Who is crowding whom in the European carbon market?

Boulis Maher Ibrahim
Accounting, Economics and Finance, Heriot Watt University, SEEC, CFI, Edinburgh, UK

Iordanis Angelos Kalaitzoglou
Audencia Nantes School of Management, CFRM, ReMA, Nantes, France

Abstract

This paper investigates whether the upstairs and downstairs markets in the two most active European carbon-trading venues crowd themselves or each other, and the implications this might have on price formation. We use a trade volume imbalance measure to proxy for crowdedness, and a recent asymmetric information microstructure model to extract price and risk components of information and liquidity for upstairs and downstairs markets. We then analyse these markets separately in terms of uncertainty resolution period, over-reaction and under-reaction, and information and liquidity risk premia by the relative degree of their crowdedness and by the intensity of upstairs trading activity. The results suggest that the carbon market is crowded, overreacts and underreacts more with higher crowdedness, and that the upstairs market prices information and liquidity differently, and resolves information faster, than the downstairs market. However, a negative correlation between their relative degree of crowdedness suggests that these two markets are complementary.

JEL classifications: C41; C58;G15
Keywords: Carbon market; Microstructure; Duration model; Ultra-high-frequency data

1. Introduction.

In his 2009 presidential address of the American Finance Association, Stein (2009) argues that although arbitrage money may eliminate arbitrage profits, an average narrowing of the gap between prices and fundamentals or a reduction in nonfundamental volatility may not necessary ensue. He introduces a possible 'crowded-trade effect,' which is a misalignment of prices to fundamental values due to the coordination problem arising from the inability of traders to condition their behaviour on current market-wide arbitrage capacity. He also mentions leverage as a possible mechanism through which arbitrageurs may inflict externalities on one another. With regard to the crowded trade issue, Stein assumes that each arbitrageur makes an unbiased estimate of the number of others that are active in the market at any point in time and trades accordingly. However, errors in this estimate can cause crowdedness. His analysis suggests an "ambiguous answer": "Because of the externalities associated with crowding and leverage, there is no clear theoretical presumption that – absent any policy intervention - the rise of sophisticated investors should necessarily be beneficial to market efficiency." If we assume that arbitrageurs with larger capital capacity for trading implement their strategies predominantly in the upstairs (Over-the-counter) market, the question that naturally arises, and is of interest to us in this paper, is whether these arbitrageurs introduce externalities to themselves in upstairs trading or to the downstairs (screen) market, or to both markets. Does the upstairs OTC market crowd the downstairs (screen) market? And if so, what are the implications on risk premia?

Within the context of carbon trading, this paper investigates whether the upstairs (OTC) market crowds itself, the downstairs (screen) market, or both, and whether this reflects in different pricing of information versus liquidity risk for OTC trades versus screen trades. This links to the usual wisdom, as mentioned by Stein (2009), that if one adopts the view that individuals are naïve investors, most probably trading on screen, and institutions are rational arbitrageurs, most probably trading over-the-counter, then the rise of hedge funds and institutional holdings (in equity markets) would suggest that smart-money players trade intensively with one another, with the dumb money playing a much-diminished role. Accordingly, one would expect the upstairs market to profit from the downstairs market if they are more informed, or are more 'astute' (or have higher capacity) towards information acquisition, arbitrage identification, and risk pricing. There is also an added liquidity dimension to these arguments. If the upstairs market trades large quantities then this may have consequences on liquidity issues that could affect either or both markets (upstairs and downstairs). This, for example, can occur in shallow markets if large traders are able to predict, at least partially, the liquidity perturbations that their trading is likely to cause. If such cases arise, would the premium on liquidity risk vary more, or be different, than that on information risk? Would the resulting variations in liquidity overwhelm information considerations, especially in a cap-and-trade system such as that of the European carbon market? Further, would these issues be different between the upstairs market and the downstairs market? These questions motive this study.

With an eye on the wide changing implications on international carbon pricing and trading that have emanated from the Paris Agreement of December 2015, and the resulting recent sprouting of numerous carbon trading and tax mechanisms around the globe, the European Union Emission Trading Scheme (EU ETS) remains to be the largest market for trading carbon. This market, of carbon credits and European emission allowance (EUA) futures, has passed through two phases of development and currently is in its third phase of trading. Phase I is a pilot period (2005-2007) modelled according to the cap-and-trade sulphur dioxide (acid rain) experience of the United States, Phase II is a more intense trading period that coincided with the Koyoto commitment period (2008-2012) where installations have to match emissions to allowances, and the current Phase III (2013-2020) is a regulatory enhanced evaluation period. The total supply of emission permits is capped at progressively decreasing levels determined by climate change politics aimed at emission abatement targets relative to pre-1990 levels. Emitting installations within EU member states (initially allocated a number of allowances and, more recently, partially purchased them in auctions) must match their annual emissions to permits by end of April the following year. Intermittent surpluses and deficits can be sold or purchased in a number of overlapping trading venues. The largest two are the European Climate Exchange (ECX), trading through the Intercontinental Exchange (ICE) platform in London, and Nord Pool (NP) trading through OMX, though the latter is much smaller. We focus on these venues in our analysis, which we conduct at a microstructural intraday level.

Prior research on the microstructure of the EU ETS sets the background to our intended analysis. Benz and Hengelbrock (2008) present evidence that EUA futures trading in ECX leads that in NP, and Mizrach and Otsubo (2014) show that ECX leads price discovery. Rittler (2012) finds that the futures market predominantly leads the spot market and price discovery occurs first in futures. Other evidence points towards the futures price being the reference price for carbon. For example, Mizrach (2012) confirms that although spot and futures prices are co-integrated, the futures curve evolves independently. Daskalakis et al. (2009) also conclude that the cost of carry model stands only on intraphase basis. In terms of maturity, Lucia et al. (2015), for example, find that speculative activity concentrates in the maturing contract (most heavily traded). These studies provide rationale for our focus on ECX and NP, futures and forward prices, maturing contracts, and intra-phase analysis.

No prior study however, measures the degree of crowdedness in the carbon market and any possible consequential over- or under-reaction or differential effects on information and liquidity risk premia. The most related studies are Rannou and Barneto (2014), Rannou (2017), and Ibrahim and Kalaitzoglou (2016). Using Granger causality analysis, GARCH specifications and volume and volatility regressions Rannou and Barneto (2014) analyse market efficiency in Phase II while differentiating between OTC and screen trading. They consider daily frequency and note a dominance of forward OTC trading on futures screen trading, and a causality from OTC to screen volumes. They advocate that this is "mainly driven by heterogeneous investor beliefs" providing an indication of how information is dispersed and held. They further conclude that leverage volatility effects indicate that EUA daily prices overreact to bad news. Specifically, the OTC market underreacts to information flows conveyed by more lagged volume effects on futures volatility. Rannou (2017) analyses the effect of order book measures on unexpected returns, realised return volatility and the autocorrelation between improvements in bid (or ask) quotes and the location of the next trade (at the bid or the ask). His analysis is carried out at 30-minute frequency and covers Phase II of the market. He reports that order imbalance has a moderate predictive power on returns (corroborating Mizrach and Otsubo, 2014), order book slope has some predictive power on return volatility, and immediacy costs (bid-ask spread) are related to volatility and autocorrelation of trade direction. Rannou and Barneto (2014) and Rannou (2017), however, do not take into account the time dimension of trades (i.e., duration), even though Rannou (2014) include a ratio of order arrivals to departures. They also do not analyse crowdedness or its possible differential effects on risk premia, and note that the coarse daily or 30-minute frequency could mask other effects. Using Autocorrelated Duration (ACD) models, Ibrahim and Kalaitzoglou (2013a) analyse agent identification in the carbon market at transaction level (Phase I and the first year of Phase II) while taking into account the time dimension of trades. Beside agent identification, their intradaily analysis finds OTC trades to be fewer in number, larger in size and transacted in a timed manner with a preference towards the end of the day in ECX (they analyse only one contract and in ECX only). They further find that there is a higher proportion of informed trades amongst OTC transactions than amongst screen trades. Although not mentioned in their paper, this points towards a difference in information acquisition and impounding into prices between the OTC market and the screen market. These studies motivate our closer investigation between these two markets at transaction level and in event time, rather than calendar time.

In this paper we analyse, at transaction level, whether price discovery is affected by OTC trades differently than by screen trades. We differentiate the contribution of these trades in their capacity to resolve uncertainty and to their respective pricing of liquidity and information risk. We do so by analysing whether the observed intradaily and intra-phase trading patterns of OTC trades, which differ from those of screen trades, crowd the market. We then investigate whether crowdedness, if it exists, is confined to the upstairs market or is inflicted on the downstairs market. Further, we investigate whether the upstairs market prices information and liquidity differently than the downstairs market, and whether crowdedness causes these two markets to over-react or under-react. This is carried out by separating the price of information from that of liquidity by market type, at different levels of crowdedness emanating from either market, and by the intensity and presence of OTC trades. The results suggest that the EU carbon market is crowded. In the first two phases, the upstairs market is more crowded than the downstairs market, while in the third phase, the reverse holds. Information is priced differently than liquidity, and these prices vary by level of crowdedness. This causes over-reaction and under-reaction detected in the autocorrelations of both returns and trade direction. However, crowdedness in the upstairs market is negatively correlated to that in the downstairs market, and the magnitude of this negative correlation increases as OTC trading intensifies. This suggests that both markets are complementary, even if uncertainty is resolved faster when trading upstairs is intense.

The remainder of this paper is organised as follows: Section 2 describes the microstructure methodology used to model price change, duration and volume of trades, and the measures of crowdedness, OTC intensity, illiquidity per unit of volume, and information and liquidity risk premia; Section 3 describes the data used and sample selection; Section 4 presents the analysis and the results; and Section 5 summarises and concludes.

2. Methodology

2.1 Autoregressive Conditional Weighted Duration (ACWD*)*

This study employs an enhanced specification of an ACWD model called the Smooth Transition Mixture of Weibull ACWD (STM-ACWD) (see Kalaitzoglou and Ibrahim, 2015). This specification provides a richer measure of intraday liquidity by incorporating volume into duration analysis and allows for smooth transition between states of weighted duration with the errors within each state following a different Weibull distribution.[1]

First, let $d_i = t_i - t_{i-1}$ be the (raw) duration of event/transaction i, measured by the time (t) elapsed since the preceding event at i-1, and x_i is the diurnally adjusted duration. Let v_i be the number of contracts (i.e., volume), associated with each event/transaction, with mean $\bar{v}$ and variance σ_v^2. The diurnally adjusted duration is then transformed into weighted duration, S_i, using the scaling factor $K(u_i) = exp\left(\frac{-u_i}{2}\right)$, where $u_i = (v_i - \bar{v})/\sigma_v$. The new variable, $S_i = x_i * K(u_i)$, is an explicit measure of 'trading intensity,' defined as the waiting time for a single contract to be traded. Trading intensity S_i is then modelled by the following STM-ACWD specification:

$$S_i = \Theta_i \varepsilon_i \tag{1}$$

[1] Extending on previous literature (Bredin et al, 2014; Kalaitzoglou and Ibrahim, 2015) transaction size is introduced to better identify a high trading intensity regime. This is a significant enhancement compared to previous studies (e.g., Kalaitoglou and Ibrahim, 2013b) that focus solely on duration. If transaction size is ignored, then OTC trades affect duration only and therefore, the estimate of the threshold value that distinguishes higher from lower trading intensity would be biased, because it does not consider the distinctively larger size of OTC trades. In contrast, this paper shifts the focus on a richer measure of liquidity, i.e. the arrival rate of contracts, which enables OTC trades to have an impact on both duration and size.

$$\Theta_i = E(S_i|F_{i-1};\varphi) = \Theta(S_{i-1},\ldots,S_1;\varphi_1) \quad (2)$$

$$\varepsilon_i|J_i \sim i.i.d., with\ density\ f(\varepsilon_i|J_i;\varphi_2) \text{ and } E(\varepsilon_i|J_i;\varphi_2) = E(\varepsilon_i) = 1 \quad (3)$$

where Θ_i is the conditional expected trading intensity; F is the econometrician's information set, and $\varepsilon_i = S_i/\Theta_i$ is the standardized trading intensity. Θ_i is modelled using the following linear specification:

$$\Theta_i = \omega + \sum_{j=1}^{m}(a_j + \zeta_j * I_{i-j})S_{i-j} + \sum_{p=1}^{q}\beta_p\Theta_{i-p} \quad (4)$$

where ζ_j (zeta) is a parameter that captures a potentially different effect of OTC transactions indicated by the dummy I, which takes the value of 1 for OTC transactions and zero otherwise. OTC transactions are allowed to revise the coefficient a_j of the past values of trading intensity, which in this case represents the impact of non-OTC trades. Zeta is revised when the transaction at *i-j* is an OTC trade. Positive values, ζ_j>0, indicate longer expected duration and/or lower expected transaction size. When ζ_j=0, the specification in Equation (4) (dubbed ACWD-OTC) collapses to the ARMA specification of ACWD.

The conditional density function, *f*, is assumed to be a smooth transition mixture of Weibull distributions in three regimes of trading intensity:

$$f(S_i|J_i;\tau) = (h_i(\tau)/S_i)[S_i\Gamma(1+1/h_i(\tau))/\Theta_i]^{h_i(\tau)}exp\left(-[S_i\Gamma(1+1/h_i(\tau))/\Theta_i]^{h_i(\tau)}\right) \quad (7)$$

$$h(J_i:\tau) = \gamma_1 + (\gamma_2-\gamma_1) * G_1(J_i:g_1,j_1) + (\gamma_3-\gamma_2) * G_2(J_i:g_2,j_2) \quad (8)$$

$$G_k(S_i:g_k,j_k) = (1+exp\{-g_k * (J_i - j_k)\})^{-1} \quad (9)$$

where the overall shape parameter of the Weibull distribution, $h(J_i:\boldsymbol{\tau})$, is a function of a threshold variable, J_i, represented by the first lag of S, and a vector of parameter coefficients $\boldsymbol{\tau} = (\gamma_1\ \ \gamma_2\ \ \gamma_3\ \ g_1\ \ g_2\ \ j_1\ \ j_2)$. The parameters, γ_1, γ_2, and γ_3 are the shape parameters of the Weibull distributions in the respective three regimes assumed for S, determined by the threshold variable J_i. The smoothness parameter g_k is either g_1 or g_2. For every observation/transaction the overall shape parameter, $h(J_i : \boldsymbol{\tau})$, is the weighted average of γ_1, γ_2 and γ_3, and the weights are determined by the two smooth transition functions, G_1 and G_2. Finally, Γ is the gamma function.

Consequently, the hazard function of 'volume' weighted durations is revised after every transaction, providing a measure of instantaneous liquidity and, thus, can be considered a conditional 'liquidity heat' measure. Note that the appearance of an OTC trade is allowed, through expected duration, to indirectly affect the rate at which a single contract is traded.[2] This is captured by the shape of the conditional hazard function $\lambda_i \equiv \frac{f(S_i|J_i;\tau)}{1-F(S_i|J_i;\tau)}$, where $F(S_i|J_i;\tau)$ is the conditional CDF of weighted durations. A decreasing shape indicates an acceleration of trading, an increasing shape a deceleration of trading, and a flat shape a continuation of the current level of liquidity.[3]

2.2 Market sentiment, illiquidity and risk premia

The second part of the analysis focuses on the longer-term impact of OTC trades on price discovery. Considering that liquidity has been a major issue during the early stages of the market (see Alberola et al., 2008 and Ellerman et al., 2014), both organized venues that are investigated in this study had OTC trade reporting and clearing (Rannou and Barneto, 2014) and, at least at the start of the carbon market, OTC trades were considered a pricing and a

[2] This model can be viewed as providing an empirical conditional liquidity 'heat' measure based on the instantaneous probability of trade arrival. This liquidity 'heat' measure is conditional on the trading intensity regimes and on the appearance of OTC trades. It is employed to distinguish the impact of high versus low relative liquidity (see Kalaitzoglou and Ibrahim, 2015) and/or the impact of OTC trades on liquidity, volatility and uncertainty resolution time. The Sequential Arrival Information Hypothesis (discussed in Rannou and Barneto, 2014) or the Mixture of Distribution Hypothesis (discussed in Bredin et al., 2014) do not explicitly account for the effect of time duration and the fact that trade arrival times convey information.

[3] If γ_i = 1, then the associated distribution is Exponential and the hazard function is flat, indicating random arrival of a single contract at a rate that does not change over time. If $\gamma_i < 1$, then the distribution and the hazard function have a monotonically decreasing slope indicating increased probability of a single contract to be traded closer to the realization of the last event. If $\gamma_i > 1$, then the distribution is bell-shaped and the hazard function has an upward slope. This indicates that the probability of a single contract arrival increases over time (see Kalaitzoglou and Ibrahim, 2013a).

liquidity benchmark (Benz and Hengelbrock, 2008; Mansanet-Bataller and Pardo, 2009).[4] Accordingly, relative OTC trade characteristics and impact are our primary focus.

