. The second paper that considers migration of price discovery, Melick and Thomas (1997), uses this approach to infer the implied probability density function for crude oil futures prices from American crude oil option prices.

We extend these papers in a number of ways. First, unlike Evans and Mahoney (1996), we take account of the fact that the futures options are in fact American-style. The early exercise premium turns out to be non-trivial and, unfortunately, requires us to abandon the simple put-call-parity as a tool to 'back out' the implied futures price. Melick and Thomas develop a method which combines a mixture of lognormal distributions with no-arbitrage bounds for American options to infer the implied martingale equivalent density of futures prices. Data limitations prevent us from replicating this approach. Instead we use the single lognormal Barone-Adesi Whaley approximation for American options. Second, unlike Evans and Mahoney (1996), but similar to Melick and Thomas, we extract the implied futures density, not just the implied futures price. Thus we can compare implied against observed (limit-locked) futures prices, as well as the implied volatility behaviour surrounding price limit moves. Melick and Thomas allow more flexible densities than we do, but once again, we believe that most empirical applications do not allow for this level of sophistication. Third, we apply our analysis to intraday futures and futures options data, instead of using end-of-day settlement prices as in the Evans and Mahoney (1996) and Melick and Thomas papers. We claim more accurate identification of the price limit distortion, which is essentially an intraday phenomenon. This allows us to expand our sample size (by including intraday limit episodes) while avoiding non-synchronicities in the futures options and futures data. Fourth, we carefully construct control samples to avoid spurious conclusions, based on either model misspecification or non-synchronous option prices. These control samples are taken when the futures price is still within the limits on the same day as the limit move occurrence. Obviously, we do need intraday transactions data to achieve this. Finally, both other studies are based on short sample periods, nine months in Melick and Thomas, respectively one month in Evans and Mahoney (1 month). Whereas Melick and Thomas choose this sample to focus on an unusual episode (i.e., the Persian Gulf Crisis in 1990/91), Evans and Mahoney's sample seems to be unnecessarily restrictive. Our sample spans 7 years of data including two years when the futures price limits were officially lifted. This creates an additional control sample where we investigate so-called phantom (or pseudo) limits.

The remainder of the paper is organised as follows. In Section 2 of the paper we briefly elaborate on the method chosen to extract the implied futures prices density function from traded futures option prices. We also discuss the (dis)advantages of a number of alternative methods that have recently appeared in the literature. Section 3 provides an empirical application to commodity futures and futures options traded at the New York Board of Trade. Sample selection issues are discussed in some detail. We conclude with Section 4.

2 Estimating implied futures prices from futures option prices

We consider futures and matching futures options contracts where the futures contracts are regulated by price limits, while the futures options contracts are allowed to trade without impediment. We therefore take the market prices for the futures options as 'equilibrium' prices and estimate the implied pricing kernel that is consistent with current market valuation of the underlying futures asset. As the futures options we consider are American-style, we will have to allow for early exercise[1]. The Barone-Adesi and Whaley (1987) approximation is a common approach adopted in the literature to price American options. Alternative, more accurate, methods to check for early exercise of American futures options are based on the binomial tree methodology (see e.g., Rubinstein, 1994, Jackwerth and Rubinstein, 1996)[2]. Starting at the expiration date, a unique risk-neutral stochastic process is identified by recursively calculating the up-move probabilities at each node in the recombining tree.

A novel (continuous time) implied density estimation methodology is the so-called mixture model, which accounts for non-normality of the returns and possibly asymmetric features of the data. Ritchey (1990) assumes that the risk-neutral distribution is a mixture of lognormals and provides an implied density estimator for European options. Söderlind and Svensson (1997) and Melick and Thomas (1997) extend the mixture model to cope with American options. To account for the early exercise premium, they treat the observed option prices as weighted averages of upper and lower no-arbitrage bounds. This combination of bounds, weights and parameters of the mixture distribution can then be estimated by non-linear least squares. Their methodology allows for skewness and excess kurtosis in the implied *pdf* of the underlying futures price returns. Time-varying volatility, and asymmetric response to price innovations are well-known empirical phenomena for futures price returns which may cause these distortions to normality in the *pdf*.

Unfortunately, these methods require a sufficient number of 'relevant' observations, i.e., a sufficiently large range of exercise price series, to robustly identify the tails (and hence the higher order moments) of the implied futures *pdf*. Our data series and focus on price limit episodes do typically not allow for this precision. Since we only want to include options that

[1] Since our futures options will typically have at least four months to maturity, the early exercise adjustment is non-trivial unlike in Söderlind and Svensson's (1997) examples.

actually trade in price limit intervals, we are typically restricted to those options that are close to being *at-the-money*. There does not seem to be much point in trying to identify implied fat-tailedness if the tails are not sufficiently represented in the data. Melick and Thomas' use of settlement prices (and bid/ask prices in the absence of actual transactions for certain exercise price series) artificially avoids this problem.

Our approach is therefore summarized as follows. For a representative limit episode, we observe a sample of N traded option prices (puts and calls), $C_i^M, i ? 1,..,N$. We assume that there is also an unobservable arbitrage-free price $C_i ? ?, i ? 1,..,N$ for each traded option, which is a function of a set of parameters, $?$. We would expect the market price to equal the arbitrage-free theoretical price, but Jacquier and Jarrow (2000) suggest that there may be two sources of pricing errors. The model, though theoretically correct, still depends on a set of parameters that will typically be estimated with error. Given that our options are American, and the model is only an approximation to the true arbitrage free equilibrium price, we can expect further errors. These are classified as model errors. In addition to these model errors, it is also possible that the market option prices are observed with error, or the market may make occasional mistakes. These are classified as market errors. In the next section, when we discuss sample selection, we explain how we attempt to control for these errors. For now, the errors are combined, such that

$$C_i^M ? C_i ? ?? ? ?_i \tag{1}$$

and we assume normally distributed pricing errors: $?_i$ ~ N(0,$?_?^2$). An important consideration is given by the fact that the error distribution is truncated due to the existence of no-arbitrage boundary conditions on the traded option prices, see Martin, Forbes and Martin (2001). The (relevant) truncation bound for an American futures call option is given by $lb_i ? \max ?0, f_i ? e^{?rt} X_i ?$, the truncation bound for an American futures put option is given by $lb_i ? \max ?0, X_i ? f_{t,i} ?$[3]. The log likelihood to maximize with respect to the parameters $?$ is then given by

$$L ? ?? ? ?_?^{?N} \prod_{i?1}^{N} \exp ? ? \frac{1}{2?_?^2} ?C_i^M ? C_i ? ??^2 ? \Big/ ?1 ? ? ?lbs_i ?? \tag{2}$$

where

[2] Note that there is an extensive literature on implied density estimation from traded option prices. We only skim this literature by referring to a handful of papers that are relevant for our purpose.

[3] Since we use puts and calls simultaneously, we should also consider the put-call-parity restriction. However, since the options are American, these bounds would be less restrictive.

$$\sigma_{\varepsilon}^{2} = \frac{1}{N-1}\sum_{i=1}^{N}\left(C_i^M - C_i(\cdot)\right)^2$$

and

$lbs_i = \dfrac{lb_i - C_i(\cdot)}{\sigma_{\varepsilon}}$ evaluated at the standard normal cumulative density function, Φ.

Given the data restrictions, our benchmark specification for $C_i(\cdot)$ is the Black-Scholes model for commodity futures options, Black (1976). Of course, since the options are American, we have to allow for early exercise value and therefore choose the Barone-Adesi and Whaley (1987) approximation (BAW). The arbitrage-free price of an American call option is then

$$\begin{aligned} C_i\left(E_t(f_0),\sigma,X,t\right) &= c_i\left(E_t(f_0),\sigma,X,t\right) + A_2\left(f^*,\sigma,X,t\right)\left(\frac{E_t(f_0)}{f^*}\right)^{q_2(\sigma,t)} && \text{for } E_t(f_0) < f^* \\ C_i\left(E_t(f_0),\sigma,X,t\right) &= E_t(f_0) - X && \text{for } E_t(f_0) \geq f^* \end{aligned} \qquad (3)$$

with the following notation:

f_0 = futures price at expiration of the option,

f^* = 'critical' futures price that triggers early exercise,

X = exercise price of the option,

σ = standard deviation of futures price returns,

r = risk-free rate of return,

t = time to expiration of the option,

E_t = expectations operator, t periods prior to expiration.

and

$$c_i\left(E_t(f_0),\sigma,X,t\right) = E_t(f_0)e^{-rt}\,\Phi\left(d_1\left(E_t(f_0),\sigma,X,t\right)\right) - Xe^{-rt}\,\Phi\left(d_2\left(E_t(f_0),\sigma,X,t\right)\right) \qquad (4)$$

is the standard Black-Scholes expression for a European futures call option,

$$A_2\left(f^*,\sigma,X,t\right) = \frac{f^*}{q_2(\sigma,t)}\left\{1 - e^{-rt}\,\Phi\left(d_1\left(f^*,\sigma,X,t\right)\right)\right\}$$

$$q_2(\sigma,t) = \frac{1 + \sqrt{1 + \dfrac{8r}{\sigma^2\left(1 - \exp(-rt)\right)}}}{2}$$

$$d_1\left(E_t(f_0),\sigma,X,t\right) = \frac{\ln\left(\dfrac{E_t(f_0)}{X}\right) + \dfrac{\sigma^2 t}{2}}{\sigma\sqrt{t}}, \quad d_2\left(E_t(f_0),\sigma,X,t\right) = d_1\left(E_t(f_0),\sigma,X,t\right) - \sigma\sqrt{t}$$

and based on the assumption of a log normal distribution for the futures price

$$E_t ? f_0 ?? \exp ? ? ? \frac{?^2}{2} ? \tag{5}$$

To find the critical futures price, we solve for it implicitly through

$$f^? ? X ? c_i ? f^?, ?, X, t ?? \frac{? 1 ? e^{?rt} ? ? d_1 ? f^?, ?, X, t ??? f^?}{q_2 ? ?, t ?} \tag{6}$$

Hence, the parameter set for which we optimize (2) is $? = ?, ?, f^*$. It should be clear from the above that unlike other studies that take the underlying asset price f as given, e.g., to find the best fitting model, our aim is to find option implied values for the underlying futures price f (and for $?$). These 'parameter' values can then be used to assess the impact of the underlying futures price being locked at a price limit.

