Government Office for Science

Foresight

What has happened to UK Equity Market Quality in the last decade? An analysis of the daily data

The Future of Computer Trading in Financial Markets - Foresight Driver Review – DR 1

Contents

What has happened to UK Equity Market Quality in the last decade? An analysis of the daily data[1]

Professor Oliver Linton, London School of Economics

This review has been commissioned as part of the UK Government's Foresight Project, The Future of Computer Trading in Financial Markets. The views expressed do not represent the policy of any Government or organisation.

[1]Thanks to Irene Galbani and Steve Grob at Fidessa for making their data available to us. We are grateful to two referees and Lucas Pedace for helpful comments.

This is intended to be a background piece that investigates the properties of the UK equity market over the last decade with a view to establishing trends or otherwise in market quality. The main findings are that market quality has improved a little over the decade in terms of volume traded and liquidity, but there has been large variation over the period. In particular, very negative outcomes were experienced during the period 2008/2009, although there has since been a rebound. Overall, the findings are consistent with almost stationary behaviour of market quality over this period.

Introduction

The last decade has seen many changes in financial markets. In the early part of the new millennium, the US stock markets went through decimalization, i.e., a reduction of the common tick size to one cent, Angel [1997]. A number of major stock exchanges went through demutualization following the London Stock Exchange and Nasdaq in 2000 and the Chicago Mercantile Exchange in 2002. The market structure has also changed dramatically recently, with many new trading venues being created, e.g., BATS in Kansas city and Chi-X in London, with a variety of different clientele, Gresse [2010]. There has been a big increase in computer-based high frequency trading in equities and other financial markets, Cliff (2011). The legislative framework has also changed with regulation NMS in the US and the MiFid series in the EU, which promoted or facilitated some of these changes. Finally, the global financial crisis (GFC) in the period 2008/2009 has had a big effect on the real economy and also on financial market outcomes.

This project is focused on the issue of computer-based trading or high frequency trading (HFT) (we use the term HFT generically) in financial markets, although this is intimately related with the underlying market structure and the process of fragmentation of trading volume that has recently occurred. Many arguments have been put forward to say that HFT is causing poor market outcomes, especially since the "flash crash" of May 6th, 2010, which showed how far stock prices could move in a very short period of time. A great deal of investigation has followed, culminating in a report by the SEC/CFTC into the events of that day, which pointed to a number of explanations that include the perfect storm of bad news from the Euro debt crisis, a large sell order entered without price constraints, and disappearing HFT market maker algorithms. Since then commentators have pointed to a number of "mini-flash crashes" where markets for individual stocks have been "spooked". Specifically, in the UK on August 24th, 2010 there were very rapid changes in the prices of five LSE-listed stocks---BT Group PLC, Hays PLC, Next PLC, Northumbrian Water Group PLC (NWG) and United Utilities Group PLC (UU). In this latter case, unlike the flash crash, there appear to have been no significant news about fundamentals.

We next discuss the issue of competition between venues and systems for order flow. The first round of MiFid regulation, implemented on November 1st, 2007, eliminated the concentration rule that operated in some European countries that prohibited trading of equities off the regulated market. This has resulted in a great deal of competition between venues and systems for order flow. Some of this competition is taking place in lit markets but there has also been an expansion in trading on public and private dark pools. The fragmentation of order flow is facilitated by computer hardware and software. Fragmentation appears to be both cause and effect of developments in HFT. The development of competition is generally viewed as good and preferable over monopoly provision, but where significant investments are being made in technologies and provision of supervision and other services is expensive, there must be a natural limit to the amount of venues the trading market can sustain over the medium term. And in the short term, we may have concerns over the interoperability of the market as a whole and whether the market outcomes achieved are better or worse as the degree of competition increases. It may also matter whether the competition is taking place in a lit or dark environment. Gresse (2010) finds improvements in bid-ask spreads and volatility following the reform in a large sample of European blue chip stocks. O'Hara and Ye (2009) found similar benefits from competition between trading venue in the United States. London Economics (2010) has claimed that these changes were hugely beneficial for equity market quality in Europe. So much so that this has led to a permanent increase in European GDP of 0.8%!

There has been a recent empirical literature that has tried to determine the validity of some of these claims. Much of the empirical work that has addressed these issues has been based on high frequency data including not just transactions but also the order book. This type of data gives an apparently complete picture of the market and trading activities. Some of the questions that have been addressed include: (1) Are HFT homogenous and less diverse in terms of their trading strategies than non HFT? (2) Is more frequent HFT associated with higher volatility in the market, and is this relationship causal? (3) How and when do HFT supply liquidity? (4) Have bid-ask spreads and available liquidity on average decreased or increased during the time that HFT has been increasing? (5) Is more frequent HFT associated with lower bid-ask spreads in the market, and is this relationship causal? (6) How are the lifetimes of orders related to HFT? (7) What is the profitability of HFT? (8) What type of strategies do HFT use? (9) Do they demand or supply liquidity, and does this vary with market conditions? (10) How does HFT affect price discovery? (11) How do the above effects differ in times of market stress versus more regular times? (12) How has the increase in competition between trading venues introduced in MiFid affected trading volume? (13) How has the increase in competition between trading venues introduced in MiFid affected liquidity? (14) How has the increase in competition between trading venues introduced in MiFid affected volatility?

We describe elsewhere, Linton (2010), the studies that have been undertaken, their methodology, and the outcomes they report. It is fair to say that this literature is growing rapidly and is informing our understanding of how markets actually work in the very high speed setting (in the absence of a comprehensive economic theory that describes the important interactions between participants). This methodology looks inside the sausage and examines all its constituents. One issue with this analysis is that it is very computationally demanding. Such datasets are extremely large and complex.

Consequently most studies have focused on short periods of time or a restricted set of assets. Furthermore, there are many errors in the data, both in terms of the outcomes of interest but perhaps more importantly in terms of the timing of the occurrence of events. A few milliseconds of inaccuracy may not sound like a big deal, but for this type of analysis it can lead to serious problems for the statistical analysis on which this is based. And inaccuracies of a millisecond or more are inevitable whatever the measurement system, Angel (2011). Obviously, for many of the questions outlined above this analysis is the only game in town. However, a number of questions may be addressed by other methods and indirect evidence can be constructed.

We supply an analysis based on daily time series data. We evaluate what has happened to a variety of equity market outcomes at that frequency in the UK over the last ten years. If HFT or MiFid have had bad effects, this should be evident in the lower frequency record even if it is not possible to distinguish between the explanations. We ask whether the sausage still tastes good, rather than why it tastes good. Specifically, we will address the following questions:

1) Have UK equity prices become more volatile in recent times?
2) Are these series more likely to have crash like episodes in recent times?
3) What has happened to the traded volume for UK equities?

4) Have the markets for these equities become more or less liquid over this period?
5) Have the markets become more or less efficient over this period?
6) Has there been a relative change in the behaviour of the overnight to intraday volatility?
7) Has market fragmentation affected the quality of market outcomes?
8) Does the quality of market outcomes depend on the amount of trading in different style of trading venue: lit (public exchanges with visible order book), dark (public exchanges with invisible order book), otc (over the counter), and si (systematic internalizer)?

We will try to answer these questions for the FTSEALL share index and a handful of individual stocks. The data we will use to answer questions 1-6 will consist of daily open and close price, the intraday high price, the intraday low price, and the total trading volume (number of shares), which are publicly available. To address the last two questions we will supplement our analysis with further work that makes use of a data series provided to us by Fidessa (2011) that breaks out the location of trading volume into different venues. We do not use a lot of formal technique, but rather present the data in various ways to make our points. We have also not provided an exhaustive reference list.

The general conclusion seems to be that 2008/2009 was a particularly bad time for these series in terms of volatility, large negative price moves, and low volume, but that since then things have settled down to yield lower volatility, more volume, and less frequent large price changes. The performance metrics we examine do not look significantly different than longer term averages over the time we consider.

Claessens, Kose, and Terrones (2011) investigated financial cycles over the period 1960:1 to 2007:4 around the world. They found evidence that such cycles can be long and deep and related to housing cycles and credit cycles and to be correlated across economies. We do not address the causal factors behind the downturn but rather confirm that it is more consistent with the cyclical interpretation of these authors rather than being part of a secular trend.

In section 2 we introduce some definitions regarding the quantities of interest. In section 3 we present the quantities of interest in time series graphs and tables. In section 4 we investigate the properties of the key market quality variables in more detail. In section 5 we present results about fragmentation and dark trading. Section 6 concludes and speculates about the future.

Data and Measurement

We take the daily opening price, the closing price, the high price and the low price and the volume transacted for each series from Bloomberg and Yahoo.

Returns

We compute the standardized daily return, or capital gain, of the series. This is defined as

$$return = \log(close) - \log(close_{-1}) \approx \frac{close - close_{-1}}{close_{-1}},$$

where $close_{-1}$ denotes the closing price from the previous day. The standardized return series is transformed to have (within sample) mean zero and variance one, so that return is measured in units of standard deviation, which can be benchmarked against a standard normal distribution. Although the normal distribution is known to be a poor approximation for stock returns since Mandelbrot (1963), it is a useful benchmark to provide interpretation. For a normal distribution, three standard deviations contains 99.73% of the data, six standard deviations contains practically all the data. Since the work of Vilfredo Pareto on income distributions at the end of the

19th century economists have modelled heavy tailed distributions by the distributions that take his name. We measure the extremity of the distribution by the so-called *tail thickness parameter*, which can be motivated from this distribution, and which determines how often we see large events. We will measure the tail thickness of the whole series treated as a common entity and each year separately; we also distinguish between the upper and lower tail thickness parameters, so that we can determine if large negative events are more likely than large positive ones. We use the methodology developed in Gabaix and Ibragimov (2011) to estimate the tail thickness parameter.

Volatility

We measure the intraday volatility by the scaled intraday price range

$$vol = \frac{high - low}{low},$$

where high denotes the intraday high price, low denotes the intraday low price. This is a widely used measure of ex-post volatility that is proportional to a standard deviation type measure under some special circumstances, Alizadeh, Brandt and Diebold (2002). Note that most authors use a different denominator, like (high+low)/2. With our choice of denominator, vol can also measure the profit rate of an omniscient daily buy and hold investor who buys (sells) at the low (high) and sells (buys) at the high (low) within a day. The realized range has the significant advantage that one can find the daily value in the newspapers for a variety of financial instruments, and so one has a readily available volatility measure without recourse to analysis of the intraday price path. Alizadeh et al. (2002) also argue that the method is relatively robust to measurement error, specifically of the bid-ask bounce variety, since the intraday maximum is likely to be at the ask price and the daily minimum at the bid price of a single quote and so one expects a bias corresponding only to an average spread, which is not so great. By contrast, in computing the realized variance (reference) one can be cumulating these biases over many small periods, thereby greatly expanding the total effect. The high low measure does detect extreme intraday movements but it cannot differentiate between a day where the stock visited the high and then the low just once versus a day where the high and low were visited many times in succession. The later day would be considered more volatile, and this would be picked up by high frequency data. The absolute value of the close to close return series can also be considered a (noisy) volatility measure.

