Time Series Properties of Liquidation Discount

F. Chan [a], **J. Gould** [a], **R. Singh** [a] **and J.W. Yang** [b]

[a]*School of Economics and Finance, Curtin University, GPO BOX U1987, Perth, Western Australia, 6845*
[b]*School of Accounting and Finance, University of Western Australia, Perth, Western Australia, 6009*
Email: *ranjodh.singh@curtin.edu.au*

Abstract: This paper proposes an approach for quantifying liquidity risk. Urgent liquidation of a portfolio will entail a liquidation discount. This is the market impact discount in value yielded by the immediate sale of the portfolio relative to its in hand market value calculated from the prevailing market conditions.The proposed approach is to firstly construct the log liquidation discount rate using stock market data available from the order book. The behaviour of this empirical time series is modelled and subsequently used to predict future behaviour of the liquidity risk associated with the portfolio. This is achieved by constructing eight different sized portfolios, each corresponding to different numbers of shares from N stocks over two time periods (morning and afternoon). Each stock is to be liquidated on a daily basis. The bid side order book is used to price the immediate sale of a given stock at time t. The price differential between the bid value and market value of the stock is defined as the liquidation discount rate of the stock at time t. Replicating this process for N stocks produces a time series of portfolio liquidation discount rates. Specifically, there are total of eight time series which based on eight different scenarios, each consisting of a different number of shares for a given stock. These scenarios are represented by α which denotes differing proportions of all shares on issue for a given stock. A log transform is applied to the series and these are further segmented into two time periods to investigate liquidity behaviour over time.

This paper proposes to model the time series properties of the log liquidation discount rate using the Autoregressive Fractional Integrated Moving Average (ARFIMA) - Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model. The mean component of the series is modelled using the ARFIMA(r, d, s) model and contains both ARMA(r, 0, s) and ARIMA(r, 1, s) as special cases ($d = 0$ and $d = 1$ respectively). The GARCH(p, q) model is used to model the variance component. A number of models are tested under varying lag structure i.e. different values of p and q. Model performance is based on a model's ability to forecast future values of log liquidation discount rate. The forecast accuracy is measured using the mean square error (MSE). Optimal models resulting from differing values of α over two time periods are identified. The results indicate that the ARFIMA(p, d, q) - GARCH(p, q) model consistently produces the most accurate forecasts over both time periods. For practical purposes a simpler model (in terms of lag structure) is proposed. This model offers a more intuitive interpretation with only a marginal loss in performance. The parameter estimates pertaining to each model are averaged over all values of α for each time period. This produces a two final models each corresponding to a time period. Using these models one can forecast (over n horizons) the variance of the log liquidation discount rate. This forecast is interpreted as the future liquidity risk associated with the portfolio. The empirical results suggest that the variance converges to its long run value at a faster rate in the morning compared to the afternoon.

Keywords: Liquidity risk, fractional differencing, GARCH, time of day effect

1 INTRODUCTION

Stock market investors are most obviously subject to market risk. This is the risk that future market prices will be less favourable than today. Additionally, investors will also be subject to liquidity risk. This is the risk that the volume of the shares to be traded cannot be immediately transacted at prevailing market prices. The degree of absence of stock market liquidity i.e. illiquidity poses a real cost to investors.

Previous studies on liquidity risk include Brennan and Subrahmanyam (1996) who investigate the empirical relation between monthly stock returns and measures of illiquidity obtained from intradaily data. Bertsimas and Lo (1998) derived portfolio transaction strategies that optimize the trade-off between market impact cost and the risk associated with price volatility. Engel et al. (2006) introduced *liquidation value at risk* to assess the price risk versus market impact cost trade-off under various liquidation strategies. Aitken and Comerton-Forde (2003) and Goyenko et al. (2009) provide a summary and comparison of a wide variety of many approaches to measuring liquidity.

This paper proposes an approach for quantifying liquidity risk; it does so by taking a perspective of a stock investor with a long spot position in a portfolio. Urgent liquidation of a portfolio will entail a liquidation discount. This is the market impact discount in value yielded by the immediate sale of the portfolio relative to its in hand market value calculated from the prevailing market conditions. The proposed approach is to firstly model the empirical time series behaviour of the log liquidation discount rate and subsequently use this to predict future behaviour of the liquidity risk associated with the portfolio. This is a step towards improved recognition of liquidity risk within analytical models of general portfolio risk. In fact, Settlements (2001) promotes the development of 'risk assessments that take account of market liquidity' and consideration of 'how such measures could be used in the disclosure of market risk'.