Given typical empirical characteristics like fat tailedness and/or skewedness of futures returns, the theory-to-market fit in (1) may not be optimal when assuming a single log normal distribution for the futures price as in BAW. More complicated mixtures of log normal distributions (discussed above) or more flexibly parameterized distributions have been found to outperform the BAW approximation. That may be less of a problem for our specific purpose since we are primarily interested in the first two moments of the distribution which are not necessarily distorted by the fat-tailed phenomenon. In any case, data considerations prevent sophisticated analysis beyond the single lognormal BAW approximation for most traded commodity options.

3 NYBOT Commodity Futures – an Application

We investigate limit occurrences for commodity futures contracts traded on the New York Board of Trade (NYBOT). NYBOT was created in June 1998 to become the parent company of the Coffee, Sugar and Cocoa Exchange (CSCE) and the New York Cotton Exchange (NYCE). Price limits existed for CSCE's cocoa, coffee "C", and sugar-11 futures contracts until December 1997, as of today they still exist for NYBOT's cotton and frozen concentrated orange juice futures contracts[4]. To facilitate a "before/after" comparison (i.e., with and without price limits), we restrict our application to the cocoa, coffee, and sugar futures contracts. All three futures contracts have frequently encountered price limit moves during our sample period. Because our investigations require a special circumstance, futures locked at limit, in order to generate an adequate sample size we require as much history as possible and for that reason our samples are based on intraday data taken from the period 1993 through

[4] NYBOT also uses limits for the NYSE Composite and Russell 1000® Index futures contracts.

1999. The futures and futures options transactions data for this study are obtained directly from the NYBOT trade records. Treasury bill rates for maturities matching the option series expiration as close as possible, have been obtained from Datastream.

INSERT TABLE 1

Table 1 gives descriptive statistics for the continuously compounded daily futures returns during our sample period, 1994-1999[5]. A further split has been made for the maturity of the futures contracts. A few surprising results appear. Unlike many empirical finance studies, there is little evidence of non-normality. Skewness is limited and generally insignificant, kurtosis is excessive for the coffee contract, but to a much a smaller extent than is commonly found. The scale of these empirical return distributions is rather excessive in comparison with financial asset returns. The coffee futures returns have an annualized standard deviation in the range 45-50% across maturities, sugar and cocoa futures returns have a more modest annualized standard deviation in the range of 18-28%, and 22-26% respectively. A comparison of maturities suggests that the empirical distributions are reasonably similar for sugar and cocoa futures returns. However, the distributions seem to "narrow" for longer maturities. This reduction in standard deviation also appears in coffee futures returns. These, however, also display significant non-normality (in terms of skewness and excess kurtosis) for the nearest and next-to-nearest maturities. The longer maturities' distributions appear more normal.

If we compute the standard deviation of daily futures price changes for the nearest maturity contracts, we find US$cts 4.85 (coffee), US$cts 0.17 (sugar), and US$20.95 (cocoa). Now compare this to the daily price limits, which can be found in the contract specifications in the Appendix. It takes 1.2 daily standard deviations to hit the coffee limits, 3 daily standard deviations to hit the sugar limits, and 4.2 daily standard deviations. Hence, fairly extreme events for sugar and cocoa, but a rather frequent event for coffee. Not surprising therefore that we observe 193 limit days (out of 1245 trading days, which is about 15%) for coffee, and only 25 limit days for sugar, respectively 3 limit days for cocoa.

INSERT TABLE 2

---

[5] Unfortunately, we did not have access to the 1993 daily futures prices.

Table 2 illustrates this impact of higher volatility (this causality statement is made tongue-in-cheek) on the frequency of limit moves. It seems obvious that we focus our attention on the coffee futures contract. For the 193 limit days, we observe 429 contracts that locked limit. Of these 429 contracts, only 15 were nearest-maturity or 2nd nearest-maturity, 322 were 3rd , 4th or 5th nearest-maturity. This is somewhat interesting if we refer back to Table 1, where we concluded that the 1st and 2nd nearest-maturity contract displayed higher volatility and fatter tails than the almost normally distributed further-out maturities.

Table 2 also gives a breakdown of the direction of the limit moves and their dating. The number of up and down limit days has been fairly balanced with 115 up moves and 107 down moves across the three futures contracts. The number of up and down limit contracts are similarly balanced at 257 up moves, respectively 225 down moves across the three futures contracts. The intertemporal spread has not been so balanced with 1994 and 1997 accounting for 75% of limit days, and 81% of limit contracts. When we compute the daily standard deviation for annual samples (nearest maturity contracts for coffee) we find for 1994: US$cts6.26, 1995: US$cts3.15, 1996: US$cts2.63, 1997: US$cts7.83, 1998: US$cts3.39, and 1999: US$cts3.68. Hence, the peak limit years coincide with a standard deviation about twice as high as in the remaining years.

INSERT FIGURE 1

This clustering of limit moves is clearly visible in Figure 1 for coffee “C”, and to some extent also for sugar-11. The 1994 and 1997 limit clusters for coffee and the 1995 limit clusters for sugar coincide with the coffee “C” futures price and the sugar-11 futures price both being at peak levels. Despite this illusion of the futures price being “up,” we still find almost as many down limit moves as up limit moves during those episodes.

3.1 *Limit and control sample selection*

Among the many limit moves recorded in Table 2, a substantial number occurred intraday and sometimes lasted only for a short time interval. Prices would therefore not necessarily lock limit for the full trading day or even close at the limit. The latter occasions would not appear in end-of-day settlement data used in other studies, including Melick and Thomas (1997) and Evans and Mahoney (1996). They argue that the settlement procedure avoids the problems typically associated with non-synchronous quotes inherent in transactions data. This non-synchronicity could arise e.g., when the option prices, from which implied futures prices are inferred, were based on morning transactions when the futures prices only locked at the limit in the afternoon. We argue that this problem is not restricted to intraday transactions data but will equally affect end-of-day settlement prices. The bid/ask quotes used to derive the option settlement prices are frequently found to be 'stale' for options that do not trade intraday, and therefore may give a misleading impression of synchronicity. In fact, by using intraday futures and options transactions data we can actually control much more carefully for the (non)synchronicity of these quotes. Of course, by construction the settlement data tend to be less "noisy" (they reduce e.g., the bid-ask noise by taking the midpoint) than transactions data. By including carefully constructed control samples, we intend to minimize this disadvantage of transactions data.

Including (temporary) intraday limit move episodes – in addition to limit-close days – considerably increases (almost doubles) our sample size. However, we select samples of effective limit intervals only if there is a sufficient number of traded options (with a sufficient spread in exercise prices) during the limit-lock period. Furthermore, these options have to be traded sufficiently 'close-in-time' to each other to guarantee that they reflect a unique implied density. This rather restrictive choice limits our sample size, but is expected to enhance the reliability of our findings. The number of available samples that qualify is indicated between parentheses in Table 2, column labelled "Total." Out of 193 coffee limit days (429 limit contracts), we count 62 (101) useful samples. Of these 101 individual samples, 45 were temporary intraday limit samples, 56 locked limit at the close of trading.

INSERT FIGURE 2

Figure 2 compares the duration of the selected temporary, and locked limit episodes with the duration for the full sample of temporary and locked limit episodes. The duration patterns are reasonably similar but, for both temporary and locked limit intervals, the selected samples tend to be biased towards longer durations. The reason for this is obvious. Longer durations provide more opportunity to observe a sufficient number of traded options. The majority of the temporary limit intervals last very briefly. In fact, seventy-five percent of selected samples last less than 50 minutes (this is ninety percent of the full sample). However, twenty percent of selected temporary intervals still last for more than two hours without locking at the limit for the day. The pattern is distinctly different for the locked limit intervals. There are very few locked limit intervals that commence shortly before market close. The shortest selected locked limit interval lasted 21 minutes. Seventy-five percent last over one hour.

Our next step is to select control samples. Jacquier and Jarrow (2001) indicate that there are two potential error sources combined in ? in equation (1) driving a wedge between market option prices and 'equilibrium' prices: model error and market error. To avoid drawing spurious conclusions from implied futures prices that deviate from the limit lock futures price purely because of these errors, we generate two sets of control samples. The first set of control samples consists of so-called 'checking' intervals. We select a limit-free intraday interval on the same day as the effective limit interval. Just as for the effective limit intervals we need to compile a sufficient number of traded options with different exercise prices, preferably traded as close in time to each other as possible. However, unlike the effective limit intervals, we now also have to guarantee that the futures price remains constant (or nearly so) to get a meaningful implied expectation. In case we were unable to find a suitable control sample on the same day as the limit interval, we used the first available control interval on the subsequent trading day. This occurred 38 times, predominantly for down limit samples and, not surprisingly, mostly for limit lock days. For a few cases (and longer maturities), the gap between control and effective sample could be a week. For two (out of 101) effective samples, we were unable to find a suitable control sample. Though not displayed in Figure 2, the control samples tend to be short-lived due to the selection criteria.

The second set of control samples consists of 'phantom' limit intervals. Prior to December 1997, CSCE price limits were lifted from the nearest- and next-to-nearest delivery futures contract two business days prior to the delivery month of the nearest-delivery contract. We have already seen that our effective limit sample consists mainly of 3rd- to 5th-nearest maturity futures contracts. This provides an ideal control sample, allowing us to investigate exceedences of the inactive price limits for the 1st- and 2nd-nearest futures contract when longer-maturity contracts were constrained at their price limits. As of December 15, 1997, the CSCE removed price limits from its coffee, sugar and cocoa futures contracts altogether. We treat the latter years (post December 1997) in our sample as if the price limits were still in place and select intraday intervals during which these non-existent limits were exceeded. Selection of traded option prices during phantom limit intervals had to satisfy the same requirements as for the checking intervals. The required combination of a 'fairly' stable futures price to match with a sufficiently dispersed option sample does not leave us with many suitable intervals. Table 2 indicates that post-limit we can only use 5 out of the potential 257 phantom limit contracts. The selection of phantom intervals for the limit years is a little more fruitful. We find 56 phantom limit contracts (46 phantom limit days), 48 (39) of which occur in 1997.