One could argue that since high frequency traders trade actively during the day and have a propensity to end the day with zero open position, their activities would affect only the intraday series in terms of its volatility etc. To investigate this we will consider a measure of the overnight volatility to the intraday volatility. We do this in two ways. The close to close return can be decomposed into the overnight return (close to open) $retco,$ and the intraday return (open to close), $retoc$. We will compare $|retco|$ with $|retoc|,$ in fact we will compute weekly average (the median) of both quantities before comparison. The second method is to compare the weekly average of $|retco|$ with the weekly average of our measure of intraday volatility vol defined above. The first method has the advantage that the two quantities are directly comparable, while the second approach retains consistency with our method of measuring volatility elsewhere in the paper. We find some uneven coverage of the index opening price and so we present the results just on some individual equities.

Volumes

The measure we have is the value of the shares transacted in the case of individual stocks and a weighted version of this for the index. There are some problems with large and small volumes in the series. For example, the volume on Christmas Eve is typically very small, in large part because this is a half (or less) trading day.

Liquidity

There are a number of ways of measuring liquidity. Sarr and Lybek (2002) review the literature and divide measures into four categories: Transaction cost measures; Volume-based Measures; Price-based measures; market impact measures. Their focus was more on fixed income markets than equities and their study preceded the recent developments in high frequency trading. Goyenko et. al. (2009) review the state of the art in this literature and compare measures computed from high frequency order book information with those computed indirectly or implicitly from transaction data at a lower frequency. We consider some standard measures computable from the daily data we have. Roll (1984) proposes the following measure based on the serial covariance of closing prices of an individual stock

$$S = 2\sqrt{-\mathrm{cov}(close - close_{-1}, close_{-1} - close_{-2})}$$

where cov means covariance. The justification of this comes from a very simple model where the fundamental value Val evolves as a random walk and the observed trade price is $P = Val + S \times Q/2,$ where Q is a buy/sell indicator for the last trade that equals +1 for a buy and -1 for a sell, which are assumed to be equally likely. The quantity can therefore be interpreted as a spread between the ask price and the bid price. In some cases, the covariance inside the square root can be positive and so the quantity is not defined. A number of refinements to this measure that address this issue have been suggested, see for example Hasbrouck (2006). Amihud (2002) develops a liquidity measure that captures the daily price response associated with one unit of trading volume. This is

$$Liq = Average\frac{|return|}{V},$$

where V is the volume, and average means that one averages this quantity over a longer period like a week or a month. Large values of this measure indicate an illiquid market where small amounts of volume can generate big price moves. It is considered a good proxy for the theoretically founded Kyle's price impact coefficient. In the current environment, a plausible alternative to close to close return is to use the intraday high minus low return, since there can be a great deal of intraday movement in the price that ends in no change at the end of the day. We compute the raw measure so that the average in (liq) is over one day only.

Price Discovery and Efficient Markets

The efficient market hypothesis has been a central plank of finance theory and subject to much empirical investigation. In the simplest version, stock prices are martingales. The most common interpretation of this is to say that stock returns are not predictable. Questions arise as to what information can be used and what method can be used on the data to form a prediction. Hendershott (2011) gives a discussion of this notion and how it can be interpreted in a high frequency world. We shall take a rather simple approach and base our measure of predictability on just the price series itself and confine our attention to linear methods. In this world, efficiency or lack thereof, can be measured by the degree of autocorrelation in the stock return series, that is, by the function of lag j

$$ro(j) = \frac{\mathrm{cov}(return, return_{-j})}{\mathrm{variance}(return)},$$

where the variance and covariance have to be computed over a certain time period and the subscript $-j$ denotes return from j -periods ago. In particular, we shall focus on the first order autocorrelation $ro(1)$ which we denote by ro. Under the efficient markets hypothesis this quantity should be zero. An alternative measure of inefficiency is based on the variance ratio test statistic of Lo and MacKinlay (1988). This is

$$vr_h = \frac{var_h}{h \times var_1},$$

where var_h denotes the variance of h-period returns. When returns are computed in the logarithmic way, the return over h periods is the sum of the returns over all the one period sub intervals. If returns were uncorrelated (independent), then the variance of the sum would be the sum of the variances. Under this hypothesis, the variance ratio should be one for all horizons. However, when there is positive dependence, the variance ratio will be larger than one, while if there is negative dependence, the variance ratio is less than one. From a statistical point of view this measure is not well motivated, (as a statistical test it is optimal against a rather strange alternative hypothesis, Faust (1992) and Deo and Richardson (2003)), but it is widely used in finance studies. We consider the weekly case with $h = 5$, so that the variance ratio is the ratio of the variance of weekly returns to the variance of daily returns. We compute this measure each year, so that we see how the market conditions have changed over the decade.

Evolution of UK Equity Market Quality

FTSE ALL Share

The FTSE All-Share Index, originally known as the FTSE Actuaries All Share Index, is a capitalisation-weighted index, comprising around 600 of more than 2,000 companies traded on the London Stock Exchange (LSE). It aims to represent at least 98% of the full capital value of all UK companies that qualify as eligible for inclusion. The index base date is 10 April 1962 with a base level of 100. To qualify, companies must have a full listing on the London Stock Exchange with a Sterling or Euro dominated price on the LSE trading systems SETS or SETSmm or a firm quotation on SEAQ or SEATS, and must meet a number of other eligibility requirements. FTSE All-Share is the aggregation of the FTSE 100 Index, FTSE 250 Index and FTSE SmallCap Index.

In Figure 1 we present the data: the closing price, the volatility measure computed from intraday high and low prices, the traded volume, and the standardized daily return series. These data are from 1/4/1997 to 10/5/2011. During this time there have been several different trends up and down in the price index. The index stands at 3109.26 on March 1st, 2011 and was at 2975.87 at January 4th, 2000. For comparison, the UK CPI was 92.1 in January, 2000 and 118.1 in March, 2011. This was a period of modest inflation both in consumer prices and asset prices. The volatility series shows typical variation with a peak in 2009 and a decline since then. The volume series is quite erratic and reached a peak sometime in 2008 and has since suffered a decline to levels consistent with the early 2000s. Note however, that the day to day variation in volume is very large and any apparent trend or not must be put into the context of this extreme variation. The marginal distribution of traded volume is well known to have very heavy tails, Gabaix et al. (2005). The standardized return series paints a similar picture to the volatility series. It displays the typical volatility clustering, meaning there are quiet periods e.g., 2004-2007, and then very noisy periods like the financial crisis period and the turn of the millennium. The time series plots confirms the common view that daily returns are not independent and identically normally distributed.

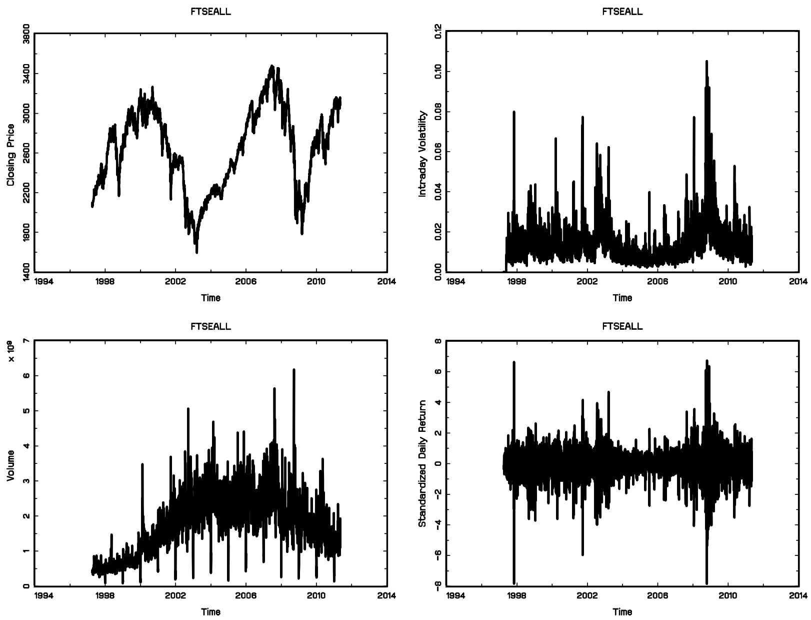

Volatility

The close to close return volatility, which can be discerned from the last panel of the above figure, is similar to the intraday volatility measure: both show an increase during the GFC and a subsequent decline.

Crash Frequency

We next address the question of whether crashes or dashes have become more frequent and/or larger in size in recent times. In Figure 2 we show all close to close returns in excess of 3 standard deviations. Clearly, there were a large number of big events during the crisis period of 2008/2009, but since mid 2009, there have been relatively few such large events. The earlier period around the change of the millennium also saw a number of large price moves. The occurrence of the extremes are clustered, which is not consistent with the hypothesis that returns are independent over time. The number of events greater than three standard deviations is also incompatible with a normal distribution, but the tails are not so extreme compared to other financial time series. Actually, there is a lot of research to suggest that the heaviest tails (largest price movements in relative terms) occur not in equity returns but in electricity prices and some currency series. First, electricity prices are well known to have very long tails, i.e., occasional very large changes in prices. These are not generally attributed to high frequency trading, but rather to occasional spikes in demand that cannot be met. Emerging market currencies, like the Russian ruble, the Thai Baht, and so on are also known to have very heavy tails, Gabaix and Ibragimov (2001), meaning some extreme rate changes take place. These "crashes" are not usually attributed to high frequency trading, but rather to changing views of fundamentals, i.e., government policy. We might investigate whether the large moves in the index are due to fundamental information, and many of them may be, but we caution that even in times gone by absent high frequency trading, authors have found it sometimes difficult to pin down what information has moved asset prices. Cutler, Poterba, and Summers (1989) and Fair (2002) investigated large price moves on the S&P500 (daily and intradaily in the latter case) and tried to match the movements up with news stories reported the following day: many large movements were associated with monetary policy, but there remained many significant movements that they could not find explanations for.

