Who gains from market fragmentation? Evidence from the early stages of the EU carbon market*

Estelle Cantillon† and Aurélie Slechten‡

March 2023

Abstract

We document the impact of market fragmentation during the first phase of the EU emissions trading scheme on the terms that traders were able to get. We observe the universe of over-the-counter (OTC) and exchange transactions and the transaction prices associated with four of the 11 exchanges that were active during that period. We define a measure of price advantage based on the difference between the transaction price and the median market-wide price that day. We decompose price advantage on exchange, counterparty and trader characteristics and show that *where* traders traded and *how connected* they and their counterparties were with the rest of the market strongly impacted the terms they were able to obtain. Such features are expected to characterize OTC transactions but not exchange transactions. The high level of market fragmentation during the first phase, which was a policy choice, hampered information aggregation about the overall balance between supply and demand in the market, and put small and non-energy compliance traders at a large disadvantage.

Keywords: Trading networks, price formation, market frictions.
JEL codes: D47, D85, G12, Q58.

1 Introduction

The European Union Emissions Trading Scheme (EU ETS) is the largest carbon emissions market in the world. Now a fairly mature market, it had a bumpy start, in part due to the laissez-faire approach that the European Commission took to trading in allowances. The view then was that private actors would naturally step in to offer trading services and that, as a result, "the price of allowances [would] be determined by supply and demand as in any other market" (European Commission, 2005, p. 14). In practice, trading picked up slowly and remained highly fragmented for a long time.

*We thank seminar audiences at AERNA-CE3 Durham UAM, Lancaster, HEC Liège for their comments and suggestions. Stefan Bergheimer provided excellent research assistance. Financial support from FNRS (CDR grant J.0029.22) and the British Academy is gratefully acknowledged.

†FNRS, Université Libre de Bruxelles, and CEPR. Email: Estelle.Cantillon@ulb.be
‡Lancaster University. Email: a.slechten@lancaster.ac.uk

We build on the recent literature in finance on fragmented markets, and on transaction data during the first phase of the EU ETS to document the impact of market fragmentation on the price faced by market participants. A key advantage of our data is that we observe trader identities. The law of one price fails generically in this market. Prices on different trading venues reflect local supply and demand conditions, but also the position of the exchange in the network formed by the transactions between market participants. The trading terms that traders on exchanges get depend both on their connection to the rest of the market and the connections of the counterparty side has. Small compliance traders and compliance traders from non-energy sectors get a penalty. These results suggest that the fragmentation of the EU carbon market hampered price aggregration and put less connected, small and/or industrial market participants at a disadvantage.

Very few securities today trade in a single place, let alone on a single centralized exchange. This has long puzzled economists. Market fragmentation raises a number of normative questions such as its impact on information aggregation, allocative efficiency and redistribution in the presence of heterogeneous traders. The first phase (2005-07) of the EU carbon market provides a valuable setting to explore these questions. During that crucial period, trading was spread across 11 exchanges and over-the-counter. The market brought together compliance firms, who had to ensure they could cover their emissions over the past year by allowances, financial intermediaries and other smaller market participants. We know who traded, when, with whom, on which platform if any, and, for most exchange-mediated transactions, at what price.

Our dataset contain the universe of inter-firm transactions during the first phase. During that period about 56% of transactions were carried out over-the-counter, the rest on exchanges. Each trading venue attracted a very different pool of traders. Compliance traders made up 91.3% of market participants and mostly traded over-the-counter. 8% of compliance traders used exchanges and even there, diversity prevailed: some exchanges were better at attracting compliance traders, and the balance between supply and demand of allowances varied widely across exchanges. Overall, traders on exchanges tended to use more venues and were better connected to the rest of the market. Transaction sizes and prices tended to be higher over-the-counter.

We explore how the trading terms received by a trader in a particular transaction co-varies with exchange, counterparty and trader-specific factors. We focus on spot exchange-based transactions from June 2005 to May 2007, which represent about 36% of transactions, as these are the only transactions to which we can associate a transaction-specific price. Our main outcome of interest is the price advantage from which a transaction benefits, which we define, for a buy order, as the difference between the median market-wide price that day and the price obtained by the buyer, normalized by the median market-wide price. A positive price advantage means that the buyer obtained a better price than the median price prevailing on the market that day. In a frictionless centralized market, price advantage is solely driven by intraday variation in prices. In a fragmented market, deviations from market-wide prices can also be driven by exchange-specific factors, such as local supply and demand conditions, or counterparty and trader-specific characteristics.

We regress transaction-level buy-side (respectively, sell-side) market advantage on trader and exchange characteristics. Following the recent theoretical and empirical literature on over-the-counter markets, we use the network formed by the transactions between participants to proxy for their connectivity with the rest of the market. Every market participant is a node and two nodes are connected if they traded together within the last 12 months. We control for the centrality of the buyer (respectively, seller) in the network formed by market participants and exchanges, as well as the average centrality of counterparties on the exchange that day, on top of other more traditional trader characteristics. To account for local market conditions, we control for the connectedness of the exchange to the rest of the market and for the degree to which the profile of surplus allowances of exchange participants differs from the market-wide profile.