This paper has the following structure: Section 2 describes the construction of the liquidation discount rate. Section 3 outlines the modelling technique applied to the liquidation discount rate as well as the forecasting methodology used in the paper. Section 4 states the data sources used in this study and discusses the results of thfrom a practical perspective. Lastly, section 5 summarizes the major findings and limitations of this study.

2 LIQUIDATION DISCOUNT RATE

A portfolio consists of N company stocks at time t where $x = 1, 2, 3, .., N$. Let S_t^x denote the total number of ordinary shares on issue belonging to company x at time t. It is assumed that this portfolio contains αS_t^x shares of company x at time t. The α is a proportionate holding factor and as such represents the fraction of all shares on issue.

An order can be thought of as an intention to buy or sell a quantity of shares of given company stock at a specified price. An order book ranks the buy orders (bids) from the highest to lowest bidding prices and the sell orders (asks) are ranked from the lowest to highest asking prices. Consequently, the bids are matched with the asks for a given company stock in an order-driven market. This process occurs instantaneously throughout the market's operating hours.

Let $q_{t,i}^x$ represent the quantity of shares for a buy order at a bid price $b_{t,i}^x$ for company x at time t. At $i = 1$, $q_{t,1}^x$ is the quantity of shares for a buy order at the highest bid price $b_{t,1}^x$ for company x at time t. Similarly, when $i = 2$, $q_{t,2}^x$ is the quantity of shares for a buy order at the second highest bid price $b_{t,2}^x$ for company x at time t. When $i = m$ (the maximum value of i), $q_{t,m}^x$ is the quantity of shares for a buy order at the lowest bid price $b_{t,m}^x$ for company x at time t. The liquidation of αS_t^x shares will expend the bid side order book to a certain *depth*. In other words, the order is processed starting from the best price and quantity combination available through to the ith best combination until the number of shares to be liquidated at time t is obtained. This is represented by the following expression:

$$\sum_{i=0}^{d_t^{\alpha,x}} q_{t,i}^x \leq \alpha S_t^x \leq \sum_{i=1}^{d_t^{\alpha,x}+1} q_{t,i}^x \tag{1}$$

where $d_t^{\alpha,x} \in (0, 1, 2, 3.., m)$ denotes the depth of the bid side order book for a given proportion (α) of x shares at time t. Note that the value of α directly influences the depth i.e. a large number of liquidated shares will deplete the bid order book to a greater depth than a small number of shares. In cases where the number of shares to be liquidated exceeds the quantity of available shares, that is the depth of the bid side order book is exhausted it is assumed that the $m + 1$th best bid price of the excess shares is zero. This is an extremely rare

occurrence in the empirical data. However, this assumption is conservative and as such represents the worst case scenario. The quantity of x excess shares at time t is defined as

$$q_{t,m+1}^x = \alpha S_t^x - \sum_{i=1}^m q_{t,i}^x. \tag{2}$$

The liquidated value of αS_t^x shares is given by $lv_t^{\alpha,x}$ where

$$lv_t^{\alpha,x} = \sum_{i=0}^{d_t^{\alpha,x}} q_{t,i}^x b_{t,i}^x + (\alpha S_t^x - \sum_{i=0}^{d_t^{\alpha,x}} q_{t,i}^x) b_{t,d_t^{\alpha,x}+1}^x. \tag{3}$$

The market value of αS_t^x shares is denoted as $mv_t^{\alpha,x}$ and is assumed to be equal to the highest bid price plus a small premium p. This premium is derived from the empirical data itself.

$$mv_t^{\alpha,x} = \alpha S_t^x (b_{t,1}^x + p). \tag{4}$$

For a given value of α at time t, the portfolio liquidation discount rate is expressed as

$$L_t^\alpha = 1 - \frac{\displaystyle\sum_{x=1}^N lv_t^{\alpha,x}}{\displaystyle\sum_{x=1}^N mv_t^{\alpha,x}}. \tag{5}$$

This is simply the discount rate between the aggregated liquidation value (equation (3)) and the aggregated market value (equation (4)). For the purpose of analysis a log transformation is applied to equation (5). This transformed quantity is denoted by ℓ_t^α.