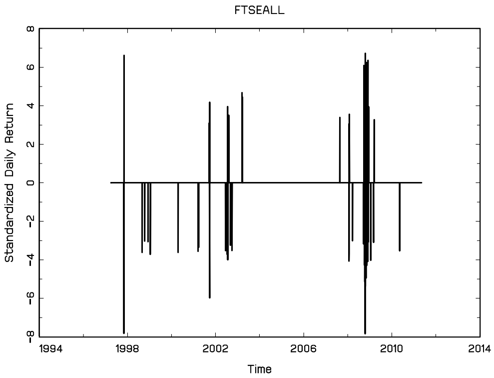

We next provide an alternative way of describing the frequency of crashes by reporting estimates of the tail thickness parameters. The estimated tail thickness parameters (upper and lower tail) by year are given in the table below: larger numbers imply thinner tails, i.e., less frequent crashes. The normal distribution has upper and lower tails equal to infinity. There is some variation in these numbers from year to year (the full sample estimates are 4.0257377 and 4.6352597, respectively) and 2009 in particular had heavy lower tails (large probability of crashes). There does not seem to be a systematic trend in this parameter according to this analysis.

Year	**Lower**	**Upper**
1997	3.7674817	2.8362299
1998	3.7057575	3.6721875
1999	6.0356831	4.3334682
2000	7.3532365	3.9120501
2001	4.3839674	2.5569542
2002	3.1661038	2.7544897
2003	2.9684990	3.0111824
2004	3.3302933	2.8868847
2005	6.9605714	2.9729499
2006	2.5277004	1.8534189
2007	3.3808208	2.2316994
2008	3.0940768	3.2003109
2009	1.9511090	2.3525702
2010	4.8150107	2.8751737
2011	3.4741242	3.9885074

Table 1. Tail thickness parameter by year

We next looked at the twenty largest events in terms of intraday variation (we just show the top five for space limitations). Five of the top twenty events were in the 2008/2009 period, but there were also very volatile days in the turn of the millennium period. These extreme days occurred before the fragmentation induced by MiFiD and before the advent of high frequency trading. Furthermore, note that the mini flash crash was not such an extreme day in terms of either high minus low price of close to close return. This is partly due to the fact that the five shares affected by this crash constitute only a small part of the index, but more importantly, the actual impact of the sell-off was limited, because the LSE's automatic circuit breakers kicked in when the losses in these stocks neared 10% and trading in them was suspended for five minutes after which the markets resumed in a relatively orderly fashion.

Date	**High**	**Low**	**Volume**	*vol*
14/05/2009	2226.00	2187.32	1844241408	0.103849179
30/11/2001	2518.30	2490.10	1838659840	0.090088841
27/03/2002	2535.75	2520.00	2273746176	0.086472563
05/12/2002	1973.94	1936.44	2441989376	0.08518024
10/01/2001	2946.57	2913.87	2155657216	0.080615892

Table 2. Top five days in terms of intraday volatility

Volume

The volume in Figure 1(c) shows an increase up until 2008 and then a decrease. We next order the days by traded volume. None of the largest twenty days occur after 8/10/2008.

Date	**High**	**Low**	**Volume**	*vo*
19/09/2008	2726.37	2496.53	6178696704	0.029020988
09/08/2007	3307.94	3222.84	5639005184	0.063321223
20/09/2002	1923.88	1830.35	5061952000	0.046900969
10/08/2007	3244.62	3128.72	4942967808	0.009711315
17/08/2007	3160.75	3014.01	4764757504	0.034700573

Table 3. Top five days in terms of volume

We next consider the recent LSE outage day, 25th February, 2011, and the days around the mini flash crash. These appear to be not particularly extreme days in terms of their trading volume and price variation when judged against the last ten years (none is in the top twenty).

Date	High	Low	Volume
25/02/2011	3112.69	3066.32	1133202048
24/08/2010	2700.96	2639.99	1241928192
25/08/2010	2668.39	2619.93	1282310784
26/08/2010	2666.25	2638.11	1021461568

Table 4. Price and volume on recent LSE outage day and around the mini flash crash

Liquidity

We show the Amihud illiquidity measures computed either with close to close return or high minus low return. We report the raw daily figures here rather than average them, so there are some extreme outcomes. Both graphs show some variation in liquidity: the measure improved considerable up until 2004 where it was flat until 2008 whence it increased considerably (although still much less illiquid than in 2000). It has since come down from this short term high. Compare this with Hendershott, Jones, and Menkveld (2011) who found that algorithmic trading causally improved liquidity over the period 2001-2005.

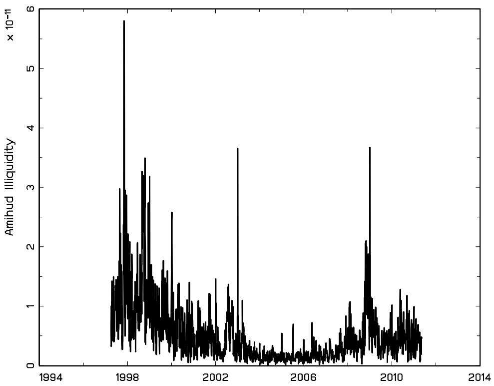

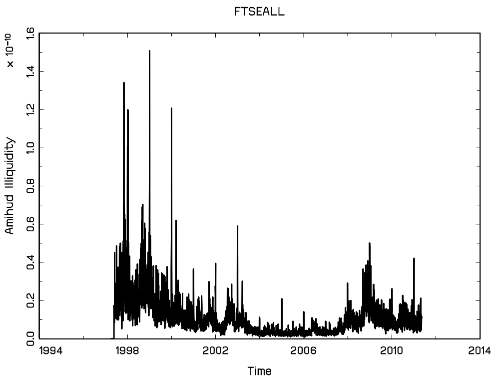

Price Discovery and Market Efficiency

The next table shows the first order autocorrelation and the five day variance ratio test by year along with t-statistics (for the autocorrelation the implicit null hypothesis is that $ro = 0$ and for the variance ratio the implicit null hypothesis is $vr = 1$). The largest amount of predictability (highest $|ro|$) was in 2008. Recall that the square root of negative auto-covariance can be used as a liquidity measure (assuming it is negative). This shows that 2008-2010 were relatively illiquid times, but that liquidity has come back down in the last year. There are some periods where the autocorrelation is statistically significant either positively or negatively, but in most periods it is not. The variance ratio test is never significant at the 5% level.

Year	ro	**t-stat**	vr	**t-stat**
1997	0.17534655	2.69942580	1.19928930	1.08470890
1998	0.14888306	2.29202550	1.04924830	0.26805281
1999	0.053816752	0.82849837	1.05755320	0.31325551
2000	0.025546638	0.39328549	1.25102780	1.36631560
2001	0.018182745	0.27991980	1.07684890	0.41827981
2002	-0.052105925	-0.80216053	1.16992870	0.92490252
2003	-0.062228323	-0.95799285	0.82380599	-0.95900378
2004	-0.16512733	-2.5421030	0.74293817	-1.3991580
2005	0.018843075	0.29008544	1.08837550	0.48101766
2006	-0.17381109	-2.6757878	0.88138544	-0.64560544
2007	0.064910134	0.99927881	0.88905414	-0.60386558
2008	-0.19485934	-2.9998214	0.80658788	-1.05272000
2009	-0.020627437	-0.31755535	0.85233903	-0.80370170
2010	-0.080197063	-1.23461810	1.17364580	0.94513416
2011	0.068291048	1.05132730	1.16472130	0.89655905

Table 5. First order autocorrelations and variance ratio test statistics by year

The Mini Flash Crashettes

We also examine the behaviour of the five firms that were most affected by the events of August 24th, 2010. If the argument is that HFT caused these bad outcomes, then most likely HFT must be trading these equities actively. We may therefore find that their performance has suffered more than the performance of the equity universe as a whole. We show a selection of the statistics we showed for the index.

First, we show the price and volume data for the three days around the mini flash crash. The day in question seemed to have a lot of intraday volatility and trading volume compared with the previous day for all equities. However, we computed the top twenty days in terms of intraday volatility and the top twenty days in terms of volume for each equity as for the index. None of these equity days shown in Table 6 were in the top twenty according to either of these criteria, so overall the day was not particularly special.

Equity	Date	Open	High	Low	Close	Volume
	25/08/2010	90.1	91.3	89.05	89.15	8083800
Hays	24/08/2010	90.8	93.7	82.50	90.50	9033000
	23/08/2010	91.8	92.1	90.85	91.25	1848900
	25/08/2010	1938	1949	1914	1932	510200
Next	24/08/2010	1964	1978	1817	1944	1564500
	23/08/2010	1967	1983	1949	1970	471100
	25/08/2010	319.5	328.9	319.5	327.9	795100
NWG	24/08/2010	317.0	336.1	313.0	320.6	1699800
	23/08/2010	313.3	320.0	313.3	318.8	828500
	25/08/2010	134.5	135.5	132.0	132.6	20436800
BT	24/08/2010	135.0	136.0	123.0	135.2	37954100
	23/08/2010	135.0	136.7	133.5	135.5	17247800
	25/08/2010	568.0	571.5	564.0	565.5	2284900
UU	24/08/2010	564.5	620.0	550.5	564.5	2944200
	23/08/2010	564.5	568.5	560.0	567.0	1759700

Table 6. Price and volume data around the mini flash crash day

Hays

Hays PLC is a FTSE250 company first listed on LSE on 26/10/89 with market capitalization on April 12th, 2011 of £1564.21 million. The data is from 1/1/2003 to 12/4/2011. During this period, the equity experienced a modest price appreciation, somewhat better than the appreciation of the all share index. The time series of closing price, intraday volatility, trading volume, and standardized return are shown next. The intraday volatility takes a higher average than for the index, and has reached 20%. Volatility increased during the GFC and has since declined. On standardization of the series though we see that the extreme events for this price series have not been as extreme as for the index Figure 6(d) and Figure 7.