We proceed by first developing a measure of the degree of presence and intensity of OTC trades. This can be looked upon as a 'Market Sentiment Index' (MSI_i), or a measure of intensity of the upstairs (OTC) market. Note that OTC trades tend to be large in size and, hence, may carry private information and/or have liquidity implications. The proposed MSI_i is:

$$MSI_i = \underbrace{\left(1 + \frac{\sum_{j=1}^{n} I_{i-j}}{n}\right)\left(1 + \frac{\sum_{j=1}^{n} D_{i-j}}{n}\right)} \tag{10}$$

where, D_i is a dummy variable that takes the value of 1 when λ_i is a decreasing function (high intensity trades), I is the OTC indicator dummy previously defined, and n is the number of transactions during the last 15 minutes. The term $\sum_{j=1}^{n} I_{i-j}/n$ is the proportion of OTC trades in the last 15 minutes, measuring the intensity of OTC presence (large trade and possibly private information) or upstairs market activity. The term $\sum_{j=1}^{n} D_{i-j}/n$ is the proportion of high intensity trades in the last 15 minutes, measuring the degree of prevailing market (upstairs and downstairs) liquidity and perhaps volatility. This index combines these two effects and quantifies the prevailing condition of the market with respect to the intensity (liquidity) and presence of OTC trades (possibly information). Given that OTC are large and possibly informed trades, the MSI can also be looked upon as a measure of upstairs market intensity of activity (big players demanding or supplying liquidity quickly). This provides an estimate of how "active" the market is in terms of overall liquidity and OTC activity and, hence, may also represent market sentiment. The index varies from a minimum value of 1 when there are no OTC and high trading intensity trades during the last 15 minutes, to a

[4] In NP, OTC trades are reported directly to the exchange. In ECX, two facilities, namely Exchange of Physical (EFP) and Exchange for Swap (EFS) have been introduced, and OTC trades are reported up to 30 minutes after the normal trading hours. For more detailed information please refer to "Guidance ICE Futures EFP EFS Policy" available at <https://www.theice.com/publicdocs/futures/ICE_Futures_%20EFP_EFS_Policy.pdf>.

maximum value of 4 when all trades over the last 15 minutes are both OTC and high intensity trades.

Second, and with an eye on the asset pricing literature, a measure of illiquidity premium ($IL_{i,l}$) of the market is investigated across different levels of MSI_i. We construct a version of the Amihud's (2002) illiquidity premium computed over three different intraday intervals (*l*) that follow a transaction (*i*):

$$IL_{i,l} = 100 \times \left(p_{i+n_l} - p_i\right)^2 \Big/ \sum_{j=1}^{n_l} v_{i+j} \tag{11}$$

where p_i is the transaction price, $l = (5', 15', 60')$ is the period in minutes following event/transaction i, and n_l is the number of transactions during this period. This measures how much prices change for every unit of volume during each interval. Higher values indicate that price changes are more sensitive to liquidity fluctuations, which, according to Amihud (2002), is a sign of market inefficiency.

We also analyse the price impact of OTC trades by focusing on their asymmetric information risk premium (*AIP*) and their liquidity risk premium (*LP*), separately. Each premium is measured as the component of transaction price change attributed to each type of risk, divided by the contribution of each type of risk to the total variance. This way, LP and AIP measure the monetary reward required for every unit of a specific type of risk. Specifically,

$$LP_i = \varphi_i / LIQ_i \tag{12}$$

$$AIP_i = \theta_i / AS_i \tag{13}$$

where, φ_i is the liquidity component of price change (and half the spread) and θ_i is the asymmetric information component. $AS_i = (1 - \rho^2)\theta_i^2$, where ρ is the first order autocorrelation in order flow, is the component of variance of price change attributed to asymmetric information, and $LIQ_i = 2(1-\rho)\varphi_i^2$ is the component of variance attributed to liquidity.[5]

Both price components, φ_i and θ_i, are estimated from transaction prices using a modified version of the asymmetric information microstructure pricing model (DJM) of Ibrahim and Kalaitzoglou (2016):

$$\Delta p_i = \theta_i(q_i - \rho q_{i-1}) + \Delta\varphi_i q_i + \varepsilon_i + \Delta\xi_i \tag{14}$$

where Δp_i is the change in price from transaction $i-1$ to transaction *i*; q_i is the order flow variable that takes a value of +1 if trade *i* is buyer initiated and -1 if it is seller initiated (this variable is assumed to follow a simple Markov process with ρ being the first order autocorrelation of q_i); $(q_i - \rho q_{i-1})$ is the surprise in order flow, which captures private information; θ_i is the price response to, or cost of, asymmetric information; φ_i is the price response to, or cost of, liquidity; ε_i is public information; and ξ_i is price discreteness. The price responses to information and liquidity, θ_i and φ_i, are updated after every transaction according to the revision in beliefs/expectations of the dealers and limit-order traders:

$$\theta_i = \theta_1 + \theta_2\Theta_\iota^{-1} + \theta_3 E[P_\iota | H_{\iota-1}] \tag{15}$$

$$\varphi_\iota = \varphi_1 + \varphi_2\Theta_\iota^{-1} + \varphi_3 E[P_\iota | H_{\iota-1}] \tag{16}$$

where, θ_m and φ_m, $m = 1,2,3$, are parameters to be estimated; $E[P_\iota | H_{\iota-1}]$ is the expected risk conditions at event/time ι, conditional on the information set H as of time *i*-1; and P_ι is a Poisson random variable that captures shocks in volatility and in the presence of OTC trades which, in estimation, we measure by the proportion of OTC trades during the fifteen minutes that precede and include trade *i*-1. This model provides a rich decomposition of conditional variance to its structural components: asymmetric information, AS; liquidity, LIQ;

[5] In order to avoid erratic changes in LP_i and AIP_i, modified versions are computed, where the exponential functions of θ_i and φ_i are used to ensure positive support.

interactions between information and liquidity, INT; public information, PI; and price discreteness, PD. This decomposition is:

$$E\left[\Delta p_i - E[\Delta p_i]\right]^2$$

$$= \underbrace{(1-\rho^2)E\left[{\theta_i}^2\right]}_{AS_i} + \underbrace{2(1-\rho)E[{\varphi_i}^2]}_{LIQ_i} + \underbrace{2(1-\rho^2)E[\theta_i\varphi_i]}_{Int_i} \tag{17}$$

$$+ \underbrace{\sigma_\varepsilon^2}_{PI} + \underbrace{2\sigma_\xi^2}_{PD}$$

This model is estimated with an iterative GMM procedure using an appropriate set of moment conditions (see Ibrahim and Kalaitzoglou, 2016).

3. Data Sample

3.1 Sample selection

We consider the most liquid EUA futures contracts (December maturity) in the two major organised venues of the European carbon market, namely the European Climate Exchange (ECX) and Nord Pool (NP), even though the latter is much smaller. The intra-day data employed consists of date, time stamp (in seconds), volume (# of contracts), trade sign (buy or sell), and indicators that identify OTC trades, of all transactions in the December contracts from 22 April 2005 to 30 April 2015.[6] The comparative nature of the analysis between ECX and NP requires the construction of overlapping trading periods. In order to ensure sufficient observations to conduct an intraday analysis the following three periods are covered. Phase I of the market (1/1/2005 to 31/12/2007) is represented by a continuous time series constructed from the contracts that mature in Phase I, i.e., the December 2005, 2006 and 2007 contracts.

[6] The data for ECX is provided by the exchange and includes trade direction and trade type indicators. We define as OTC trades that are block, EFP (Exchange for Physical) or EFS (Exchange for Swap). EFP, EFS and most block trades are OTC trades registered through special ICE reporting facilities. The Exchange allows reporting of these trades after trading and up to 30 minutes after the closing of the market. The data from NP is extracted from Thomson Reuters Tick History (TRTH), where no trade sign is provided. This data set consists of all trades and revisions of best bid-ask quotes. All transactions occur at the best bid or the best ask, and trade direction is assigned according to the prevailing best bid-ask prices at the time of the transaction (or prior to it). OTC trades are reported directly and indicated by specific market flags.

The viable period for analysis for this phase is 22/4/2005-3/12/2007, which covers almost the entire 'pilot period' of the market and is considered as the Phase I sample period. Phase II series is constructed similarly from December contracts that mature in Phase II, i.e., the December 2008, 2009, 2010, 2011 and 2012 contracts. The viable sample period for this phase series is 14/3/2006-14/12/2012. This covers the entire 'trade period' of Phase II. Finally, Phase III series is constructed similarly from the December 2013, 2014 and 2015 contracts and the viable sample period is 4/4/2011-30/4/2015, which accounts for the first 2 years and 4 months of Phase III.

Prices in both exchanges are quoted in Euros (€) and the minimum tick is €0.01 (or €0.05 in ECX prior to 27 March 2007). Trading is continuous from Monday through Friday from 8:00-18:00 hrs Central European Time (CET).[7] There is a 15-minutes pre-open session in both markets, during which market participants can register their interests by submitting limit orders, but no trade occurs prior to the official start of the main session.

To represent trading and prices in the three phases separately, and for comparability and consistency reasons with the earlier literature, the following treatment is applied to the data. First, volume rollover is used for the construction of the continuous series in each phase.[8] Second, some transactions are omitted either due to very thin trading or in order to match maturity. These are: all transactions before 22/4/2005 and after 3/12/2007 for contracts with maturity in Phase I; transactions before 14/3/2006 and after 14/12/2012 for contracts with maturity in Phase II; and transactions before 4/4/2011 for contracts with maturity in Phase III. Third, all transactions out of the official trading hours are excluded because the recording of

[7] Nord Pool opening hours were from 8:00:00 to 15:30 until 2/2/2009. More information https://cns.omxgroup.com/cdsPublic/viewDisclosure.action?disclosureId=405432&messageId=488565 and http://www. nasdaqomxnordic.com/tradinghours for Nord Pool and theice.com for ECX.

[8] We rollover when daily volume of the nearest contract exceeds that of the maturing contract consistently for two days. In ECX the rollover dates for Phase I are 21/11/2005 and 23/11/2006; for Phase II are 9/12/2008, 4/12/2009, 17/12/2010 and 14/12/2011; and for Phase III are 13/12 2013 and 12/12/2014. In NP the rollover dates for Phase I are 24/11/2005 and 27/11/2006; for Phase II are 28/11/2008, 27/11,2009, 29/11/2010 and 19/12/2011; and for Phase III are 16/12/2013 and 15/12/2014.

OTC trades up to 30 minutes after the closing of the market in ECX would distort the estimates of durations. Fourth, time stamps are adjusted for daylight savings and durations are calculated in seconds, which is the highest frequency in the data provided by ECX. This reveals many observations with the same time stamp (zero duration).[9] Trades that have the same time stamp, trade direction/type and same price are considered as possible strategic placements/large broken trades, and since they cause no change in price or order flow we aggregate their volume and consider them as one trade.[10] The remaining zero duration trades that show a change in either price or order flow are left as unique individual trades and their volume is not aggregated. Fifth, the focus is on 'normal' market conditions (including known market crashes) and, therefore, trades with price changes greater than 20 times the standard deviation are considered as outliers and are omitted. Finally, to deal with intraday seasonality, durations are diurnally adjusted as in Engle (2000) (see technical appendix).

3.2 Data descriptive statistics

Table 1 presents descriptive statistics of the duration, volume, price, and returns (change in price) of the three phase series in each venue (ECX and NP). These are presented for 'All' and 'OTC' trades separately. In terms of general liquidity, the results show that ECX is far more liquid, with substantially higher number of trades and lower duration than NP, confirming ECX as the main trading venue for EUA futures, which corroborates prior studies (see Benz and Hengelbrock). Average prices are consistently lower in ECX (e.g. €15.68 in Phase I) than in NP (€17.55 in Phase I), pointing to less significant market frictions such as

---

[9] In ECX the proportion of trades with the same time stamp is 9% in Phase I, 29% in Phase II and 45% in Phase III. In NP the proportion of trades with the same time stamp is 3% in Phase I, 4% in Phase II and 1% in Phase III.

[10] The inclusion of these trades as individual trades would significantly affect the magnitude and the significance of the ARMA coefficients, introducing bias to the estimates. The methodological focus of this paper is on trading intensity measured in seconds per contract, and by aggregating volume there is no loss of information since the time per traded contract remains the same at the specific time stamps at which these trades occur.

spreads (Medina et al., 2014) and higher operational efficiency in ECX (Kalaitzoglou and Ibrahim, 2013a, 2015). Average volume (trade size) increases over the phases, especially in NP and for OTC trades, while the median volume decreases only for ECX screen trades. Average duration decreases substantially over the phases in ECX, but rises substantially in the last phase in NP. It is also shorter in Phase II in both markets.

OTC trades (upstairs market) exhibit comparable trends in different venues and phases, but appear to be radically different from non-OTC trades (downstairs market). The average duration of OTC trades (e.g., 149.70s in ECX II and 183.91 s in ECX III) is consistently considerably longer compared to that of all trades as well as, by implication, screen trades (e.g., 59.40s in ECX II and 57.11s in ECX III). They are also distinctively larger (e.g., 43.72 in ECX II and 212.09 in ECX III) than all trades (e.g., 13.03 in ECX II and 18.61 in ECX III). This is probably why they are associated with higher price changes (e.g., standard deviation of returns is 0.45 versus 0.17 in ECX II, and 0.66 versus 0.39 in NP II). Thus, on average, OTC trades are larger, slower, more expensive (priced higher) and riskier than screen trades. This points towards different trading intensities, price discovery contributions and perhaps information and liquidity risk premia between the upstairs and downstairs markets.

4. Results and Analyses

To conserve space the estimation results of the STM-ACD model of Equations (1) to (9) and the DJM model of Equations (14) to (16) are presented and discussed in the online appendix.

4.1 Intradaily patterns

First we confirm whether the trading patterns over the day of OTC trades mentioned in Section 1 above four our sample phase series are similar to those reported by Kalaitzoglou and Ibrahim (2013) for the December 2008 contract. Figure 1 plots the number of all trades, high intensity trades (identified by the DJM model), and OTC trades by hour of the day in both market venues in each of three market phases. Against the background pattern of trading, OTC trades increase throughout the day in ECX but roughly follow the pattern of hourly trading in NP (even a decreasing trend in Phase III). This confirms that the observations of Kalaitzoglou and Ibrahim (2013) made through analysis of one contract extend to the three phase series used here. However, Figure 1 adds that the number of high intensity trades (whether OTC or screen) generally follows the background hourly pattern of trading (i.e., the number of all trades). Thus, in addition to OTC trades being larger, slower, more expensive and riskier than screen trades, they increase in number throughout the day in ECX but decrease in NP.

Given these patterns of general trading, we further investigate other properties of OTC trades. Figure 2 plots the averages within each hour of the trading day of price change volatility (realised volatility) five transactions prior (t-5) to ten transactions (t+10) following OTC trades. The obvious patterns are: a) volatility is high when OTC trades occur, b) volatility experiences a hump around OTC trades and the width of this hump differs by phase and trading venue, and c) the two effects just described are larger towards the beginning of the trading day in ECX, but larger towards the end of the trading day in NP.[11] Thus, OTC trades cause shocks in volatility, and these shocks start before the appearance of an OTC trade and lasts for a number of transactions after the OTC trade. This implies that OTC trades carry price-sensitive information or liquidity variations that take time to resolve or absorb. The figure also shows that the magnitude of the volatility shocks differs by time of day and

[11] In the online appendix we plot two versions of Figure 1. One shows the volatility around high intensity OTC trades, and the other shows volatility around non-high intensity OTC trades. The patterns are similar. Thus, volatility reacts to OTC trades in similar patterns regardless of whether or not these trades are of high intensity.

trading venue. This implies that the arrival of information or liquidity consequences of OTC trades are different during different hours of the day, and that this pattern is opposite between ECX and NP. It also points out that the uncertainty resolution period (time it takes for high relative volatility to die out) might well be different between OTC and screen trades and across venues. This is investigated further below.

One way of telling whether these patterns are due to information arrival or liquidity variations is to check whether trading intensity (volume weighted duration) as a proxy for (il)liquidity, exhibits similar patterns. We use the ACWD model, which distinguishes trades by three regimes of trading intensity, to separate OTC trades into those of high intensity from those of medium and low intensity. Figure 3 plots the trading intensity (lower values indicate higher trading intensity) of the 25 transactions around medium and low intensity OTC trades (t-5 to t+20). In general, there is a mild lower-intensity trading surrounding these trades (slightly discernible humps) in ECX, but no discernible similar behaviour in NP. Figure 4 is a similar plot but around high intensity OTC trades instead. There is a clear dip at t followed by a clear hump. Thus, trading becomes markedly less intense for a period that extends over the subsequent 7 to 13 transactions that follow high intensity OTC trades. This is not the case around medium and low intensity OTC trades. Accordingly, OTC trades in general cause a depletion of liquidity, and high intensity OTC trades cause a clear illiquidity shock (a more substantial depletion) that persists for up to 13 transactions. To see whether other high intensity trades also cause such a shock, Figure 5 plots trading intensity around high intensity screen trades. A clear dip occurs at t but is followed by a milder rise (lower liquidity) extending for a longer period in both venues and all three phases (except NP Phase III, which has very few high intensity observations; see Table 1). Thus, high intensity OTC trades cause an illiquidity shock from which the market recovers its intensity faster, while other high intensity trades cause a slower recovery in market trading intensity. In other words, liquidity is replenished faster after high intensity OTC trades than after high intensity screen trades. We next investigate the implications of this on the illiquidity premium and the uncertainty resolution period.

4.2 Illiquidity premium and resolution of uncertainty

Given the above results that OTC trades cause distinctly different liquidity perturbations in the market, we investigate whether they also have different implications on uncertainty resolution period and the illiquidity premium.

Figure 6 plots the time (in minutes) it takes for price change volatility to revert back to prior levels following an OTC trade, a high intensity (HI) trade, and a high intensity OTC (OTC+HI) trade, i.e., the time takes for volatility shocks introduced by each of these trade types to be absorbed. This can be considered as period taken for uncertainty to be resolved, and is presented by averages within each hour of the day. All graphs exhibit hump during the middle hours of the day, implying a generally slower resolution of uncertainty around the lunch hour. The graphs also show that, in general, uncertainty resolution is slowest after OTC trades and fastest after high intensity OTC trades (the OTC graph encapsulates the HI and the OTC+HI graphs during most hours). This means that low and medium intensity OTC trades introduce information and liquidity shocks that take longer for the market to resolve, but despite the fact that high intensity OTC trades are larger and faster, the information and liquidity shocks that accompany them are resolved fastest.