Figure 2 also compares the duration of the selected phantom limit episodes with the duration for the full sample of phantom limit episodes. The duration patterns are now distinctly different. As for the control samples, the selected phantom limit samples are short-lived with no intervals exceeding 28 minutes. Hence, we have selected the phantom limit episodes from the bottom 80 percent of full sample durations.

With the samples selected, we next investigate the impact of partial or incomplete price limit regulation from three perspectives. First, we consider the migration of volume from the restricted market to the related unrestricted market. Second, we investigate the consequences on price discovery. Third, we analyze whether price limits have a beneficial or adversary impact on price volatility.

3.2 *Effects on trading volume*

When the futures market locks limit either temporarily or for the day, it has been suggested that trading volume migrates to related but unrestricted markets. Subrahmanyam (1994), for example, develops a theoretical model where triggering a circuit breaker (a temporary price limit) on the dominant market causes trading volume to shift to the satellite market (without a circuit breaker). Berkman and Steenbeek (1998) find empirical evidence for this phenomenon for the Nikkei stock index futures contract which is simultaneously traded on the Osaka Securities Exchange (OSE) and the Singapore International Monetary Exchange (SIMEX). Price limits on the OSE cause a migration of volume to the unconstrained SIMEX. Subrahmanyam (1994) suggests that this migration might be impeded by switching costs. We conjecture that it is considerably cheaper to switch between derivatives on a single market than it is to switch between markets. Migration of volume should therefore be straightforward from futures to futures options at the NYBOT. Evans and Mahoney (1997) provide evidence of such volume migration by plotting the fraction of trading day in limit against futures contracts traded for both markets. First they observe that daily futures volume declines significantly with the duration of the limit interval. It seems somewhat trivial to conclude this. A more interesting question would be whether futures volume ‘compensates’ for the trade suspension post limit interval. More interesting is their finding of a significant increase in trading volume in the options market during futures limit lock. When standardizing option volume into futures-equivalent volume, they find that total volume (futures and options combined) is virtually unchanged.

Instead of looking at daily volume on limit days, we investigate the exact intraday limit intervals and compare those with the phantom limit control episodes. Obviously, when the limits are inactive, the futures price can change without bound after it crosses a phantom limit. In comparison with effective limit intervals this implies that the phantom limit intervals tend to last much longer. Taking this into account, controlling for the duration of a limit interval, we not only compute nominal volume but also ‘standardized’ per-minute volume.[6] The results are given in Table 3.

INSERT TABLE 3

The effective limits results are based on 614 intervals. Note that this is substantially more than the 429 limit contracts in Table 2. The difference occurs because of multiple effective limit spells that are counted as a single limit contract in Table 2, but counted separately for the purposes of Table 3. Interestingly, the lower left part of Table 3 indicates that futures volume does not completely disappear during effective limit episodes. It does, however, dry up to such an extent that we needed a different count indicator to distinguish it from the phantom limit interval count. Fourty-five percent of the effective limit intervals have at least two futures transactions at the limit price. Ten percent have more than five futures transactions at the limit price. The mean number of futures transactions is 2.61. That compares to a mean of 52.87 futures transactions for the phantom limit intervals (of which there are 3,596). If we standardize for the duration, we find that the effective limit mean is 1.12 transactions per minute against the phantom limit mean of almost 10 transactions per minute. In light of the skewness of these count distributions, the median is even more informative. The standardized effective limit median is 0.14 per minute against the standardized phantom limit median of 5.69 per minute. Hence, a considerable drop in futures volume!

As to the question whether option volume picks up where futures volume drops off, consider the upper part of Table 3. Most strikingly, unlike futures volume, we do not need an ‘expanded’ count measure for phantom limit options volume. In fact, the count distributions look fairly similar. For both effective limit and phantom limit intervals, sixty percent have zero option transactions. Eighty-five percent of effective limit intervals have five or less option transactions. This is eighty-four percent for phantom limit intervals. Standardized by duration, there is some evidence that option volume is more (!) dispersed for phantom limit

[6] Note that volume is here defined as the number of transactions that occur within a particular time period. The size of the transaction is not taken into account. We find very little evidence of a change in the transaction size conditional on a limit move. We do, however, find significant changes in the number of transactions.

intervals than for effective limit intervals. This is reflected in the mean option volume measures. If we look at the median option volume, however, there is little between them!

We conclude that there is little (or rather, no) evidence to suggest a migration of volume from the constrained to the unconstrained market. This is somewhat unfortunate for our next exercise: implied futures price derivation. Given that we need two or more options at the very least to compute the implied density measures this leaves us with very few useful limit episodes as already discussed above in the sample selection section.

3.3 *Effects on price discovery*

Most academics agree that price discovery is (at the very least temporarily) obstructed by price limits. Brennan (1994) investigates to what extent 'external' (noisy) price signals alleviate the price discovery problem in a constrained futures market. High signal to noise ratios, as used in his model simulations, suggest that Brennan believes these signals to be typically strong (in fact he calls a correlation between signal and equilibrium futures price of 0.75 moderate!). Interestingly, Brennan suggests that these signals come predominantly from the underlying cash/spot market. Convergence of cash and futures price close to maturity then explains why limits are lifted in the delivery month (they become obsolete in Brennan's cost minimizing contract design) of the nearest contract. Unfortunately, the researcher typically has no access to a high-frequency cash commodity price for agricultural futures contracts. We argue that even for futures market practitioners, this quest for a unique underlying spot price may prove elusive. Instead, we believe that a 'stronger' signal may emerge from the unconstrained futures options price.

That said, given the observed lack of depth in the options market, we would not expect to get overly strong price signals indicating market direction during limit lock. We therefore selected limit episodes that contain a sufficient number of traded options to identify the implied futures price and implied volatility. It should be clear that our stringent sample selection criteria limit our sample size rather considerably. Ultimately, we end up with 101 useful effective limit intervals. We separate this sample into temporary limit intervals and locked limit intervals. Of the 45 temporary limit intervals, 12 are down limits and 33 are up limits. Of the 56 locked limit intervals, 19 are down limits and 37 are up limits. Implied futures pricing errors for the locked limit sample are given in Figure 3. These pricing errors are defined as the difference between the options implied futures price and the limit price. We expect positive pricing errors for up limits, and negative pricing errors for down limits.

INSERT FIGURE 3

The left-hand side of Figure 3 displays the down limit episodes (? markers on the lower axis). Options implied futures prices are given by square symbols (?) for the effective limit samples, and by triangle symbols (?) for the control samples. Note first that 15 (out of 19) effective down limit samples have a negative pricing error, and 36 (out of 37) effective up limit samples have a positive pricing error. The direction of market expectations seems to make sense. For the control sample we find 7 (out of 17) negative pricing errors for down limits, respectively 23 (out of 37) positive pricing errors for up limits. However, the up limit control pricing errors are generally (31 out of 37) smaller than the effective limit pricing errors. For the down limit sample, the control pricing errors are less distinguishable from the effective limit pricing errors. We only find a significant limit effect for the up limit episodes.

Implied futures pricing errors for the temporary limit sample are given in Figure 4. The fact that these were temporary trade disruptions suggests that the market subsequently resumed trading within the allowed price range. We would not (necessarily) expect a positive (negative) pricing error for up (down) limit episodes.

INSERT FIGURE 4

Thus, we find most implied pricing errors to be insignificantly different from zero (and frequently smaller than the control pricing errors). None of the down limit samples give any price direction signal. Only four (out of 33) up limit samples suggest an upward price expectation. The remainder do not give any price direction signal. This implies that there is not much evidence of a rebound/reversion in the futures price either. Under that scenario we would expect a negative (positive) pricing error for up (down) limit episodes.

Implied futures pricing errors for the phantom limit sample are given in Figure 5. For the up limit intervals, there is little evidence of a direction effect with an even spread of positive/negative pricing errors. For the down limit intervals there seems to be a rebound direction effect, with most pricing errors positive. Few of these tend to be significantly different from zero.

INSERT FIGURE 5

As mentioned above, we would a priori expect up (down) limit moves to be associated with positive (negative) pricing errors. However, from Figure 4 we concluded that prices sometimes revert inside the allowable price range. If the option traders correctly assess this probability, we are likely to find opposite pricing errors to those a priori expected. We next investigate the fit of the futures price expectation to the next available 'free' futures price. A scatterplot of implied futures price changes against next futures price changes is given in Figure 6.

INSERT FIGURE 6

For the temporary limit intervals we take the next off limit price inside the allowed price range on the same trading day. For the locked limit intervals we take the first traded price on the next trading day (there were no occasions when this price was immediately at the next limit). Figure 6a suggests that there is little evidence of a relationship between implied and observed futures price change (the relationship is clearly distorted by the single outlier) for temporary limit intervals. Figure 6b, on the other hand, suggests a reasonably strong positive relationship between implied and observed futures price change for locked limit intervals.

At the risk of presenting anecdotal evidence, it seems worthwhile to investigate the intraday price discovery in some detail. Figure 7 gives plots for three limit days with a mixture of successive temporary and locked limit intervals.

INSERT FIGURE 7

On the 17th of May 1994, the upper price limit was hit early after the opening of trading. The first temporary limit interval gives an implied futures price well above the limit price. However trading resumed within the price range and after a while the upper limit is hit for a second time. The second temporary limit interval now delivers an implied futures price just below the limit price. Subsequently, trading resumes at prices well within the price range, but the upper price limit is still hit on a number of occasions. The implied futures price successively converges to the limit price. Clearly, option traders' market expectations are in line with the observed price movements. At the end of the trading occurs at prices substantially below the upper limit price.

