SYRTO

SYstemic Risk TOmography
Signals, Measurements, Transmission Channels, and Policy Interventions

Is there light in “dark trading"? A GARCH analysis of transactions in dark pools

Philippe De Peretti, Oren J. Tapiero

SYRTO WORKING PAPER SERIES
Working paper n. 1 | 2014

This project has received funding from the European Union’s Seventh Framework Programme for research, technological development and demonstration under grant agreement n° 320270.

This documents reflects only the author's view. The European Union is not liable for any use that may be made of the information contained therein.

Is there light in "dark trading" ? A GARCH analysis of transactions in dark pools

Philippe de Peretti[1]

Maison des Sciences Economiques, Centre d'Economie de la Sorbonne (CES) Université Paris 1 Panthéon-Sorbonne, 106 Avenue de l'Hopital, 75634 Paris CEDEX 13 France

Oren J. Tapiero[2]

Maison des Sciences Economiques, Centre d'Economie de la Sorbonne (CES) Université Paris 1 Panthéon-Sorbonne, 106 Avenue de l'Hopital, 75634 Paris CEDEX 13 France

---

Abstract

The ability to trade in dark pools without publicly announcing trading orders, concerns regulators and market participants alike. This paper analyses the information contribution of contemporary transaction in dark pools on future prices. The analysis is conducted by performing a nested and non-nested pairwise comparisons between GARCH models, with GED distributed disturbances. We compare several models using (log) scores based on the one-step ahead density forecasts. In each model, different proxies of activity in dark pools (volume and number of transactions) are included in the variance equation. Results indicate that the activity in dark pools conveys meaningful information for the one-step-ahead log-returns density forecast. Furthermore, it is to some extent meaningful to the price discovery. At last, traded volume and transactions in dark pools each provide distinct information.

Keywords: Dark pools, Price Discovery, One-step-ahead density forecast.

☆The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7-SSH/2007-2013) under grant agreement n 320270 SYRTO.

Email addresses: `Philippe.de-Peretti@univ-paris1.fr` (Philippe de Peretti), `oren.tapiero@univ-paris1.fr` (Oren J. Tapiero)

[1]The Systemic Risk Tomography: Signals, Measurements, Transmission Channels, and Policy Interventions (SYRTO)

[2]Laboratoire dexcellence - régulation financière (Labex-Réfi)

Introduction

The ability to trade in dark pools, without publicly announcing trading orders concerns market participants and regulators alike. In 2009, SEC Chairman and head of division on trading and markets Mary Schapiro and James Brigagliano expressed their concern indicating that the trading activity in dark pools may impair the price discovery process. In an article on the New-York Times (March 31^{st}, 2013), regulators have further expressed their concern that such an impairment would eventually drive ordinary investors away from the markets. Therefore, regulators suspect that dark pools may negatively affect trading liquidity. To address these concerns, some countries have taken regulatory measures over dark trading activity. Canada, for instance, heavily regulates this activity by allowing these kinds of trades only if there is a significant price improvement relative to executions on public exchanges. While, in Australia regulators have recently[3] proposed to impose a minimum threshold for orders in dark pools. Another potential concern for regulators is that it may be also a potential venue for price manipulations. For example, a trader may push up the price on the public exchange (by issuing multiple buy orders) while simultaneously selling in the dark pool. However, Kratz and Schöneborn (2012) overrules this possibility.

There are several incentives for institutional investors to trade in dark pools. First they are not obliged to make their intentions public. This implies that an institutional investor is able to execute large orders with fewer trades and without significantly affecting market impact risk. Boni et al. (2012) support this claim by indicating improved execution quality for large trades carried in dark pools. Thus, combined with mid-quote pricing, overall transaction costs paid by the institutional investor decreases. However, an investor engaging in this activity faces an execution risk because the dark pool does not guarantee

[3] Reuters web site (April 9^{th}, 2013): http://www.reuters.com/article/2013/04/09/us-exchanges-sec-darkpool- idUSBRE93818520130409

trading executions. This may imply that in moments of high intraday price volatility, the institutional investor will prefer to trade in public exchanges. Another incentive to trade in dark pools relates to the possibility of information asymmetry that may exists in public exchanges. Zhu (2012) state that dark pools enable investors to avoid trading against an informed order-flow. Moreover, both medias and regulators assert that dark pool trading activity has been increasing almost in tandem with the high frequency trading activity. In other words, the increasing activity in dark pools may be reflecting institutional investors distrust of public exchanges due to high frequency trading activity[4]. Provided this is true and provided institutional investors are able to detect high frequency trading activity, trading in dark pools may coincide with the latter trading activity.

While regulators and CEOs of public exchanges[5] have expressed their concerns, academic papers indicate some of the potential benefits and questionable aspects of dark pool trading activity. Buti et al. (2011) indicate that dark pool trading activity is higher on days with high share volume, low intraday volatility and high depth. O'Hara and Ye (2011) find that market fragmentation (in general) does not impair overall market quality. Moreover they find that while short-term volatility has increased, price dynamics has become closer to the random walk (implying greater market efficiency). At last they find that overall executions are faster and transaction costs are lower. Nevertheless, Ye (2010) indicate that introducing a dark pool does negatively affect price discovery on the public exchange while improving overall liquidity. He notes that his improvement is explained by less informed trading taking place on the exchange. Weaver (2011) also finds a negative relationship between increased dark pool activity and market quality (i.e.: price discovery) by indicating the positive effect it has on the measures of bid-ask spread. On the other hand, Zhu (2012) indicates that while

[4]Boni et al. (2012) indicate that some dark pools are specifically designed for institutional investors and discourage other participants such as high frequency traders.

[5]see: footnote 3

price discovery is improved by the presence of a dark pool, trading liquidity on the public exchange is reduced. Ready (2010) analyses volume in dark pools to finds that lower stock spreads (in dollars term) coincide with reduced dark pool activity, which conforms Zhu (2012) prediction. Nimalendran and Ray (2011) mitigate the "price discovery impairment" argument by indicating the possibility that informed traders may also trade in dark pools and therefore "spilling" information into the quotes that are seen in public exchanges. Nevertheless, two years later, the same authors (using propriety data) find increased quoted spreads on public exchanges following dark pool transactions (Nimalendran and Ray, 2013). Moreover, they find that informed traders may be concurrently trading in the "light" and in the "dark".

To compete with dark pools public exchanges (e.g.: Euronext-Paris, BATS, NASDAQ, NYSE and others) have started to allow traders to hide some or all of their order size. Bessembinder et al. (2009) (using data from Euronext-Paris) find that hidden orders take more time to be executed and that there is some execution risk associated with these orders. However, they also find that allowing hidden orders does not drive away defensive investors from the exchange. Buti and Rindi (2013) find that allowing hidden orders on public exchanges benefits large traders, while small traders are beneficial only when the tick size is large. Furthermore, they find that internal spreads widens with the presence of hidden orders. Therefore, overall, it seems that the effect of hidden orders on trading is to an extent similar to effect of dark pools.