However, are these trades more costly because of the higher information asymmetry and liquidity risk that most likely accompanies them?[12] Table 2 tabulates the version (*IL*) of

[12] Ibrahim and Kalaitzoglou (2016) show that larger and faster trades (high intensity) can sometimes increase or decrease the spread (trading costs). If information asymmetry dominates immediacy improvements with high intensity trades then costs increase, otherwise they decrease or remain unchanged. On average, OTC trades are

Amihud's illiquidity premium introduced in Equation (11) above. Values are tabulated for the 5-miuntes, 15-minutes and 1-hour intervals that follow the 15-minutes over which MSI is measured and which precede each transaction. As MSI measures the degree of presence of high intensity OTC trades (measured for each transaction), and hence the intensity of recent trading activity in the upstairs market, the table provides illiquidity premia estimates for different levels of high-intensity OTC presence (MSI bands). This is tabulated by each phase of each venue. The last three columns present the percentage ratio of this measure in ECX over NP. A clear pattern in the first six columns is the decline in the illiquidity premium as MSI increases. This pattern is largely consistent throughout, except perhaps for the smallest two or three bands of MSI. This implies a generally declining cost of liquidity per contract following high intensity of OTC presence, or intense activity of the upstairs market. Together with the results reported in the previous Section 4.1, this indicates that although high intensity OTC trades cause short lived episodes of liquidity depletion, their volatility shocks are absorbed faster and they are generally followed by lower liquidity premia. On face value, this may seem counter-intuitive, since one would expect a higher illiquidity premium with greater liquidity depletion. However, the *IL* measure does not differentiate between asymmetric information and liquidity costs since it is based on price changes that include both information and liquidity components. It also ignores the increased risk (volatility) that accompanies high intensity OTC trades (as shown in the previous sub-section). Thus, the risk dimension ought to be taken into account in proper measurement of risk premia. Before we discuss risk adjustments (in Section 4.5 below) we first investigate whether OTC trades crowd themselves or other trades in the market. This is carried out next.

---

larger but not faster than screen trades and hence, they may have different implications on uncertainty resolution and information and liquidity risk premia.

4.3 Crowding for Carbon

The increasing number of OTC trades throughout the day in ECX, and the opposite pattern in NP, calls for an investigation of whether the larger but slower OTC trades crowd themselves or screen trades during different times of the day, or when their intensity increases. To this end we compute a measure of crowdedness based on trade imbalance (whether buys are more than sells or vice versa). Define C_{OTC} as the sum of signed OTC trade volume during the 15 minutes prior to a trade,

$$C_{OTC} = \sum_{i=0}^{n} OTC\ indicator_i \times trade\ direction\ indicator_i \times Volume_i, \qquad (18)$$

and define C_{screen} equivalently for screen trade volume prior to each trade. These measure the OTC and screen trade imbalances, respectively, prior to each trade. Positive values indicate more buys than sells, and negative values indicate more sells than buys of each type of trade during the 15 minutes that precede a trade. To first investigate patterns in the magnitude of trade imbalance we initially consider averages of absolute values. Table 3 reports the averages of the absolute values of these measures during every hour of the day and within different MSI bands for the three phases and the two market venues. In general, the tabulated values show that the magnitude of both OTC and screen crowdedness increases throughout the day in ECX, but decreases in NP. This means that crowdedness is linked to the hourly pattern in the number of OTC trades shown in Figure 1 (i.e., higher towards the end of the day in ECX, but lower in NP). Thus, the higher the number of OTC trades the higher is the crowdedness in the markets (for both OTC and screen trades). This is also clear from the increasing values of the measures of crowdedness at higher MSI bands (bottom panel of Table 3). Accordingly, as the intensity of OTC trades increases (higher MSI), their larger size crowds each other as well as screen trades.

Note that, during trading hours, the magnitude of screen crowdedness is higher than OTC crowdedness in all three phases in ECX, but lower in NP, again in all three phases. This means that trade imbalance in the ECX downstairs market is higher that in the upstairs market, especially in Phase III. The opposite, however, is true in NP, where the imbalance in the upstairs market is higher than that in the downstairs market. Note also that the disparity in magnitude of crowdedness between these two markets is higher during the morning session than the afternoon session. This could be explained by intradaily movements in market share across venues, or by segregation of customer type (perhaps geographically).

Although this shows the relative magnitude of OTC versus screen crowdedness, it does not necessarily reveal whether the upstairs (OTC) market crowds the downstairs (screen) market, as not all 15-minute intervals have OTC trades (i.e., OTC trades are not uniformly distributed over time intervals or hours of the day). This is clearer when we analyse crowdedness across MSI bands. The MSI measures the joint intensity and presence of OTC trades over the 15 minutes that precede a trade. From the bottom panel of Table 3, the 1-1.3 band shows no presence of OTC trades in any venue or phase, and in Phase III the highest band is 2.5-2.8 in ECX and 2.2-2.5 in NP. In general, when there are OTC trades (populated MSI bands) the tabulated values show that OTC crowdedness is higher than screen crowdedness in Phases I and II, in both ECX and NP. This is reversed, however, in Phase III in both markets, where screen crowdedness is higher than OTC crowdedness. Thus, despite the decreased presence of high intensity OTC trades in Phase III, the substantially increased size of OTC trades in that phase may have contributed to higher screen crowdedness. Alternatively, more informed large trading may have moved to the downstairs market in Phase III. This alternative explanation is consistent with the fact the number of identical zero-duration trades has increased over the phases, and is highest by a large margin in Phase III (these are trades that have identical time stamp, price, trade direction, and trade type indicator and, hence, are suspected to be large broken up strategic trades).

To further investigate the relationship between OTC and screen crowdedness, Figure 7 plots their average correlation by hour of day (top graph) and MSI bands (middle graph). Further, the bottom graph depicts average screen crowdedness across bins of OTC crowdedness. From the top graph, and by hour of day, the correlations generally range from -0.06 to 0.03 in ECX and -0.25 to 0.05 in NP. This graph shows a very small positive correlation in ECX during the first three hours of the day (first five hours in Phase II), which then becomes negative. The general pattern in ECX is a decline in correlation throughout the day. In NP, however, the correlations are mainly negative but increase throughout the day and become positive only during the afternoon session in Phase II and the last two hours in Phase III. Given that OTC trades increase towards the end of the day in ECX, but decrease in NP, the generally negative correlation means that OTC and screen crowdedness mildly compensate each other, and when one rises the other falls (both on a generally rising drift throughout the day as shown in Table 3). This alternation, or 'dance,' increases in intensity towards the end of the day in ECX and decreases in NP.

The middle graph of Figure 7 shows that when crowdedness is measured relative to degrees of presence of intense OTC trades, the correlations between OTC and screen crowdedness decreases substantially as OTC intensity rises in both venues and all phases. Thus, relative trade imbalance between the OTC and screen markets alternates ('dances') more vigorously when the upstairs OTC market is more active.

The bottom graph of Figure 7 shows that the more negative the OTC trade imbalance (i.e., more sells than buys) is, the more positive the screen imbalance (i.e., more buys than sells) is, and vice versa. This is consistent across both venues and over the three phases. It is clear evidence, therefore, that a larger relative imbalance in one market (say upstairs) on average coincides with an opposite imbalance in the other market (downstairs). When there are more OTC buys, there are more screen sells, and when there are more OTC sells, there are more screen buys. This compensatory mechanism operates in a similar manner in both ECX and NP, and in all three phases.

In summary, the carbon market is, on average, imbalanced. Towards the end of the trading day, OTC trades increase in number and trade imbalance in ECX, but decrease in number and trade imbalance in NP. Variations in trade imbalance alternate between the OTC market and the screen market. The degree of alternation increases towards the end of the day in ECX, but decreases in NP. However, an increasingly selling screen crowd faces an increasingly buying OTC crowd, and a buying screen crowd meets a selling OTC crowd. Hence, a demand pressure in the upstairs market seems to coincide with a supply pressure in the downstairs market, and vice versa. This is consistent in both market venues and over the three phases.

4.4 The effect of crowdedness

Given these patterns in crowdedness, we further analyse their effect in terms of over-reaction or under-reaction on returns and trade direction. Normally, a buying pressure should lead to price increases, and a selling pressure should lead to price decreases (if supply is constant). Would the cap-and-trade carbon market, however, over-react or under-react to these imbalances? We look at the autcorrelations in returns (change in price) and in trade direction over the five-minute period that follows every screen or OTC transaction (each binned by relative degree of OTC or screen crowdedness). The top three graphs of Figure 8 plot these autocorrelations for different bins of OTC crowdedness, and the bottom three graphs plot the same by bins of screen crowdedness, both for the ECX venue. In general, the observations fan out from the thinner centre (at zero crowdedness). Thus, correlations are stronger as crowdedness becomes more positive or negative. Both of returns and trade direction are more autocorrelated following greater buying or selling pressure, indicating more aligned reaction and higher predictability of direction. Note also that both negative and positive autocorrelations are observed at high levels of crowdedness (as one moves away from the centre to the left or to the right towards the edges of the x-axes). This means that both over-reaction and under-reaction occur.

One main observation from Figure 8 is that the autocorrelations of returns are lower than those of trade direction, while both are negative on average. This indicates more reversals in returns than in trade direction over the five minutes that follow a trade, on average. Thus, in terms of relative magnitude, returns swing more than trade direction. Since most return autocorrelations are negative, there are more reversals in returns than continuations regardless of the direction of trade imbalance – sign of crowdedness (i.e., whether buys are more than sells, or vice versa). Further, the degree of this reversal increases with the level of crowdedness. Accordingly, there are larger or more frequent price adjustments over the next five minutes as crowdedness increases. This can be due to either more vigorous resolution of uncertainty because of higher degrees of asymmetric information or due to longer time requirements to replenish depleted liquidity. In either case, the higher levels of crowdedness convey a lack of market depth.

Note also, that the observations of positive autocorrelation of returns do not occur except at high levels of crowdedness (although this varies by phase). Thus, there are cases in which the returns have runs in the same direction, and judging by the size of the positive autocorrelation over the next five minutes these runs can be highly aligned in direction. This is a sign of overreaction. In terms of phases, note that this effect is more pronounced after screen crowdedness (bottom graphs) than after OTC crowdedness (top graphs) within each phase. Thus, overreaction in the same direction (i.e., a continuation of buys or of sells) seems to be more pronounced in the downstairs market at high levels of screen crowdedness.

Another observation is that autocorrelations by OTC crowdedness are more dispersed than those by screen crowdedness, at all levels of crowdedness. Autocorrelations are more bunched up at low to moderate levels of screen crowdedness than at similar levels of OTC crowdedness. At low to medium levels of crowdedness, there are larger changes in autocorrelations following OTC crowdedness than following screen crowdedness. This implies that relatively low to moderate levels of screen crowdedness do not cause as much of over-reaction or under-reaction as similar levels of OTC crowdedness do. However, at higher levels of screen crowdedness (roughly outside the -500 to 500 range) the dispersion of autocorrelations is comparable to, or even exceeds (e.g., in Phase III), that of OTC crowdedness. This pattern is the same for returns and trade direction. Thus, there seems to be threshold levels of screen crowdedness beyond which the degree of over-reaction or under-reaction matches or exceeds that of OTC crowdedness.

These patterns are roughly similar in NP (Figure 9) bearing in mind that NP's share of the carbon market is far smaller than that of ECX. The median return autocorrelation (horizontal level around which autocorrelations fan out) is different in NP than in ECX and is distinctly less negative in Phase II. Autocorrelations are also less bunched up at the centre and fan out faster (i.e., have a wider range at comparable levels of crowdedness) than in ECX. These indicate greater levels of over-reaction or under-reaction. Thus, NP over-reacts or under-reacts more than ECX, especially in Phase II.

In summary, the markets are more crowded following OTC trades than screen trades in Phases I and II, but the reverse is true in Phase III. Autocorrelations of returns and trade direction (indicators of overreaction and under-reaction) increase in magnitude with the level of crowdedness. This effect is more pronounced after OTC crowdedness than after screen crowdedness for low to medium levels of crowdedness. Thus, there is more over-reaction and under-reaction following OTC trade imbalances of medium size. At higher levels of crowdedness this effect is stronger on screen (i.e., in the downstairs market). Over-reaction or under-reaction is more pronounced in NP than in ECX, especially in Phase II.

We next investigate whether the intensity of OTC trades and crowdedness have implications on the information and liquidity risk premia.

4.5 Illiquidity and information risk premia

Crowdedness could be motivated by information arrival or liquidity shocks. It is of interest to see whether the above patterns in the number of trades, volatility, liquidity premium, uncertainty resolution, crowdedness and autocorrelations of returns and trade direction coincide with a decrease or an increase in the risk premia of information and liquidity across the upstairs and the downstairs markets. In theory, as liquidity rises one expects a lower risk premium, but since information arrival is usually accompanied by an increasing risk premium, these two opposing forces might cancel out if they are equal and if risk premia are measured on the basis of total risk (i.e., price change volatility). To investigate the above effects on the information risk premium separately from the liquidity risk premium, we use the version of Ibrahim and Kalaitzoglou's (2016) microstructure model in Equations (14) to (16) to extract the components of return and variance that are due to information and those that are due to liquidity. We then study the ratios of these components, i.e., the *AIP* and *LIQ* measures described in Equations (13) and (14) above. *AIP* measures the percentage change in price due to information per unit of 'information' risk, and *LIQ* measures the percentage change in price due to liquidity per unit of 'liquidity' risk.

Figure 10 plots hourly averages of *AIP* and *LIQ* for OTC and screen trades. First, the relative scale of the two *y*-axes shows that *AIP* is larger than *LIQ*, as expected. The variations shown in the graphs are quite small, but significant nonetheless.[13] In other words, the information risk premium is larger than the liquidity risk premium, on average. Thus, in

[13] Differences between values at the beginning of the day versus those at the end of the day are tested in risk premia and across MSI in Table A2 in the online Appendix A.

general, information risk demands a higher compensation than liquidity risk, even if the latter is sizable.

More importantly, there are discernible patterns in risk premia over the trading day. In Phases I and II of ECX there is a marked decline in the information risk premium for OTC trades during the first few hours of the day. In Phase III it oscillates around an almost flat trend. For screen trades there is a milder declining trend throughout the day in all phases, and in the last two phases the information risk premium of screen trades is lower than that of OTC trades. In NP the patterns are almost opposite: the information risk premium for OTC trades oscillates around a distinctly rising trend, but that of screen trades trends upwards in Phase I but less so in Phase II and trends downwards in Phase III, with the relative magnitude in the last two phases being lower than that of OTC trades. Thus, the information risk premium of OTC trades is larger than that of screen trades in the last two phases of the market, which are more active. It declines throughout the day in ECX (or stays flat) but rises in NP. Taken together with the hourly patterns in the number of OTC and screen trades discussed previously, these results may mean that as OTC trades increase in number, they crowd themselves and the screen market but the information risk premium for these trades decreases as price discovery is faster (uncertainty resolution is also faster), while the information premium for screen trades is not affected as much. Thus, in terms of compensation for information risk, it seems that OTC trades ‘crowd’ themselves more than they crowd screen trades. Although their demand or supply imbalance on average is met by an opposite imbalance from screen trades, their competitiveness as far as information is concerned reduces the risk adjusted reward for themselves more than for screen trades, even though overall information costs is higher for these trades than for screen trades.

Figure 10 also plots the hourly average liquidity risk premia for OTC and screen trades. These show patterns that are a mirror reflection of those just discussed for the information risk premia. Using the same dataset Ibrahim and Kalaitzoglou (2016) show that a reciprocal relationship between information and liquidity components of price change exists. Figure 10 adds that this reciprocal relationship also extends to risk premia. The liquidity premium per unit of liquidity risk for OTC trades rises throughout the day in ECX and falls in NP in Phases I and II. In Phase III it follows a very mild downtrend for OTC trades, and an uptrend for screen trades in ECX, but falls dramatically for OTC trades and oscillates around a roughly stable daily average for screen trades in NP. In general, screen trades demand a higher liquidity premium than OTC trades, especially in the last two phases. Thus, as the competitiveness of OTC trades decreases, their information risk adjusted rewards decrease in ECX and increase in NP. The opposite is true for liquidity risk adjusted rewards.

To see variations in rewards as OTC intensity increases we plot the risk premia by MSI bands in Figure 11. This is shown for OTC, screen, and all trades separately. The observed patterns are clear. Information risk premia decline and liquidity risk premia rise at higher MSI bands. The exception being the information rewards for OTC trades in NP Phase III (but there are only 46 OTC trades in NPIII). Thus, as OTC intensity increases, the rewards on information erode out, but those on liquidity are enhanced. However, the erosion in information risk adjusted rewards is more than the enhancement in those of liquidity. Accordingly, in terms of risk premia, as OTC trading intensifies the OTC market ends up crowding itself as well as the downstairs market.