We find that the prices that traders get on exchanges depend on the local conditions (balance between supply and demand) prevailing on these exchanges, unless the exchange is well connected to the rest of the market. Counterparty characteristics also matter: better connected counterparties and a higher seller-to-buyer on an exchange significantly reduce the price advantage that a seller can get (the same result holds, *mutatis mutandis*, for buyers). Most interesting and surprising, a trader's connectivity with the rest of the market matters, even after controlling for all exchange and counterparty characteristics. The advantage obtained by individual connectivity is one order of magnitude lower than the advantage obtained from exchange and counterparty characteristics but is still in the range of effects found in over-the-counter market (of the order of 0.5 percentage point). Finally, we find that small compliance traders and compliance traders from non-energy sectors received significantly worse terms than other traders.

There are several ways to organize a market for carbon allowances. In Europe, a laissez-faire solution was chosen. Though the level of fragmentation observed in phase 1 has reduced over time, with many of the small exchanges exiting, trading in the secondary market remains a multi-venue affair with no single price. In Korea, market-makers are tasked to ensure sufficient liquidity at all times, thereby maintaining a level-playing field. In China, only compliance traders are allowed to participate in the physical market. Our paper suggests that these choices have consequences for market participants. The chosen market design for EU ETS created an unequal playing field and excess transaction costs for less central traders.

Our paper provides empirical evidence for recent theories on price formation in fragmented markets. Formally, two main approaches are used to model price formation in fragmented markets. One approach is to view them as OTC markets and model price formation as the result of random unstructured interactions between traders (in the spirit of search models) or structured interactions based on a network of relationships between traders. The search model approach views transactions as between atomistic dealers with non persistent links (Duffie et al., 2005). Network-based models are characterized by long-lived relations between traders (Babus and Kondor, 2018). An alternative approach is to explicitly model competing exchanges and traders' participation to them. A common finding of either approach approach is that the law of one price fails unless markets are "sufficiently connected".

We contribute to this literature by documenting the importance of trader centrality to determine the price advantage obtained by a trader on exchanges –a finding that is consistent with the predictions of network-based models of price formation in over-the-counter markets but are typically not expected in exchanges where trading is centralized and pricing anonymous. The pattern of transactions in the EU ETS captures some of the key elements of these models. Links between traders and exchanges are persistent. If we observe a transfer of allowances between two traders in 2005, the probability that we observe the same pair of traders in 2006 is around 70%. Moreover, 75% of the traders in exchange-based transactions only trade on one trading venue. In addition, the ranking in terms of centrality score for the exchanges and the main financial intermediaries is relatively constant during our sample period. There is a growing empirical literature on trading costs and centrality premium (Hollifield et al, 2017; Di Maggio et al. 2017; Li and Schürhoff, 2019; Kondor and Pinter, 2022, among others). We extend prior studies by focusing on exchange transactions and a newly created market with features that significantly differ from other financial markets (presence of highly heterogeneous traders, abatement costs, necessity to buy allowances for compliance purposes). Our results suggest that trading on exchanges does not differ that much from trading over-the-counter and, in particular, that individual trader and counterparty connectivities matter for the prices traders face.

Finally, we are not the first to exploit the transaction log of the EU ETS. The existing literature has documented several specificities of the market. First, transactions are highly seasonal and concentrated in April, the month when allowances need to be surrendered for compliance, and December (Martino and Trotignon, 2013). Second, participation by regulated firms is highly heterogeneous (Martino and Trotignon, 2013, Zaklan, 2013, Betz and Schmidt, 2016, Jaraité-Kažukauské and Kažukauskas, 2015, Abrell et al., 2022). Some firms (and industrial sectors) are very active while others barely interact or do not interact at all with the market. Reasons for limited participation include the design of the market that allowed firms to borrow the equivalent of one year of allowances, limited incentives for firms with surplus allowances to sell them, and prohibitive transaction costs faced by small firms with limited trading experience. Third, financial intermediaries and other non-compliance traders play an important role in this market (Martino and Trotignon, 2013, Borghesi and Flori, 2018). Fourth, large compliance traders tend to use exchange or banks for their transactions whereas small compliance traders tend to use brokers (Cludius and Betz, 2020). In other words, the *fragmentation* of the EU carbon market is well established. Because we are able to match the exchange-based transactions with a transaction-specific prices, we can go one step further and *quantify* the impact of this fragmentation on market participants.

2 The EU carbon market

The setting for our analysis is the first phase of the EU emissions trading scheme (EU ETS), which was established by the European Union as part of its commitment under the Kyoto Protocol. The EU ETS was officially launched in 2005 and is, to date, the largest emissions market in the world. The first phase of the EU ETS covered emissions in 2005-07.[1] During the first phase, close to 11,000 installations from the most energy intensive sectors in the economy

[1]The second phase covered 2008-12 and corresponded to the Kyoto Commitment. We are currently in the fourth phase.