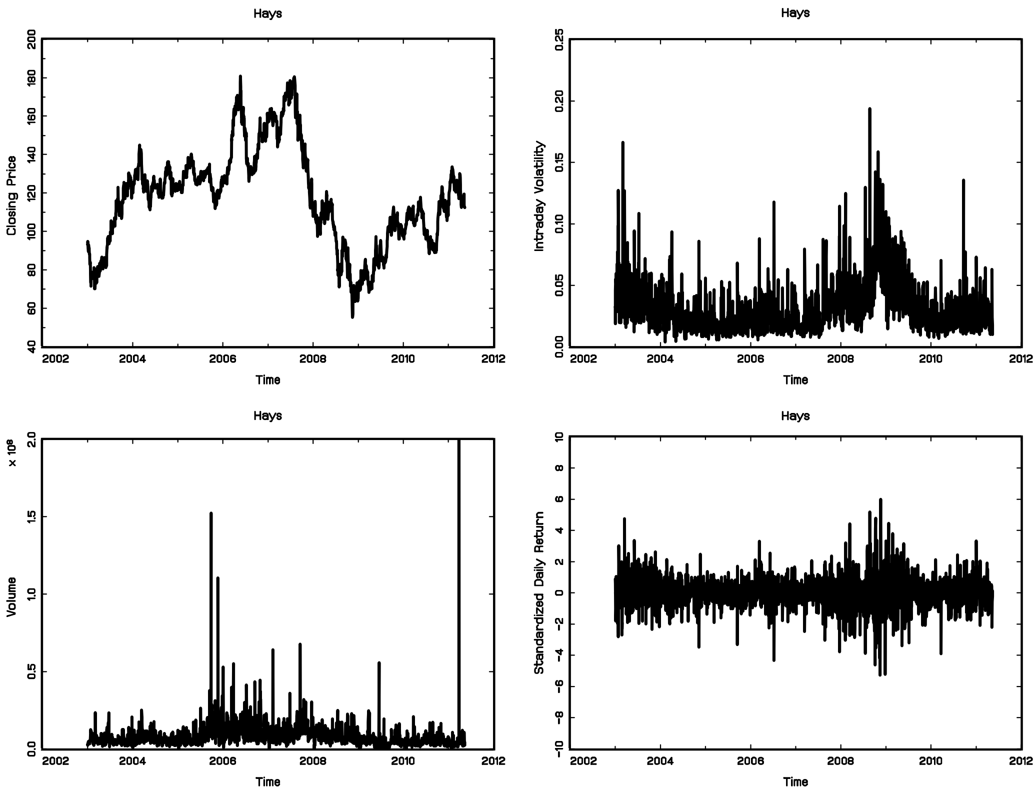

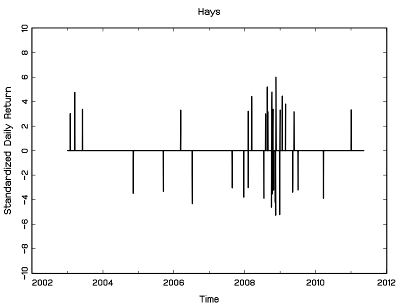

We next show the raw liquidity measure, which appears to go up at the end of the period, certainly relative to the 2004-2007 period. The degree of illiquidity is obviously much greater than for the index. The autocorrelations and variance ratio tests are not statistically significant in any year and are not shown here.

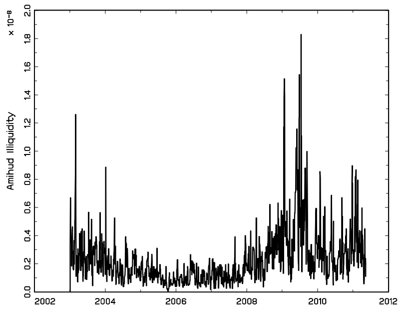

Date	Open	High	Low	Close	Volume	*vo*
05/08/2008	81.25	97.00	81.25	90.00	17332200	0.193846154
04/03/2003	77	80.50	69	73.75	23427000	0.166666667
10/10/2008	63	73.00	63	71.5	14105700	0.158730159
18/09/2008	79.75	86.25	75.5	83.25	21545100	0.142384106
24/10/2008	59	62.5	55	55.25	17213600	0.136363636

Table 7. Five largest days in terms of intraday volatility

Next

Next PLC is a FTSE100 company first listed at the LSE on 12/3/48 with market capitalization of £3781.04 million. The data is from 1/1/2003 to 12/4/2011. Over this period there was a considerable increase of nearly six-fold in the share price, although the period 2007-2009 saw consistent declines.

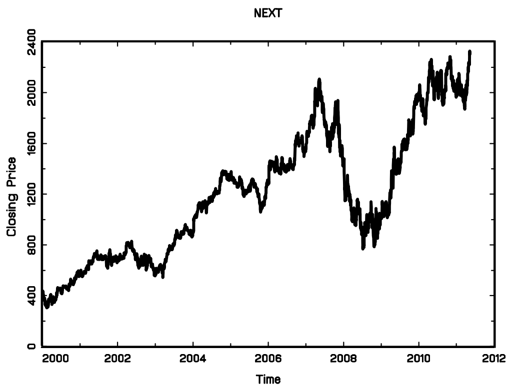

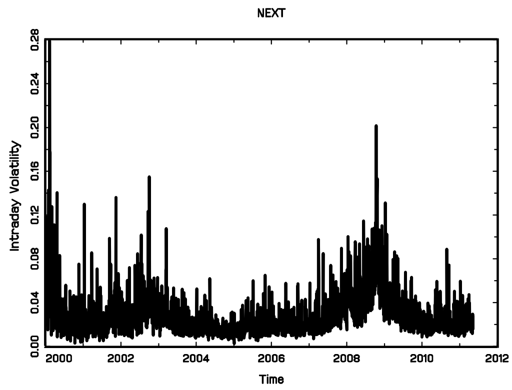

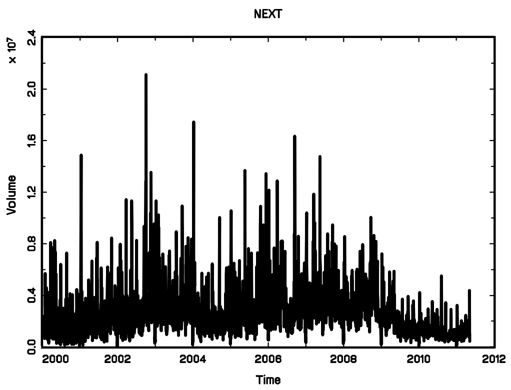

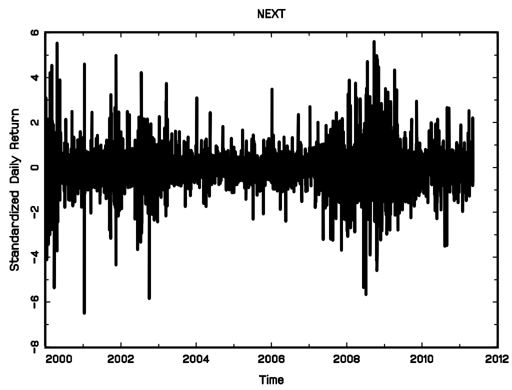

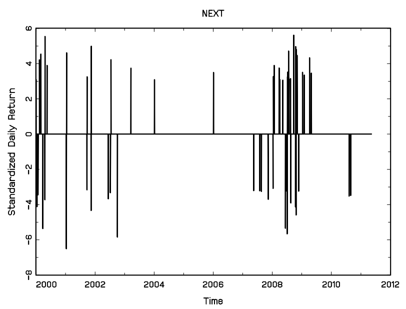

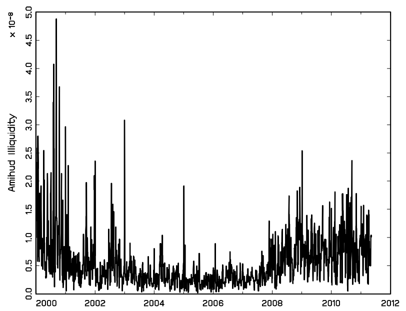

The autocorrelations and variance ratio tests are not statistically significant in any year. The five largest days in terms of intraday volatility were all in 2008 and beginning of 2009.

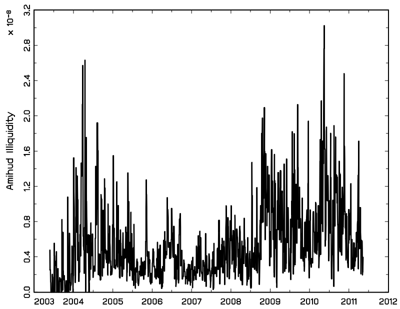

The next graph shows the ratio of (the weekly average of) $|retco|$ to (the weekly average of) $|retoc|$. The graph shows no discernible trend, the ratio seems to be fairly constant throughout the decade, and is generally less than one (the average over the whole sample is 0.44), but there are some days where this ratio is very large. We remark that we have not distinguished between normal close to open periods and weekend or holiday close to open periods, and the large values of the ratio generally come from the latter class of days. The point remains however that this ratio has not significantly trended up over the decade (as can be confirmed by taking a longer term average).

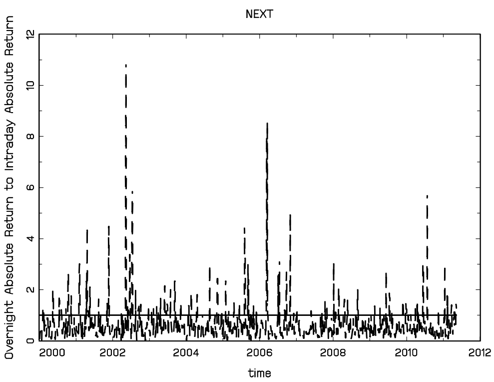

Northumbrian Water

Northumbrian Water PLC is a FTSE250 company first listed on 23/9/2003 with a market capitalization of £1677.75 million. The data is from 23/5/2003 to 12/4/2011. Over this period the share price increased nearly five-fold, although the period 2008/2009 saw the price almost half.

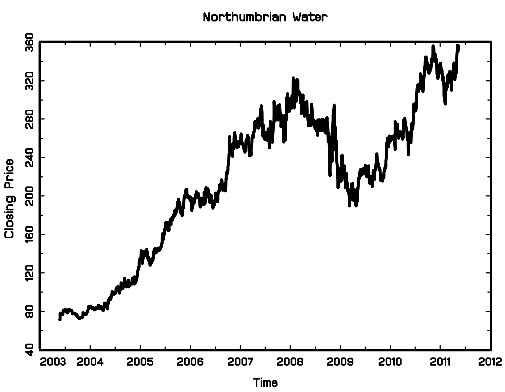

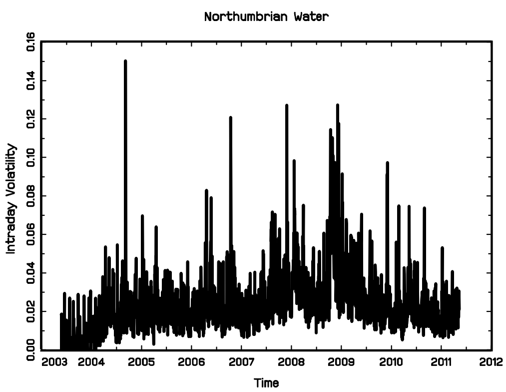

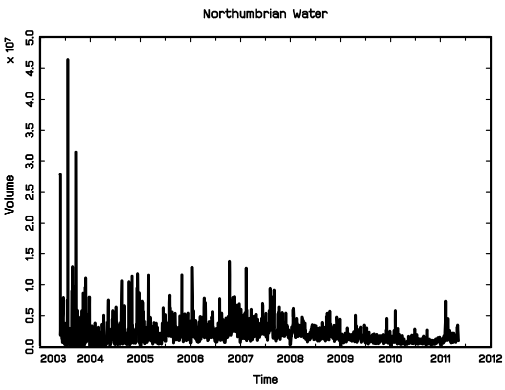

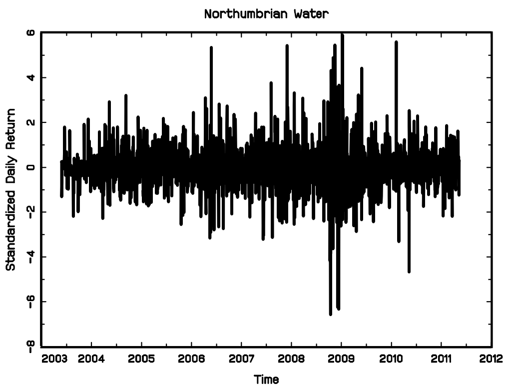

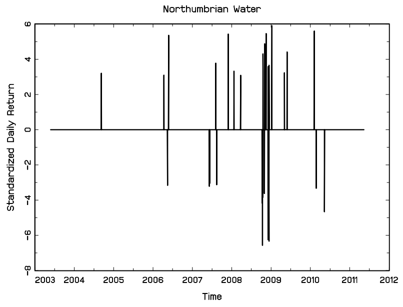

The autocorrelations and variance ratio tests are not statistically significant in any year.