The goal of this paper is to analyse the effects of trading in dark pools on future log-returns. To that end, the analysis is based on the predictive content of dark trades concerning future prices, modelled using an Auto-Regressive model with a conditional variance following a Generalised Autoregressive Conditional Heteroskedastic model (AR-GARCH). To take into account excess kurtosis, we set the law of the innovations as the Generalised Error Distribution (GED). We then perform pairwise comparisons over various models including:

1. A benchmark model, with no information on trading activity in dark pools,
2. models in which the proportion of volume traded in dark pools is included in the variance equation (linearly and non-linearly),
3. models in which the proportion of transactions made in dark pools is included in the variance equation (linearly and non-linearly).

We follow Amisano and Giacomini (2007) to compare between evaluated models. This involves using a Weighted Likelihood Ratio test, based on log-scores computed from log-returns one-step ahead density forecast. It enables to clarify the importance of dark pool trading activity on the accuracy of log-returns forecasts. This test is to be implemented on the entire or selected areas of log-returns forecast density. For practitioners and regulators, the center and tails of the distribution may be of most significance. This test is implemented on the Trade and Quotes (TAQ) transaction data . The data, with milliseconds timestamp, is aggregated at 5-minute intervals and reflects the trading period of the first six months of 2013 for IBM, Microsoft and General Electric.

For Microsoft and General Electric we find that the proxy for proportion of dark trades over-performs the proportion of dark volume (both linearly and non-linearly) in predicting future intraday returns. This, regardless of the highlighted region of the one-step-ahead forecast density. In the case of IBM, we find that the proportion of dark volume has a superior information contribution. Nevertheless, overall, the inclusion of either of these proxies in the AR(1)-GARCH(1,1) estimation framework over-performs the benchmark model in forecasting both the center and tails of the distribution. These results highlight the informational content that transactions on dark pools may contain. They also highlight that for the market, the size of dark trades (in monetary terms) is not as important as the frequency at which they occur. This is especially important when considering extreme one-step ahead realisations of returns for which the number of dark-trades is more predictive. Furthermore, as our results indicate, it becomes even more important (as in the case of Microsoft and General Electric) when considering the non-linear effects that dark-trading activity may have on the returns process.

This paper is divided into four sections. The first section describes the dataset used for this work. The second section describes the empirical methodology used in this paper. The third section presents and discusses the empirical results and the last section concludes this paper.

1. Data description and analysis

We retrieved data from the Trades and Quotes (TAQ) database. The database consists of two distinct files, one indicating quotes and another the transactions. The data set is time stamped to the milliseconds and reflects the transactions made within active markets hours, i.e.: 9:30:00:000 to 16:00:00:000, though we discard the first and last 15 minutes of the trading session. The transaction file contains all the transactions made in the existing trading venues. For our sample period we choose the period that starts on January 2013 and ends on June 2013. Moreover, we choose the actively traded and liquid stocks of IBM, Microsoft and General Electric as reference.

Information on dark-trading is indicated by the letter 'D' in the 'Exchange' column of the transaction dataset. More precisely, the designated letter 'D' indicates all trades reported by the Financial Industry Regulatory Authority (FINRA) that oversees trades executed in other Trade Reporting Facilities (TRF) including dark pools. We indicate that this variable has been used as a proxy for dark pool trading in Boni et al. (2012) and Weaver (2011). Furthermore, Weaver (2011) indicates that 90% of all TRF trades are executed in dark pools. Therefore, we assume that the Exchange variable provides a plausible proxy for dark pools trading activity. The transactions file also contain information on trade size (volume) and the condition at which it was executed. Thus, we are able to have an approximation of the proportion of dark pool trading both in terms of traded volume and number of transactions.

Using the transactions data we compute 5 minutes log-returns, $r_t = \ln(P_t/P_{t-1})$, using only prices reported on public exchanges. Where the price (P_t) used to calculate log-return is the last observed price within a predetermined time interval. We choose a 5 minutes time interval in order to avoid microstructural noise. Then, using the 'Exchange' variable and the indicating letter 'D' in the TAQ transaction data, we compute our two proxies. The first is the proportion of volume traded in the dark designated by the letter $V_t^D \left(= Q_s^D/Q_s\right)$, while the second is the proportion of dark-trades designated by the letter $N_t^D \left(= TN_t^D/TN_t\right)$. Where: $Q_s(Q_s^D)$ is the quantity traded in the s'th transaction carried in public (dark) exchange and $TN_t(TN_t^D)$ total number of public (dark) transaction within a pre-specified time interval.

The relationship between absolute returns, proportion of dark-trades (N_t^D) and the proportion dark volume (V_t^D) is plotted (for each stock) in three (figures 1-3) and two (figures 4-6) dimensional figures. A-priori, in the case of Microsoft and IBM, the three dimensional figures indicate that the absolute value of log-returns decreases with respect to the proportion of dark-trades and the proportion dark volume. While, in the case of General Electric, figure 3 indicates a possibly concave relationship between the absolute value of log-return and the two variables. The two dimensional figures provide a closer examination, however. It reveals a possibly concave relationship (for all stocks) between absolute log-returns and the two proxies (proportion of dark trades and volume). This concave relationship is likely to be related to dark trading occurring when volatility is low due to the execution risk associated with trading in dark pools.

Figure 1: **3D Distribution of $|r_t|, N_t^D$ and V_t^D (Microsoft)**

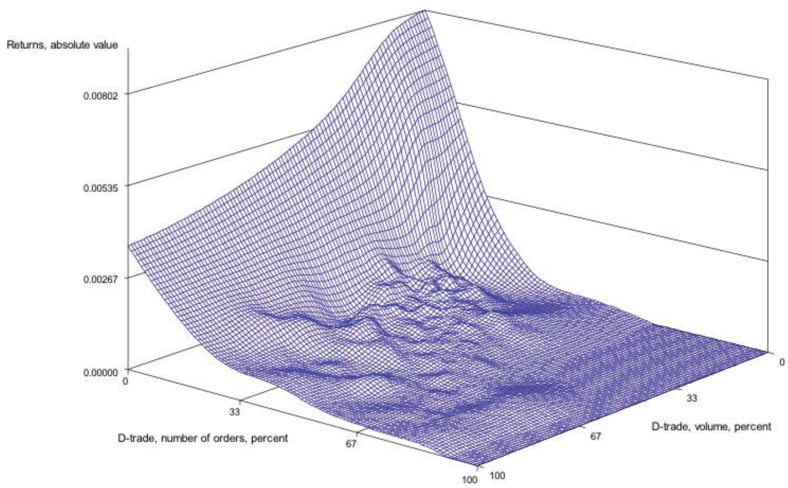

2. Empirical Methodology

Let $\{r_t\}_{t=1}^T$ be a series of returns, and assume that the data admits the following data generating process[6]:

$$\begin{aligned} r_t &= \rho r_{t-1} + \alpha + \varepsilon_t \sqrt{h_t} \\ h_t &= \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta h_{t-1}^2 + \varphi X_t \end{aligned} \tag{1}$$

Where:

- ρ, α, α_0, α_1, and β are parameters to be estimated.
- $\alpha_0 > 0$, $\alpha_1 > 0$, $\beta > 0$ and $\alpha_1 + \beta < 1$.
- φ is a $(1 \times k)$ vector of parameters associated with the $(k \times n)$ matrix of exogenous variables X_t.
- $\varepsilon \sim L(\cdot)$

[6]Before choosing the AR(1)-GARCH(1,1) model, we have estimated various models, also with different laws for the residuals (Student, Skew-Student, Skew-GED). Clearly the AR(1)-GARCH(1,1)-GED performs best.