3   MODELLING FRAMEWORK

This study attempts to investigate the time series properties of the log liquidation discount rate. Preliminary analysis indicates that the time series exhibits a slow decaying autocorrelation function. This could potentially indicate that the series is fractionally integrated. Consider the following model:

$$\phi_r(L)(1-L)^d \ell_{t,\alpha} = \mu + \theta_s(L)\epsilon_t \tag{6}$$

$$\epsilon_t = \eta_t \sqrt{h_t} \quad \eta_t \sim \text{iid}(0,1) \tag{7}$$

$$h_t = \omega + \sum_{i=1}^p \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^q \beta_i h_{t-i}. \tag{8}$$

Equation (6) represents the ARFIMA(r, d, s) model where r and s denote the order of autoregressive and moving average parts of the model respectively. The d parameter represents the fractional difference term in the model and determines the long run behaviour of the time series. It is conveniently related to the Hurst index ($d = H - \frac{1}{2}$) which is a measure of the long term memory/dependence in a time series (refer to Mandelbrot and Van Ness (1968) for more details). The d parameter is estimated in the modelling process. However, fixing the value of d to 0 or 1 produces the ARMA and ARIMA models respectively. Granger and Joyeux (1980) and Hosking (1981) introduced this model. It attempts to capture the dynamics of the mean of the process. In equation (6), μ represents the drift term and $\phi_r(L) = 1 - \phi_1(L) - \phi_2(L) - ... - \phi_r(L)$ is the autoregressive operator where L is the lag operator such that $Ly_t = y_{t-1}$. Similarly, $\theta_s(L) = 1 + \theta_1(L) + \theta_2(L) + ... + \theta_s(L)$ represents the moving average operator. Both these polynomials have their unit roots outside the unit circle and share no common roots. The fractional differencing operator $(1-L)^d$ can be rewritten as

$$(1-L)^d = 1 - dL - \frac{d(1-d)}{2!}L^2 - \frac{d(1-d)(2-d)}{3!}L^3 + ... \tag{9}$$

$$= 1 + \sum_{j=1}^\infty \frac{\Gamma(j-d)}{\Gamma(j+1)\Gamma(-d)}L^j \tag{10}$$

where $\Gamma(.)$ is the Gamma function. This model assumes that the conditional variance is constant over time. However, estimating the ARFIMA model for the log liquidation discount rate shows that the residuals are not constant over time. Therefore, the variance of the process is modelled using a GARCH(p, q) model (equation (8)) where p and q represent the order of the autoregressive and moving average parts of the model respectively. Bollerslev (1986) introduced the GARCH model whilst extending the work of Engle (1982) on ARCH models. This allows one to model a time varying conditional variance. Hence, the ARFIMA(r, d, s) - GARCH(p, q) model is used to model the log liquidation discount rate. When $\omega > 0$, $\alpha_1, \alpha_2, ..., \alpha_p > 0$, $\beta_1, \beta_2, ..., \beta_q \geq 0$, $\sum_{i=1}^{p} \alpha_i + \sum_{i=1}^{q} \beta_i < 1$ and the estimated value of $d \in (-1, 0.5)$, the proposed model is stationary as well as invertible. Specifically, when $d \in (0, 0.5)$ the model exhibits long memory. In a pioneering study, Baillie et al. (1996) apply this model to analyze the inflation rate time series. A latter study by Ling (2003) shows that the US consumer price index inflation series is fractionally integrated.

The AFRIMA(r, d, s) - GARCH(p, q) model is used to capture the time series behaviour of log liquidation discount rate for a given α and time period. This is achieved by assessing the model's performance under varying lag structures. For example, let r and s take on values of 0,1,2 and 3. Similarly, p and q are assigned values of 1,2 and 3. The d parameter is either set to a fixed values of 0 and 1 or estimated. Given that α takes eight values (0.0001% - 0.0008%) for each time period (morning and afternoon), the number of models that are fitted is 6912 (4 r values, 3 d values, 4 s values, 3 p values, 3 q values, 8 α values and 2 time periods). From this set, models are selected based on their ability to forecast future values of log liquidation discount rates. Given that the data has 1283 observations for each value of α for the morning period, model parameters are estimated on the first 1000 observations. These estimates are then used to generate 283 forecasts. These forecasts are compared to the 283 actual observations. Let $d_{t,\alpha}$ denote the difference between the actual value and the forecast on tth day for a given α i.e.