Figure 1: Network of transactions (May 2005-April 2006)

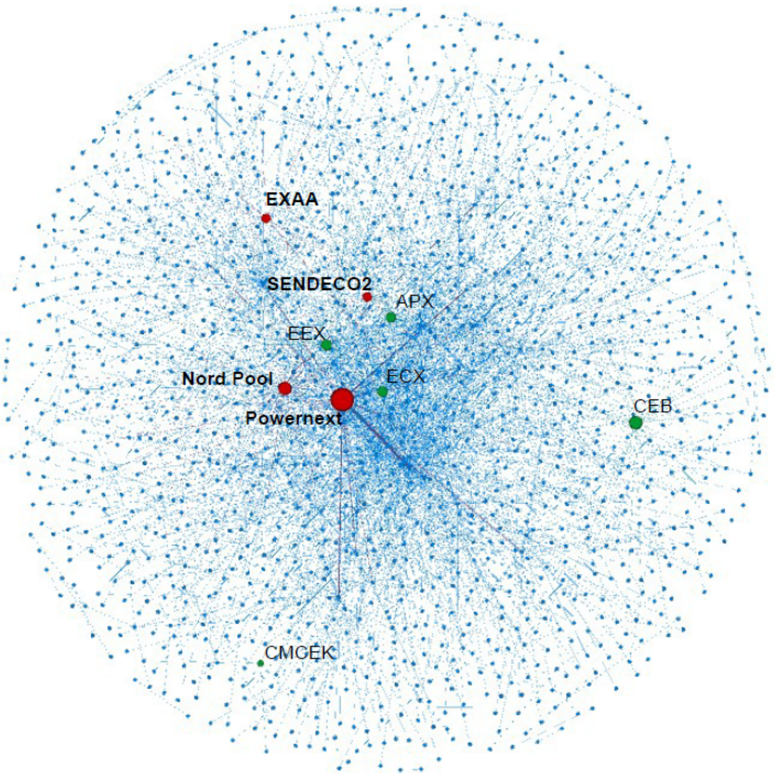

Notes: This graph represents the network of transactions as of April 2006. A trader is considered as active on the market if they traded at least once over the last 12 months. Each trader is a node and two nodes are connected if they have at least one transaction in common. The thickness of the edges depends on the number of transactions between the two nodes. All nodes representing traders are in blue and have the same size. The exchanges used in our analysis are in red, while the exchanges for which we do not have transaction-specific prices are in green. The size of the nodes representing exchanges is proportional to the number of transactions that happened on those exchanges between May 2005 and April 2006.

(electricity generation, basic chemistry, cement, steel, glass and ceramics, pulp and paper,...) received allowances to cover their emissions during the year, with the obligation to buy allowances on the market to cover any excess. National registries were set up to record ownership and transfers of these allowances.

Allocations of allowances for the whole phase were decided at the beginning of the phase, but allowances were actually distributed in three installments at the end of February of each year. Firms had until April 30 to surrender the allowances corresponding to the emissions of the previous year. Unused allowances could be banked for future years within the phase.

The market benefited from very little support, beyond the creation of allowances and the registries. As a result, a diverse set of financial intermediaries - brokers, dealers and exchanges - entered what was promising to be a major new market. During our sample period, 11 exchanges entered the market. Most of these exchanges were incumbent power exchanges, already offering trading services for the largest segment of compliance firms, namely electricity producers.[2] These include Amsterdam-based APX, Leipzig-based EEX, Oslo-based Nord Pool, Rome-based GME, Paris-based Powernext (later called Bluenext), Vienna-based EXAA, and Warsaw-based POLPX. The Czech (Czech Moravian Commodity Exchange Klano - CMCEK)

[2]Electricity producers represented approximately 60% of the emissions covered by the ETS at the time.

and Slovak (Commodity Exchange Bratislava - CEB) commodity exchanges also entered. Additionally, the market attracted new entrants. Spain-based SENDECO2 served the compliance needs of non-energy firms, with a focus on Southern Europe. The European Climate Exchange (ECX) offered trading in allowances futures. Some large financial institutions and even energy companies (for example, Electrabel, Shell, Statkraft) set up dedicated intermediation services to serve the nascent market.

The result was a highly fragmented market, weaving together centralized exchanges, dealers, brokers, other financial intermediaries and compliance firms along geographical and sector lines. Figure 1 represents the network of transactions over the 12 months period between May 2005 and April 2006. A node is a market participant and two nodes are connected on the graph if the two market participants transacted during the May 2005 - June 2006 period. The 9 exchanges present during that period are indicated by large dots, proportional to their volume of transactions during that period. The red dots correspond to the exchanges for which we have transaction-level prices and which are therefore the focus of our main analysis in Section 4. The graph was generated using the force-directed Fruchterman-Reingold algorithm which seeks to place connected nodes together and minimize the number of crossings among edges (Fruchterman and Reingold, 1991). This means that central market participants tend to be located closer to the center of the graph. Figure 1 confirms the fragmentation of the market and the absence of clear central market participants. Instead, exchanges (except for the smaller ones, EXAA, CEB and CMCEK) share the central spots with many other market participants.