5. Summary and Conclusion

This paper analyses differences in trading patterns and price formation between the upstairs and downstairs markets of two main trading venues of the European Union Emission Trading System. It investigates whether the upstairs market crowds itself or the downstairs market, or vice versa, and the implications this might have on uncertainty resolution, over-reaction and under-reaction, and information and liquidity risk premia. Through intraday analysis of transaction data covering the period from market inception to 30 April 2015 of the European Climate Exchange and NordPool trading venues, the following results are reported. Upstairs OTC trades have distinctly different characteristics than downstairs screen trades. On average, they are larger, slower, more expensive and riskier than screen trades, and increase in number throughout the day in ECX but decrease in NP. They also introduce a volatility and trading intensity shock that are larger at the beginning of the day in ECX and at the end of the day in NP. High intensity OTC trades are accompanied by a larger shock in trading intensity and cause discernible short-lived liquidity depletion episodes and shocks that takes the market up to 13 transactions to recover from. However, trading intensity recovers, and liquidity is replenished, faster after high intensity OTC trades than after high intensity screen trades. Uncertainty resolution is also fastest after high intensity OTC trades and slowest after low and medium intensity OTC trades, but the Amihud's illiquidity premium is lower following high intensity OTC trades.

On microstructural scale, the carbon market is on average imbalanced and crowded. Crowdedness increases throughout the day in ECX and decreases in NP, in line with the general hourly pattern in the number of OTC trades. These venues could be dominated by different customer base with different liquidity demand preferences over the day. In ECX, the upstairs market is more active towards the end of the day, and in NP towards the beginning of the day. Crowdedness also increases with the presence of high intensity OTC trades. Thus, as OTC intensity increases so does the magnitude of crowdedness and the larger sized OTC trades crowd the upstairs as well as the downstairs markets. By hour of the day, the downstairs market (screen trades) is more crowded than the upstairs market, on average. However, by degree of presence of high intensity OTC trades (upstairs high activity) OTC crowdedness is higher than screen crowdedness in Phases I and II, but lower in Phase III.

Thus, despite more intense trades upstairs in Phase III the substantially higher size of OTC trades in that phase may have contributed to crowdedness downstairs. An alternative rationale is that more informed large trading may have moved downstairs in Phase III. This would be consistent with the increased number of zero-druation identical trades (large broken up strategic trades) over the phases and especially in Phase III.

Over-reaction and under-reaction over the five minute interval that follows a trade increase in magnitude with crowdedness. There are more return reversals than trade direction reversals. Consequently, there is lack of market depth with higher levels of crowdedness. Distinctly positive autocorrelations in return are observed only at extreme levels of crowdedness, indicating decidedly high over-reaction. Relatively low to moderate levels of screen crowdedness do not cause as much over-reaction and under-reaction as similar levels of OTC crowdedness, but high screen crowdedness beyond the -500 to 500 threshold levels can cause greater over-reaction and under-reaction than OTC trades. There are differences in these patterns between ECX and NP.

The information risk premium declines and the liquidity risk premium rises as the levels of intense upstairs activity rise. As OTC intensity increases rewards for asymmetric information erode, but those on liquidity are enhanced. However the former exceeds the latter. Hence, in terms of risk premia, the upstairs market ends up crowding itself as well as the downstairs market as its trading intensifies.

References

Alberola, E., J. Chevallier, and B. Chèze, B. 2008. Price drivers and structural breaks in European carbon prices 2005-07. *Energy Policy* 36:787–797.

Amihud, Y. 2002. Illiquidity and stock returns: Cross-section and time series effects, *Journal of Financial Markets* 5:31–56.

Benz, E. A. and J. Hengelbrock. 2008. Price discovery and liquidity in the European CO_2 futures market: An intraday analysis. Working Paper, The Bonn Graduate School of Economics.

Bredin, D., S. Hyde, and C. Muckley. 2014. A microstructure analysis of the carbon finance market. *International Review of Financial Analysis* 34:222–234.

Daskalakis, G., Psychoyios, D., and R. Markellos. 2009. Modeling CO2 emissions allowance prices and derivatives: Evidence from the European Trading Scheme. *Journal of Banking and Finance* 33:1230–1241.

Ellerman, D., C. Marcantonini, and A. Zaklan. 2014. The EU ETS: Eight years and counting. EUI Working Paper RSCAS 2014/04, Robert Schuman Centre for Advanced Studies, Climate Policy Research Unit, European University Institute.

Engle, R. F. 2000. The econometrics of ultra–high–frequency data. *Econometrica* 68:1–22.

Ibrahim, B. M. and I. Kalaitzoglou. 2016. Why do carbon prices and price volatility change? *Journal of Banking and Finance*

Kalaitzoglou, I., and B. M. Ibrahim. 2013a. Does order flow in the European carbon allowances market reveal information? *Journal of Financial Markets* 16:604–635.

Kalaitzoglou, I., and B. M. Ibrahim. 2013b. Trading patterns in the European carbon market: The role of trading intensity and OTC transactions. *Quarterly Review of Economics and Finance* 53:402–416.

Kalaitzoglou, I., and B. M. Ibrahim. 2015. Liquidity and resolution of uncertainty in the European futures market. *International Review of Financial Analysis* 37:89–102.

Lucia, J., M. Mansanet-Bataller, and A. Pardo. 2015. Speculative and hedging activities in the European carbon market. *Energy Policy* 82:342–351.

Mizrach, B. 2012. Integration of the global carbon markets, *Energy Economics* 34:335-349.

Mizrach, B., and Y. Otsubo. 2014. The market microstructure of the European climate exchange. *Journal of Banking and Finance* 39:107–116.

Rannou, Y. 2017. Liquidity, information, strategic trading in an electronic order book: New insights from the European carbon markets. *Research in International Business and Finance* 39: 779–808.

Rannou, Y. and P. Barneto. 2014. Futures trading with information asymmetry and OTC predominance: Another look at the volume/volatility relations in the European carbon markets. *Energy Economics* 53: 159–174.

Rittler, D. 2012. Price discovery and volatility spillovers in the European Union Emissions Trading Scheme: A high-frequency analysis. *Journal of Banking and* Finance 36:774–785.

Stein, J. 2009. Presidential address: Sophisticated investors and market efficiency, *Journal of Finance* 64:1517-1548.

Table 1 Descriptive Statistics

			Phase I					Phase II					Phase III				
			#	Duration	Volume	Price	Return	#	Duration	Volume	Price	Return	#	Duration	Volume	Price	Return
ECX	All trades	Mean	26,854	864.50	13.08	15.68	0.00	1,043,043	59.40	13.03	13.79	0.00	654,324	57.11	18.61	5.87	0.00
		Median		210.00	10.00	16.25	0.00		9.00	5.00	13.70	0.00		6.00	5.00	5.58	0.00
		Maximum		62054.00	600.00	31.00	17.75		49200.00	5500.00	33.70	23.77		16431.00	9000.00	20.22	4.11
		Minimum		0.00	1.00	0.01	-19.40		0.00	1.00	0.00	-23.47		0.00	1.00	2.46	-4.11
		Std. Dev.		2431.81	19.04	8.79	0.93		273.87	40.79	5.03	0.17		211.04	81.06	1.96	0.03
		Skewness		9.95	9.43	-0.24	-0.75		47.86	46.63	0.64	0.45		15.46	40.41	2.61	0.29
		Kurtosis		145.47	151.35	1.96	113.70		4300.34	4154.81	3.11	2554.56		494.19	2870.15	15.26	3574.70
	OTC	Mean	8,626	898.56	19.23	16.72	0.06	91,361	149.70	43.72	16.04	0.01	9,208	183.91	212.09	6.51	0.00
		Median		240.00	10.00	16.50	0.00		49.00	25.00	14.97	0.00		56.00	90.00	5.88	0.00
		Maximum		62054.00	600.00	31.00	17.75		49200.00	5500.00	33.50	23.77		11696.00	9000.00	19.70	3.41
		Minimum		0.00	1.00	0.03	-19.40		0.00	1.00	2.82	-8.50		0.00	1.00	2.48	-4.11
		Std. Dev.		2378.90	28.47	8.07	1.37		647.48	115.08	5.23	0.45		427.17	525.89	2.84	0.16
		Skewness		11.40	7.28	-0.40	0.08		26.53	21.94	0.36	6.05		8.22	8.19	2.28	-1.79
		Kurtosis		193.27	82.97	2.29	57.37		1150.50	717.70	2.43	346.01		124.77	96.05	9.50	163.82
NP	All trades	Mean	5,070	2667.00	12.00	17.55	0.00	17,910	2241.00	18.57	17.01	0.00	2,823	6105.69	21.60	7.22	0.00
		Median		1020.00	10.00	17.20	0.00		705.00	10.00	16.20	0.00		3007.00	10.00	6.65	0.00
		Maximum		27000.00	550.00	51.15	34.30		35868.00	100000.00	37.25	21.50		35925.00	7000.00	19.62	6.11
		Minimum		0.00	1.00	0.01	-36.00		0.00	0.00	1.00	-36.25		0.00	1.00	2.04	-5.92
		Std. Dev.		3991.52	16.34	8.41	0.75		3829.75	753.58	5.61	0.39		7523.54	263.11	2.97	0.22
		Skewness		2.72	11.95	-0.48	-3.13		3.28	130.58	-0.10	-33.16		1.62	26.45	1.88	0.45
		Kurtosis		12.22	285.24	2.38	1946.01		17.03	17305.01	1.99	4973.34		5.16	701.65	7.33	389.06
	OTC	Mean	2,489	2752.75	15.11	17.58	-0.01	5,390	2108.20	40.08	19.22	0.00	46	7148.87	628.57	5.46	0.03
		Median		1140.00	10.00	17.00	0.00		780.00	10.00	20.05	0.00		4266.00	13.00	5.03	0.00
		Maximum		26940.00	550.00	51.15	34.30		29893.00	100000.00	37.25	21.50		22880.00	7000.00	8.70	0.88
		Minimum		0.00	1.00	0.04	-36.00		0.00	1.00	1.00	-36.25		61.00	3.00	3.15	-0.16
		Std. Dev.		3976.37	20.48	8.43	1.04		3428.51	1373.44	4.58	0.66		6891.60	1988.30	1.27	0.17
		Skewness		2.71	11.09	-0.39	-2.32		3.31	71.64	-0.28	-22.17		1.01	2.93	0.36	3.21
		Kurtosis		12.26	223.51	2.26	1045.65		17.27	5208.24	2.39	1909.70		2.76	9.59	2.28	15.74

Table 1 presents the descriptive statistics of trade duration, volume, price and price change of all trades versus OTC trades (OTC, Block, EFP/EFS) for ECX and NP in Phase I, II and III.

Table 2 Illiquidity Premium across Market Sentiment Levels

		ECX I	ECX II	ECX III	NP I	NP II	NP III	Phase I	Phase II	Phase III
5'	1-1.3	2.5474	0.0964	0.0051	10.5399	0.7879	1.3296	24.17%	12.24%	0.39%
	1.3-1.6	2.4115	0.1205	0.0025	7.6171	0.2968	0.3172	31.66%	40.62%	0.78%
	1.6-1.9	1.8586	0.1428	0.0034	3.2611	0.6139	0.2432	56.99%	23.26%	1.41%
	1.9-2.2	2.0579	0.1272	0.0025	3.2679	0.5409	0.3642	62.97%	23.51%	0.70%
	2.2-2.5	1.4002	0.1028	0.0018	2.5139	0.2660	-	55.70%	38.63%	-
	2.5-2.8	1.3109	0.0293	-	2.1973	0.3214	-	59.66%	9.12%	-
	2.8-3.1	1.2561	0.0286	-	2.1624	0.0957	-	58.09%	29.89%	-
	3.1-3.4	1.0579	0.0188	-	2.5913	-	-	40.83%	-	-
	3.4-3.7	-	0.0072	-	-	-	-	-	-	-
	3.7-4	-	0.0332	-	-	-	-	-	-	-
15'	1-1.3	1.1584	0.0721	0.0042	2.6199	0.5349	0.8890	44.21%	13.48%	0.48%
	1.3-1.6	0.8929	0.0837	0.0013	2.0871	0.1443	0.1303	42.78%	58.03%	0.98%
	1.6-1.9	0.9421	0.0856	0.0021	1.4858	0.4661	0.1391	63.41%	18.36%	1.53%
	1.9-2.2	1.0768	0.0502	0.0022	1.5971	0.4113	0.3815	67.42%	12.20%	0.57%
	2.2-2.5	0.4471	0.0421	0.0012	1.3655	0.1056	-	32.74%	39.92%	-
	2.5-2.8	0.1818	0.0101	-	1.2480	0.0690	-	14.57%	14.57%	-
	2.8-3.1	0.2760	0.0069	-	1.4737	0.0483	-	18.73%	14.25%	-
	3.1-3.4	0.2039	0.0065	-	0.7196	-	-	28.34%	-	-
	3.4-3.7	-	0.0030	-	-	-	-	-	-	-
	3.7-4	-	0.0108	-	-	-	-	-	-	-
1h	1-1.3	0.4385	0.0217	0.0034	1.1309	0.4887	0.7305	38.78%	4.44%	0.46%
	1.3-1.6	0.4286	0.0276	0.0010	1.1518	0.1400	0.1243	37.21%	19.70%	0.80%
	1.6-1.9	0.5402	0.0282	0.0010	0.7659	0.1505	0.1138	70.53%	18.74%	0.88%
	1.9-2.2	0.4511	0.0251	0.0013	0.5660	0.1894	0.1531	79.70%	13.24%	0.82%
	2.2-2.5	0.2708	0.0139	0.0009	0.6913	0.1057	-	39.17%	13.14%	-
	2.5-2.8	0.1368	0.0038	-	0.4225	0.2010	-	32.38%	1.90%	-
	2.8-3.1	0.1069	0.0023	-	0.2431	0.0329	-	44.00%	7.14%	-
	3.1-3.4	0.0913	0.0027	-	0.1613	-	-	56.58%	-	-
	3.4-3.7	-	0.0023	-	-	-	-	-	-	-
	3.7-4	-	0.0041	-	-	-	-	-	-	-

Table 2 presents the illiquidity premium, $IL = 100 \times (P_{i+s} - P_i)^2 / \sum_{j=1}^{n} v_{i+j}$, for a period of 5 minutes, 15 minutes and 1 hour across different levels of MSI, in both markets in all three phases. The last panel of the table presents the relative illiquidity premium in ECX as a percentage of the equivalent in NP, i.e., IL_{ECX}/IL_{NP}.

Table 3 Average Magnitude of Crowdedness by Hour of Day and OTC Intensity Levels

		ECX I		ECX II		ECX III		NP I		NP II		NP III	
		OTC	Screen	OTC	Screen	OTC	Screen	OTC	Screen	OTC	Screen	OTC	Screen
Hour of Day	08	33.23	39.19	13.66	67.81	13.35	115.82	57.40	57.23	199.98	95.43	192.38	102.10
	09	29.32	45.66	14.35	76.44	18.59	123.24	50.08	42.84	179.56	75.99	133.67	90.23
	10	31.28	53.95	26.62	79.92	33.65	124.42	48.47	37.42	156.41	65.88	143.29	93.10
	11	35.26	55.97	33.20	83.83	37.31	139.01	49.41	41.11	161.78	62.37	113.63	70.19
	12	40.67	50.85	39.52	80.56	36.26	135.10	45.86	38.45	124.33	37.66	98.04	76.80
	13	41.45	52.38	46.36	79.24	38.75	134.59	43.80	39.61	62.01	42.50	101.93	88.32
	14	40.41	51.20	38.61	79.48	49.27	139.71	43.47	34.92	69.49	37.31	60.87	71.96
	15	46.76	49.80	43.39	81.80	43.47	143.29	38.28	41.90	50.92	40.83	71.82	60.01
	16	49.25	58.23	56.47	86.13	66.15	143.50			49.72	31.91	62.76	65.05
	17	53.79	61.44	64.57	93.25	55.20	155.56			44.17	27.43	53.71	49.09
MSI	1-1.3												
	1.3-1.6	14.90	2.91	24.83	2.77	36.21	46.85	13.08	13.98	22.78	18.35	6.15	43.50
	1.6-1.9	34.98	9.50	59.80	6.88	73.34	132.52	25.90	13.25	56.62	16.80	30.82	52.08
	1.9-2.2	50.00	16.09	70.73	19.36	80.47	172.27	46.29	19.51	63.42	28.03	43.63	67.17
	2.2-2.5	60.33	22.69	70.34	31.85	100.60	162.53	64.82	25.16	73.47	36.22	65.00	65.59
	2.5-2.8	57.09	30.07	89.87	41.29	164.31	189.05	80.76	30.18	71.95	41.56		
	2.8-3.1	61.79	45.61	110.52	43.27			85.25	43.72	117.53	51.57		
	3.1-3.4	66.50	58.06	130.22	55.75			140.78	53.94				
	3.4-3.7	71.21	78.24	134.39	78.12			195.98	65.13				
	3.7-4			157.71	84.37								

Table 3 presents average absolute crowdedness for OTC and screen trades computed as $OTC = \sum_{i=0}^{n} \left| C_{OTC,i} \right| / n$ and $Screen = \sum_{i=0}^{n} \left| C_{Screen,i} \right| / n$, where n is the number of trades in each hour (top panel), or in each band of MSI (bottom panel),
$C_{OTC} = \sum_{i=0}^{n} OTC\ Indicator_i \times Trade\ Direction\ Indicator_i \times Volume_i$, and
$C_{Screen} = \sum_{i=0}^{n} (1 - OTC\ Indicator_i) \times Trade\ Direction\ Indicator_i \times Volume_i)$