On the 13th of June 1994, the first traded futures price occurs at the upper limit. The options implied futures price, surprisingly, is somewhat below the upper limit price. Subsequently, trading resumes at, and then below, the upper limit price until the limit is hit for a third[7] time at noon. This last episode is a locked limit interval for which the implied futures price is well above the limit price. The next day trading resumes at a price even higher than this implied price.

The trading pattern on the 28th of May 1997, looks similar to the previously discussed limit day. The futures price hits the upper limit early morning (with an implied futures price just below the limit price). It then reverts back inside the price range but very quickly hits the upper limit again. The second implied futures price is now well above the limit price, and trading is interrupted for well over an hour. Then, trading resumes for a little while just inside the upper limit, before locking limit for the rest of the day. This resumption of trade leads to a downwards adjustment of the implied futures price, though it is still well above the limit price. The next day trading resumes at a price between the second and third implied futures prices.

Combining these pieces of evidence, we tentatively conclude that the futures options market does indeed provide (reasonably) accurate futures price signals when the futures market is constrained by price limits. Recall that this is despite the fact that options volume does not significantly increase. Of course, this may imply that the uncertainty surrounding the implied futures price expectation increases when the futures price locks limit. Also consider that we compute the risk-neutral expectation, not the ‘true’ expectation. Since we cannot observe whether there is a risk premium, or whether this risk premium depends on limit moves, we have to be somewhat cautious in drawing conclusions from our findings.

3.4 *Effects on volatility*

Most of the price limit literature has in fact focused on the impact on volatility. Ma, Rao and Sears (1989) find a moderating impact on treasury bond futures volatility. Lee, Ready and Seguin (1994) find evidence of excessive volatility subsequent to a trading halt on the New York Stock Exchange. McMillan (1991) finds similar evidence subsequent to a circuit breaker in the S&P500 index futures market. Of course, one has to be careful in drawing conclusions from a post-limit increase in observed volatility. Part of this increase reflects the spillover of unresolved volatility when the price limits were hit. Carefully constructed control samples (as

[7] Since the second limit episode did not meet our selection criteria, there is no implied futures price.

in Lee, Ready and Seguin who use “pseudohalts”) are needed to unravel this intertemporal distortion in volatility from the underlying change in volatility. We follow a somewhat similar procedure. We compare the options implied volatility during effective limit intervals with the options implied volatility during the matching control intervals. NYBOT publishes daily options implied volatility measured across At-The-Money, In-The-Money and Out-of-The-Money options. Due to the volatility ‘smile, ’ this average tends to be biased upwards. Since we restrict ourselves to traded options, we are automatically restricting the implied volatility computation to ATM or nearly-ATM options.

Figure 8 displays implied volatility for the temporary limit sample. All volatility series are measured as the annualized percentage standard deviation of futures returns. The (?) diamond markers indicate historical volatility as computed by NYBOT over the past 30 days of (nearest-maturity) futures settlement prices by NYBOT. In comparison with the implied volatility derived from the limit sample (?) as well as implied volatility derived from the control sample (?), historical volatility is substantially higher. More importantly for our purposes, limit and control volatilities are typically very close, with a few exceptions.

INSERT FIGURE 8

Figure 9 displays implied volatility for the locked limit sample. Historical volatility still exceeds the implied volatility measures. Now, however, control and limit volatilities are clearly different. Locked limit implied volatility consistently exceeds control implied volatility.

INSERT FIGURE 9

Figure 10 displays implied volatility for the phantom limit sample. We compare these measures with historical volatility and implied volatility as computed and published by NYBOT. There is a fairly constant ‘gap’ between NYBOT’s daily measure and ours which is based on the phantom limit sample only. We suspect that this gap reflects the upward bias in NYBOT’s measure.

INSERT FIGURE 10

We conclude that volatility is affected by price limits, but a qualification applies. When the trade interruption is temporary, we observe no increase in volatility. When the futures price locks limit for the remainder of the trading day, we do however, observe increased volatility.

4 Conclusions

This paper considers the migration of price discovery from a price constrained derivatives market to a related, unconstrained, derivatives market. The typical research questions regarding price limits consider the obstruction of price discovery and their impact on the volatility of price changes. Based on implied futures price density estimates from traded options, we directly address both questions. Data restriction focus our application on Coffee "C" futures trading on the New York Board of Trade. When the futures price hits either limit, we derive the implied futures price and implied volatility from traded futures options prices that are not constrained. We separate our sample into temporary limit intervals and locked limit intervals. For the former, trading resumes on the same trading at futures prices within the allowed range. For the latter trading is interrupted for the day and trading only resumes the next trading day when the limits have moved. This distinction is important since we find that only for locked limit intervals the implied futures prices indicate that price discovery indeed shifts to the futures options market. This is particularly interesting since we find little evidence of an increase in traded options volume during futures limit lock. Another interesting feature of our results is that options implied volatility seems to increase only for the limit lock intervals. Implied volatility during temporary trade interruptions is indistinguishable from implied volatility during control periods when the market can observe the underlying futures price.

Two previous papers have dealt with this issue and found that both volume and price discovery migrate to the unconstrained market when limits are invoked on the constrained market. Whereas these two studies were based on event-like samples (a single day for one study), our analysis is more comprehensive in that it covers a five year sample of constrained futures trading as well as a subsequent period of two years when the limits were lifted. This provides an ideal before/after setting. Another important contribution of this study is its use of options transactions data, instead of the more common end-of-day settlement data. This allows to avoid non-synchronicities between futures and futures options prices as well as between different futures options series.

References

Barone-Adesi, Giovanni, and Robert Whaley, 1987, Efficient analytic approximation of American option values, *Journal of Finance*, 42, 301-320.

Berkman, Henk, and Onno W. Steenbeek, 1998, Daily price limits and trading in Nikkei futures, *Journal of Futures Markets* 18, 265-279.

Black, Fisher, 1976, The pricing of commodity contracts, *Journal of Financial Economics*, 3, 167-179.

Evans, Joan, E., and James M. Mahoney, 1996, The effects of daily price limits on cotton futures and options trading, Research Paper no.9627, Federal Reserve Bank of New York: New York.

Evans, Joan, E., and James M. Mahoney, 1997, The effects of price limits on trading volume: a study of the cotton futures market, *Current Issues in Economics and Finance*, vol.3, no.2 Federal Reserve Bank of New York.

Hall, Anthony D., and Paul Kofman, 2001, Regulatory tools and price changes in futures markets, *Australian Economic Papers*, forthcoming.

Jackwerth, Jens Carsten, and Mark Rubinstein, 1996, Recovering probability distributions from option prices, *Journal of Finance*, 51, 1611-1631.

Jacquier, Eric, and Robert Jarrow, 2000, Bayesian analysis of contingent claim model error, *Journal of Econometrics* 94, 145-180.

Lee, Charles M.C., Mark J. Ready, and Paul J. Seguin, 1994, Volume, volatility, and New York Stock Exchange trading halts, *Journal of Finance* 49, 183-214.

Ma, Christopher K., Ramesh P. Rao and R. Stephen Sears, 1989, Limit moves and price resolution: the case of Treasury Bond futures, *Journal of Futures Markets* 9, 321-335.

Martin, Gael M., Catherine S. Forbes, and Vance L. Martin, 2001, Implicit Bayesian inference using option prices, Working Paper, Monash University.

McMillan, Henry, 1991, Circuit breakers in the S&P 500 futures market: their effect on volatility and price discovery in October 1989, *Review of Futures Markets* 10, 248-274.

Melick, William R., and Charles P. Thomas, 1997, Recovering an asset's implied pdf from option prices: an application to crude oil during the Gulf crisis, *Journal of Financial and Quantitative Analysis*, 32, 91-115.

Ritchey, R.J., 1990, Call option valuation for discrete normal mixtures, *Journal of Financial Research* 13, 285-296.

Söderlind, Paul, and Lars Svensson, 1997, New techniques to extract market expectations from financial instruments, *Journal of Monetary Economics* 40, 383-429.

Rubinstein, Mark, 1994, Implied binomial trees, *Journal of Finance*, 49, 771-818.

Subrahmanyam, Avanidhar, 1994, Circuit breakers and market volatility: a theoretical Perspective, *Journal of Finance* 49, 237-254.

Whaley, Robert, 1986, Valuation of American futures options: Theory and empirical tests, *Journal of Finance*, 41, 127-150.

Appendix: CSCE Futures and Futures Options Contract Specs

COFFEE "C"

Price limits are set 6 cents above and below the previous settlement price (lifted on 15/12/1997).
<u>KC Futures</u>
Trading time 9:15am-1:32pm New York time
Price quotation in cents per pound
Delivery months: March, May, July, September, December
Minimum price fluctuation: 5/100 cent/lb.
Last trading day: one business day prior to last notice day
Last notice day: seven business days prior to last business day of delivery month
<u>KC Futures Option</u>
Trading time 9:15am-1:30pm New York time
Price quotation in cents per pound
Delivery months: March, May, July, September, December
Minimum price fluctuation: 1/100 cent/lb.
First trading day: first trading day following the last trading day of any expiring regular option month
Last trading day: second Friday of the calendar month (minimum of 4 trading days between last trading day and the first notice day of the expiring future)
Expiration date/time: 9:00pm New York Time on the last trading day.

SUGAR-11

Price limits are set 0.5 cents above and below the previous settlement price (lifted on 15/12/1997).
<u>SB Futures</u>
Trading time 9:30am-1:20pm New York time
Price quotation in cents per pound
Delivery months: March, May, July, September, December
Minimum price fluctuation: 1/100 cent/lb.
Last trading day: last business day of the month preceding delivery month.
Notice day: first business day after the last trading day.
<u>SB Futures Option</u>
Trading time 9:30am-1:18pm New York time
Price quotation in cents per pound
Delivery months: March, May, July, September, December
Minimum price fluctuation: 1/100 cent/lb.
First trading day: first trading day following the last trading day of any expiring regular option month
Last trading day: second Friday of the calendar month (minimum of 4 trading days between last trading day and the first notice day of the expiring future)
Expiration date/time: 9:00pm New York Time on the last trading day.