Many of the exchanges that entered in Phase I remained small. By the end of the phase I, CMCEK and POLPX had left, Powernext had become the leading exchange for spot allowances, and ECX the leading exchange for futures. Much of trading in spot allowances remained over-the-counter.

3 Data and preliminary evidence

Our analysis covers spot transactions that took place during the first compliance phase of the EU ETS. Following Hintermann (2010) and Ballietti (2016), we restrict attention to transactions before May 2007 to avoid the period of very low prices at the end of Phase I.[3]

We use three sources of data. The first source is the Community Independent Transaction Log (CITL) which records every physical transaction that took place between market participants in the EU ETS. This dataset contains information about the identity of the buyer and the seller, a time stamp and the number of allowances exchanged. The second dataset is the national accounts dataset. Every market participant must hold an account to be able to buy and sell allowances. By default, every regulated installation is associated with a separate account but individuals or companies could easily open an account for trading. The accounts dataset provides information on the account holder, whether it is a compliance trader and if so, the associated installation, its sector, the number of free allowances received, its verified emissions and the number of allowances

[3]Given the non bankability of allowances into phase II and the revealed surplus in the market, prices dropped below 0.30 EUR/ton after May 2007 and never recovered.

Table 1: Trading Venue Characteristics

	Nb. Transactions	Volume (mtCO2)	Nb. Traders	Compliance Traders (%)	Net Surplus (mtCO2)	Venue Centrality	HHI (buy-side)
Powernext	397.2	4.47	35.3	54.3	27.63	1.01	0.41
EXAA	7.3	0.02	10.0	49.8	1.11	0.16	0.87
Nord Pool	43.8	3.05	41.3	65.9	19.46	0.76	0.68
SENDECO2	20.8	0.24	47.9	93.0	-2.75	0.13	0.73
CEB	10.5	0.14	24.0	77.4	0.56	0.07	0.94
CMCEK	2.3	0.00	3.3	43.8	1.30	0.01	0.36
APX	22.6	0.57	16.4	69.7	1.39	0.14	0.51
ECX	14.3	7.40	46.8	63.3	30.41	1.21	0.56
EEX	17.9	0.60	37.6	75.1	8.46	0.94	0.25
GME	5.3	0.12	4.7	77.8	-0.92	0.04	0.55
POLPX	1.8	0.03	3.4	49.0	0.04	0.00	0.56
OTC	576.6	25.93	1,187.8	91.2	94.18	-	0.32
Total	1,019.6	37.49	1,242.0	91.3	95.55	-	0.29

Notes: The unit of observation for this table is a *trading venue*$\times$*month* observation and the numbers correspond to averages over the sample period. The number of traders active on a trading venue is computed on the basis of traders who have traded on that trading venue in the past 12 months. Net surplus is computed as the accumulated net surplus of active traders. Venue centrality is measured by the eigenvector centrality of the exchange. The HHI for the buy-side is computed as the percentage of allowances purchased by each buyer on a trading venue during a specific day, squared and then summed across all buyers on that trading venue during that day. It takes value between $1/n$, where n is the number of active buyers that day (least concentrated) and 1 (most concentrated) (the HHI for the sell-side of the market presents a very similar pattern). Monthly averages are adjusted for the time of operations during our sample period: Powernext (2005m6-2007m5), CEB (2006m1-2007m5), CMCEK (2006m3-2006m6), APX (2005m6-2007m5), ECX (2005m12-2007m5), EEX (2005m4-2007m5), EXAA (2005m6-2007m5), GME (2007m3-2007m5), Nord Pool (2005m10-2007m5), POLPX (2006m9-2007m2), and SENDECO2 (2005m12-2007m5).

surrendered for compliance. The third source of data that we use are transaction-level price data provided by the exchanges and a daily price index for OTC transactions and for SENDECO2 transactions. The transaction and accounts datasets are public. The price data were public at the time (Powernext, EXAA, ECX, GME, Nord Pool, SENDECO2) or commercially available (Point Carbon, EEX).

To construct our final dataset, we aggregate accounts at the level of ownership to ensure we focus on transactions between independent companies rather than on internal transfers.[4] We also remove transactions in the CITL that correspond to initial allocations and surrenders of allowances for compliance.

We match transactions to transaction-specific prices wherever possible (specifically, transactions on Powernext, EXAA, ECX and Nord Pool). Transactions on SENDECO2 are associated with a SENDECO2-specific daily price. OTC transactions and transactions on POLPX and CEB are associated with the daily Point Carbon index. APX, CMCEK, EEX and GME required traders to deposit allowances prior to trading and transactions on these exchanges cannot, therefore, be matched to a price.