Figure 2: **3D Distribution of $|r_t|, N_t^D$ and V_t^D (IBM)**

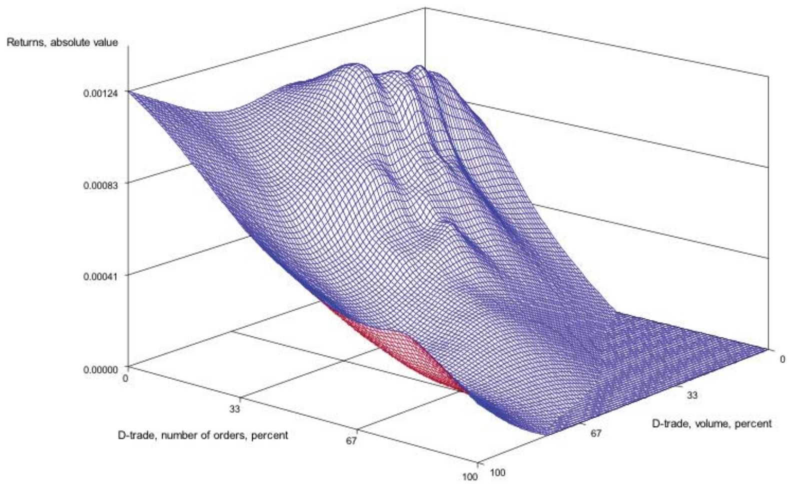

Figure 3: **3D Distribution of $|r_t|, N_t^D$ and V_t^D (General Electric)**

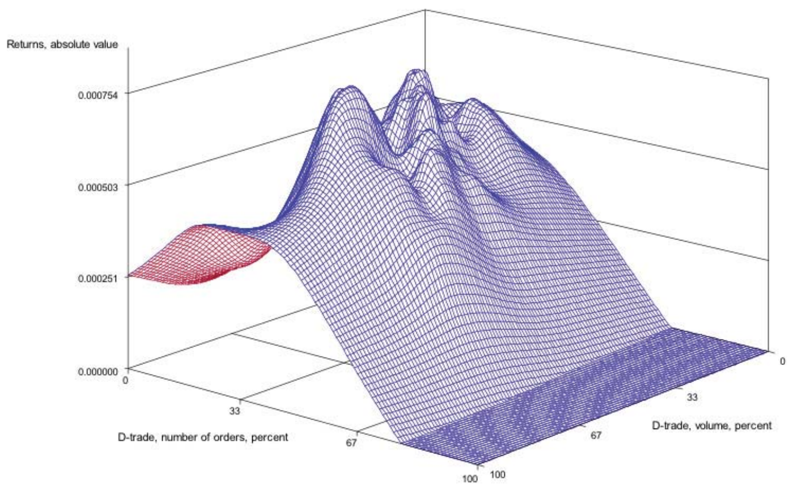

Figure 4: **3D Distribution of $|r_t|, N_t^D$ and V_t^D (Microsoft)**

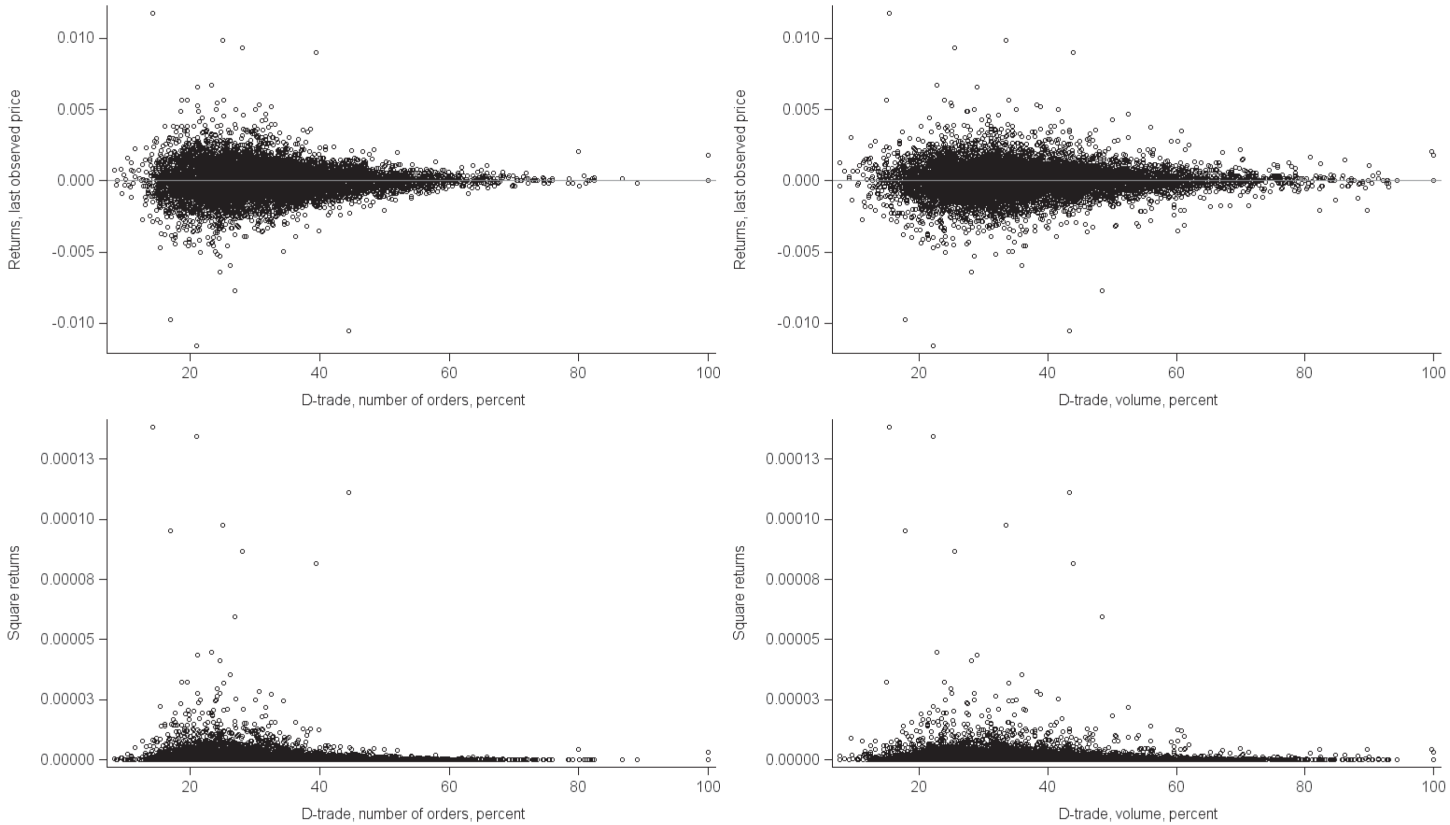

Figure 5: **3D Distribution of $|r_t|, N_t^D$ and V_t^D (IBM)**

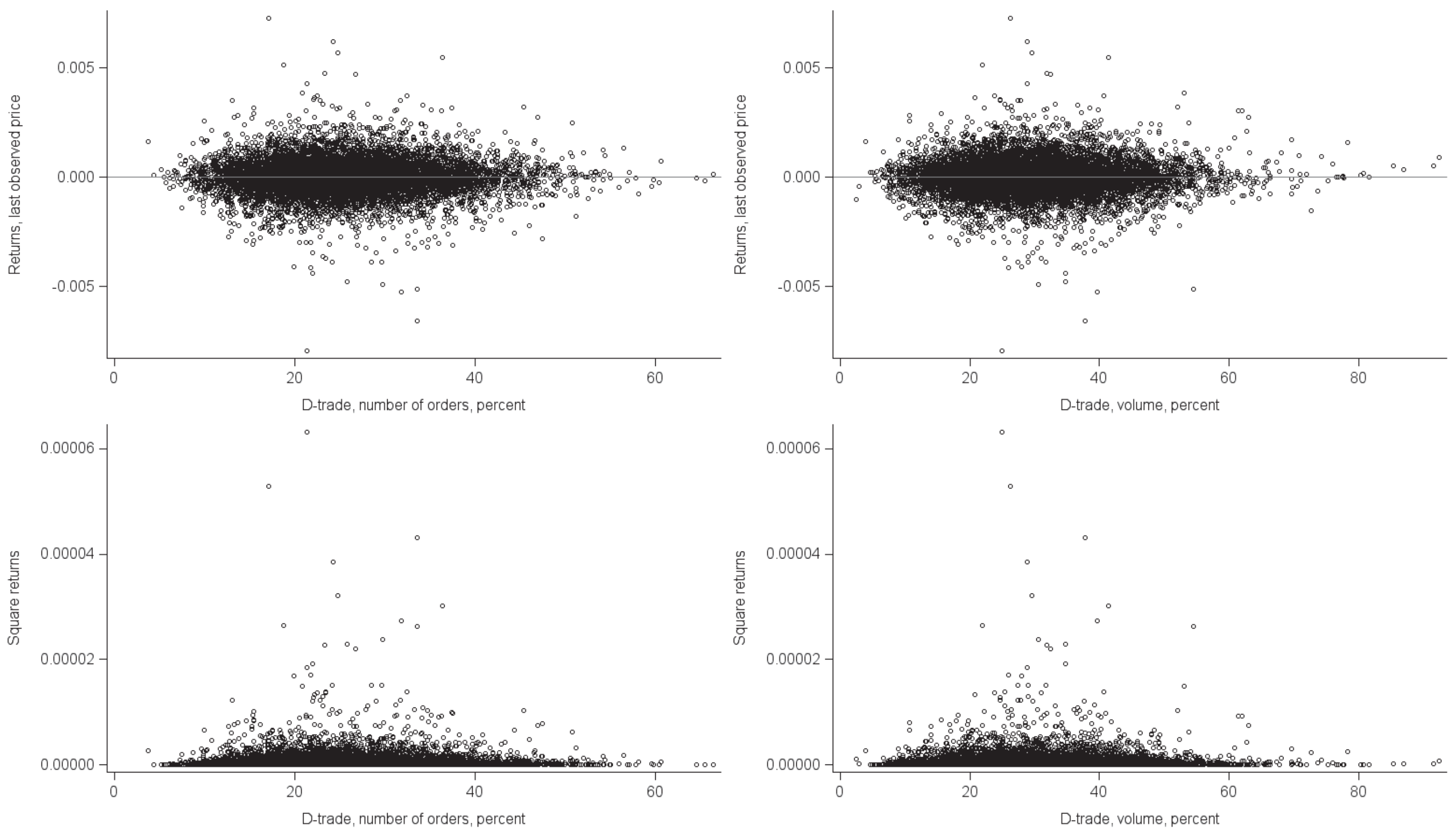

Figure 6: **3D Distribution of $|r_t|, N_t^D$ and V_t^D (General Electric)**

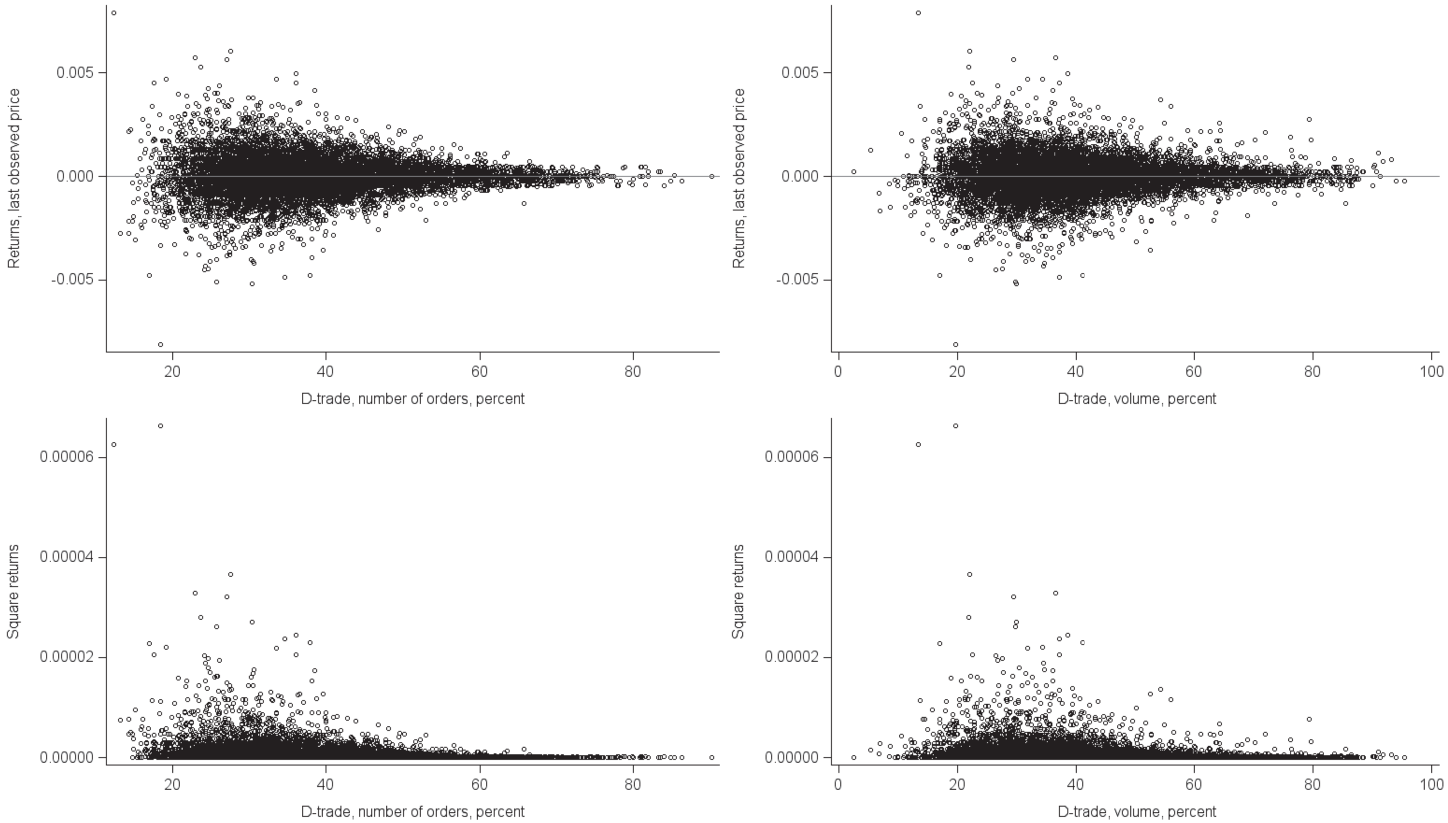

Moreover, to capture excess kurtosis in intraday returns, define $L(\cdot)$ as the Generalised Error Distribution (GED) law with ν degrees of freedom (Nelson, 1991). We shall refer to the above model as the AR(1)-GARCHX(1,1)-GED model (X_t standing for included external explanatory variables), which is reduced to the AR(1)-GARCH(1,1)-GED model if $\varphi = 0$.

To analyse the informational contents of dark-trades on the conditional variance (h_t) and then on returns, we focus on the predictive accuracy of competing AR(1)-GARCHX(1,1)-GED models. Where each one is differing by the variables included in the matrix X_t. Especially, we focus on pairwise comparisons based on the accuracy of out-of-sample one-step-ahead density forecasts.

The use of density forecasts for comparing both nested and non-nested models is popular in finance and economics (e.g.: Tay and Wallis, 2002). This approach bears several interesting features. First, the comparison takes place over either the full or a particular region of the one-step-ahead density forecast. In particular, distribution center and tails. Such an approach enables to investigate on the significance of dark trading activity in predicting extreme future price realisations. For practitioners and regulators, it may be bear important information on the measure of value-at-risk. Second, the competing models are allowed to be only a rough approximation of the true underlying data generating process. In other words, a certain degree of misspecification is allowed. Third, tests are designed to deal with heterogeneous data. Fourth, for two nested models, the suggested approach allows to analyse the marginal influence of a given exogenous explanatory variable in terms of predictive content. Thus, providing information that is different from the one returned by the standard Student t-statistics