COCOA

Price limits are set 88 dollars above and below the previous settlement price (lifted on 15/12/1997).
<u>CC Futures</u>
Trading time 8:30am-1:30pm New York time
Price quotation in dollars per metric ton
Delivery months: March, May, July, September, December
Minimum price fluctuation: $1/mt.
Last trading day: one business day prior to last notice day
Last notice day: ten business days prior to last business day of delivery month
<u>CC Futures Option</u>
Trading time 9:00am-1:25pm New York time
Price quotation in dollars per metric ton
Delivery months: March, May, July, September, December
Minimum price fluctuation: $1/mt.
First trading day: first trading day following the last trading day of any expiring regular option month
Last trading day: first Friday of the month preceding the contract month.
Expiration date/time: 9:00pm New York Time on the last trading day.

Table 1. Descriptive statistics of daily futures returns

	Nearest	Next-to-nearest	Third	Fourth	Fifth
COFFEE					
Mean	0.07%	0.05%	0.05%	0.05%	0.04%
Std.deviation	3.02%	2.83%	2.47%	2.38%	2.36%
Max	23.77%	23.23%	19.04%	18.28%	17.54%
Min	-15.03%	-13.89%	-13.25%	-13.45%	-15.59%
Range	38.80%	37.12%	32.29%	31.73%	33.13%
Skewness	0.58	0.67	0.04	-0.02	-0.13
Kurtosis	7.80	8.94	4.50	4.80	5.18
SUGAR					
Mean	-0.00%	-0.03%	-0.04%	-0.04%	-0.03%
Std.deviation	1.77%	1.53%	1.33%	1.24%	1.16%
Max	8.60%	8.20%	7.54%	8.11%	8.33%
Min	-9.07%	-7.23%	-5.75%	-5.82%	-5.27%
Range	17.67%	15.43%	13.28%	13.93%	13.60%
Skewness	-0.20	-0.18	-0.21	-0.26	-0.20
Kurtosis	2.96	3.42	3.22	4.18	4.61
COCOA					
Mean	-0.07%	-0.07%	-0.07%	-0.07%	-0.06%
Std.deviation	1.63%	1.55%	1.46%	1.41%	1.37%
Max	9.96%	9.87%	9.49%	9.15%	8.73%
Min	-5.32%	-6.81%	-6.53%	-6.47%	-6.49%
Range	15.28%	16.68%	16.03%	15.62%	15.23%
Skewness	0.55	0.50	0.43	0.43	0.40
Kurtosis	3.17	3.38	3.06	3.13	3.02

The table reports the sample mean, standard deviation, maximum, minimum, range, skewness and kurtosis of continuously compounded futures returns from 1993-1999. Column headers indicate futures contracts of different maturities.

Table 2. Limit occurrences of CSCE futures contracts, 1993-1999

	Limit Days					Limit Contracts				
Year	Up	Down	Exp. Up	Exp. Down	Total #	Up	Down	Exp. Up	Exp. Down	Total #
KC – Coffee "C"										
✍ 1993	1	2	0	0	3 (1)	2	3	0	0	5 (2)
✍ 1994	35	32	9	6	82 (38)	90	71	24	18	199 (70)
✍?1995	6	19	0	1	26 (1)	9	35	0	1	45 (1)
✍?1996	4	4	0	0	8 (1)	8	4	0	0	12 (1)
✍?1997	45	21	2	6	74 (21)	98	50	5	15	168 (27)
?1998*	17	22	-	-	39 (3)	36	51	-	-	87 (3)
? 1999	21	17	-	-	38 (1)	94	76	-	-	170 (2)
SB – Sugar – 11										
1993	4	4	0	0	8	6	6	0	0	12
1994	3	3	0	0	6	6	5	0	0	11
1995	0	8	0	0	8	0	16	0	0	16
1996	2	1	0	0	3	2	1	0	0	3
1997	0	0	0	0	0	0	0	0	0	0
1998*	2	4	-	-	6	4	11	-	-	15
1999	3	2	-	-	5	9	3	-	-	12
CC – Cocoa										
1993	0	0	0	0	0	0	0	0	0	0
1994	3	0	1	0	3	6	0	1	0	7
1995	0	0	0	0	0	0	0	0	0	0
1996	0	0	0	0	0	0	0	0	0	0
1997	0	0	0	0	0	0	0	0	0	0
1998*	0	0	-	-	0	0	0	-	-	0
1999	3	0	-	-	3	12	0	-	-	12

Notes:
numbers in parentheses indicate useful samples.
✍ indicates that effective limit samples could be selected during this year.
? indicates that phantom limit samples could be selected during this year.
Shaded rows indicate years when the limits were abandoned.
- indicates that samples are not relevant during the two years when limits were abandoned.
* 1998 numbers include the period December 15, 1997 until December 31, 1997 – during which limits on CSCE futures were abandoned.
Note that we also have 56 useful phantom limit samples for 1993-1997.

Table 3. Coffee "C" futures volume migration?

OPTIONS Transactions

	Effective Limits #		Phantom Limits ##	
Count	Frequency	Per minute	Frequency	Per minute
0	59.12 %	59.12 %	60.07 %	60.07 %
1	9.93 %	36.48 %	11.15 %	15.41 %
2	6.35 %	2.44 %	6.34 %	9.45 %
3	4.89 %	0.98 %	3.45 %	5.39 %
4	2.77 %	0.49 %	2.22 %	3.03 %
5	2.28 %	0.33 %	1.33 %	1.64 %
<5,10]	6.35 %	0.16 %	3.75 %	2.98 %
<10,15]	3.09 %	-	1.97 %	0.83 %
<15,20]	1.63 %	-	1.17 %	0.50 %
<20,25]	0.16 %	-	1.06 %	0.19 %
<25,30]	0.81 %	-	0.67 %	0.31 %
>30	2.61 %	-	6.81 %	0.19 %
Max	159	5	1506	60
Mean	3.60	0.17	12.41	1.18
Median	0.00	0.00	0.00	0.00
Std.dev.	11.75	0.52	59.89	3.73

FUTURES Transactions

Count	Frequency	Per minute	Frequency	Per minute	Count
1	44.95 %	81.27 %	33.34 %	13.10 %	1
2	22.48 %	8.47 %	29.84 %	32.59 %	<1,5]
3	11.24 %	3.09 %	9.23 %	23.78 %	<5,10]
4	7.00 %	0.81 %	4.92 %	12.12 %	<10,15]
5	4.40 %	1.14 %	2.61 %	6.79 %	<15,20]
6	2.77 %	0.98 %	2.09 %	2.09 %	<20,25]
7	2.44 %	0.81 %	1.25 %	5.26 %	<25,30]
8	1.14 %	0.65 %	1.42 %	0.44 %	<30,35]
9	1.14 %	0.33 %	0.81 %	0.70 %	<35,40]
10	0.33 %	0.49 %	1.08 %	0.14 %	<40,45]
<10,15]	1.30 %	0.33 %	0.56 %	0.03 %	<45,50]
>15	0.81 %	1.63 %	12.85 %	2.98 %	>50
Max	19	45	2491	120	
Mean	2.61	1.12	52.87	9.95	
Median	2.00	0.14	3.00	5.69	
Std.dev.	2.58	3.62	200.43	12.50	

The count column measures the number of futures options, respectively futures transactions. Effective limit episodes are sampled from 1993 to December 1997 (total of 614). Phantom limit episodes are sampled from 1993 through 1999 (total of 3596). The "Frequency" columns measure the number of episodes as a percentage of the total number of episodes with a particular count. The "Per minute" columns standardize the frequency for the duration of the episodes.

The mean duration of an effective limit episode is 54 minutes (median duration is 16 minutes; maximum duration is 280 minutes, minimum duration is 2 seconds).

The mean duration of a phantom limit episode is 21 minutes (median duration is 45 seconds; maximum duration is 285 minutes, minimum duration is 2 seconds).