[4]To do this, we first use fuzzy matching (Levenshtein distance) based on the names, address and parent company, after converting everything into lower case letters and removing all punctuations, spaces and accents. We then search for accounts that could serve as dedicated trading desks for firms under common ownership and merge them with the aggregated account of these firms, the idea being that the transactions of interest are the transactions between the trading desk and third parties, whereas transactions between the trading desk and the account of the firms under common ownership are just internal transfers. The online appendix provides more detail on the data cleaning and construction.

Our final dataset contains 28,548 transactions, including 10,503 spot transactions on exchanges for which we have a transaction-specific price, 16,133 OTC transactions, and 1,912 other transactions, either transactions corresponding to deposits and transfers with GME, EEX, APX and CMCEK, settlement transactions associated with futures trading, or POLPX and CEB transactions for which we do not have transaction-specific price information.[5] During our sample period, 5,499 market participants are connected to the EU ETS, including 5,254 compliance traders and 11 exchanges. The remaining market participants are non-compliance traders, most of which financial intermediaries.

For each trader, we construct a monthly measure of their accumulated net surplus. For compliance traders, this is defined as the sum of free allocations minus surrender, net settlement of future allowances, and net purchases of spot allowances up to that month. Annual free allocations and surrenders are intrapolated at the month level. Likewise, settlements of future transactions are intrapolated at the month-level over a 12-month period for contracts with a maturity date in December, and 3-month period for contracts with a maturity date in March.[6] Non-compliance traders do not get free allowances nor are subject to surrenders, so we only use their net spot transactions and intrapolated settled transactions to compute their net positions.

We define a trader as being active on an exchange in a given month if they have traded on the exchange over the past 12 months.[7] This allows us to compute a monthly measure of accumulated net surplus at the level of each exchange (by summing the accumulated net surplus of the traders active on that exchange) and at the level of the market.

A major focus of our analysis is traders' centrality in this market. For each trader and exchange, we compute a monthly measure of their centrality using the graph-theoretical concept of eigenvector centrality on the network formed by the transactions that took place over the course of the past 12 months (a node is a trader or an exchange, and two nodes are connected if a transaction took place between these two nodes in the past 12 months). Eigenvector centrality measures the connectedness of a trader or an exchange in a network by accounting both for the number of traders with which they traded and the connectedness of these trading partners.[8] An eigenvector is defined up to a constant, which implies that the centrality scores can only be used to compare nodes within the same network. While there were some small independent networks of OTC transactions during our sample period, all the exchanges and close to 95% of traders belonged to the same main network. This network is the one we use for computing exchange and trader network centrality. To provide a basis for comparison across time, we normalize the eigenvector centrality such that the sum over all nodes in the network is 100.

Table 1 provides descriptive statistics for the 11 exchanges that operated during phase I and for OTC transactions (the top panel describes the four exchanges that will be used in the core of

[5]At the settlement of futures contracts, allowances change hands and this generates a transaction in our data. However, the price associated with these futures at maturity is the same for all transactions and is therefore not informative of the trading terms that the trader received originally when the position was open.

[6]We include futures positions in the computation of accumulated net surplus because they represent future commitments to buy or sell.

[7]The 12-month window is motivated by the low frequency of trades in this market.

[8]The literature identifies a number of centrality measures for trading networks (see e.g. Hollifield et al., 2017, Li and Schürhoff, 2019, Kondor and Pinter, 2022). We chose to go for eigenvector centrality because it weighs important trading partners more.

Table 2: Transaction and trader characteristics

Panel A: Spot exchanges with transaction-specific prices							
	N	mean	0.05	0.25	0.50	0.75	0.95
Transaction characteristics							
Size (10,000tCO2)	10,503	1.15	0.20	0.50	1.00	1.00	3.00
Price (EUR/t)	10,503	14.03	0.90	6.60	14.59	21.90	26.95
Price Advantage	10,503	-0.07	-5.75	-1.04	0.00	1.01	5.33
Buyer is a compliance trader	5,319	0.68	0.00	0.00	1.00	1.00	1.00
Seller is a compliance trader	5,184	0.45	0.00	0.00	0.00	1.00	1.00
Trader-month characteristics							
Nb transactions per month	2,034	9.28	0.06	0.25	2.40	13.75	37.67
Accumulated net surplus (mtCO2)	2,034	0.34	-0.96	-0.00	0.00	0.11	2.83
Trader centrality	2,034	0.32	0.01	0.01	0.05	0.42	1.48
Trader characteristics							
Also trading OTC	197	0.56	0.00	0.00	1.00	1.00	1.00
Multi-exchange trading	197	0.25	0.00	0.00	0.00	1.00	1.00
Panel B: OTC market							
	N	mean	0.05	0.25	0.50	0.75	0.95
Transaction characteristics							
Size (10,000tCO2)	16,133	4.49	0.07	0.50	1.02	3.00	15.00
Price (EUR/t)	15,381	14.39	0.78	6.78	15.13	22.63	27.18
Buyer is a compliance trader	16,133	0.51	0.00	0.00	1.00	1.00	1.00
Seller is a compliance trader	16,133	0.74	0.00	0.00	1.00	1.00	1.00
Trader-month characteristics							
Nb transactions per month	33,228	1.13	0.05	0.10	0.20	0.50	4.00
Accumulated net surplus (mtCO2)	33,228	0.08	-0.11	-0.00	0.00	0.02	0.34
Trader centrality	33,228	0.08	0.00	0.00	0.02	0.06	0.28
Trader characteristics							
Also trading on exchanges	2,744	0.08	0.00	0.00	0.00	0.00	1.00