Following Amisano and Giacomini (2007), define $Z_t = (r_t, X'_t)'$ and let $\mathscr{F}_t = \sigma(Z_1, Z_2, .., Z_t)$ be the information set at time t. suppose we have two competing AR(1)-GARCHX(1,1)-GED models, say $f_t(Z_1, Z_2, .., Z_{t-m+1} : \varphi_1)$ and

$g_t(Z_1, Z_2, .., Z_{t-p+1} : \varphi_2)$ (where φ_1 and φ_2 are parameters to be estimated) and we want to rank these models according to their out-of-sample one-step-ahead forecast accuracy. We can either analyse point forecasts (e.g.: Clark and McCracken, 2009) or density forecasts. Since the latter represent the complete characterisation associated with the one-step-ahead forecast, it contains all the relevant information. Furthermore, let $d_t^f(\cdot)$ and $d_t^g(\cdot)$ be the two out-of-sample one-step-ahead density forecasts and let $\ln d_t^f(r_{t+1})$ and $\ln d_t^g(r_{t+1})$ be the two log-scores evaluated at the outcome r_{t+1}. Amisano and Giacomini (2007) suggest a test based on a loss function that uses these logarithmic scoring rules.

Define $\lambda \in (max(m,p), \frac{T-1}{T})$. Using a rolling scheme, one can estimate the two models on the time period $1 : t = int(\lambda T)$. Then, produce density forecasts and re-estimate the model on $1 : t = int(\lambda T) + 1$. This procedures is repeatedly carried on, which yields two sets of n log-scores: $\{\ln d_t^f(r_{t+1})\}_{t=int(\lambda T)}^{T-1}$ and $\{\ln d_t^g(r_{t+1})\}_{t=int(\lambda T)}^{T-1}$. Note that by using this scheme, we allow the models to capture structural changes in the parameters as well as in the kurtosis of the returns.

To test for null of equality of density forecasts, the following statistic is used:

$$t_n = \frac{\overline{WLR}_{\lambda T,n}}{\hat{\sigma}/\sqrt(n)} \tag{2}$$

Where:

- $\overline{WLR}_{\lambda T,n} = n^{-1} \sum_{t=int(\lambda T)}^{T-1} WLR_{\lambda T,t+1}$,