$$d_{t,\alpha} = \ell_{t,\alpha} - \hat{\ell}_{t,\alpha}. \tag{11}$$

The optimal model is the one that minimizes these differences for all t and α. The mean squared error (MSE) is used for this purpose. The MSE is defined as

$$\text{MSE}_\alpha = \frac{1}{f_j} \sum_{t=N_j-f_j+1}^{N_j} d_{t,\alpha}^2 \tag{12}$$

where f_j is the forecast horizon and N_j is the total number of observations for a given time period j. Similarly, the mean absolute deviation (MAD) is also computed for comparison purposes.

$$\text{MAD}_\alpha = \frac{1}{f_j} \sum_{t=N_j-f_j+1}^{N_j} |d_{t,\alpha} - \bar{d}_{t,\alpha}|. \tag{13}$$

where $\bar{d}_{t,\alpha}$ the mean value of the deviations for a given α. The models producing the minimum MSE_α across all α are selected. The results from both measures are consistent for most values of α. The Schwarz-Bayesian Information Criterion (SBIC) together with standard residual information is used to assess the validity of a model fit. The SBIC is defined as

$$\text{SBIC} = \frac{-2LL}{N} + \frac{k \ln N}{N} \tag{14}$$

where LL is the log likelihood value, N is the number of observations (1000) and k is the number of parameters that are estimated. The resulting optimal models (one model for each value of α and time period) are compared to a first order lag model. This is done in order to gauge the increase in forecasting ability that higher order models possess. It is expected that this comparison will assist in the model selection process. The parameter estimates of the selected model are computed across all values of α and time period. These estimates are averaged over each time period in order to yield two final models i.e. one for each time period. From a practical perspective, an investor may want to forecast (over a given horizon) the liquidity risk based on the current information. This liquidity risk is defined as the variance of the log liquidation discount rate. Assuming that the variance component is governed by GARCH(1,1), iterating equation (8) over n horizons given the current information yields the following result:

$$\mathbb{E}[h_{t+n}|I_{t-1}] = h + (\alpha_1 + \beta_1)^n (h_t - h) \tag{15}$$

where $h = \omega/(1 - \alpha_1 - \beta_1)$ and represents the unconditional variance. For further details refer to Baillie and Bollerslev (1992).

4 RESULTS

4.1 DATA

Data for this study consists of stocks listed on the Australian Securities Exchange (ASX). The ASX is a purely order-driven market. In January of each year from 2006 to 2011, a value weighted portfolio of the top $N = 10$ stocks from the ASX and S&P200 index is compiled. The number of shares on issue is extracted from the Morning DatAnalysis database. In addition, morning (10:15 a.m.) and afternoon (15:45 p.m.) snapshots of the bid side order book are obtained from the Securities Industry Research Centre of Asia-Pacific (SIRCA) AusEquity database for the period October 2006 through to October 2011. This database provides bid orders to a maximum depth of 20 i.e. $m = 20$. The value of the premium p in equation (4) is $0.005. This is the average gap between the highest bid and the current share price. The values of α are $0.0001\%, 0.0002\%, .., 0.0008\%$. Given that the maximum depth of the bid side order book is 20, increasing the value of α beyond 0.0008% more frequently produces excess shares (refer to equation (2) in section 2). Hence, due to the available depth (data) restriction the maximum value of α is set to 0.0008%.

For all combinations of α ($0.0001\% - 0.0008\%$) and $d \in [0, 1]$, the model with the minimum MSE is selected. Results are classified by the differencing scheme implemented. As a consequence, there are 24 optimal models for each time period resulting from eight values of α and three differencing options. For a given value of α and time period, each model's MSE is divided by of the minimum MSE produced by the three differencing options. This *relative* MSE allows one to rank MSE values. The results indicate that the fractional differencing option produces the minimum MSE across all values of α for the morning period. This is also the case for majority of the α values for afternoon period. The exceptions being α values of $0.0002\%, 0.0003\%$ and 0.0004%. For these cases, $d = 1$ leads by a small margin.

This may suggest that the zero differencing ($d = 0$) is not be enough (under differencing) and the first difference ($d = 1$) leads to over differencing. Hence, fractional differencing ($0 < d < 1$) is deemed to be optimal. However, the results across all differencing options are quite close together. The estimated value of d is always less 0.5 and this indicates that there is long memory present in the data. These values are shown in table 2.

For both time periods, the lag structure seems to be little erratic for first three or four values of α. The remaining values of α have an identical lag structure (with the exception of $\alpha = 0.0008\%$). This lag structure represents an ARFIMA(1, d, 0) - GARCH(2, 3) model.