Notes: The unit of observation for transaction characteristics is a transaction (exchange-based transactions have the exchange as one of the counterparties). The unit of observation for trader-month characteristics is a trader either trading on an exchange (panel A) or trading over-the-counter (panel B) in the past 12 months. The unit of observation for trader characteristics is a trader who has been active any time during our sample period on an exchange (panel A) or over-the-counter (panel B).

our analysis). Trading was fragmented: 56% of transactions and 69% of transaction volume took place over the counter. The remainder was split across 11 exchanges. The table documents large differences across trading venues in the number of transactions, trading volumes and number of active traders. Powernext was by far the largest trading venue by number of transactions but ECX, whose spot transactions correspond to settlement transactions at maturity, dominated in terms of volumes. As already suggested by Figure 1, Powernext , ECX and EEX were the exchanges best connected to the rest of the market, based on their eigenvector centrality.

Trading venues also differed in the characteristics of traders they attracted. Compliance traders made the bulk (91.3%) of market participants overall but they used exchanges less than non-compliance traders and tended to stick to a single trading venue, unlike non-compliance traders. This explains why they accounted for a lower proportion of market participants on exchanges. An exception is SENDECO2, a trading platform specifically dedicated to serve the compliance needs of non-energy traders. This specific positioning is also reflected in the net surplus numbers, which is negative for SENDECO2, unlike for other exchanges except GME. Powernext had one of the lowest fraction of compliance traders and also the lowest level of concentration.

Table 2 provides descriptive statistics on transactions and traders present in our data, distinguishing according to whether the transaction occurred on an exchange for which we have transaction-specific prices (panel A, which will be our main dataset moving forward) or over the counter (panel B) (for brievety we do not report information on the 1,912 exchange transactions for which we do not have transaction-specific prices or which correspond to settlements of futures).

The range of realized prices and of transaction sizes is larger on the OTC market than on exchanges. Realized prices are also slightly higher on the OTC market. Consistent with the existing literature (Zhu, 2014, Degryse et al., 2015), the OTC market seems to attract less experienced (lower frequency of trades) and less informed (less connected) traders. 92% of them only trade on the OTC market. Traders on exchanges, on the other hand, "multi-home" more: a quarter trades on multiple exchanges and 56% is also active on the OTC market (and actually - not reported in the Table - are involved in 74% of OTC transactions). Trading patterns are persistent. If we observe a transfer of allowances between two traders in 2005, the probability that we observe the same pair of traders in 2006 is around 70%.

4 Determinants of price advantage

In a frictionless centralized market, we expect the law of one price to hold and transaction prices to differ at most by the bid-ask spread. This is no longer true in fragmented markets. The theoretical and empirical literature has identified a number of covariates of realized prices in fragmented markets. We explore these relationships in our sample of exchange-based spot transactions (panel A of Table 2).

Our main object of interest is the price advantage that a trader is able to obtain for their transaction, which we define (for a seller) as the difference between the price they got and the hypothetical market-wide frictionless price that day.[9] Formally, for each transaction by trader i on exchange k and day t, we define the price advantage of this transaction as:[10]

$$\text{Adv}_{ikt} = 100\frac{(p_{ikt} - \bar{p}_t)1_{i\in S} + (\bar{p}_t - p_{ikt})1_{i\in\mathcal{B}}}{\bar{p}_t}$$

where p_{ikt} denotes the transaction price and $\overline{p}_t$ denotes the hypothetical market-wide frictionless transaction price of day t. We partition exchange-based transactions according to whether the trader is on the sell side ($\mathbb{S}$) or on the buy side ($\mathcal{B}$).[11] The variable Adv takes positive values when the trader trades on favorable terms relative to the rest of the market. It takes negative values otherwise. We proxy the market-wide transaction price by the volume-weighted median transaction price of the day based on the exchange spot transactions with a transaction-specific price and OTC transactions (panels A and B sample). We normalize the price advantage by the market-wide price to account for the non-stationarity of prices over the trading phase. In our data, the observed price advantage on the exchanges typically lie within 5 percent of the market-wide median transaction price (Table 2).

The existing literature on price formation in segmented and decentralized (OTC) markets provides some indication about the way price advantage covaries with exchange and trader characteristics and motivates the following empirical specification:

$$\text{Adv}_{ikt} = \alpha_{type(i)} + \beta X_{kt} + \gamma Z_{it} + \delta W_{ikt} + \epsilon_{ikt} \tag{1}$$

where X_{kt} is a vector of exchange-specific covariates, Z_{it} contains trader-specific covariates and the remaining terms collect observable and unobservable transaction-specific covariates.