- $WLR_{\lambda T,t+1} = \omega(r_{t+1}^{st}) \Big(\ln d_t^f(r_{t+1}) - \ln d_t^g(r_{t+1}) \Big)$,
- $\hat{\sigma}$ is an heteroskedastic and autocorrelation consistent (HAC) estimator of the standard error of $WLR_{\lambda T,t+1}$ over the n considered periods.
- $\omega(r_{t+1}^{st})$ is a weighting function discussed below.
- r_{t+1}^{st} is the observed standardised returns defined as $r_{t+1}^{st} = \frac{r_{t+1} - \hat{\mu}_n}{\hat{\sigma}_n}$, where

$\hat{\mu}_n$ and $\hat{\sigma}_n$ are the unconditional mean and standard error of the n realisations of r_{t+1}.

Such a test is known as the weighted likelihood ratio test (WLR). Under the null hypothesis, t_n is distributed as a standard Normal deviate with unit variance. Notice that a large and significant positive value for t_n leads to choose $f(\cdot)$ over $g(\cdot)$. While a negative value of t_n will lead to choose $g(\cdot)$ over $f(\cdot)$. The weighting function $\omega(\cdot)$ is used to set to highlight a particular region of the density forecast. If $\omega(\cdot)$ is uniform, i.e. taking the value of one whatever r^{st}_{t+1} is, then the test highlights the entire distribution. Four other definitions of $\omega(\cdot)$ are of interest for any random variable 'y' with zero mean and unit variance:

- Center of the distribution: $\omega(y) = \phi(y)$, where $\phi(\cdot)$ is the standard normal density function.
- Tails of distribution: $\omega(y) = 1 - \frac{\phi(y)}{\phi(0)}$.
- Right tail of the distribution: $\omega(y) = \Phi(y)$, where $\Phi(\cdot)$ is the cumulative standard normal density function.
- Left tail of the distribution: $\omega(y) = 1 - \Phi(y)$.

3. Empirical results and discussion

We implement the WLR test on a given stock five minute returns. As previously mentioned, two proxies of dark trading are used in various competing models, proportion of dark-volume (V_t^D) and transactions (N_t^D). Figures 7-9 plot the two proxies for each stock, together with their trends estimated using a spline function. Clearly, for all stocks, the two series exhibit similar trends but with different volatilities.