Figure 1 Number of Trades

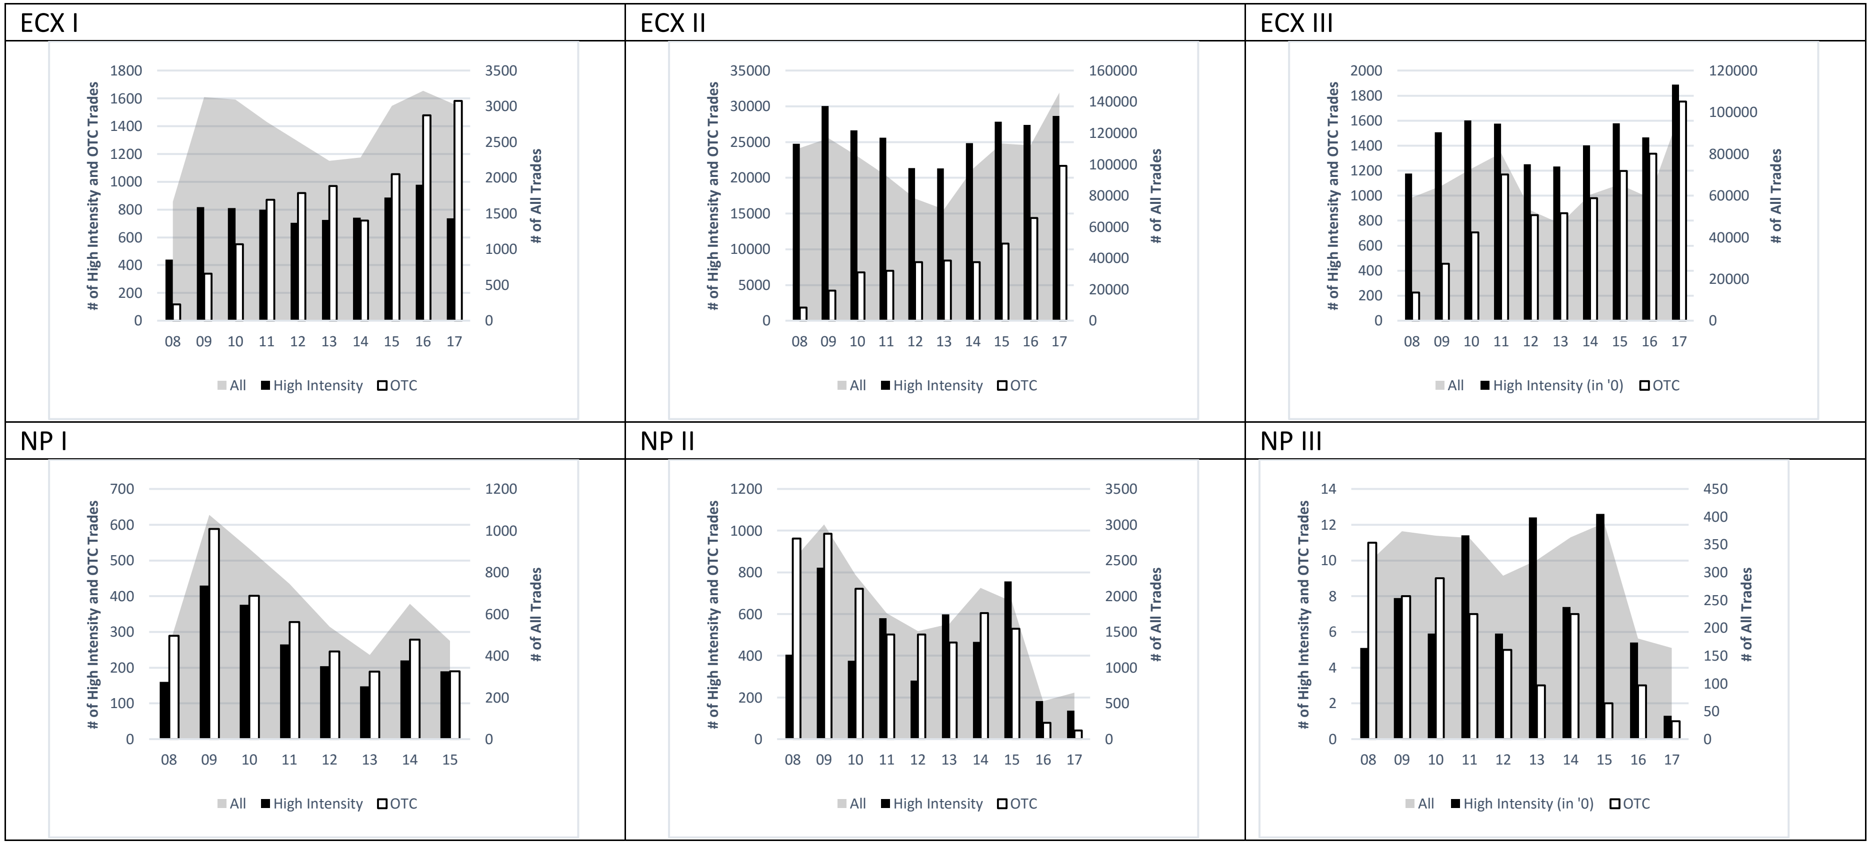

Figure 1 presents the total number of trades during the day in both ECX and NP for Phase I, II and III, along with the total number of high intensity trades, and the number of OTC trades. In Phase III the number of high intensity trades is reported in 10s in order to match the scale of OTC trades.

Figure 2 Price Change Volatility around OTC Trades Over the Day

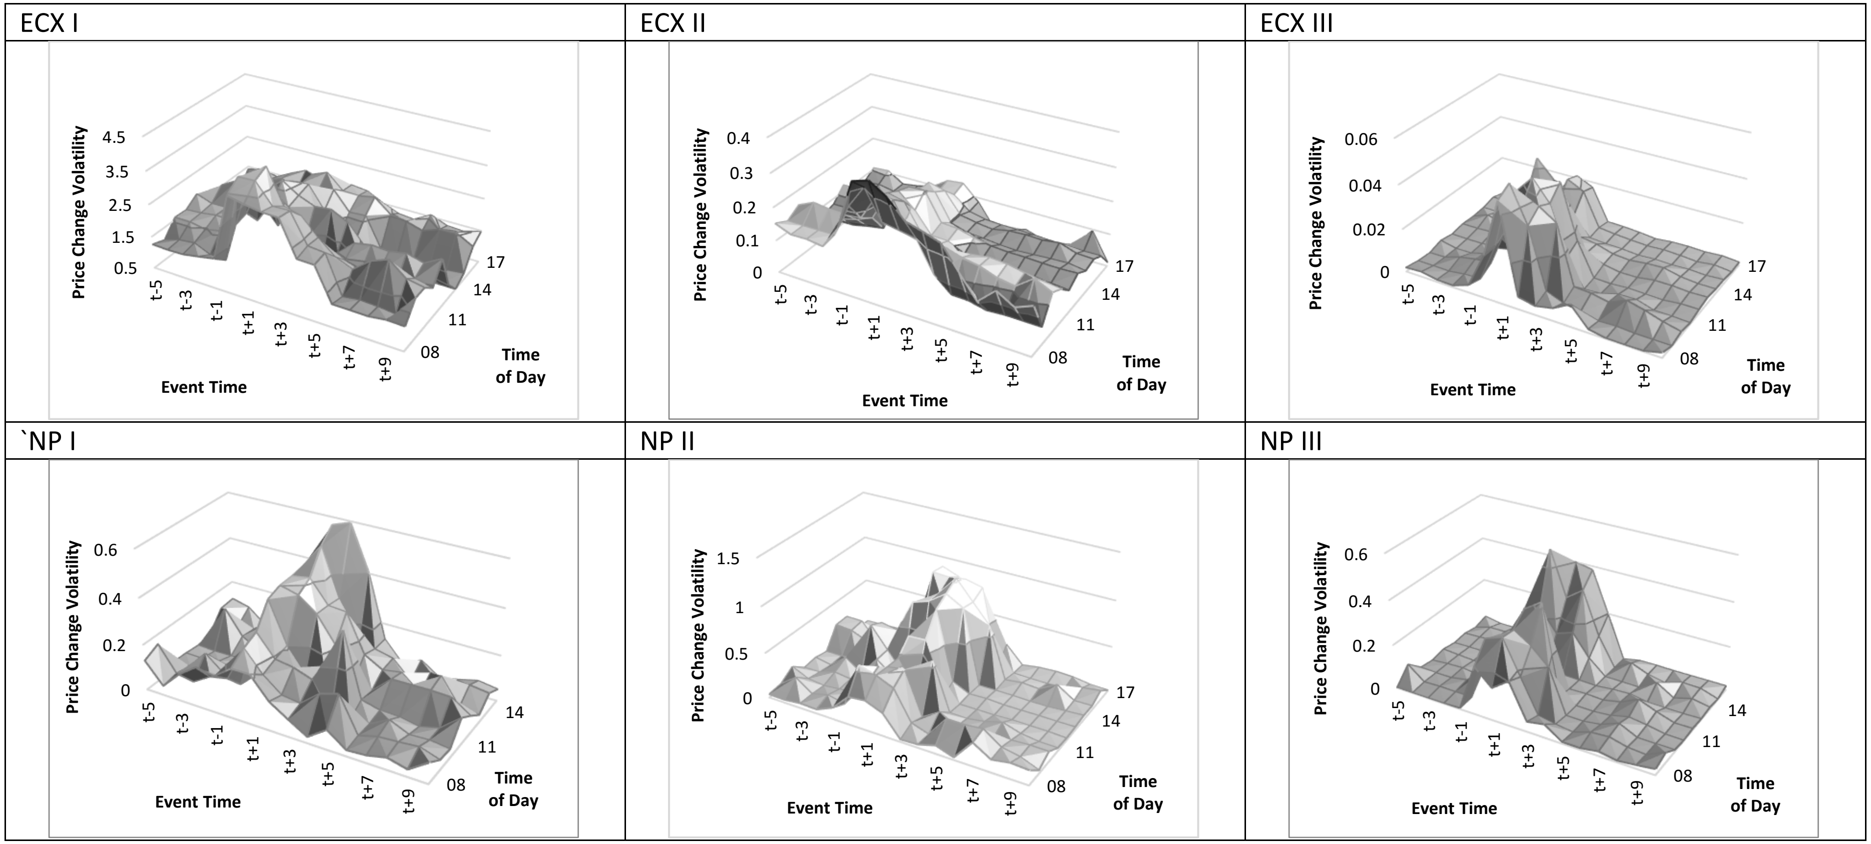

Figure 2 presents the (realised) variance of price change (vertical axis), i.e. $\Delta Price_i{}^2$, for the five transactions that precede and the ten that follow an OTC trade (horizontal axis), averaged over every hour of the trading day (depth axis) in both ECX and NP and for Phase I, II and III.

Figure 3 Trading Intensity around Low and Medium Intensity OTC Trades Over the Day

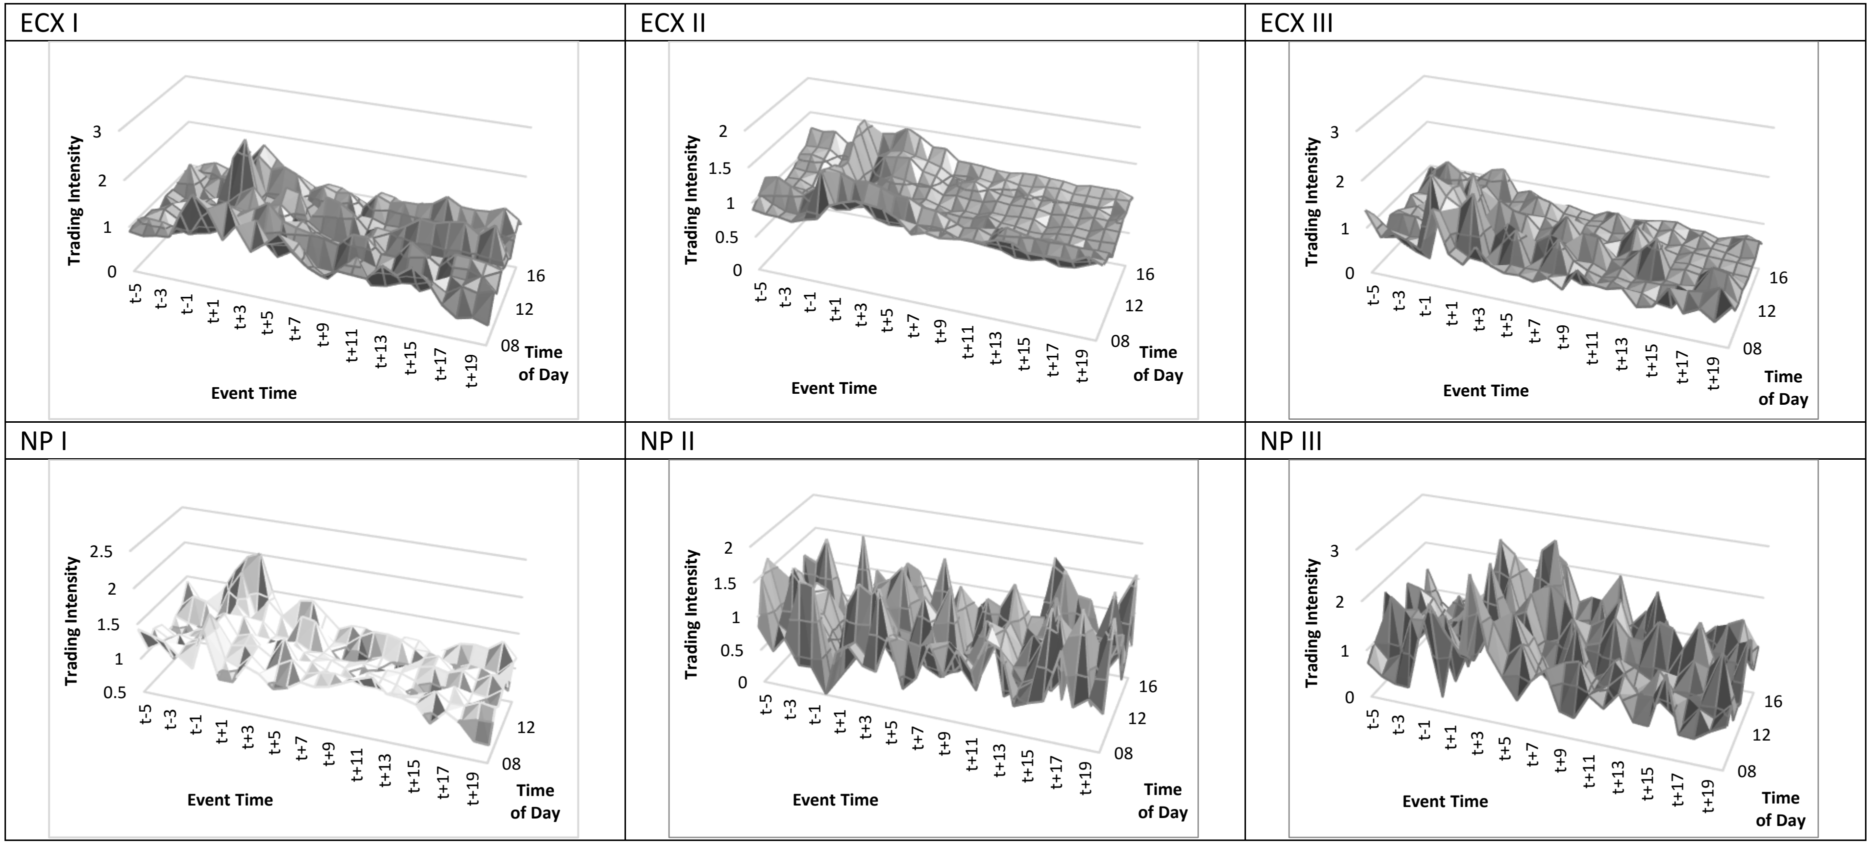

Figure 3 presents the average trading intensity (vertical axis), i.e., S_i, volume weighted duration, for the five transactions that precede and the 20 that follow (horizontal axis) a non-high intensity OTC trade, during each hour of the trading day (depth axis) in both ECX and NP for Phase I, II and III.