Figure 1a. Coffee "C" (KC nearest-to-delivery) futures price

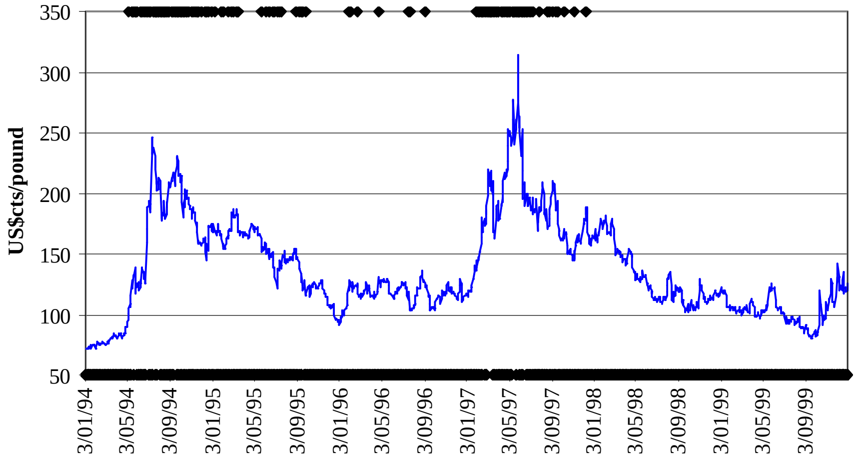

Figure 1b. Sugar-11 (SB nearest-to-delivery) futures price

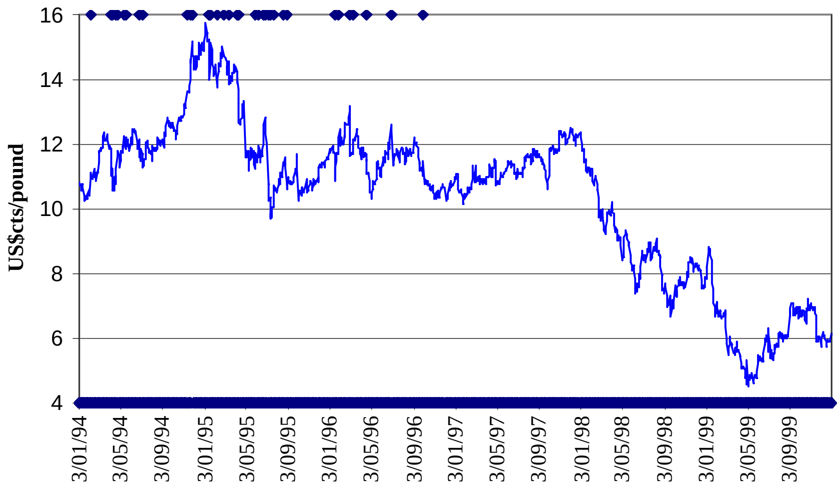

Figure 1c. Cocoa (CC nearest-to-delivery) futures price

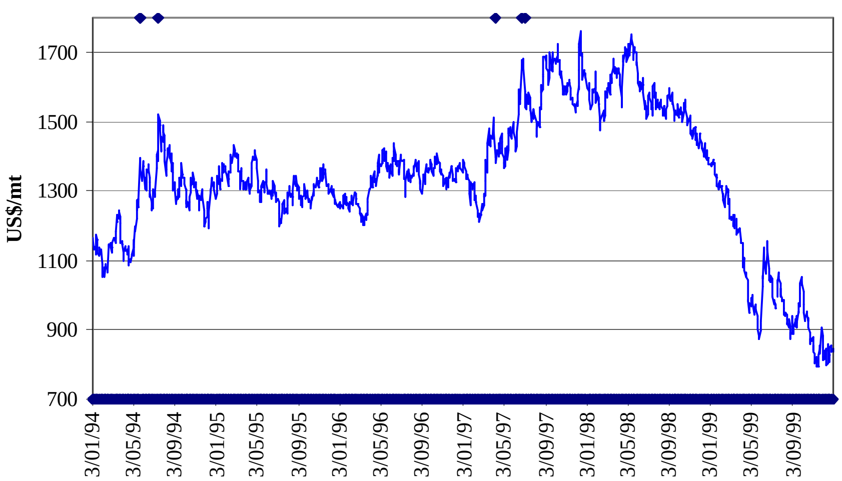

The upper ? marks indicate limit days.

Figure 2a. Duration of trade interruption for Coffee "C" futures – full sample

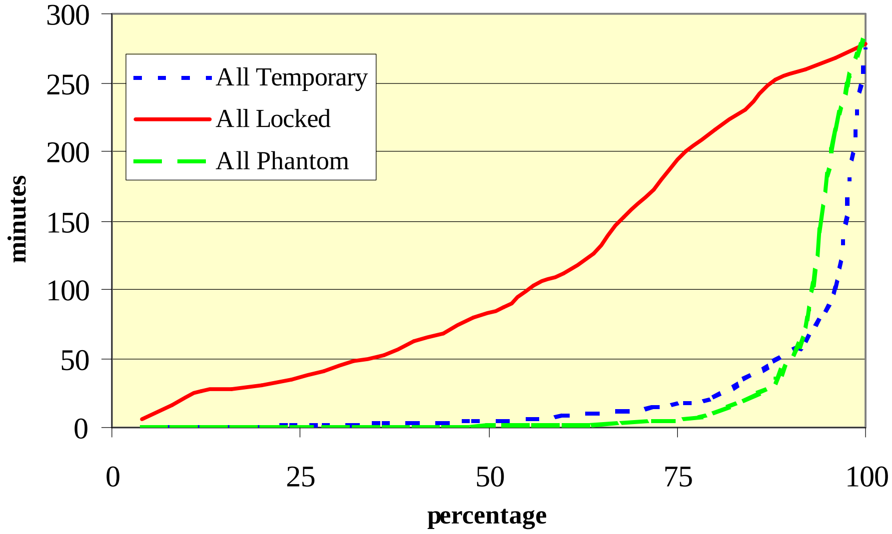

The x-axis gives the percentage of all limit episodes that occurred from 1993-1997 (until 1999 for the phantom limit sample). A limit episode starts when the price first hits the upper or lower limit and ends when trading resumes at prices within the limit range or when trading closes for the day. The y-axis gives the duration of a limit episode in minutes.

Figure 2b. Duration of trade interruption for Coffee "C" futures – selected sample

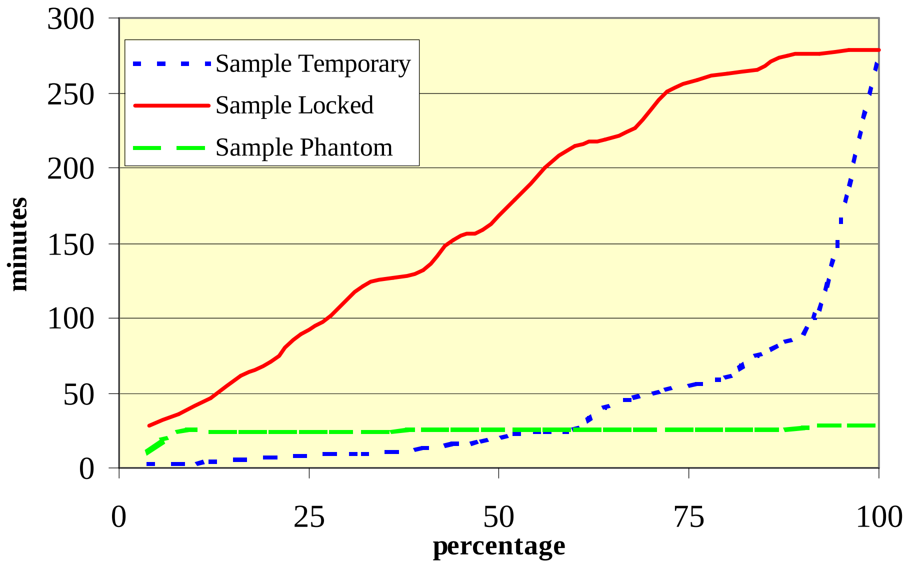

The x-axis gives the percentage of limit episodes that satisfy the selection criteria for our empirical exercise. A limit episode starts when the price first hits the upper or lower limit and ends when trading resumes at prices within the limit range or when trading closes for the day. The y-axis gives the duration of a limit episode in minutes.

Figure 3. Price discovery for Coffee "C" futures during limit locked intervals

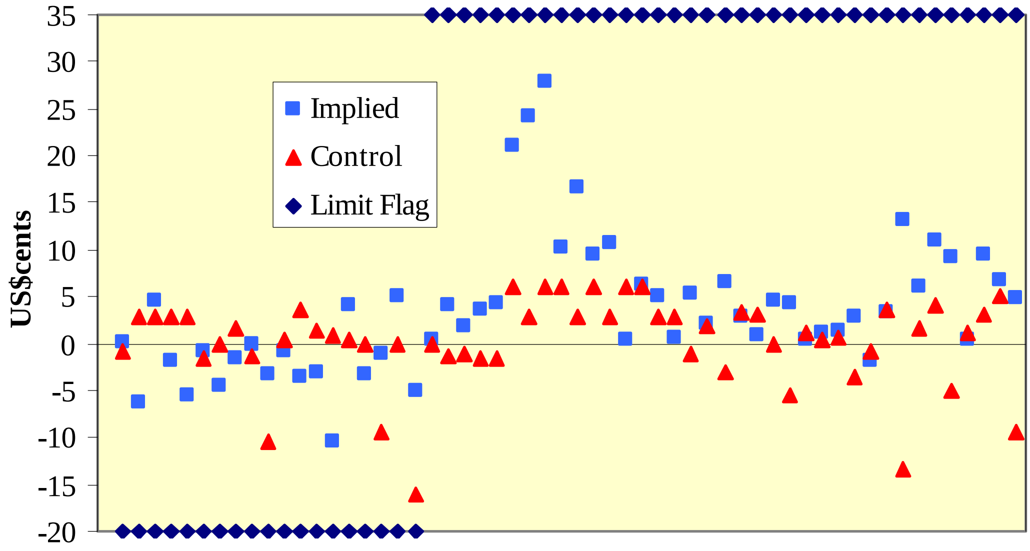

Pricing error is defined as the difference between the options implied futures price and the observed (limit) futures price. Implied (?) indicates limit episodes; Control (?) indicates control episodes when the futures price was not limited on the same day as the limit episode. The sample is partitioned into lower limit episodes (? marks on the lower axis) and upper limit episodes (? marks on the upper axis).

Figure 4. Price discovery for Coffee "C" futures during temporary limit intervals

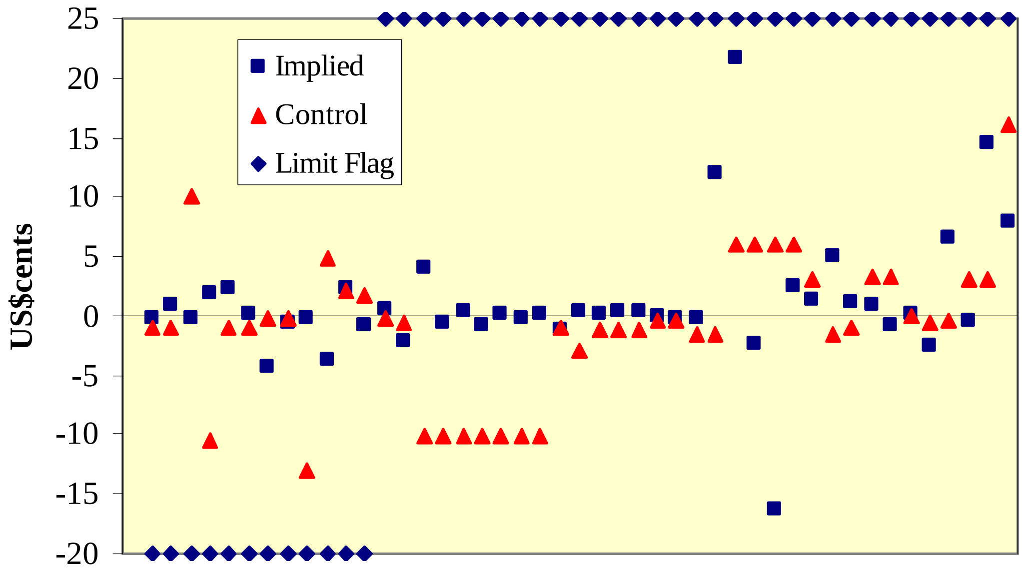

Pricing error is defined as the difference between the options implied futures price and the observed (limit) futures price. Implied (?) indicates limit episodes; Control (?) indicates control episodes when the futures price was not limited on the same day as the limit episode. The sample is partitioned into lower limit episodes (? marks on the lower axis) and upper limit episodes (? marks on the upper axis).