Table 1 reports the estimated parameters of the AR(1)-GARCHX(1,1)-GED model with $\varphi = 0$ (the benchmark model), which is estimated by implementing

Figure 7: **Time evolution of N_t^D and V_t^D (Microsoft)**

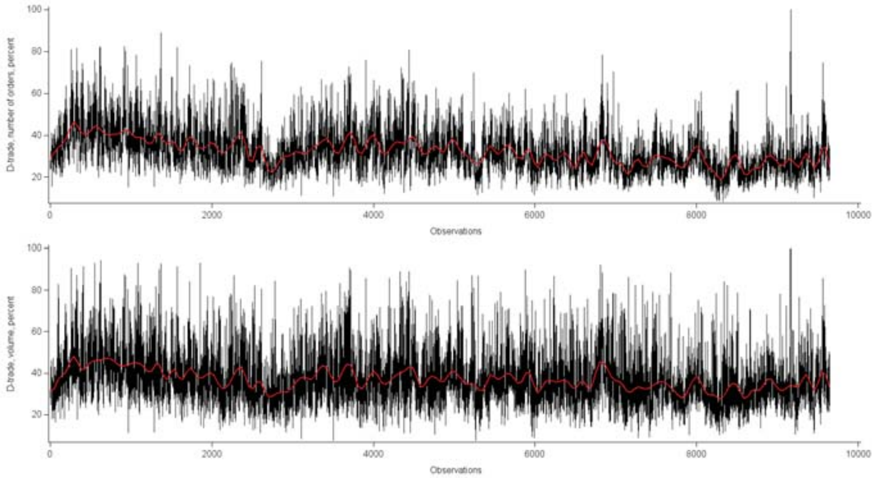

the Full Information Maximum Likelihood (FIML) framework. The autocorrelation coefficient in the mean equation is significant, and the low degree of freedom for the GED law leads to reject the normality assumption ($\nu = 2$) in favour of a fat tailed distribution. Moreover, there are no evidences for autocorrelation (Q-stat) or heteroskedasticity (ARCH-LM) (results not reported here).

We next turn to pair-wise comparisons. Table 2 summarises the five competing models used in this study. The reference model with $\varphi = 0$ is designated by M_0, whereas models M_1 to M_4 all include various proxies of dark trading. Tables 3-7 indicates the results of WLR tests. Main entries are the t_n statistics and the p-values (in the parentheses). A significant positive value for t_n indicates that model M_i (row) is to be preferred to M_j (column) and conversely. Clearly, three kinds of information are of interest:

i The information content dark trading, relative to a simple AR(1)-GARCH(1,1) model.

ii The relative information contribution to future returns of the two proxies, i.e.: proportion of dark volume (V_t^D) versus proportion of trades (N_t^D).

iii Linear versus non-linear effects of dark trading.

Figure 8: **Time evolution of N_t^D and V_t^D (IBM)**

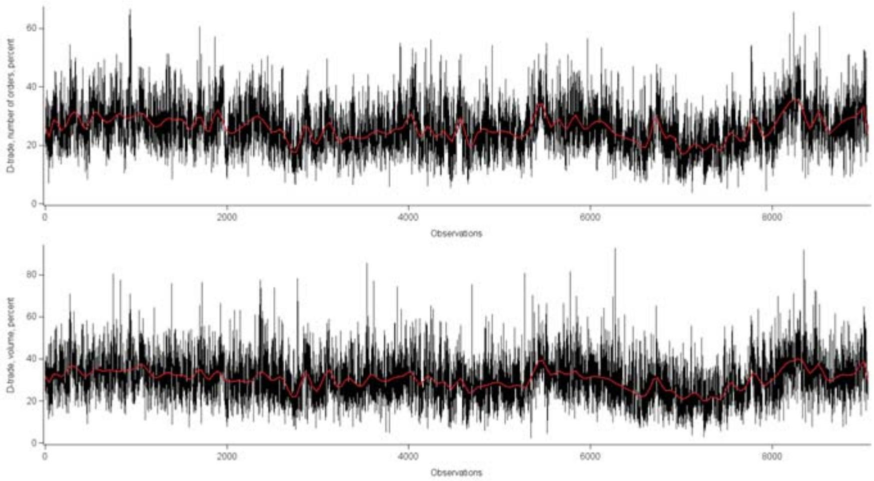

Figure 9: **Time evolution of N_t^D and V_t^D (General Electric)**

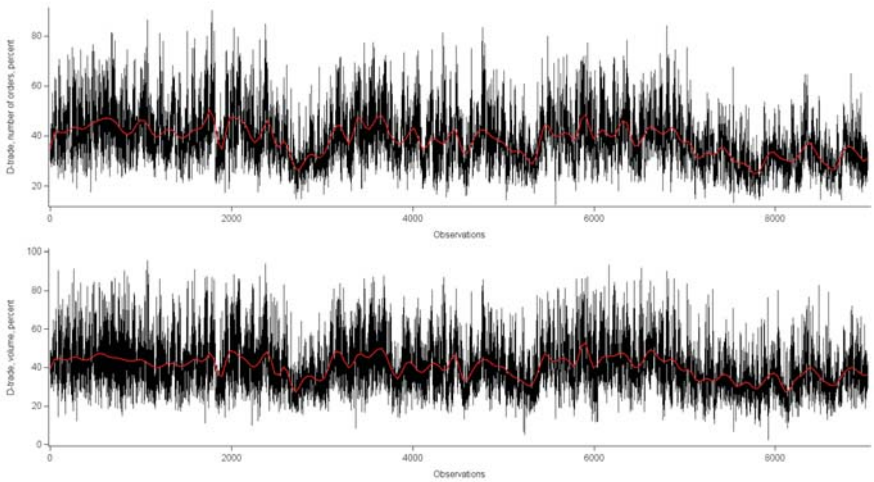

Table 1: **FIML Parameter estimates of the AR(1)-GARCHX(1,1)-GED (Benchmark Model, $\varphi = 0$)**

	Microsoft				**IBM**				**General Electric**			
	Estimate	**Std. Err.**	**t-stat**	**p-value**	**Estimate**	**Std. Err.**	**t-stat**	**p-value**	**Estimate**	**Std. Err.**	**t-stat**	**p-value**
ρ	-0.056	0.014	-4.03	$< .0001$	-0.034	0.012	-2.92	0.0035	-0.059	0.012	-4.85	$< .0001$
α	11×10^{-4}	8.1×10^{-6}	1.33	0.1823	15×10^{-4}	7×10^{-6}	2.17	0.03	12×10^{-4}	7×10^{-6}	1.72	0.0863
α_0	9.2×10^{-8}	1.1×10^{-8}	8.57	$< .0001$	4.9×10^{-8}	551×10^{-9}	8.83	$< .0001$	5.7×10^{-8}	7.3×10^{-9}	7.82	$< .0001$
α_1	0.179	0.014	12.46	$< .0001$	0.143	0.012	12.44	$< .0001$	0.131	0.011	12.02	$< .0001$
β	0.752	0.018	40.83	$< .0001$	0.789	0.015	51.60	$< .0001$	0.811	0.015	54.46	$< .0001$
ν	1.144	0.021	55.04	$< .0001$	1.269	0.023	54.24	$< .0001$	1.209	0.025	47.84	$< .0001$