British Telecom

British Telecom is a FTSE100 company first listed 3/12/84 with current market capitalization £14,638 million. The data is from 9/11/2001 to 12/4/2011. The share price has declined to almost half the level it was at the beginning of the period. There were especially rapid declines in price throughout 2008 and early 2009 followed by a rally.

There does not appear to be a secular trend in volatility; recently, volatility has fallen since the beginning of 2010. The picture in terms of close to close returns also supports this, a lot of volatility in the early 2000s and then in 2008/2009 with less substantial volatility since then. The largest single event was -9 standard deviations, which happened at the end of 2008. Volume went dramatically down throughout 2008/2009 but has recovered since the beginning of 2010.

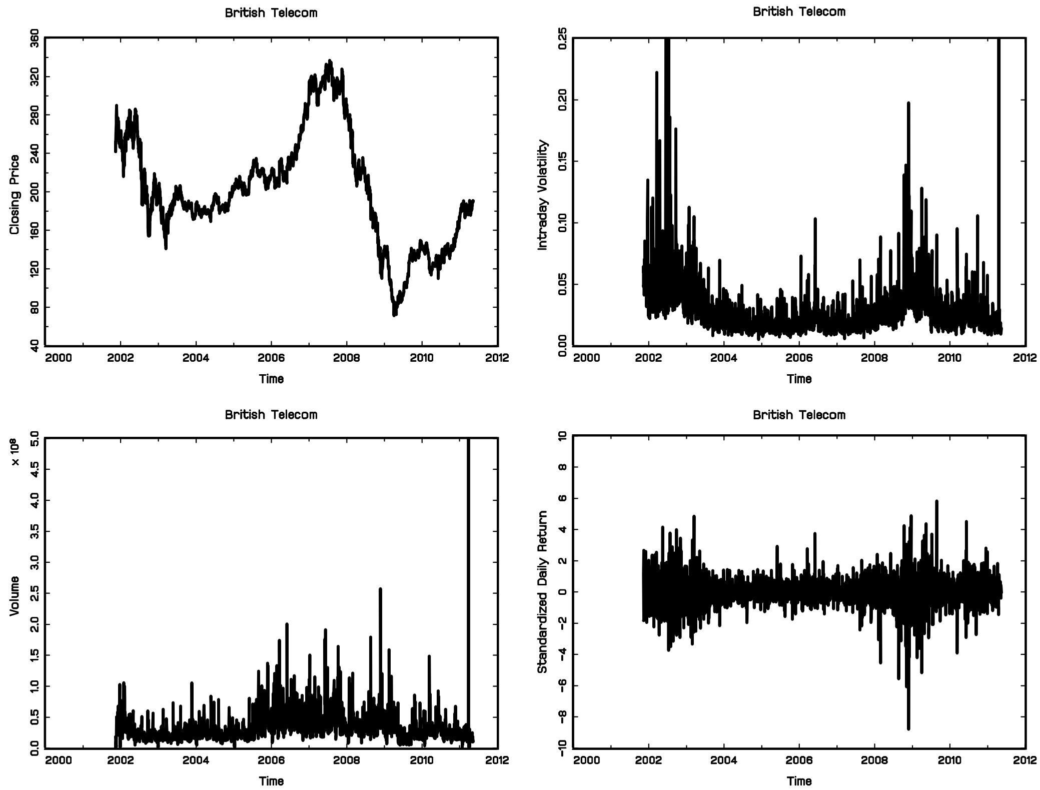

We next show the observations corresponding to the largest close to close returns over the period. Clearly, 2008/2009 were years when such large movements occurred frequently, but 2010 has been quieter. Similar results apply to the intraday variation.

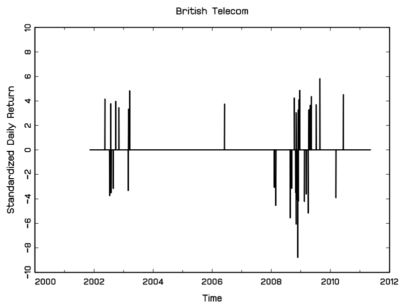

The table shows the first autocorrelation by year and the five day variance ratio test along with t-statistics. In 2009, there was significant negative autocorrelation, and the variance ratio test is significant.

Year	*ro*	t-stat	*vr*	t-stat
2002	-0.0077	-0.1179	0.7976	-1.1016
2003	-0.0862	-1.3267	1.1025	0.5580
2004	-0.0727	-1.1188	0.9390	-0.3321
2005	0.0493	0.7597	0.9777	-0.1213
2006	-0.0991	-1.5259	0.7984	-1.0971
2007	0.0016	0.0251	0.9096	-0.4920
2008	-0.1099	-1.6924	0.8897	-0.6002
2009	-0.1001	-1.5414	0.5527	-2.4343
2010	0.0090	0.1383	1.1263	0.6875
2011	0.0352	0.5414	1.2324	1.2647

Table 8. First order autocorrelations and variance ratio test statistics by year

United Utilities

United Utilities Plc is a FTSE100 company first listed on 28/7/08 with a market capitalization of £4056.89 million. The data is from 30/7/2008 to 12/4/2011. Over the shorter period we see the price declined considerable to almost half its initial value and then recovered.

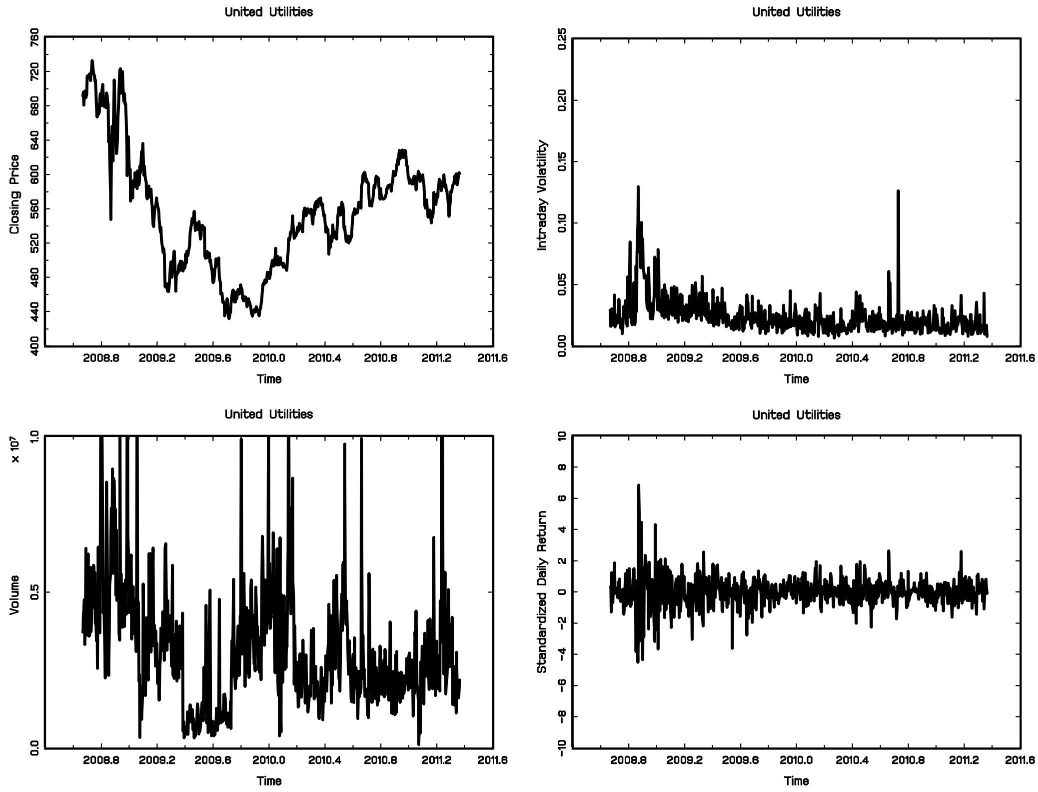

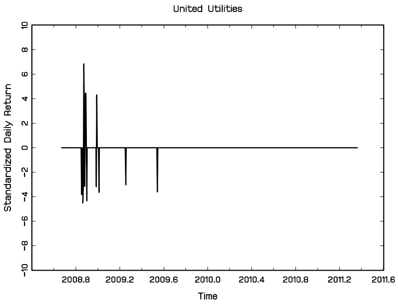

Standard and Poors 500

We show some results for the S&P500 index for benchmarking purposes. This index is extremely broad and represents a much greater value than the FTSE index. The series is shown from the first trading day of 1950 up to the present day, which gives a much longer perspective. The price series has basically been going sideways since 1999 when compared with the steady increase up until then. The return graph is truncated at ten standard deviations, so that the full scale of the 19/10/87 crash is not revealed. Note that the market on that day opened at the high and closed at the low, so the intraday volatility is essentially the same as the close to close return, both being extremely large. The longer time frame shows the cyclical variation in volatility and the upward trend in volumes. It also shows that periods with several large returns seem to occur every ten years or so. The most recent period was very bad by historical standards.