Figure 5. Price discovery for Coffee "C" futures during phantom limit intervals

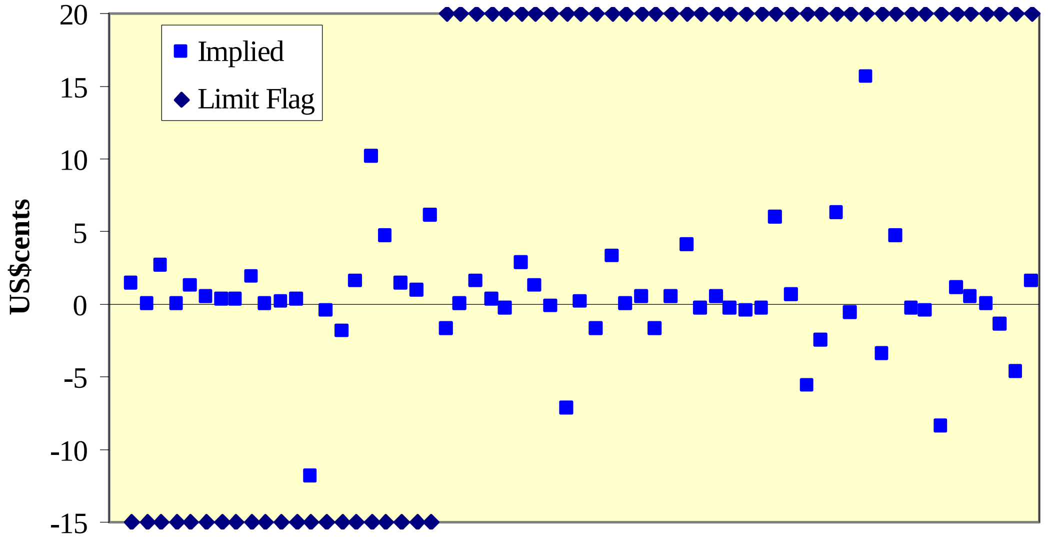

Pricing error is defined as the difference between the options implied futures price and the observed futures price. Implied (?) indicates phantom limit episodes. The sample is partitioned into lower limit episodes (? marks on the lower axis) and upper limit episodes (? marks on the upper axis).

Figure 6a. Temporary limit – Coffee "C" futures price discovery

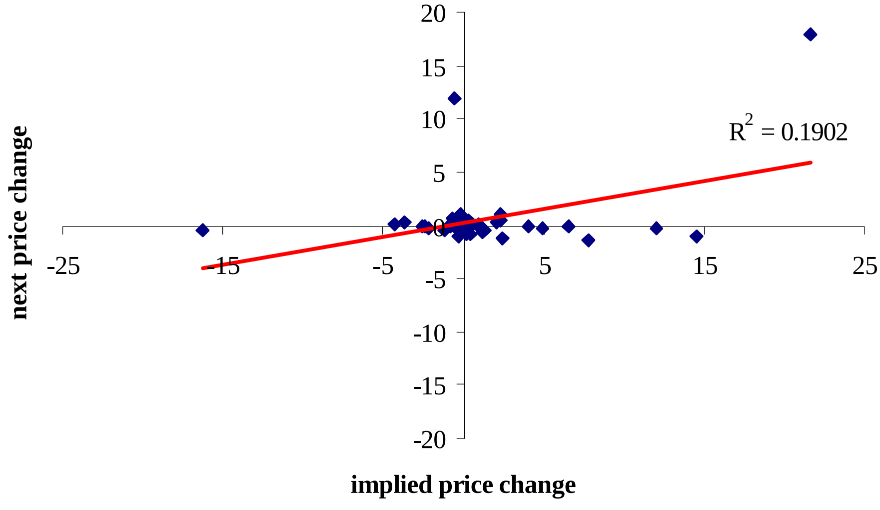

The implied price change is the futures price implied by options traded during the limit lock interval minus the limit price in US$cts. The next price change is the first available futures price which occurs inside the limits on the same trading day minus the limit price in US$cts.

Figure 6b. Locked limit – Coffee "C" futures price discovery

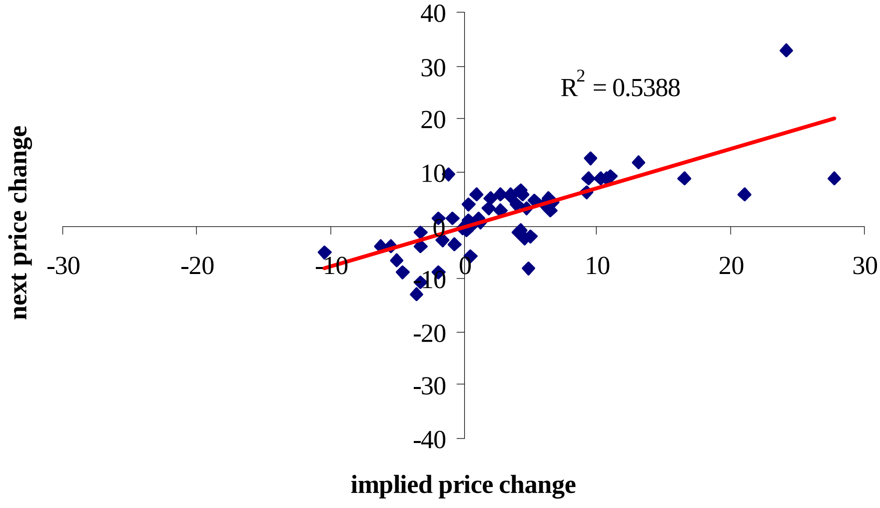

The implied price is the futures price implied by options traded during the limit lock interval minus the limit price in US$cts.. The next price is the opening price of the next trading day minus the limit price in US$cts..

Figure 7. Evolution of Coffee "C" futures price discovery

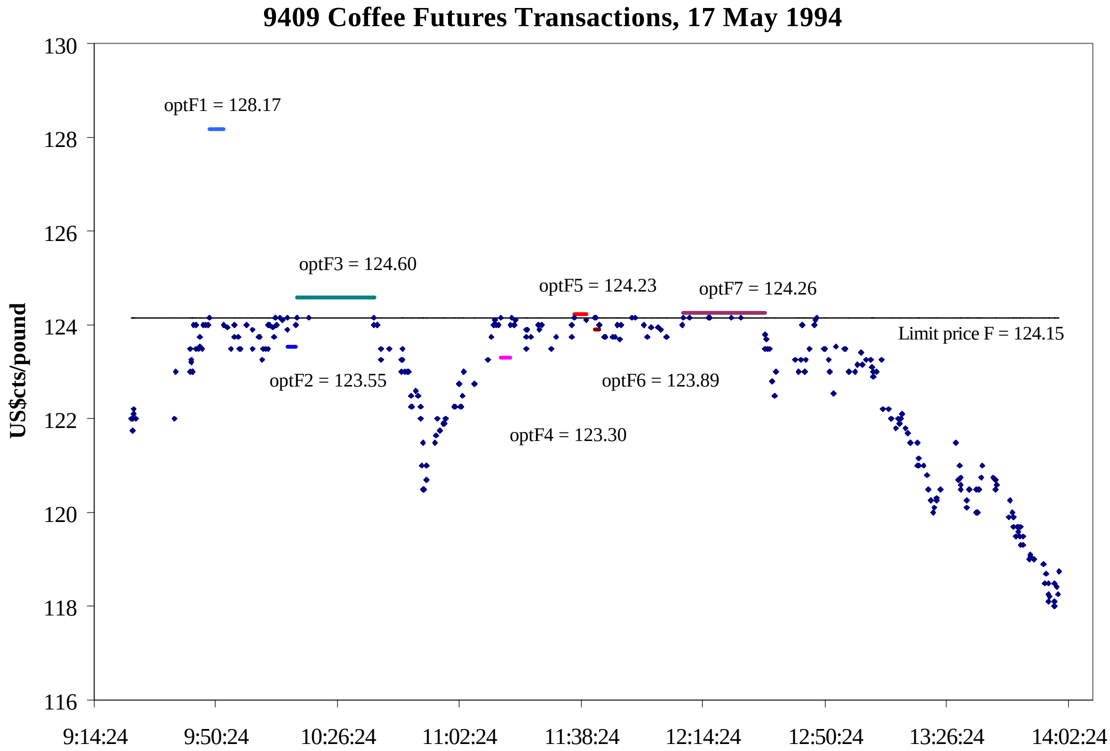

OptF1 to OptF7 indicate the options implied futures prices when the observed futures price was at its upper limit.