Table 2: **Estimated Models**

Model	*Mean Equation*	*GARCH(p,q)*	X_t	*Dist. of errors*
M_0			None	
M_1			V_t^D	
M_2	$AR(1)$	$p = 1, q = 1$	$V_t^D, \left(V_t^D\right)^2$	GED
M_3			N_t^D	
M_4			$N_t^D, \left(B_t^D\right)^2$	

3.1. The information contribution of dark trading

Results are reported in tables 3-7. Overall, they highlight the contributed information of dark pool trading activity when included in the conditional variance equation of the GARCH process. In the case of Microsoft and General Electric, this is especially true when including linear and non-linear information on dark pool activity. This results is indicated in the second and fourth columns of the respective matrices in tables 3-7. Hence, in the case of these two stocks, it is important to include non-linear (square) terms on dark pool trading activity (in terms volume and number of transactions) in order to gain predictive information over the benchmark model. In the case of IBM, however, only the proportion of volume traded on dark pools (V_t^D) adds information to the forecast density (i.e.: M_1 over-performs M_0). While conversely, the benchmark model over-performs the model where the proportion of volume traded on dark pools is also included non-linearly (i.e.: M_0 over-performs M_2). Moreover, in the case of IBM, there is no significant difference between the benchmark model and the models where information on the proportion of transactions on dark pools (N_t^D) is included (M_3 and M_4). At last, we note that in the case of Microsoft the benchmark model under-performs regardless of the information on dark pool trading activity that is included in the GARCH variance equation.

These results suggest that data on dark pool trading activity adds information to log-return forecast density.However, tables 3-7 also suggest that it also important to correctly specify the adequate model since the indicated results in these tables varies from one model specification to another. Furthermore, for the three chosen stocks, table 3-7 indicate that the obtained results are persistent regardless of which region of one-step-ahead log-return forecast density is highlighted. This provides an evidence that information on dark pool trading activity can be used to improve risk measures such as the value at risk (in the case where the left tail of forecast density is highlighted). Provided the correct model is specified.

3.2. Proportion of volume traded vs. proportion of transactions in dark pools and linear vs. non-linear effects

We recall that M_1 and M_2 designate models where the proportion of volume traded in dark pools is included in the variance equation. While M_3 and M_4 designate models where the proportion of transactions made in dark pools is included. Overall, tables 3-7 indicate that the two variables are distinct in terms of their informational content. However, one is superior to the other depending on the stock that is being studied. Furthermore, in the case of Microsoft and General Electric, including a non-linear term (either the proportion of volume or transactions made in dark pools) in the variance equation contributes additional information to future realisations. However, this is not the case for IBM.

More precisely, with respect to Microsoft and General Electric it is clear that the proportion of transactions made in dark pools bears additional information relative to the volume traded in dark pools. Though, in the case of General Electric, a non-linear (square) term has to be included (i.e. M_4 over-performs the other tested models). This result is indicated for Microsoft in the third and fourth columns of tables 3-7, while for General Electric it is indicated in the fourth column only. Furthermore, as previously mentioned, our results remain consistent regardless of the highlighted region of the one-step-ahead forecast density. With respect to IBM, the first model (M_1) over-performs all other models. Indicating that in this case, the proportion of volume traded in dark pools has additional information relative to the other variables as long as in is only linearly included

Table 3: **Weighted Likelihood Ratio tests (Entire distribution)**

Microsoft

	M_1	M_2	M_3	M_4
M_0	-5.9 (< 0.001)	-5.0 (< 0.001)	-8.7 (< 0.001)	-8.8 (< 0.001)
M_1		-4.02 (0.004)	-8.1 (0.865)	-8.6 (< 0.001)
M_2			0.9 (0.366)	-8.6 (< 0.001)
M_3				-7.6 (< 0.001)

IBM

	M_1	M_2	M_3	M_4
M_0	**-3.1 (0.002)**	3.5 (< 0.001)	-0.3 (0.7341)	1.7 (0.0964)
M_1		8.7 (< 0.001)	-6.6 (< 0.001)	10.3 (< 0.001)
M_2			-3.6 (< 0.001)	-1.7 (0.0867)
M_3				6.9 (< 0.001)

General Electric

	M_1	M_2	M_3	M_4
M_0	3.8 (< 0.001)	-1.8 (0.0721)	5.7 (< 0.001)	**-5.4** (< 0.001)
M_1		-7.9 (< 0.001)	5.2 (< 0.001)	-10.5 (< 0.001)
M_2			9.5 (< 0.001)	-5.4 (< 0.001)
M_3				-16.5 (< 0.001)

Table 4: **Weighted Likelihood Ratio tests (Center of distribution)**

Microsoft

	M_1	M_2	M_3	M_4
M_0	-6.7 (< 0.001)	-5.8 (< 0.001)	-11.4 (< 0.001)	-10.5 (< 0.001)
M_1		-4.5 (< 0.001)	-9.5 (< 0.001)	-10.4 (< 0.001)
M_2			0.2 (0.818)	-9.8 (< 0.001)
M_3				-8.9 (< 0.001)

IBM

	M_1	M_2	M_3	M_4
M_0	-3.3 (< 0.001)	4.2 (< 0.001)	-0.3 (0.764)	1.6 (0.106)
M_1		10.0 (< 0.001)	8.6 (< 0.001)	11.9 (< 0.001)
M_2			-4.6 (< 0.001)	-3.0 (0.003)
M_3				7.0 (< 0.001)

General Electric

	M_1	M_2	M_3	M_4
M_0	3.8 (< 0.001)	-0.4 (0.671)	5.3 (< 0.001)	-2.6 (0.009)
M_1		-4.8 (< 0.001)	3.3 (0.001)	-6.5 (< 0.001)
M_2			6.0 (< 0.001)	-3.3 (< 0.001)
M_3				-10.0 (< 0.001)

Table 5: **Weighted Likelihood Ratio tests (Distribution left and right tails)**

Microsoft

	M_1	M_2	M_3	M_4
M_0	-3.1 (0.002)	-3.3 (0.001)	-4.9 (< 0.001)	-5.7 (< 0.001)
M_1		-2.8 (0.004)	-5.3 (< 0.001)	-5.7 (< 0.001)
M_2			1.1 (0.271)	-6.1 (< 0.001)
M_3				-5.3 (< 0.001)

IBM

	M_1	M_2	M_3	M_4
M_0	-2.3 (0.024)	2.1 (0.039)	-0.3 (0.769)	1.3 (0.177)
M_1		5.6 (< 0.001)	4.2 (< 0.001)	7.0 (< 0.001)
M_2			-2.1 (0.033)	-0.4 (0.703)
M_3				5.3 (< 0.001)

General Electric

	M_1	M_2	M_3	M_4
M_0	2.9 (0.003)	-2.1 (0.033)	4.6 (< 0.001)	-5.7 (< 0.001)
M_1		-7.8 (< 0.001)	4.8 (< 0.001)	-10.1 (< 0.001)
M_2			9.2 (< 0.001)	-5.3 (< 0.001)
M_3				-15.9 (< 0.001)