Exchange-specific covariates of price advantage. When trades are distributed across different exchanges with no connection among them, local prices will reflect local conditions and, in particular, the existing balance between supply and demand (Jensen, 2007). We proxy local market conditions by the difference between the average accumulated net surplus of traders active on the exchange (S_{kt}) and its market-wide equivalent ($\bar{S}_t$): $\text{local\_mkt\_conditions}_{kt} = (\bar{S}_t - S_{kt})1_{i\in S} + (S_{kt} - \bar{S}_t)1_{i\in\mathcal{B}}$. An increase in this variable indicates more favorable local market conditions for traders on exchange k, relative to the rest of the market.

There are countervailing forces, however. When price information is sufficiently well distributed, either by design (consolidated tape) or because some traders multi-home and are able to arbitrage

[9] Hollifield et al. (2017), Li and Schürhoff (2019) or Di Maggio et al. (2017) use measures of dealers markups in OTC markets (differences between dealers buying and selling prices, in percentage). Our measure is closer to the one used in Kondor and Pinter (2022) who are studying clients' trading performance in the UK government bond market. Their trading performance measure has two components: anticipation component (ability to anticipate future price changes) and transaction component (difference between transaction price and the average transaction price of all dealer-client trades around the time of the transaction). Our objective is to understand whether some traders (depending on their characteristics or their chosen trading venue) were able to trade on more favorable terms within the same day (second component of the measure of Kondor and Pinter, 2022) .

[10] In principle, the same trader may make several transactions on an exchange in a given day so the triplet (i, k, t) does not uniquely define a transaction. We keep this notation in the text for expositional simplicity but do take the unique transaction level as the unit of analysis in the regressions.

[11] Exchange-based transactions in our data have the exchange on one of the side of the trade. Buy-side transactions are transactions where the trader is the buyer and the exchange appears as the counterparty. Sell-side transactions are transactions where the trader is the seller.

across the different venues, prices tend to converge across trading venues and reflect market-wide conditions (Barclay et al., 2008, Brogaard et al., 2014). We capture these ideas by controlling for exchange eigenvector centrality, a proxy for market connectedness, and local market conditions interacted with the exchange eigenvector centrality.

Counterparty-specific covariates of price advantage. The trading terms obtained by market participants on the different exchanges may also depend on the mix of traders on those exchanges, independently of the level of information fragmentation or local market conditions, for example because of market power or because of the market design of the trading venue. We account for these effects by allowing for exchange fixed efffects and controlling for the ratio of sellers to buyers, the Herfindahl Hirschmann Index and the average centrality of the counterparty side.

Trader-specific covariates of price advantage. The recent literature on OTC markets has suggested that the terms that traders get depend on their centrality in the network of all market participants and the centrality of their counterparty (Babus and Kondor, 2018). In the context of emissions markets, traders' bargaining power also depends on their commitments (emissions and allowances surrenders or settlements of futures contracts). We control for both traders' centrality and accumulated net surplus. To account for time-invariant trader characteristics, we include trader type fixed effects ($\alpha_{type(i)}$) and, specifically, distinguish between compliance traders in the energy sector (the largest and most active group), compliance traders outside of the energy sector, and non-compliance traders. We also distinguish between small compliance traders and large compliance traders, based on their initial allocation of allowances.[12]

Other controls. We control for the size of the transaction, W_{ikt}, as earlier research has found that it is correlated with the markups charged by traders (Li and and Schürhoff, 2019, Di Maggio et al., 2017). To account for market-wide drivers, we allow for month fixed effects and adjust standard errors for heteroskedasticity and clustering at the transaction day level.

Note that most of the covariates in (1) are invariant at the month level, whereas the dependent variable varies both within and across days. Our normalization of price advantage by daily prices helps account for some of the within month variation. The rest will be captured by the error term clustered at the day level. Table 4 in the Appendix provides descriptive statistics for all regression variables.

Table 3 summarizes the results separately for the buy-side and the sell-side. Most coefficients have the expected sign and, when this is not the case, they are not statistically significant. The results indicate that price advantage covaries with exchange *and* trader characteristics.

First, local market conditions and exchange centrality covary with the price advantage that traders are able to get. Favorable local market conditions are associated with a higher price advantage (first row) but this effect is neutralized if the exchange is well connected with the rest of the market (second row, interaction term, recalling that in our sample the most connected exchange has an average eigenvector centrality of 1). Exchange centrality reduces buy-side price advantage and increases sell-side price advantage (third row). This is a mechanical consequence

[12]In our main regressions, we use a cutoff of 1 million tCO2, which corresponds to the top quintile of intial allocations and approximately 95% of compliance traders' transactions in sample A.

Table 3: Regressions of the transaction price advantage on exchanges and traders characteristics.