Figure 4 Trading Intensity around High Intensity OTC Trades Over the Day

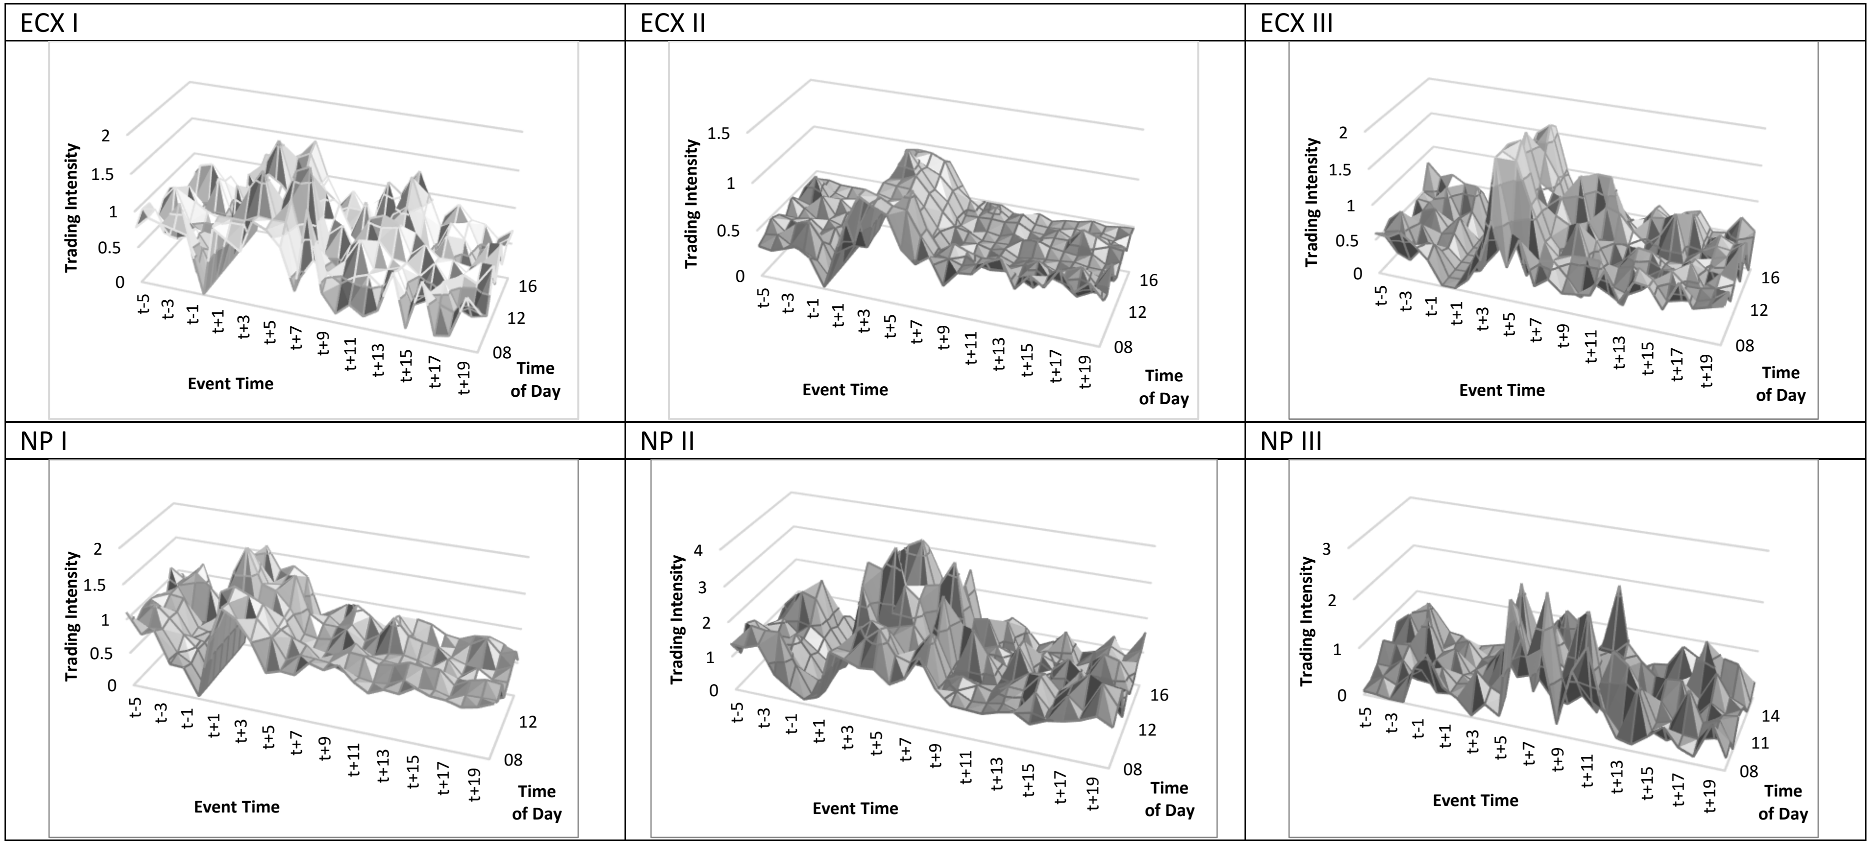

Figure 4 presents the average trading intensity (vertical axis), i.e., S_i, volume weighted duration, for the five transactions that precede and the 20 that follow (horizontal axis) a high intensity OTC trade, during every hour of the trading day (depth axis) in both ECX and NP for Phase I, II and III.

Figure 5 Trading Intensity around High Intensity Trades Over the Day

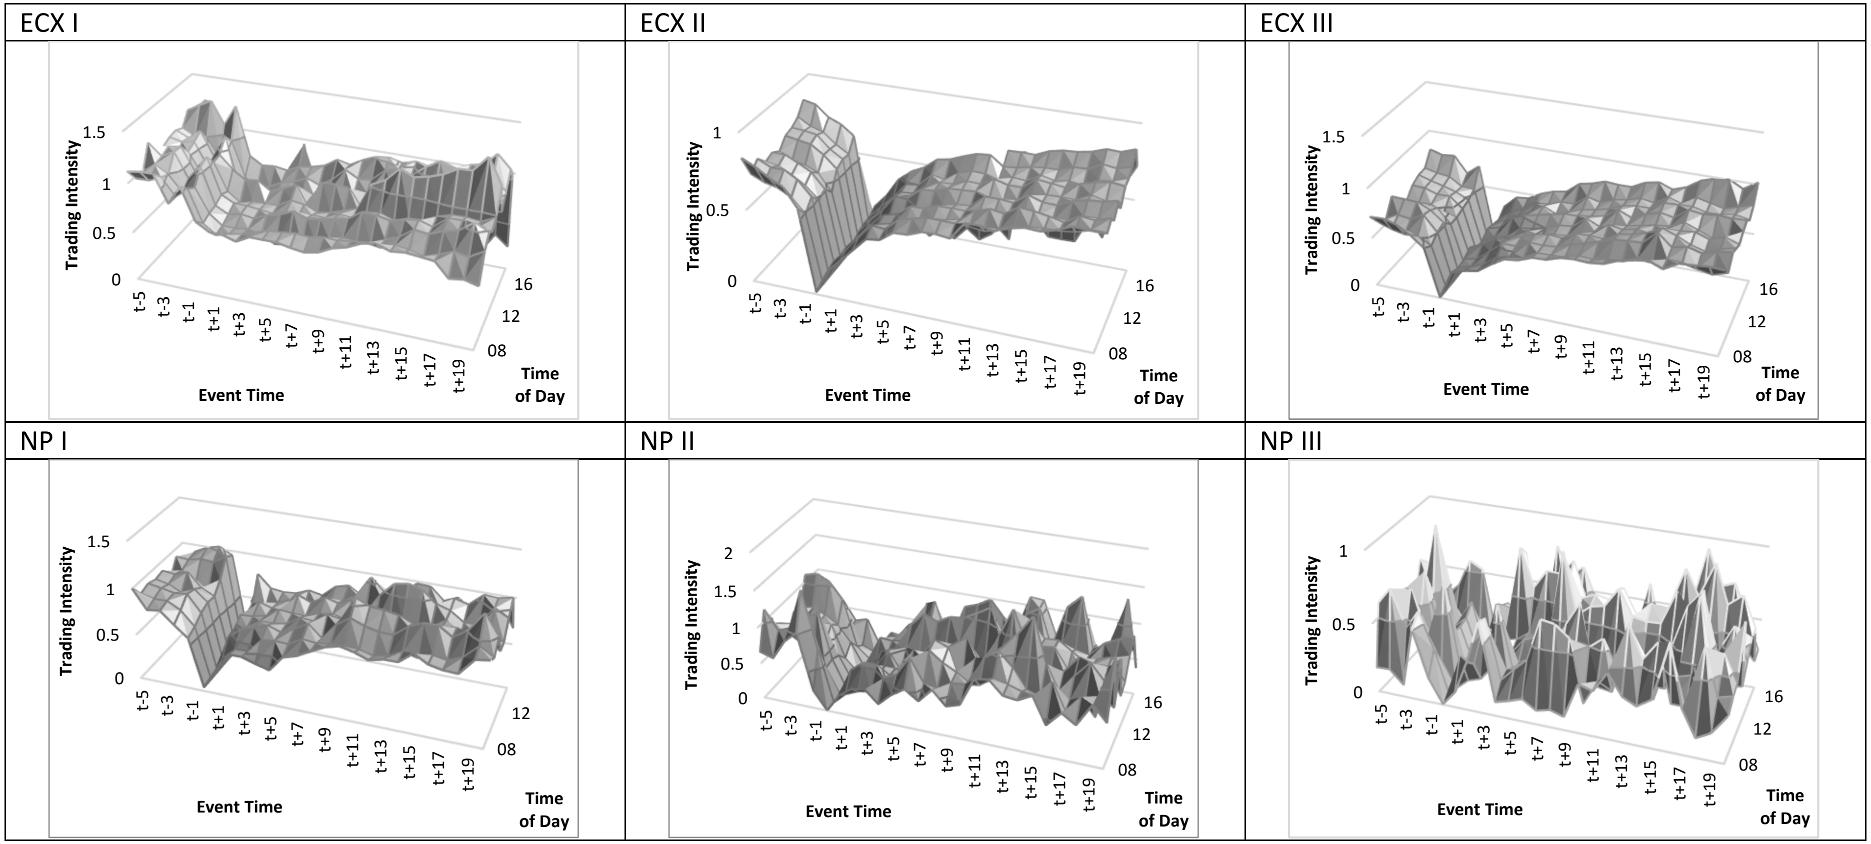

Figure 5 presents the average trading intensity (vertical axis), i.e., S_i, volume weighted duration, for the five transactions that precede and the 20 that follow (horizontal axis) a high intensity trade (whether OTC or screen), during every hour of the trading day (depth axis) in both ECX and NP for Phase I, II and III.

Figure 6 Period of Uncertainty Resolution

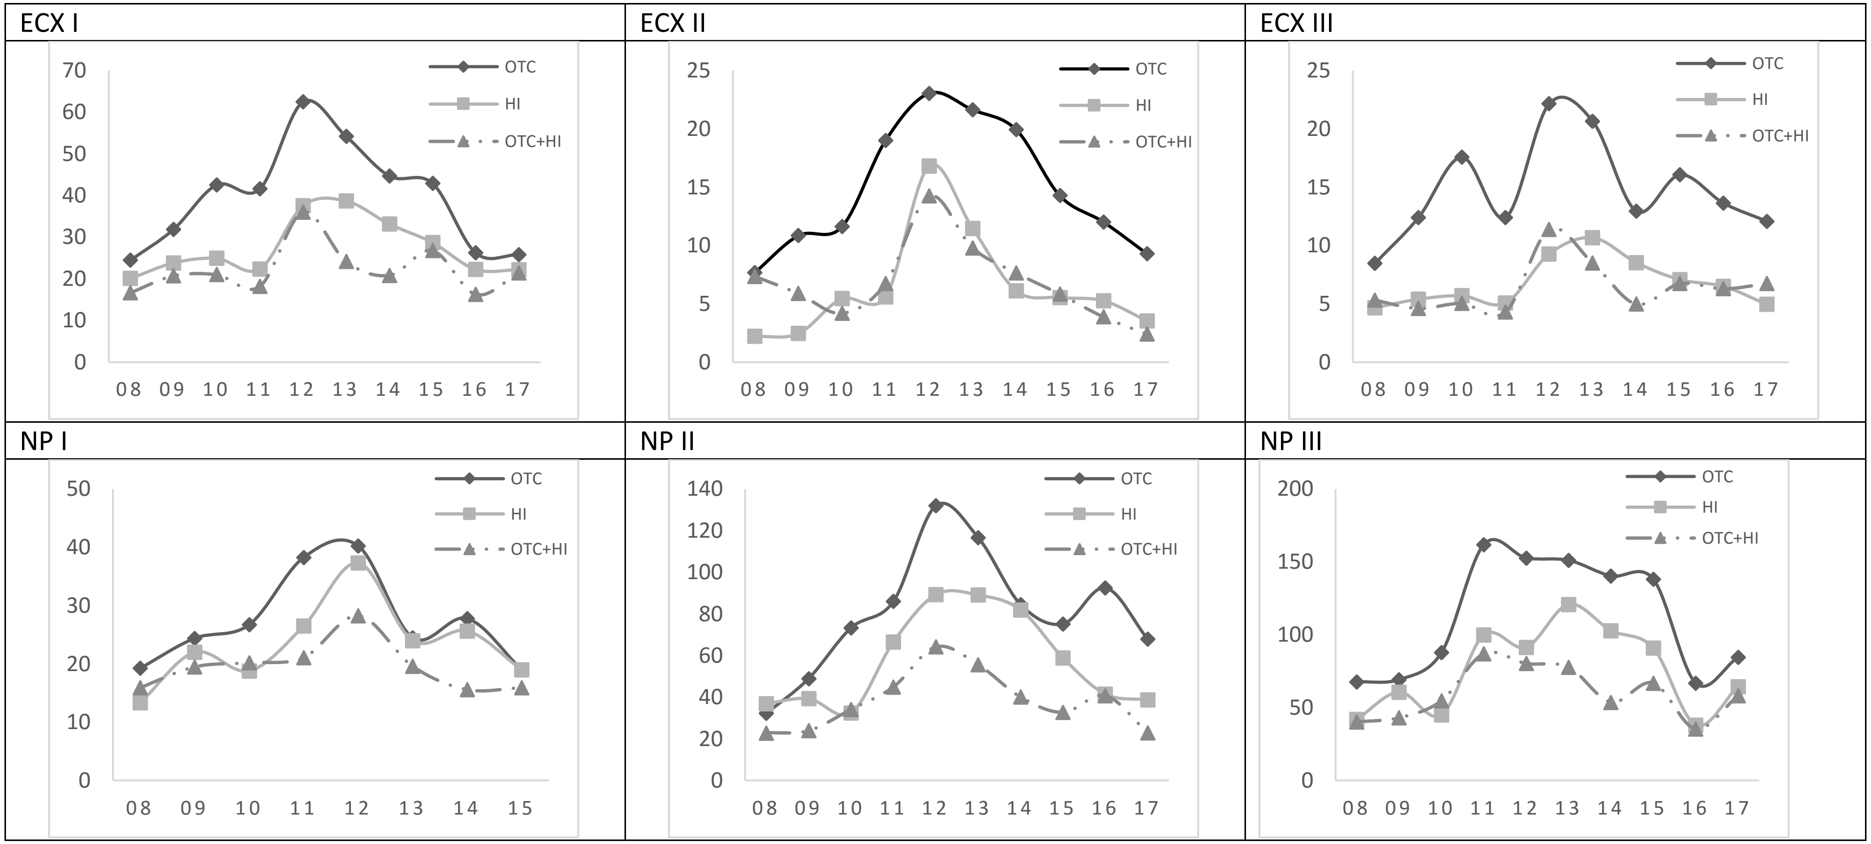

Figure 6 presents the product of average duration of the trades subsequent to high intensity and OTC trades multiplied by the number of transactions it takes for price change volatility to return to its prior levels. The product, measured in minutes (vertical axis), shows how long it takes for volatility shocks to be absorbed and is measured over one-hour intervals over the trading day (horizontal axis) in both ECX and NP for Phase I, II and III.

Figure 7 Correlation of OTC and Screen Crowdedness

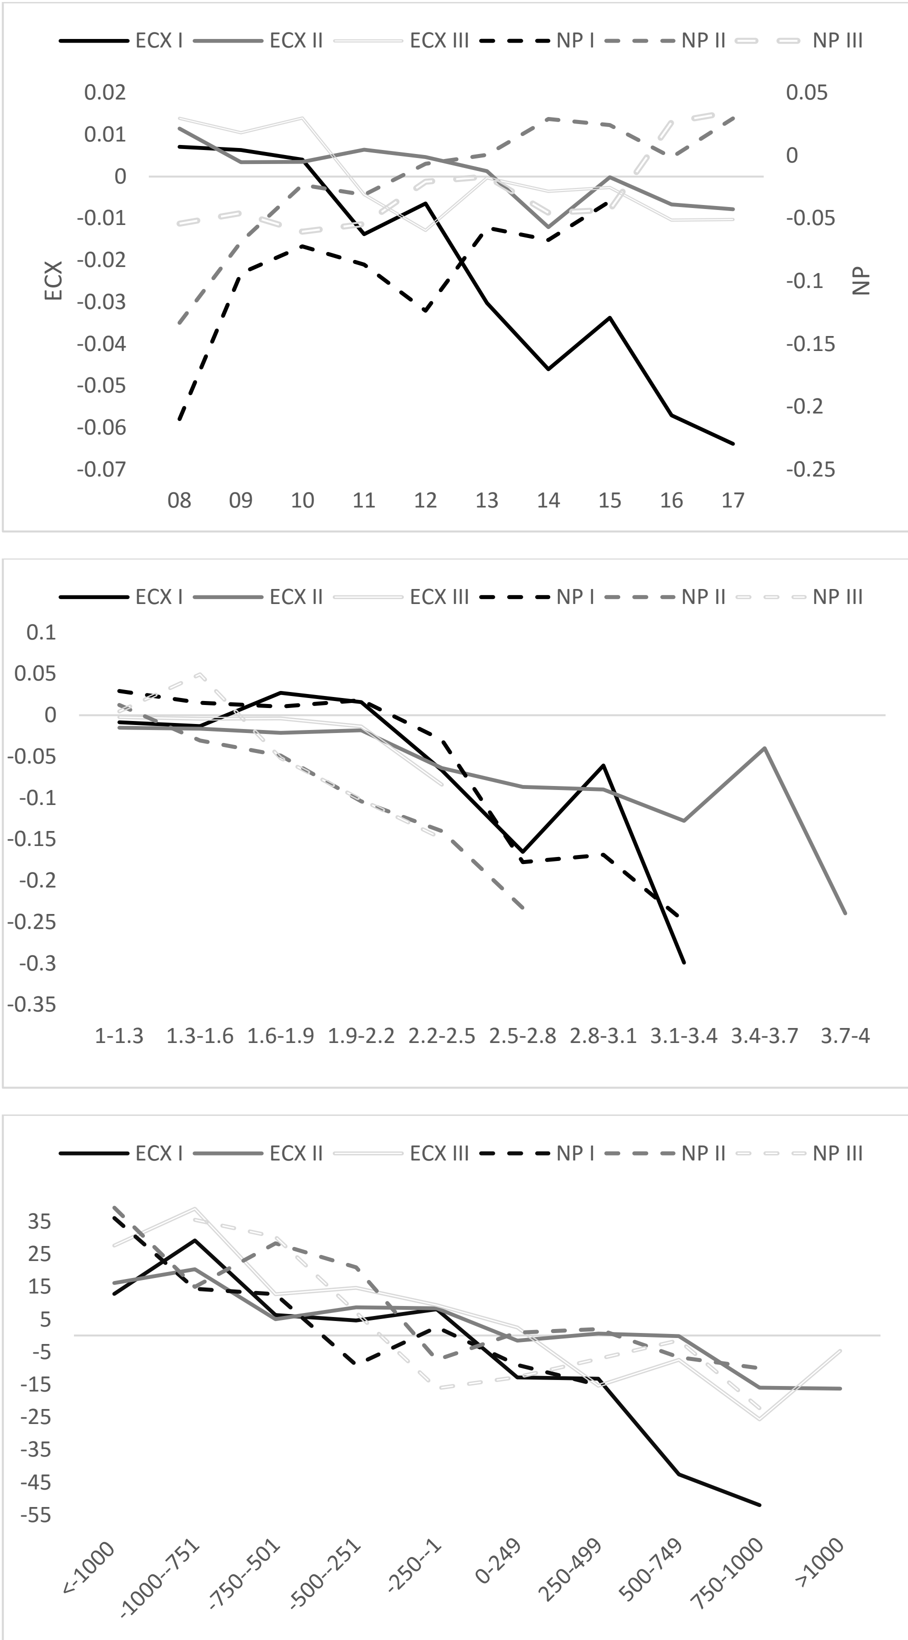

Figure 7 top graph plots the correlation between C_{OTC} and C_{Screen} (actual not absolute values $|\cdot|$) over the trading day. The middle graph plots the average correlation across MSI bands. The bottom graph plots C_{Screen} (y axis) over different levels of C_{OTC} (both in actual not absolute values$|\cdot|$).