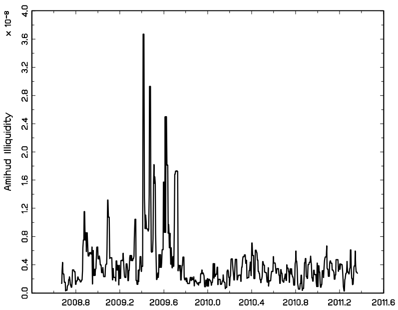

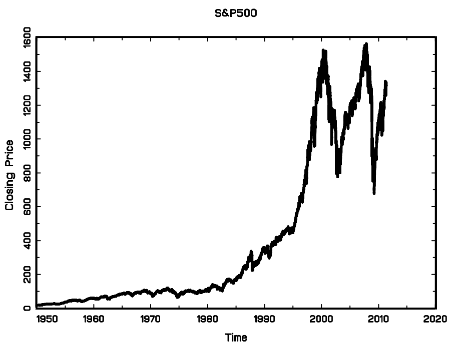

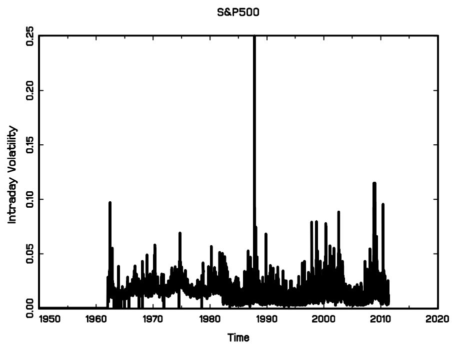

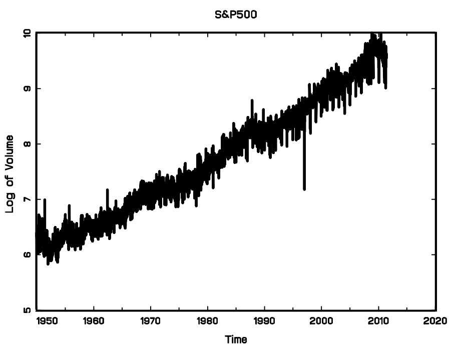

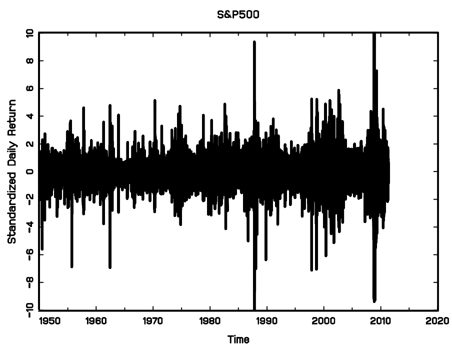

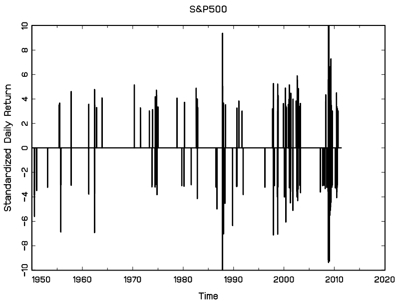

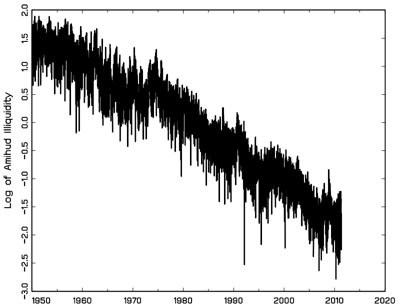

Has Market Quality Become Better or Worse in the last Decade?

We next investigate whether the market quality variables, (the logarithm of) , $volume$, and liq have experienced a trend either up or down in recent times. The graphs do not by themselves give a clear answer to this question, and may even be misleading. We need to take account of random variation in the series that can account for short term ups and downs. Let us just describe some possible models and their implications. The simplest model would be a stationary short memory linear autoregressive time series. In this case, the series fluctuates around some long run average and pulls back fairly quickly to that average after being hit by some shock. A limiting case is where the process is a random walk, i.e., non-stationary, in which case the series may never return to a steady state, and is called a stochastically trending series. A second type of non-stationary process is the deterministic trend, in which the process is driven along by some growth or decline that is determined outside the system and which eventually dominates any short term fluctuations in the series. The difference between trend stationary processes and difference stationary processes was at the heart of a lot of debate in macroeconomics thirty years ago, Nelson and Plosser (1982), especially with regard to the proper modelling framework to adopt for (the logarithm of) GDP. If GDP is trend stationary, then business cycle shocks or government policy shocks will have only a short run effect as the series is returned to trend eventually. If GDP is difference stationary, then these shocks may have permanent effects on outcomes and current losses may never be replaced. In recent times, authors have investigated an alternative possibility, somewhat intermediate, called long memory processes. These can be stationary or non-stationary, but possess the main feature that whatever the long run effects of a shock the transition to the long run may take a very long time. A common finding is that daily volatility series are very persistent to the point where standard stationary autoregressive models of short order are inadequate to describe the dependence of the series, and a long memory class provides a better fit to the data, Anderson, Bollerslev, Diebold, and Labys (2003). Rather less work has been published on volume and liquidity series, but they seem to possess similar time series properties to volatility.

We investigate the properties of the index and individual stocks. We first test for stationarity. Specifically, we perform standard statistical tests for the presence of "unit roots" in the logarithm of the three quality variables, called the Augmented Dickey Fuller test (full details are available

from the authors upon request). We use a fully automated method (in Eviews) to determine some auxiliary parameters. For volatility and liquidity we strongly reject the unit root hypothesis in favour of stationarity. For volume, the p-value for the test is just over 0.05 suggesting the unit root (non-stationarity) hypothesis cannot be rejected against the stationary alternative, but we note that the evidence is marginal. The logarithmic transformation is used because the resulting series has a marginal distribution that is much closer to a symmetric bell-shaped distribution than the untransformed series, and therefore one hopes that the linear methods we have used will be therefore better grounded.
Our investigations find similar results to Anderson, Bollerslev, Diebold, and Labys (2003). Specifically, we find that the volatility and liquidity series seem to be best described by a stationary but "long memory" model, meaning that shocks to the system take a long time to die out, but do eventually get expunged so that the series returns to some long run level. We however find some weak evidence that the volume series may even be non-stationary, meaning, it has no tendency to mean revert to some long run average. That is, a reduction in volume may never be replaced. We also estimated the same regressions for 70 of the FTSE100 stocks including BT and Next. The main findings are similar to that for the index. We find that volatility and liquidity appear to be stationary for most of the shares, whereas a significant fraction of the shares possess non-stationary volume series.

Fragmentation and Dark Trading

We next look at the issue of whether the fragmentation of equity trading and/or the increase in trading on unlit venues has affected market outcomes in a positive or negative way. MiFid round one was implemented November 1st, 2007, but fragmentation of the UK equity market began sometime before that, and by 13th July, 2007, Chi-X was actively trading all of the FTSE 100 stocks. We work with weekly UK data supplied to us by Fidessa (2011) that gives the volume traded for the FTSE-100 index, the FTSE-250 index, and the FTSE-All share index and where that volume was traded over the period 2008-2011. The data distinguish between lit (public exchanges with visible order book), dark pools (invisible order book), otc (over the counter), and si (systematic internalizer) venues. The list of lit venues includes: Bats Europe, Chi-X, Equiduct, LSE, Nasdaq Europe, Nyse Arca, and Turquoise. The list of dark pools includes: BlockCross, Instinet BlockMatch, Liquidnet, Nomura NX, Nyfix, Posit, Smartpool, and UBS MTF. The list of otc includes: Boat xoff, Chi-X OTC, Euronext OTC, LSE xoff, Plus, XOFF, and xplu/o. The list of systematic internalizers includes: Boat SI and London SI.

We first show the extent of "fragulation" (that is, the percentage of volume traded on the LSE) for the FTSE100 over the time period. The percentage of volume traded on the LSE has declined considerably. We also divide the volume into the four venue categories and show the evolution of trading volume; volume traded in the otc has increased considerably compared to the lit exchanges. Finally, we show the average size of trade by type of venue. This has decreased considerably for lit venues, but has increased a little for the otc and si venues.

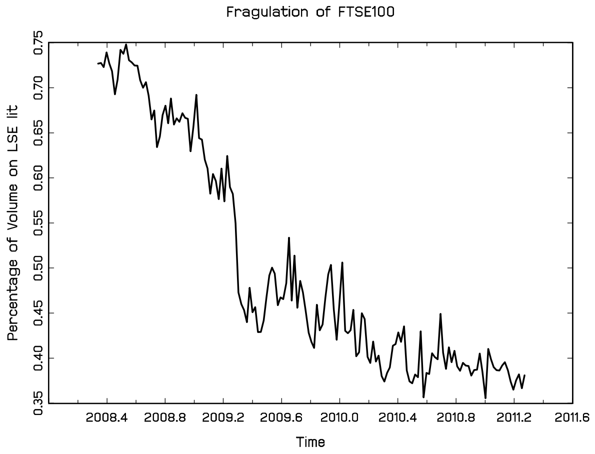

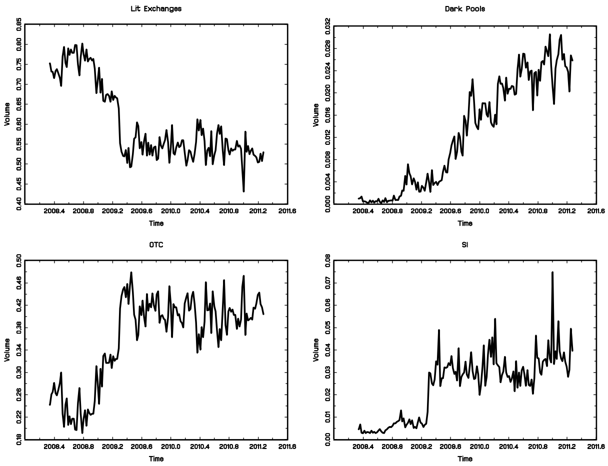

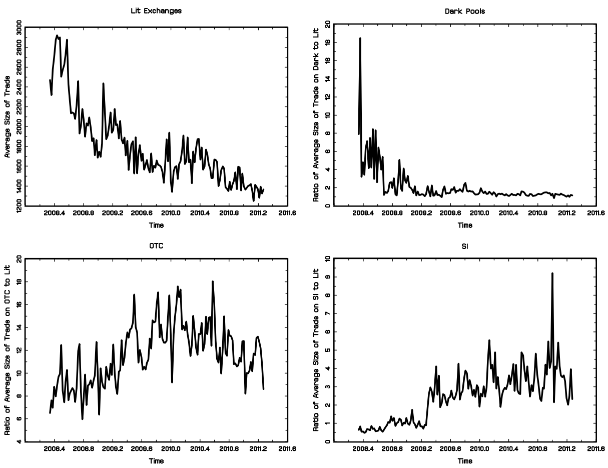

We did the same exercise for two of the individual stocks.