A natural question to ask is that how much *better* is GARCH(2,3) compared to GARCH(1,1). It can be argued that on average one cannot do better than a GARCH(1,1). To test this proposition, an ARFIMA(1, d, 0)-GARCH(1, 1) model is fitted and benchmarked against the optimal model for each value of α. The results are displayed in table 1. The MSE resulting from the proposed model is within four percent of the MSE of the optimal model. The simpler lag structure offers a more intuitive interpretation with only a marginal compromise in forecasting ability.

The parameter estimates of the proposed model are computed for all α values across both time periods as shown in table 2. It is evident from the results that these estimates are influenced by the values of α. In particular, the drift term μ, becomes marginally larger as the value of α increases. Some departures from central values can be seen in the results. These are caused by outliers in the data series. For example, the p and q parameter estimates for the morning periods at $\alpha = 0.0002\%$ are different from the remaining p and q values. The results also indicate that the fractional differencing estimate (d) is largely the same across both time periods. The computed value of d is less than a 0.5 implying that the series is stationary. Furthermore, this translates to a Hurst index of approximately 0.9 (defined in section 3) which implies that there is a relatively high degree of persistence present in the log liquidation discount rate.

Given that the parameter estimates remain relatively unchanged over different values of α, it may be ideal to propose an overall model for each time period. The overall estimates are the average values of the individual estimates across all values of α. Bates and Granger (1969) showed that combining models will produce more superior forecasts than individual models on average. The resulting models (morning and afternoon) are fitted on all values of α for a given time period. The MSE is once again used to assess the quality of the fit. From the results, it is evident that the MSE increases as the value of α increases. This is to be expected since the bid side order book starts to move towards its maximum depth as the number of shares to be liquidated increases i.e. as α increases. The afternoon models produces lower MSE values compared to the morning model. This is due to the fact that the orders are being fulfilled more smoothly. Potentially, this could imply that the market is more *stable* in the afternoon compared to the morning.

Table 1. MSE comparison between the optimal and proposed model

	Morning		Afternoon	
Optimal	alpha0001.3.0.2.1	0.02675	alpha0001.0.1.2.1	0.01117
Proposed	alpha0001.1.0.1.1	0.02713	alpha0001.1.0.1.1	0.01155
Optimal	alpha0002.1.0.2.3	0.03771	alpha0002.1.2.2.2*	0.01716
Proposed	alpha0002.1.0.1.1	0.03787	alpha0002.1.0.1.1	0.01782
Optimal	alpha0003.3.2.2.1	0.04305	alpha0003.1.2.3.2*	0.02027
Proposed	alpha0003.1.0.1.1	0.04368	alpha0003.1.0.1.1	0.02083
Optimal	alpha0004.1.0.3.2	0.04724	alpha0004.1.3.1.1*	0.02234
Proposed	alpha0004.1.0.1.1	0.04762	alpha0004.1.0.1.1	0.0227
Optimal	alpha0005.1.0.1.3	0.05078	alpha0005.1.0.1.3	0.02279
Proposed	alpha0005.1.0.1.1	0.05132	alpha0005.1.0.1.1	0.02308
Optimal	alpha0006.1.0.1.3	0.05388	alpha0006.1.0.1.3	0.02273
Proposed	alpha0006.1.0.1.1	0.0546	alpha0006.1.0.1.1	0.02306
Optimal	alpha0007.1.0.2.3	0.05646	alpha0007.1.0.2.3	0.02324
Proposed	alpha0007.1.0.1.1	0.05726	alpha0007.1.0.1.1	0.02361
Optimal	alpha0008.1.0.2.3	0.05905	alpha0008.1.0.3.3	0.02401
Proposed	alpha0008.1.0.1.1	0.05977	alpha0008.1.0.1.1	0.02442