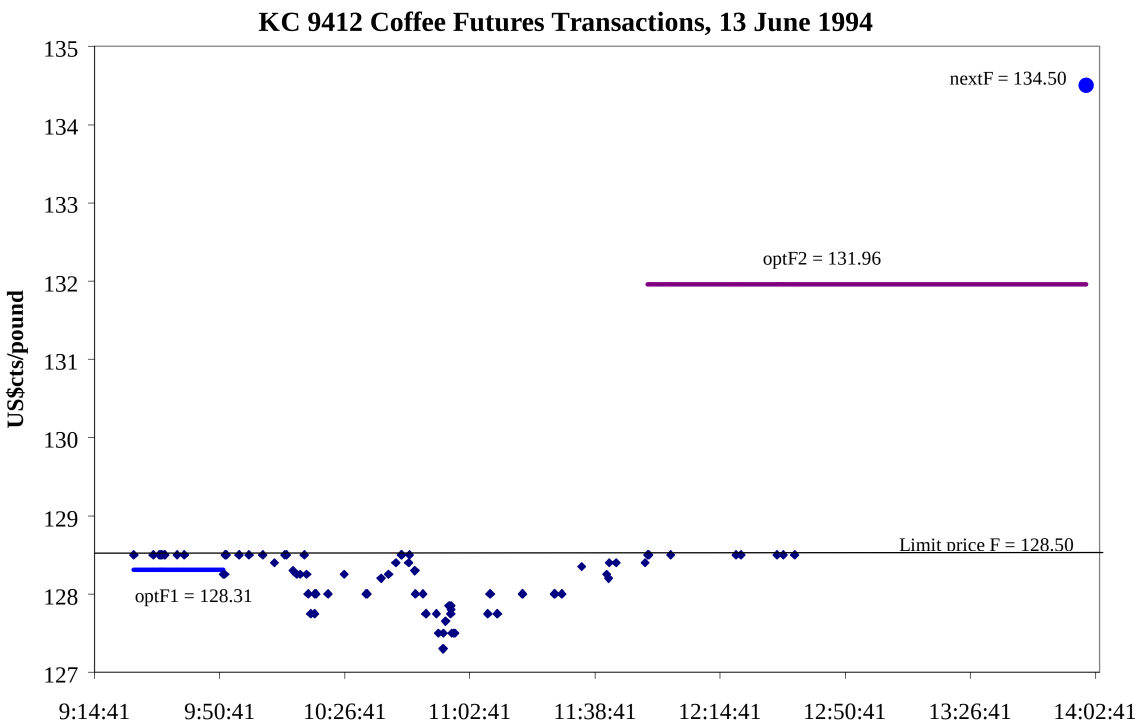

OptF1, OptF2 indicate the options implied futures price when the observed futures price was at its upper limit. NextF indicates the next available futures price that is off limit (in this case, the next day's first traded price).

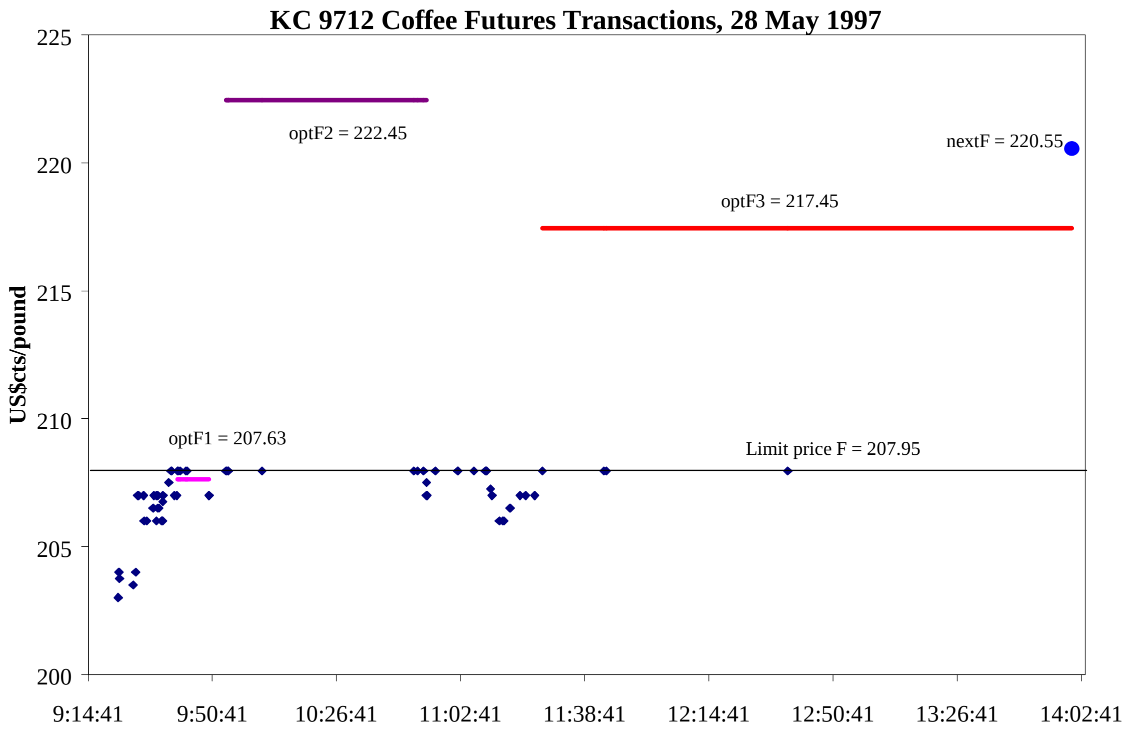

OptF1 to OptF3 indicate the options implied futures prices when the observed futures price was at its upper limit. NextF indicates the next available futures price that is off limit (in this case, the next day's first traded price).

Figure 8. Volatility during temporary limit intervals

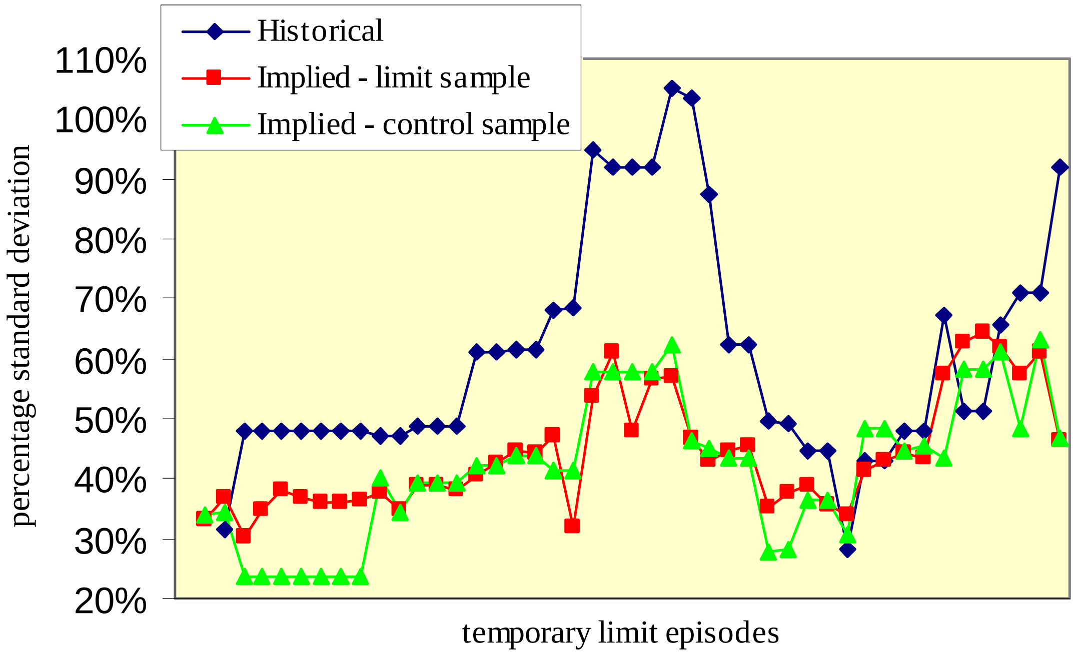

Figure 9. Volatility during locked limit intervals

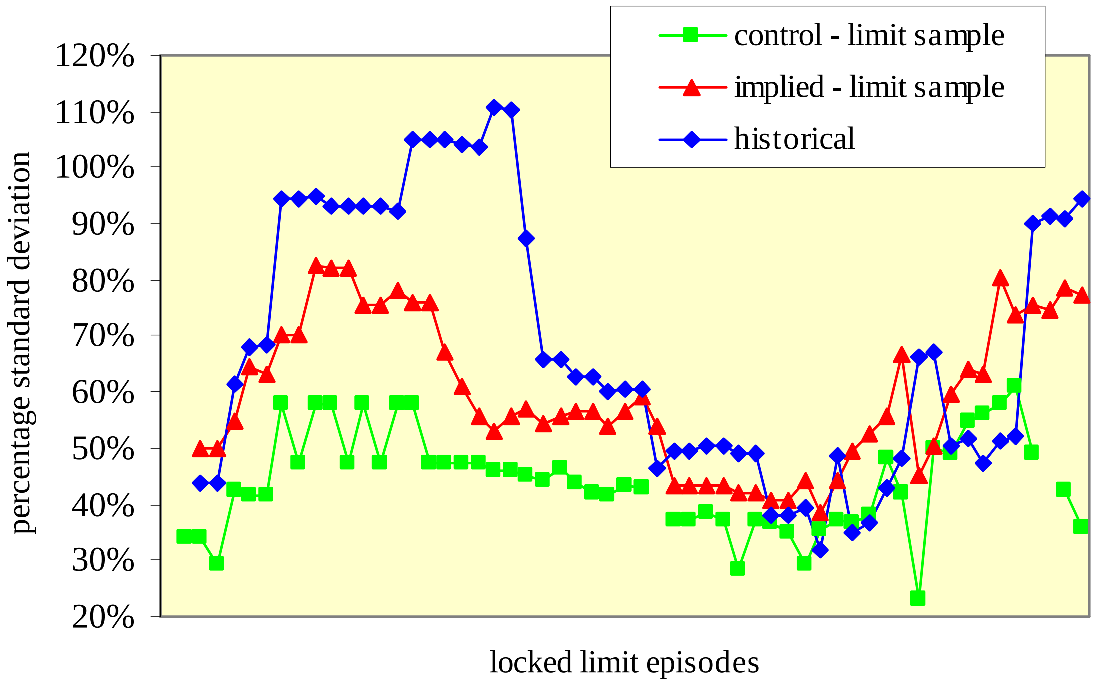

Figure 10. Volatility during phantom limit intervals