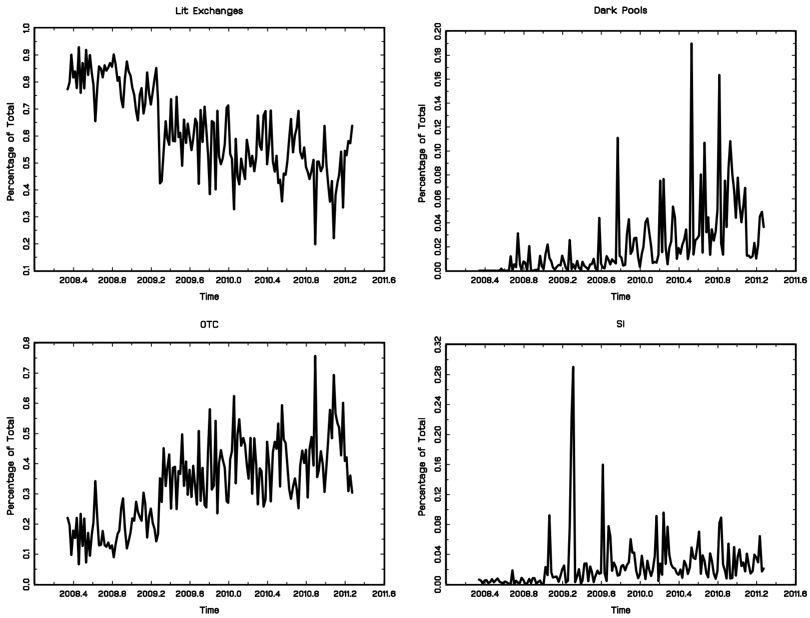

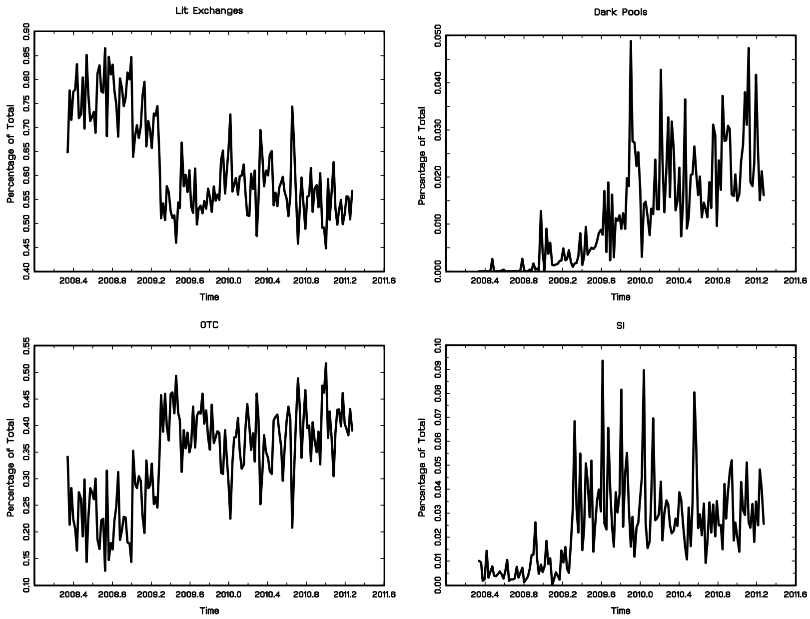

We also show the average size of trade by type of venue. Rather surprisingly we find that the average size of trades have increased on the public exchanges in the last three years, whereas the average size of trades has fallen in relative terms and even in absolute terms in some cases on the other type of venue.

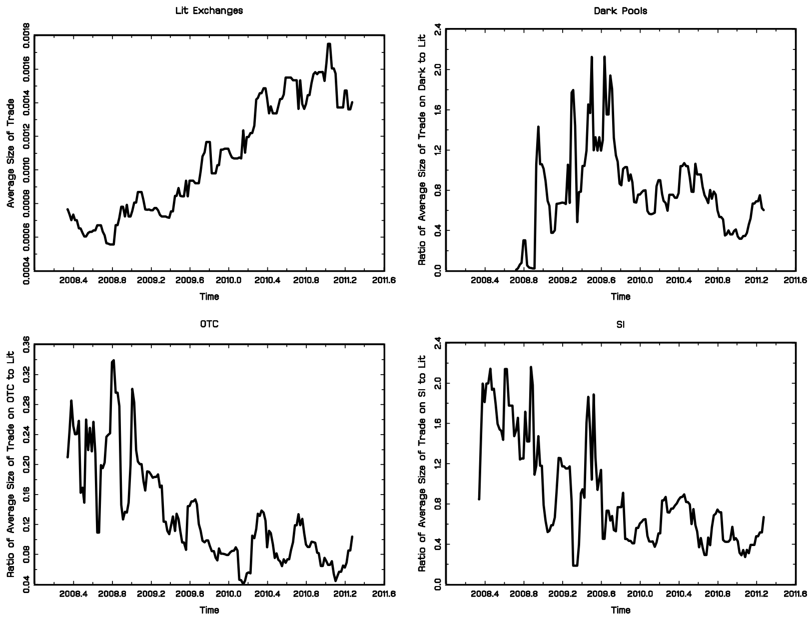

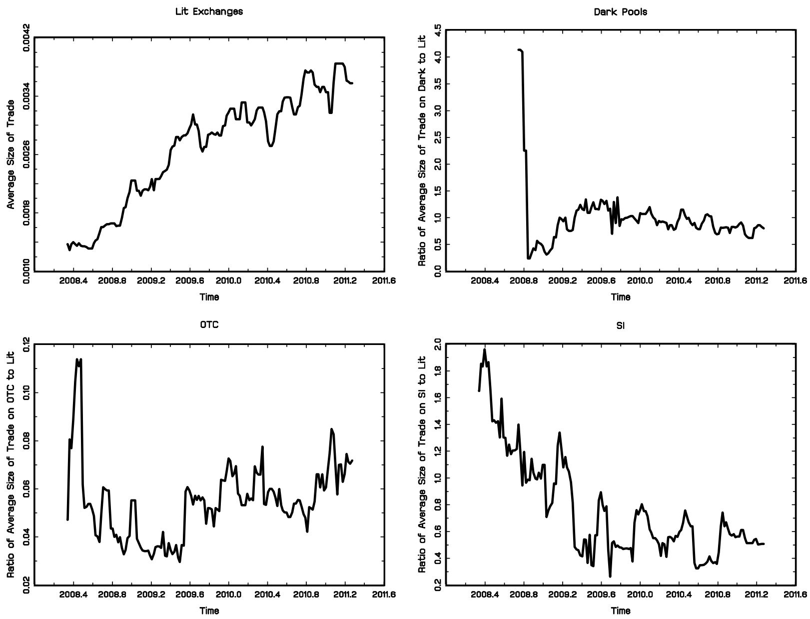

We next use this data set to ask whether the increasing market fragmentation and dark trading has been responsible for making market outcomes worse than they otherwise would have been, We look at volatility and liquidity as outcome variables, now at the weekly frequency. We consider several different measures of fragmentation or dark trading. Specifically, by index and by week: (1) the % volume on LSE-lit; (2) the % volume on dark venues; (3) the %volume on otc venues; (4) the % volume on si venues. One could just do a time series regression of the outcome variable for each series directly on the explanatory variable and read off the direct effect, but there are reasons why this will not give reliable results. Specifically, there may other variables that are changing over time in a similar pattern to the fragmentation variable and which have an effect on the outcome variables. By not explicitly including them or controlling for their effect we may be misattributing the changes to fragmentation.

To avoid this issue we will use what is called a difference in difference methodology applied to the panel data for the three indices, Card and Krueger (1994). The implicit model allows for unobservable index specific (time invariant) factors that affect the outcome variable and unobservable time varying factors that affect the outcome variables for all indices. We assume that there is a time invariant coefficient that measures the effect of fragmentation on outcome; this can vary across the indices. The model is thus more general than a direct regression and allows for unobservable time varying influences on outcome. The methodology exploits the fact that there is cross sectional variation in the degree of fragmentation across the indices (large cap indices appear to be more fragmented) as well as time series variation. It is this variation in the experimental conditions that allows one to measure the fragmentation effect by controlling for and using this variation. Specifically, we work with differences in differences, that is changes over time in the difference between volatility, say, of the FTSE-100 and the volatility of the FTSE-250. The methodology is robust to other time varying factors that affect outcomes but we don't observe or don't include explicitly in the regression. The variable lse100 is the percentage of the trading volume of the FTSE100 on LSE-lit, while lse250 is the percentage of the trading volume of the FTSE250 on LSE-lit. We find the following results

Dep	lse100	t-stat	lse250	t-stat
liq	3.440592	1.142608	-2.300237	-0.806536
Equal	0.323761	0.132819	0.323761	0.132819
vol	0.142839	1.761649	0.061199	0.796895
Equal	0.098515	1.511776	0.098515	1.511776

Table 9. Difference in difference regression of quality on fragulation

The row Equal refers to the case where the coefficients on lse100 and lse250 are restricted to be equal. In statistical significance testing, the p-value is the probability of obtaining a test statistic at least as extreme as the one that was actually observed, assuming that the null hypothesis is true. One often "rejects the null hypothesis" when the p-value is less than the significance level, which is often 0.05 or 0.01. When the null hypothesis is rejected, the result is said to be statistically significant. None of the effects in Table 9 is statistically significantly different from zero at the 5% level for the two sided test, although the (two-sided) p-value for the effect on the FTSE100 volatility is 0.078128620, so the test would reject at the 10% level, and indeed it would reject at the 5% level if our alternative hypothesis was that the coefficient was positive (one-sided test) since the p-value is then 0.039064310. The direction of the effects are as follows: positive effect of LSE volume percentage on illiquidity and volatility of FTSE100 and negative effect of LSE volume percentage on illiquidity of FTSE250 but positive effect on the volatility of FTSE250.

We next look at the issue of whether venue style affects market quality. We now divide the transaction flow according to whether the venue is lit or not, that is the variable lit100 is the percentage of the trading volume of the FTSE100 on lit venues, while lit250 is the percentage of the trading volume of the FTSE250 on lit venues. The results are similar to the LSE results; not much statistical significance and generally positive effects (meaning more volume on lit exchanges is bad for market outcomes).

Dep	lit100	t-stat	lit250	t-stat
liq	2.592981	1.103154	0.073370	0.034511
Equal	1.148788	0.635571	1.148788	0.635571
vol	0.116317	1.850374	0.046044	0.809829
Equal	0.076038	1.572593	0.076038	1.572593

Table 10. Difference in difference regression of quality on lit percentage

We remark that the limited cross-section dimension we have does not provide us with a great deal of power; furthermore, the indices are themselves averages of individual prices and therefore we may underestimate the actual level of temporary volatility and illiquidity pertaining to individual stocks. In future work we intend to carry out this analysis on the individual stocks of the FTSE250.

A further way of investigating the effects of competition is to look at the effects of the recent outages at the LSE, the Milan Borse and the Stockholm exchange. A couple of studies have shown that bid-ask spreads widened dramatically on Chi-X and BATS for example relative to their previous average values. One could argue that this shows the importance of the primary venue for price discovery. One could also argue that this shows that competition is good for liquidity. We hope to present in future completed work that investigates this question in more detail using order book information from multiple venues on and around that day.