Table 2. Parameter estimates for morning and afternoon data

Morning						
α (%)	μ	r	d	ω	p	q
0.0001	-7.93035	-0.26481	0.37842	0.00058	0.01624	0.97434
0.0002	-7.66334	-0.27665	0.40132	0.04754	0	0.37783
0.0003	-7.46893	-0.25241	0.38740	0.00055	0.01202	0.98096
0.0004	-7.29658	-0.24724	0.38464	0.00067	0.01479	0.97710
0.0005	-7.14616	-0.23499	0.38435	0.00064	0.01556	0.97667
0.0006	-7.01559	-0.22552	0.38687	0.00055	0.01443	0.97883
0.0007	-6.90785	-0.22181	0.39075	0.00057	0.01495	0.97806
0.0008	-6.82838	-0.23730	0.40638	0.01865	0.06719	0.71947
Afternoon						
α (%)	μ	r	d	ω	p	q
0.0001	-8.09875	-0.14601	0.37772	0.00020	0.01953	0.97461
0.0002	-7.96591	-0.15864	0.37295	0.00084	0.03365	0.94928
0.0003	-7.85305	-0.17299	0.37768	0.00078	0.03282	0.95332
0.0004	-7.74218	-0.17606	0.38263	0.00021	0.02183	0.97436
0.0005	-7.64380	-0.18472	0.38769	0.00011	0.02058	0.97730
0.0006	-7.55451	-0.18912	0.39187	0.00006	0.02072	0.97782
0.0007	-7.47340	-0.19409	0.39628	0.00003	0.02096	0.97802
0.0008	-7.39675	-0.19912	0.39996	0.00001	0.02056	0.97863

4.2 LIQUIDITY RISK

From equation (15) it is evident that as forecast horizon increases to some large value i.e. $n \to \infty$ the conditional variance forecast approaches the unconditional variance. The rate at which this convergence happens is dependent on quickly $(\alpha_1 + \beta_1)^n \to 0$. In this empirical study $\alpha_1 + \beta_1 \approx 0.89$ for the morning period and 0.99 for the afternoon period. This implies that the conditional variance (risk) of the morning log liquidation discount rate converges to its limiting case (unconditional variance) much faster than the afternoon discount rate. From an investment prospective, it may be best to liquidate in the afternoon since, the discount rate in the morning converges to its maximum at a faster pace. As mentioned in the previous section, the order book's operation is smoother in the afternoon compared to the morning.

5 CONCLUSION

This paper has introduced a liquidation discount measure. This being the market impact discount in value yielded by immediate sale of the portfolio relative to its in-hand market value. For differently sized portfolios of major Australian stocks, it is found that an ARFIMA(1, d, 0)-GARCH(1, 1) model offers an explanation and a consistent parametrization for log liquidation discount rate. This model is used to quantify liquidity risk in terms of the conditional variance of the log liquidation discount rate. The results suggest that this risk changes over time. However, these results are restricted to top performing stocks in the ASX over the period October 2006 through to October 2011.

REFERENCES

Aitken, M. and C. Comerton-Forde (2003). How should liquidity be measured? *Pacific-Basin Finance Journal 11*(1), 45 – 59.

Baillie, R. T. and T. Bollerslev (1992). Prediction in dynamic models with time-dependent conditional variances. *Journal of Econometrics 52*(12), 91 – 113.

Baillie, R. T., C.-F. Chung, and M. A. Tieslau (1996). Analysing inflation by the fractionally integrated ARFIMA-GARCH model. *Journal of Applied Econometrics 11*(1), 23–40.

Bates, J. M. and C. W. J. Granger (1969). The combination of forecasts. *OR 20*(4), pp. 451–468.

Bertsimas, D. and A. W. Lo (1998). Optimal control of execution costs. *Journal of Financial Markets 1*(1), 1 – 50.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307 – 327.

Brennan, M. J. and A. Subrahmanyam (1996). Market microstructure and asset pricing: On the compensation for illiquidity in stock returns. *Journal of Financial Economics 41*(3), 441 – 464.

Engel, R., R. Ferstenburg, and J. Rusell (2006). Measuring and modeling execution cost and risk. Working Paper.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica 50*(4), pp. 987–1007.

Goyenko, R. Y., C. W. Holden, and C. A. Trzcinka (2009). Do liquidity measures measure liquidity? *Journal of Financial Economics 92*(2), 153 – 181.

Granger, C. W. J. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis 1*(1), 15–29.

Hosking, J. R. M. (1981). Fractional differencing. *Biometrika 68*(1), 165–176.

Ling, S. (2003). Adaptive estimators and tests of stationary and nonstationary short- and long-memory ARFIMA-GARCH models. *Journal of the American Statistical Association 98*(464), 955–967.

Mandelbrot, B. B. and J. W. Van Ness (1968). Fractional brownian motions, fractional noises and applications. *SIAM review 10*(4), 422–437.

Settlements, B. f. I. (2001). *Final Report of Multidisciplinary Working Group on Enhanced Disclosure*. Bank for International Settlements.