Figure 8 Autocorrelation of Returns and Trade Direction by OTC and Screen Crowdedness Levels in ECX

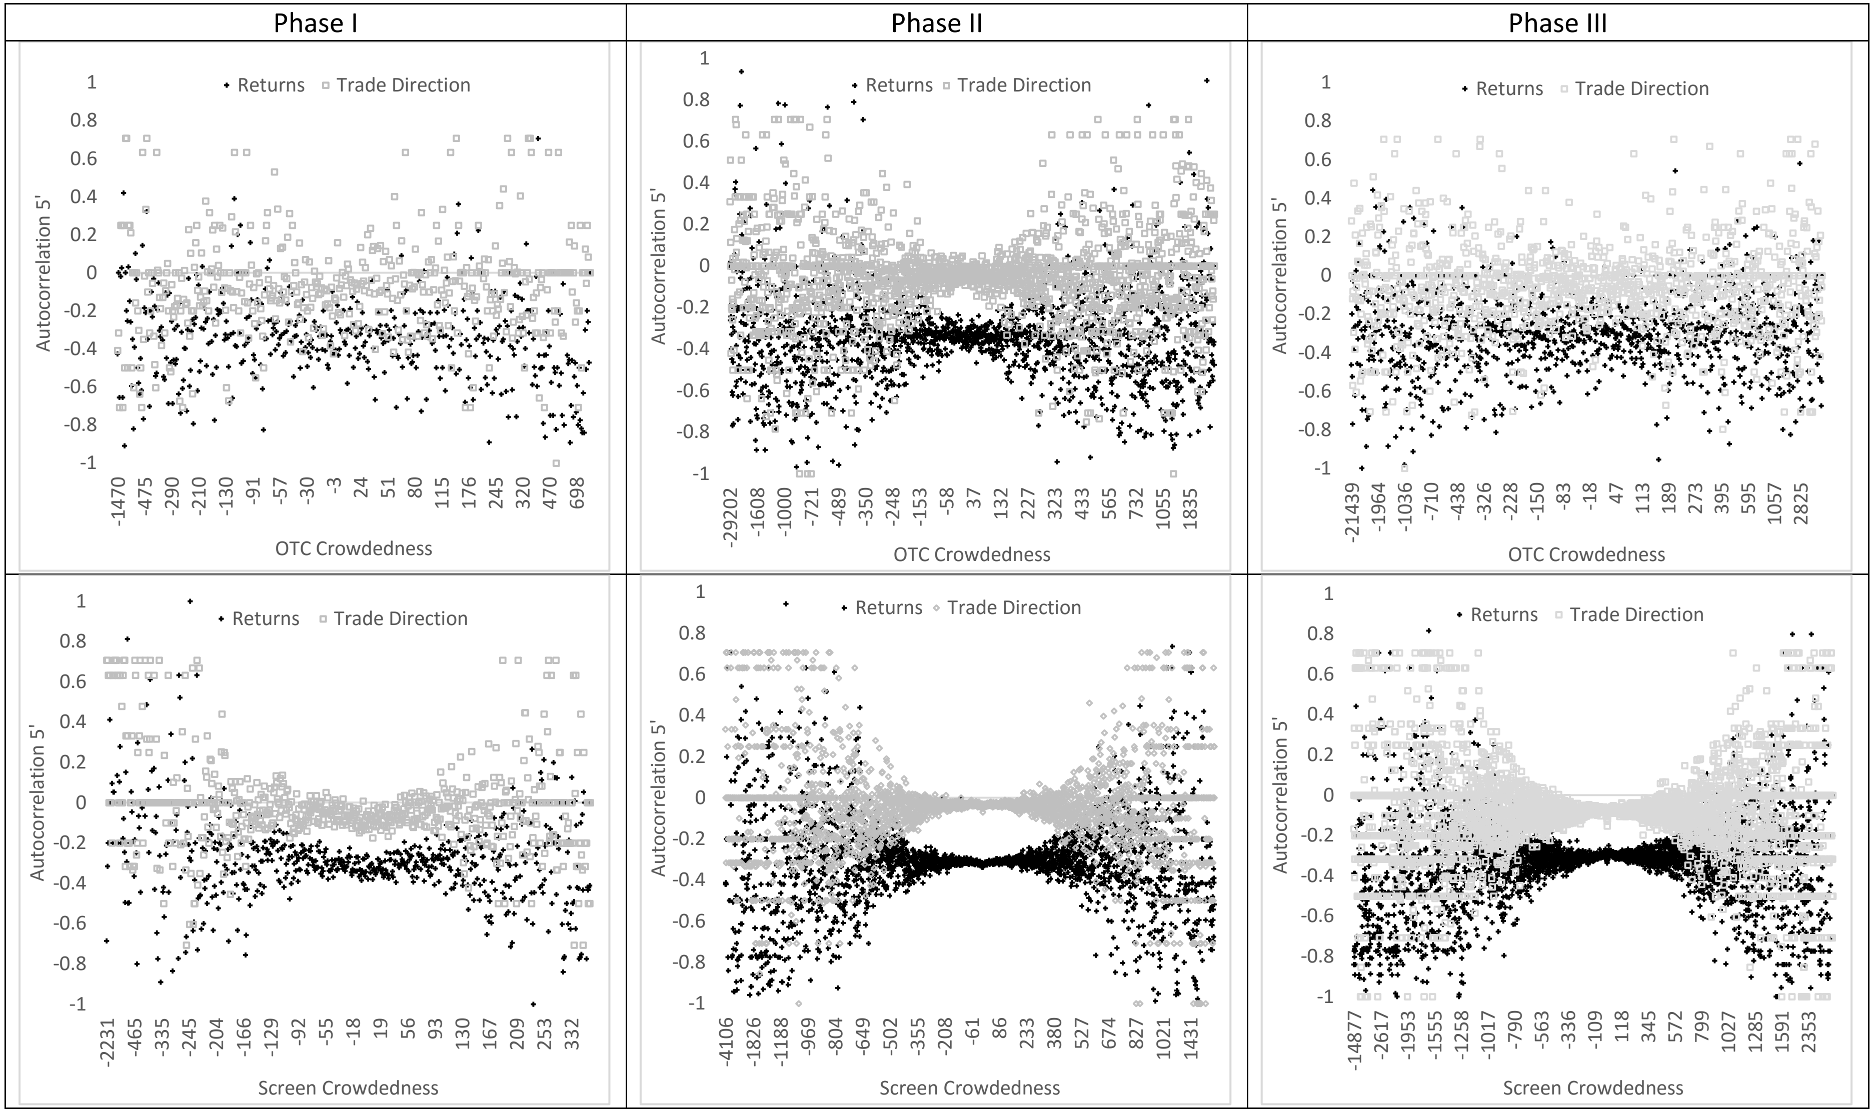

Figure 8 are scatter plots of the average autocorrelation of returns (grey boxes) and trade direction (black crosses) across C_{OTC} and C_{Screen}.

Figure 9 Autocorrelation of Returns and Trade Direction by OTC and Screen Crowdedness Levels in NP

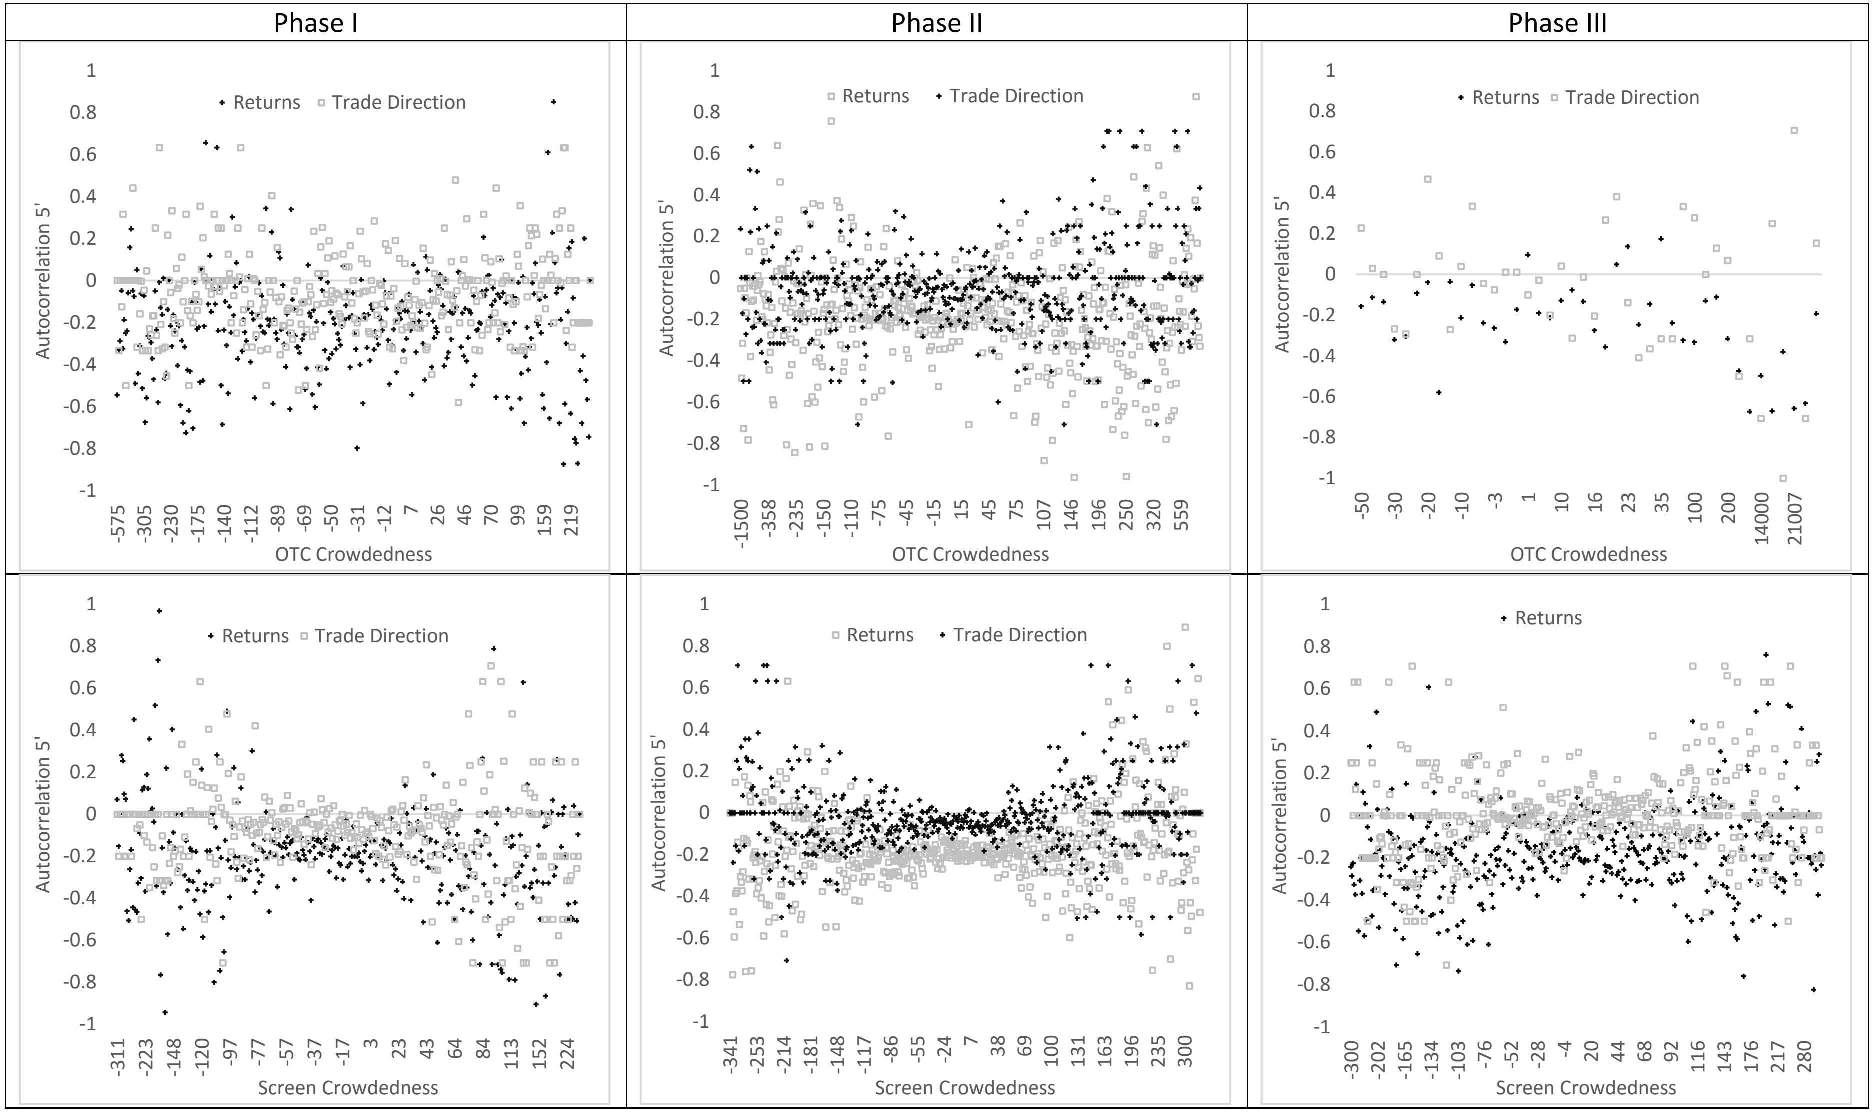

Figure 9 are scatter plots of the average autocorrelation of returns (grey boxes) and trade direction (black crosses) across C_{OTC} and C_{Screen}.