Table 6: **Weighted Likelihood Ratio tests (Right tail)**

Microsoft

	M_1	M_2	M_3	M_4
M_0	-3.4 (< 0.001)	-2.9 (0.003)	-6.4 (< 0.001)	-5.3 (< 0.001)
M_1		-2.2 (0.029)	-6.8 (< 0.001)	-5.1 (< 0.001)
M_2			-0.4 (0.650)	-5.4 (< 0.001)
M_3				-4.2 (< 0.001)

IBM

	M_1	M_2	M_3	M_4
M_0	-2.3 (0.017)	2.4 (0.014)	-0.2 (0.864)	1.3 (0.179)
M_1		6.8 (< 0.001)	5.0 (< 0.001)	7.8 (< 0.001)
M_2			-2.6 (0.010)	-1.1 (0.277)
M_3				5.2 (< 0.001)

General Electric

	M_1	M_2	M_3	M_4
M_0	2.5 (0.013)	-1.4 (0.155)	3.9 (< 0.001)	-4.8 (< 0.001)
M_1		-5.2 (< 0.001)	3.3 (< 0.001)	-7.7 (< 0.001)
M_2			5.9 (< 0.001)	-4.1 (< 0.001)
M_3				-12.4 (< 0.001)

Table 7: **Weighted Likelihood Ratio tests (Left tail)**

Microsoft

	M_1	M_2	M_3	M_4
M_0	-4.5 (< 0.001)	-4.6 (< 0.001)	-6.8 (< 0.001)	-7.8 (< 0.001)
M_1		-3.8 (< 0.001)	-5.8 (< 0.001)	-7.6 (< 0.001)
M_2			1.7 (0.093)	-7.8 (< 0.001)
M_3				-7.1 (< 0.001)

IBM

	M_1	M_2	M_3	M_4
M_0	-2.4 (0.016)	3.0 (0.002)	-0.3 (0.731)	1.2 (0.222)
M_1		6.6 (< 0.001)	5.0 (< 0.001)	7.6 (< 0.001)
M_2			-3.0 (0.003)	-1.5 (0.123)
M_3				5.2 (< 0.001)

General Electric

	M_1	M_2	M_3	M_4
M_0	3.2 (0.001)	-1.3 (0.195)	4.5 (< 0.001)	-3.4 (< 0.001)
M_1		-7.0 (< 0.001)	4.5 (< 0.001)	-7.8 (< 0.001)
M_2			8.9 (< 0.001)	-4.1 (< 0.001)
M_3				-11.9 (< 0.001)

in the variance equation. As with Microsoft and General Electric, this result holds regardless of the highlighted region of the one-step-ahead forecast density.

4. Conclusion

In this paper we have applied the GARCH estimation framework to a problem of market microstructure. More precisely, we have attempted to answer whether the activity of trading in dark pools conveys any information on the intraday return and volatility process. Our results indicate that indeed this trading activity conveys relevant information to the process determining the one-step-ahead returns. Moreover, not only it conveys information over the one-step-ahead return forecast, it also conveys important information on the entire density forecast of returns. This, with a special emphasis on the tails of this density forecast. Hence, we conclude that the trading activity in dark pools has an important role in determining intraday returns and the uncertainty that may relate to them. Hence, having important microstructural implications

Nevertheless, we emphasize that a the correct model needs to be identified in order to use the information that is conveyed by dark pool trading activity. In the case of Microsoft and General Electric, our results indicate that the proportion of transactions made in dark-polls provides more information regarding the (one-step-ahead) point and density forecast of returns. While in the case of IBM, it is the proportion of volume traded in dark pools that contributes more information. This regardless of which region of the forecast density is being highlighted. Moreover, in the case of Microsoft and General Electric we find an important non-linear effect on the conditional volatility process. While this is not the case for IBM. Nevertheless, with respect to the three stocks, we conclude that trading that takes place in dark pools bear important information. Though we did not discuss the issue of price discovery, it is obvious that dark trading has a role in the price discovery process. From our results, it seems that it may contribute to the price discovery process for the three studied stocks.

Given highlighted results, knowledge of trading in dark pools may provide valuable information to regulators and market participants alike. For regulators, dark trading might provide information over the effects of high-frequency trading, provided that dark trading activity coincides with the latter activity. Therefore, an important issue for further research is to empirically determine how trading in the dark coincides with high frequency trading. Determining this relationship may provide an important piece of information for regulators in the activity of overseeing financial markets. Another important outcome that is indicated in our results, is that dark-trading seems to be well integrated in current trading activity. Furthermore, as mentioned already, it seems that traders on public exchanges react to dark trading once it is exposed to the public.

Besides determining the relationship between high-frequency-traders (or more generally, informed trading) and trading in dark pools, further research will require to study different aspects of the intraday trading activity such as trading liquidity, volume duration and so on. Furthermore, it may be also interesting to consider if trading in dark pools coincides also with different corporate events such as stock repurchases and issuance or even earning announcements.

References

[1] Gianni Amisano and Raffaella Giacomini. Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics*, 25(2):177–190, 2007.

[2] Hendrik Bessembinder, Marios Panayides, and Kumar Venkataraman. Hidden liquidity: an analysis of order exposure strategies in electronic stock markets. *Journal of Financial Economics*, 94(3):361–383, 2009.

[3] Leslie Boni, David C Brown, and J Chris Leach. Dark pool exclusivity matters. *Available at SSRN*, 2055808, 2012.

[4] Sabrina Buti and Barbara Rindi. Undisclosed orders and optimal submission strategies in a limit order market. *Journal of Financial Economics*, 2013.

[5] Sabrina Buti, Barbara Rindi, and Ingrid Werner. Diving into dark pools. *Charles A. Dice Center Working Paper*, (2010-10), 2011.

[6] Todd E Clark and Michael W McCracken. Tests of equal predictive ability with real-time data. *Journal of Business & Economic Statistics*, 27(4): 441–454, 2009.

[7] Peter Kratz and Torsten Schöneborn. Optimal liquidation in dark pools. In *EFA 2009 Bergen Meetings Paper*, 2012.

[8] Daniel B Nelson. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, pages 347–370, 1991.

[9] Mahendrarajah Nimalendran and Sugata Ray. Informed trading in dark pools. *SSRN eLibrary*, 2011.

[10] Mahendrarajah Nimalendran and Sugata Ray. Informational linkages between dark and lit trading venues. *Journal of Financial Markets*, 2013.

[11] Maureen O'Hara and Mao Ye. Is market fragmentation harming market quality? *Journal of Financial Economics*, 100(3):459–474, 2011.

[12] Mark J Ready. Determinants of volume in dark pools. *Preprint*, 2010.

[13] Anthony S Tay and Kenneth F Wallis. Density forecasting: a survey. *Companion to Economic Forecasting*, pages 45–68, 2002.

[14] Daniel Weaver. Internalization and market quality in a fragmented market structure. *Available at SSRN 1846470*, 2011.

[15] Mao Ye. A glimpse into the dark: Price formation, transaction cost and market share of the crossing network. *Transaction Cost and Market Share of the Crossing Network (June 9, 2011)*, 2010.

[16] Haoxiang Zhu. Do dark pools harm price discovery? *Available at SSRN 1712173*, 2012.