Concluding Remarks and Speculations

Regarding the questions we aimed to address, we have found that over the last decade:

1) The volatility of the FTSEAll index shows no statistically significant deterministic trend over the decade or since the first round of MiFid. There was an increase in volatility during 2008/2009, but it has since returned to a more tranquil level; the variation is within the bounds allowed by a stationary process.
2) There were a lot of large price changes in 2008/2009 but since then the frequency of large scale price moves has reduced. Over the longer time frame, there appears to be bursts of large price movements approximately every ten years or so, and so it is difficult to judge whether the last few years have seen an increase in the frequency of market crashes or dashes. The LSE circuit breakers evidently prevented massive price changes on the day of the mini flash crash, so that this day does not even show up as an extreme event. This is different from what transpired during the US flash crash and is due partly to the scale of what was going on there and partly perhaps to the different institutional structure in the UK.
3) There may be some concerns about traded volume, although it is hard to make a firm conclusion about the statistical significance of the recent declines in volume for the index. The process driving traded volume of UK equities seems to be highly persistent, which means that bad shocks to volume, like that which occurred in 2008/2009 can take a long time to correct.
4) There does not seem to be a statistically significant trend in liquidity according to the daily measures. This is true both for the index and individual firms.
5) The markets do not appear to have changed in terms of their predictability (market efficiency) both for the index and individual stocks
6) The contribution to volatility from the intraday period seems to have remained constant over the last decade, at least for the index
7) There seems to be a slight positive relation between competition and market quality based on our index regressions, but the statistical significance of our evidence is low.
8) There seems to be a slight negative relation between illuminated trading and market quality based on our index regressions, but the statistical significance of our evidence is low.

We have made no adjustment for dividends in this work, although we doubt this would make much difference. In fact, there is some evidence that dividend payouts have been on the decrease in the last ten years.
It is always possible to over interpret the data. If one just worked with data from 2000-2009 and compared the pre MiFid period with the post miFid period one could say that the latter period was terrible and one might then blame it on the MiFid changes or on HFT. However, taking into

account the longer term variation of these series makes one see that this conclusion is unwarranted.

We are not "flash-crash deniers": we repeat again the caveat that our analysis does not say anything about the micro-picture, what is going on in the millisecond by millisecond environment of today's equity markets. We may not be able to identify the improvements in execution speed and transaction costs from our indirect measures, perhaps because the relation between the low frequency measures we work with and the high frequency real time measures has changed. Likewise, regular market disruption at the very high frequency could be taking place and we may not see the negative effects in the daily record because of the regular employment of circuit breakers that limit their effect.

Circuit breakers and their ilk are not without costs. After the 1987 crash, many major exchanges including the NYSE instituted several circuit breakers to halt or limit trading in times of market stress. However, it is not clear whether the existence of circuit breakers indeed stabilized price movements. Existing studies show mixed results. For example, Yoon (1994) investigated the Korean Stock exchange, which at the time had an extensive system of trading halts and circuit breakers with quite narrow limits which were triggered on approximately 13% of the days. He found some evidence of price overshooting and bad effects on price discovery. There is some worry that over reliance on trading halts to calm markets may lead to poor market outcomes, over the longer term.

The analysis we have conducted is based on historical data, including the most recent history. We now ask how this can be used to say what will happen in the future. The general finding we have arrived at is that the last decade has been one of "sideways" developments in which the broadest index of UK stocks has seen little or no improvement in the level. Furthermore, the market quality variables like liquidity, volatility, and trading volume have also shown only slight improvements over the decade. The same is true of other developed economy stock market indices like the S&P500, which has been in a similar holding pattern. We may judge this against previous decades where we have seen secular improvements in the price level of such broad indices, and considerable improvement of the market quality variables like liquidity and trading volume. The reasons for this stagnation have been widely discussed elsewhere. The global security situation has obviously created lots of anxiety and raised certain risk perceptions. The tax and regulation burden, the pension funding crisis, the banking crisis, and the sovereign debt crisis have all added to the costs of investing. At the same time, many new investment vehicles have appeared that have competed for investor attention like ETF's, ETP's, spread betting, etc. Also, emerging market economies and their stock markets have seen substantial growth. According to the World Federation of Stock exchanges, the 10 biggest stock markets in the world by market capitalization in (USD millions) at the end of 1999 and 2010 were

Rank	Exchange 1999	Dollar	Exchange 2010	Value
1	NYSE	11,437,597.3	NYSE Euronext (US)	13,394,081.8
2	Nasdaq	5,204,620.4	NASDAQ OMX	3,889,369.9
3	Tokyo	4,463,297.8	Tokyo SE Group	3,827,774.2
4	London	2,855,351.2	London SE Group	3,613,064.0
5	Paris	1,496,938.0	NYSE Euronext (Europe)	2,930,072.4
6	Deutsche Börse	1,432,167.0	Shanghai SE	2,716,470.2
7	Toronto	789,179.5	Hong Kong Exchanges	2,711,316.2
8	Italy	728,240.4	TSX Group	2,170,432.7
9	Amsterdam	695,196.0	Bombay SE	1,631,829.5
10	Switzerland	693,133.0	National Stock Exchange India	1,596,625.3
			BM&FBOVESPA	1,545,565.7

The top four positions have not changed (ignoring the name branding changes), although both Nasdaq and Tokyo have seen declines in market capitalization over the decade and NYSE and London have both seen only relatively modest increases in the market value. The striking feature of the 2010 picture is that positions 6-10 have been taken by emerging economy stock markets, like China, India, and Brazil, and the smaller European ones have been replaced by these larger capitalized overseas exchanges, which have evidently grown enormously throughout the decade. The growth of these exchanges has been due to the increase in market capitalization of their domestic firms, and this is likely to continue for the foreseeable future. It is not clear what role high frequency trading will have in the growth of these markets: the Shanghai market for example is subject to strict government controls and rules about trading that would ceteris paribus place it at a competitive disadvantage with regards to the more sophisticated European or American venues. It is despite the developments in high frequency trading that Shanghai (and to a lesser extent the other exchanges in this group) is growing, not because of it. One role that computer-based trading may have in the future is to facilitate the international flow of investment funds from savers to where they will make the greatest risk-adjusted, tax-adjusted, trading-cost adjusted returns. International investors may not currently choose to put a large fraction of their wealth under the jurisdiction of the People's Liberation Army, but this may change in the future depending on how the relative risk return profiles evolve. In conclusion, we think the next ten years is unlikely to yield rapid improvements in the price level of the main UK indices nor are we likely to find big improvements in the liquidity of the market or the amount of traded volume. We think the level of volatility in the market is likely to go up and down over this horizon but without establishing a persistent good or bad trend. There will come future crises in the equity markets but we cannot say what the future cause is likely to be or whether computer-based trading will be implicated as a primary factor, but at least we can say that so far the evidence is that it has not caused persistent negative outcomes in the markets we looked at.

References

Alizadeh, S., M.W. Brandt, and F.X. Diebold. Range-based estimation of stochastic volatility models. Journal of Finance, 57:1047--1091, 2002.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time series effects. Journal of Financial Markets 5, 31-56.

Anderson, T., T. Bollerslev, F. Diebold, and P. Labys (2003). Modelling and Forecasting Realized Volatility, Econometrica 71, 579-625.

Angel, J.J. (1997). Tick Size, Share Prices, and Stock Splits. The Journal of Finance Vol. 52, No. 2, pp. 655-681

Angel, J.J. (2011) Impact of Special Relativity on Securities Regulation. Georgetown University.

Card, D. and A. B. Krueger (1994), "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania," American Economic Review, v. 84, n. 4, pp. 774--775

Claessens, S., M.A. Kose, and M.E. Terrones (2011). Financial Cycles: What? How? When? IMF WP 11/76.

Cutler, D.M. Poterba, J.M. and L.H. Summers (1989). What Moves Stock Prices? Journal of Portfolio Management

Deo, R. and Richardson, M. (2003). "On the Asymptotic power of the variance ratio test", Econometric Theory, 19, 231-239.

Fair, R.C. (2002). Events that shook the market. Journal of Business 75, 713-731.

Faust, J. (1992). When Are Variance Ratio Tests For Serial Dependence Optimal?", Econometrica, 60, 1215-1226.

Fidessa (2011). http://fragmentation.fidessa.com/author/steve-grob/

Gabaix, X., and R. Ibragimov (2011). Rank-1/2: A simple way to improve the OLS estimation of tail exponents. Journal of Business and Economic Statistics 29, 24-39.

Gabaix, X., P. Gopikrishnan, V. Plerou, and H.E. Stanley (2005). Institutional investors and stock market volatility. NBER working paper w11722

Goodhart., C. and M. O'Hara. High frequency data in financial markets: Issues and applications. Journal of Empirical Finance, 4:73--114, 1997.

Goyenko, R.Y., C.W. Holden, and C.A. Trzcinka (2009). Do liquidity measures measure liquidity? Journal of Financial Economics 92, 153-181.

Gresse, C. (2010). Effects of the Competition between Multiple Trading Platforms on Market Liquidity: Evidence from the MiFID Experience.

Hendershott, T. (2011). High-Frequency Trading and Price Efficiency. University of California at Berkeley.

Hendershott, T., C.M. Jones, and A.J. Menkveld (2011). Does Algorithmic trading improve liquidity? The Journal of Finance 66,

Linton, O. (2010). Academic studies on the empirical effects of algorithmic trading on market quality. Mimeo, LSE.

Lo, AW and AC MacKinlay (1988). Stock market prices do not follow random walks: evidence from a simple specification test. Rev. Financ. Stud. (1988) 1(1): 41-66 doi:10.1093/rfs/1.1.41

London Economics (2010). Understanding the impact of MiFiD. www.cityoflondon.gov.uk/economicresearch

Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. The Journal of Business, Vol. 36, No. 4, pp. 394-419

O'Hara, M., and M. Ye (2009). Is fragmentation harming market quality? Cornell University.

Nelson, C.R. and C. I. Plosser (1982) Trends and random walks in macroeconomic time series: Some evidence and implications. Journal of Monetary Economics, vol. 10, issue 2, pages 139-162

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39, 1127-1139.

Sarr, A., and T. Lybek (2002). Measuring Liquidity in Financial Markets. IMF working paper 02/232.

Wang, F., S-J. Shieh, S. Havlin, and H.E. Stanley (2009). Statistical analysis of the overnight and daytime return. Physical Review E. 79, 056109-1-9.

World Federation of Exchanges. http://www.world-exchanges.org/statistics/annual/2010/equity-markets/domestic-market-capitalization

Yoon, J-W (1994). Circuit Breakers and Price Discovery: Theory and Evidence. PhD thesis, UCLA.

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk
URN: 11/1220