Figure 10 Risk Premia Over The Trading Day

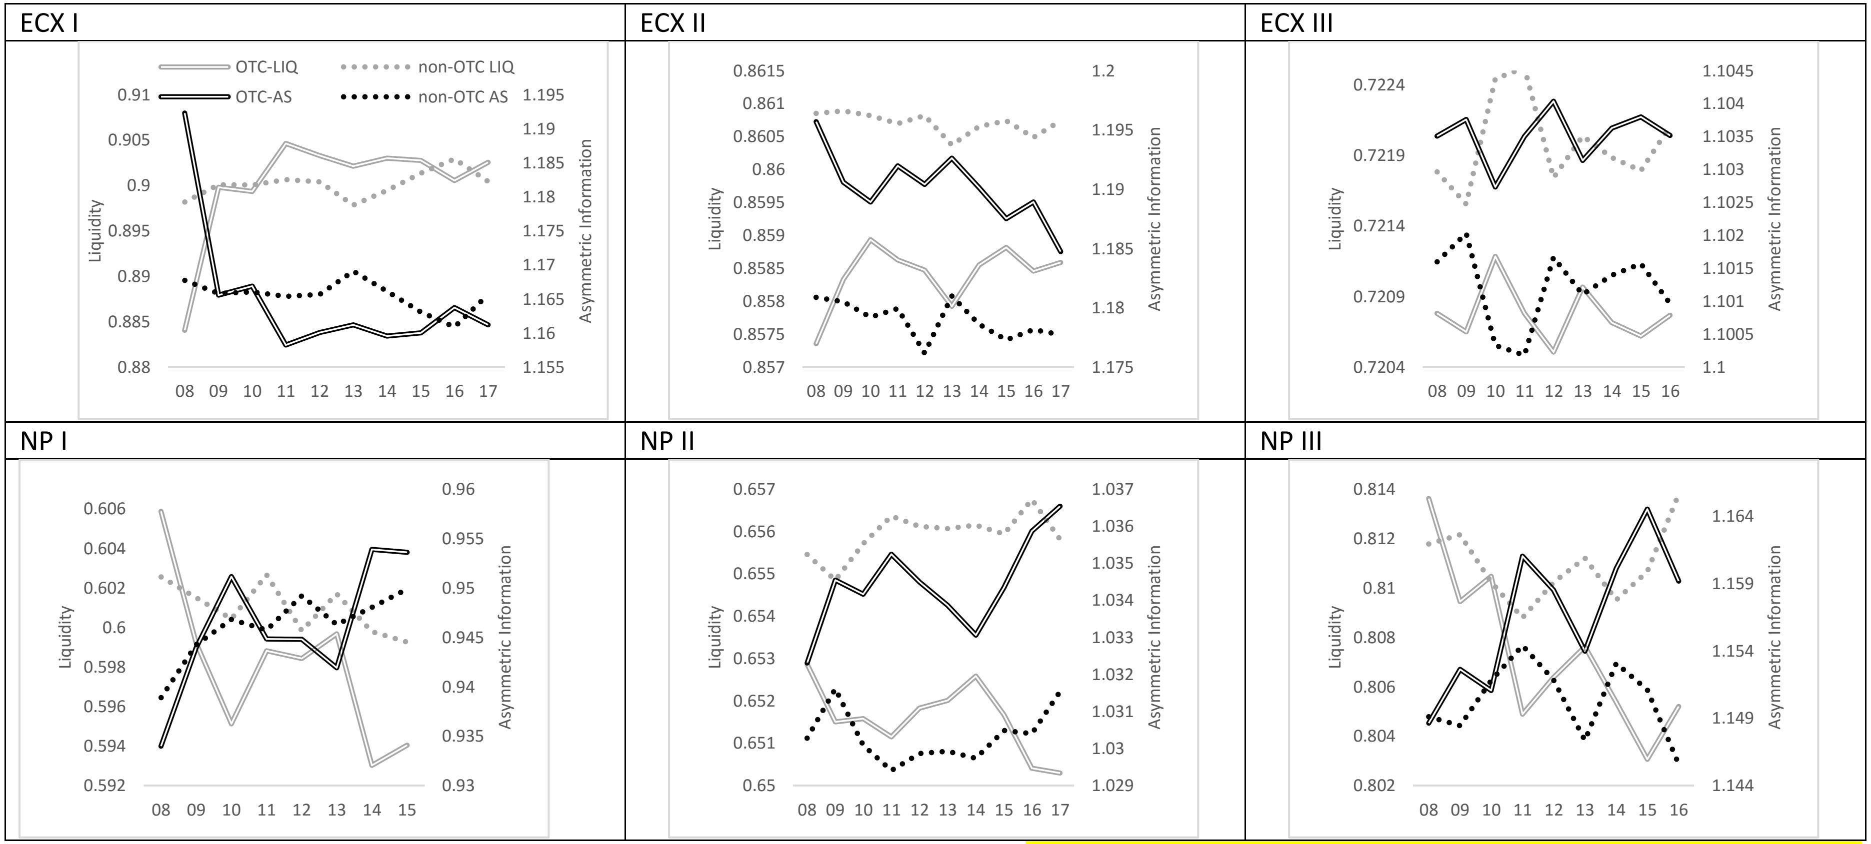

Figure 10 presents the hourly averages of the asymmetric information risk premium (*AIP*) and the liquidity risk premium (*LP*). These are denoted by *AS* and *LIQ* in the graphs, respectively. The legend in the first graph applies to all six. The black lines refer to asymmetric information (right axis), while the grey lines refer to liquidity (left axis).

Figure 11 Risk Premia across MSI for OTC, Non-OTC and All Trades

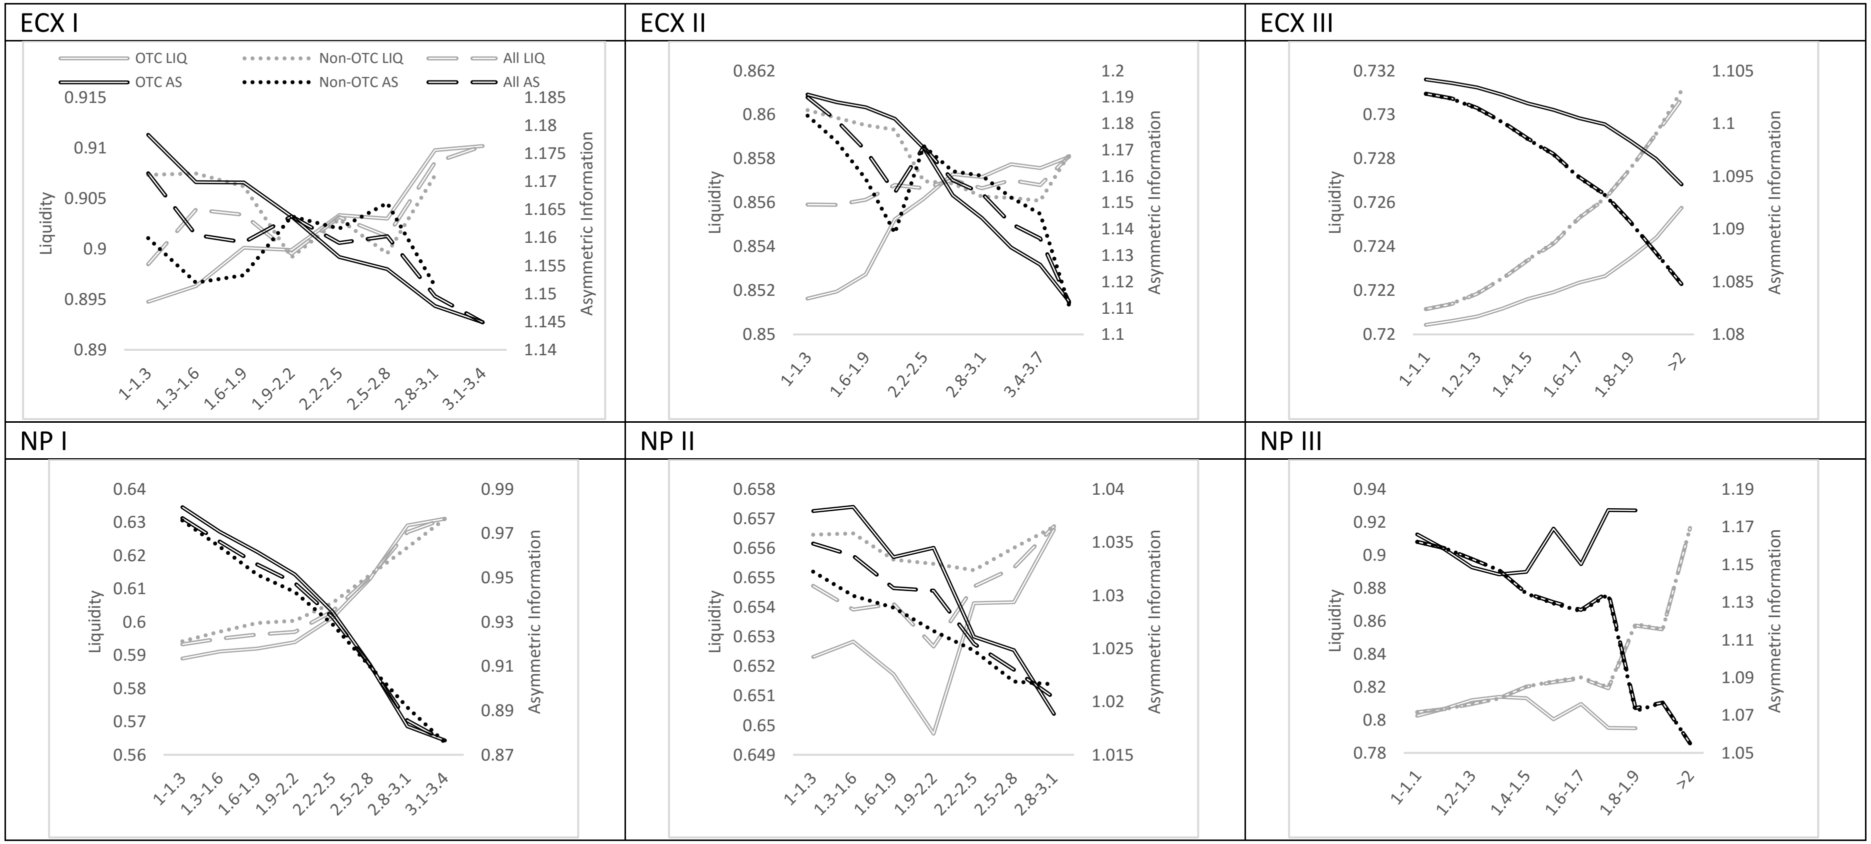

Figure 11 plots the average information risk premium (*AIP*) and the liquidity (*LP*) risk premium for OTC (solid), non-OTC (dotted) and all (dashed) trades, across different levels of MSI. The black lines refer to asymmetric information (right axis), while the grey lines refer to liquidity (left axis).

Appendix A (Online Supplement) – Not for publication

Diurnal Adjustment

First, each trading day is divided into five time intervals, each of two hours long. The nodes (time benchmarks) used are 09:00, 11:00, 13:00, 15:00 and 17:00. Second, to obtain $E\left(d_t \middle| f(t)\right)$, the raw duration ($d_t$) is regressed on the cubic-spline time function $f(t) = \beta_0 + \sum_{i=1}^{3} \sum_{j=1}^{5} \left(\beta_j \left(t - n_j\right)^i\right)$, where, j = 1, ..., 5 stands for the five nodes used, i = 1, 2, 3, and n_j's are five dummy variables calculated as $n_j = \begin{pmatrix} t - k_j, & if\ k_{j-1} < t < k_j \\ 0, & elsewhere \end{pmatrix}$. Third, given the estimates of β_j, the expected durations ($E\left(d_t \middle| f(t)\right)$) are calculated. Finally, the durations are normalised (diurnally adjusted) using the following ratios: $x_t \equiv d_t / E\left(d_t \middle| f(t)\right)$, where, x_t is the diurnally adjusted duration and $E\left(x_t \middle| f(t)\right)$ is the expected duration derived by regressing the raw durations on the cubic spline.

Estimation

Estimation is carried out by maximising the log-likelihood function using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) (1970) optimisation algorithm with numerical derivatives and robust errors to account for heteroscedasticity. The market gained complexity and liquidity gradually (Bredin et al., 2014; Medina et al., 2014). Consequently, the values of the trading intensity measure used consistently decrease over the years, introducing heteroscedasticity. Although, see Ibrahim and Kalaitzoglou (2016) for how the DJM model explains the majority of the GARCH effect in price change.

The estimation results of the STM-ACD and the version of the DJM model used in this paper are presented in Table A1 below. The top part of the table tabulates the estimation results of the STM-ACD model in both ECX and NP and in each phase. The specification is guided by Kalaitzoglou and Ibrahim (2013). Given the high number of observations, the parameters are significant as expected. We thus focus interpretation on relative values. First, estimates of *omega*, *alpha* and *beta* (which are similar to ARMA coefficients), confirm the stationary but high persistence and autocorrelation of duration. Decreasing sums of *alpha* and *beta* over the phases in ECX (NP) indicate slightly reducing (increasing) persistence over the phases. The differential effect of OTC trades, despite their lower number, is confirmed by significant values of *zeta* (see hypothesis test of zeta=0 in middle section of the table) that increase in magnitude from Phase I to Phase II in both markets, but declines in Phase III (negative in NP Phase III). The relative values of g_1 and g_2 over the phases (larger g_1 than g_2 in Phase I, but smaller in Phases II and III) confirm Kalaitzoglou and Ibrahim (2013) findings in Phases I and II (for only one contract) and extend it to all contracts and to Phase III. These values indicate a shift in the learning speed amongst agents across phases. The learning speed of discretionary liquidity traders is slow in Phase I but faster in Phases II and III. Kalaitzoglou and Ibrahim (2013) argue that this is likely to be the reason that forces informed traders to increase their trade size in the latter phases. The differences in estimates of the smoothness parameters s_1 and s_2, especially in Phases II and III, indicate a clear distinction between the three regimes of trading intensity that categorises trades into informed, discretionary liquidity traders, and uninformed. The significant differences (see hypothesis tests in the middle section of the table) in estimates of γ_1<1, γ_2>1 and γ_3=1 indicate a clear distinction in the behaviour of the agents and trichotomy of regimes. Note that in this analysis γ_3 stands for the shape parameter of the distribution for low intensity trades.

The bottom section of the table presents estimates of the version of the DJM model used in this paper. Estimates of θ_1 are insignificant, and hence not reported. θ_3 is significant only in Phase I, which indicates significant asymmetric information response to shocks in volatility in that phase only. Estimates of φ_1 are constant liquidity costs (e.g., order processing). Negative estimates of φ_2 indicate a negative effect of time varying liquidity on liquidity costs (i.e., lower immediacy costs with larger and faster trades). Estimates of φ_3 indicate that volatility shocks affected liquidity costs in Phase I only. Estimates of ρ, the trade-by-trade autocorrelation in order flow, indicate an average reversal of order flow from one trade to the next, and as estimates of ρ decrease in ECX but increase in NP from Phase I to Phase III, the probability of continuation (buys follow buys and sells follow sells) decreases in ECX but increases in NP. This reflects increasing reversals in ECX but decreasing reversals in NP.

Table A1 Estimation Results

	ECX			NP		
	I	II	III	I	II	III
omega	0.0108	0.0351	0.0122	0.0603	0.0405	0.0263
	(8.28)	(16.88)	(19.34)	(5.37)	(17.47)	(21.53)
alpha	0.0777	0.1873	0.2348	0.1148	0.0489	0.1086
	(17.81)	(40.29)	(29.72)	(10.80)	(25.13)	(2.63)
zeta	0.0148	0.0454	0.0054	0.0198	0.0454	-0.1656
	(6.33)	(16.73)	(2.67)	(2.60)	(14.16)	(-1.76)
beta	0.9075	0.7673	0.7706	0.8223	0.8963	0.8785
	(68.84)	(95.11)	(86.03)	(27.39)	(26.45)	(19.06)
g_1	1.4720	0.7428	0.8346	1.2211	0.6180	0.9508
	(15.96)	(11.91)	(13.21)	(16.36)	(19.05)	(14.07)
g_2	1.0011	1.0875	1.0449	0.7508	1.7520	1.3048
	(14.70)	(20.07)	(15.10)	(17.88)	(12.86)	(13.05)
s_1	0.5044	0.0441	0.3925	0.5208	0.4999	0.4232
	(16.24)	(22.33)	(21.73)	(12.66)	(29.69)	(19.90)
s_2	0.5812	0.3259	0.7043	0.6347	0.6998	0.6977
	(18.37)	(15.97)	(18.28)	(17.04)	(18.96)	(12.33)
γ_1	0.0027	0.6798	0.7387	0.2344	0.2412	0.4155
	(22.24)	(21.87)	(25.41)	(13.76)	(19.84)	(13.40)
γ_2	4.2813	1.9093	1.3934	4.1556	2.5024	3.3078
	(33.81)	(26.73)	(23.04)	(15.26)	(28.82)	(21.44)
γ_3	0.9797	0.9824	0.9716	0.9431	0.9044	0.8744
	(18.67)	(26.47)	(17.19)	(15.89)	(16.34)	18.3662
L	-12380.1	-375888.7	-183614.1	-4368.4	-15821.2	-1423.8
zeta=0	40.05	1644.62	58.84	12.57	32.23	3.09
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.08)
$\gamma_1=1$	2063.0	5812.7	34550.1	245.4	1421.4	99.3
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$\gamma_2=1$	110545.7	176906.9	20980.4	13430.5	28476.0	870.5
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
$\gamma_3=1$	3.86	5.06	3.53	7.94	8.75	16.26
	(0.05)	(0.03)	(0.06)	(0.01)	(0.00)	(0.00)
$\gamma_1=\gamma_2=\gamma_3=1$	1381332.2	261049.5	641122.2	30370.6	361633.4	14112.6
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
ϑ_2	0.0372	0.0136	0.0012	0.0836	0.0246	0.0161
	(6.34)	8.26	(3.26)	(2.53)	(3.59)	(2.88)
ϑ_3	0.0032			0.0057		
	(4.68)			(2.01)		
φ_1	0.0595	0.0038	0.0074	0.1050	0.0435	0.0305
	(4.12)	(9.58)	(5.31)	(3.31)	(3.77)	(2.79)
φ_2	-0.0261	-0.0013	-0.0010	-0.0742	-0.0217	-0.0149
	(-5.66)	(-7.87)	(-5.99)	(-2.39)	(3.38)	(-3.04)
φ_3	0.0102			0.0410		
	(2.95)			(2.95)		
ρ	0.4699	0.4225	0.3137	0.2098	0.2503	0.3879
	(66.84)	(298.77)	(160.01)	(27.85)	(33.52)	(25.50)

Table A1 presents coefficient estimates and *t*-stats for the estimates of the STM-ACD model, in both ECX and NP and in Phase I, II and III. The second panel lists the maximum likelihood value (L), as well as statistics and *p*-values of hypothesis tests on relevant coefficient values. The bottom panel reports the estimates of the DJM model.

Table A2 Tests on Differences in Premia between Low and High MSI

		Phase I	Phase II	Phase III
ECX	AIP	38.04	76.20	36.47
	LIQ	29.44	30.49	35.57
NP	AIP	24.59	7.055	20.35
	LIQ	14.28	1.99	32.72

Table A2 presents *t*-statistics testing the difference between premia at low versus high MSI