	Trader is on the					
	Buy-side			Sell-side		
	(1)	(2)	(3)	(4)	(5)	(6)
Exchange characteristics						
Local Mkt Conditions	18.113***	18.211***	17.990***	14.300**	11.194**	11.186**
	(5.368)	(5.018)	(5.062)	(5.777)	(5.319)	(5.289)
Local Mkt Cond. × centrality	-19.333***	-19.445***	-19.141***	-14.293**	-11.055*	-11.074*
	(6.103)	(5.687)	(5.731)	(6.359)	(5.874)	(5.837)
Exchange centrality	-10.048***	-8.902***	-8.870***	6.383**	5.182**	5.029**
	(2.604)	(2.317)	(2.290)	(2.711)	(2.413)	(2.398)
Counterparty characteristics						
Nb. Sellers / Nb. Buyers	8.262***	6.333***	6.013***	-3.502*	-2.229	-2.057
	(2.225)	(2.076)	(2.024)	(1.862)	(1.709)	(1.677)
Counterparty average centrality	-1.159*	-0.930	-1.015*	-1.602***	-1.342***	-1.302***
	(0.594)	(0.601)	(0.522)	(0.423)	(0.399)	(0.405)
counterparty HHI			0.998			-1.047
			(1.374)			(1.054)
Powernext	5.933*	5.454*	5.213	-2.773	-5.647**	-5.058**
	(3.210)	(3.181)	(3.384)	(2.252)	(2.457)	(2.467)
EXAA	-0.846	0.881	-0.022	1.880	-1.397	-0.384
	(2.321)	(2.549)	(3.429)	(1.386)	(1.790)	(1.948)
Nord Pool	3.375	3.719	2.963	-1.149	-4.399*	-3.386
	(3.555)	(3.540)	(4.218)	(2.330)	(2.606)	(2.754)
SENDECO2	-10.355***	-7.629**	-8.450**	6.378**	4.394	5.365*
	(3.389)	(3.390)	(3.997)	(2.930)	(2.914)	(3.034)
Trader characteristics						
Trader centrality		0.422***	0.423***		0.438***	0.445***
		(0.099)	(0.099)		(0.158)	(0.161)
Trader Surplus (mtCO2)		0.016	0.015		0.067*	0.070*
		(0.027)	(0.027)		(0.040)	(0.040)
Small compliance trader		-2.298***	-2.264***		-1.865**	-1.830**
		(0.820)	(0.818)		(0.730)	(0.724)
Energy Sector		0.963	0.994*		1.635*	1.591*
		(0.594)	(0.601)		(0.877)	(0.869)
Non-compliance traders		0.748	0.786		1.806**	1.759**
		(0.600)	(0.603)		(0.894)	(0.886)
Transaction characteristics						
Transaction Vol. (log)	-0.075	-0.167	-0.170	0.092	0.079	0.071
	(0.129)	(0.123)	(0.121)	(0.127)	(0.127)	(0.127)
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,319	5,319	5,319	5,184	5,184	5,184
R-squared	0.158	0.168	0.169	0.120	0.130	0.131

Notes: This table presents the results of the estimation of equation (1) for exchange-based transactions where the trader is on the buy-side (columns 1-3) or the sell-side (columns 4-6) between June 2005 and May 2007. Robust standard errors clustered at the day level are shown in parentheses. Asterisks denote significance levels: *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

of the fact that prices on exchanges tend to be lower than on the OTC market, and exchange connectedness brings prices across trading venues closer to one another: the resulting higher prices reduce exchange buyers' advantage (negative coefficient) and increase sellers' advantage (positive coefficient).

Both effects are economically significant. Holding exchange centrality fixed at its mean sample value, an improvement of one standard deviation in local market conditions is associated with a 4.25 percentage point (p.p.) increase in buyer's advantage. This advantage is larger for less connected exchanges. Likewise, holding local market conditions fixed at their mean sample value, buyers on less connected exchanges (eigenvector centrality around 0.15) benefit from an additional price advantage of the order of 3 p.p. relative to the better connected exchange (eigenvector centrality of 1).

Second, the mix of traders on an exchange matters beyond their aggregate net surplus, and is suggestive of the presence of market power: A higher seller-to-buyer ratio is advantageous for buyers (the effect for sellers is negative but not statistically significant) and more central counterparties are associated with smaller price advantages. These effects are also economically significant. An increase of one standard deviation in the seller-to-buyer ratio is associated with a 2 p.p. increase in the buyer's advantage. Trading on Powernext is associated with an additional buyer advantage and seller disadvantage of the order of 5 p.p. This may reflect the fact that traders on Powernext had on average a large net accumulated surplus (Table 1) which enabled them to be more strategic about when and at what price to buy. Reversely, trading on SENDECO2 is associated with a large buyer disadvantage (between 7.6 to 10.4 p.p. depending on the specification) and a seller advantage, possibly reflecting the high fraction of non-energy compliance traders on SENDECO2.

Looking at trader characteristics reinforces the picture that traders' relative position matters for the trading terms they get on exchanges. First, trader centrality is statistically and economically significant. A trader with an eigenvector centrality one standard deviation above the mean, is associated with a 0.22 p.p. improvement in price advantage. The top 5% traders in terms of eigenvector centrality get a 0.49 p.p. price improvement. This is consistent with Hollifield et al. (2017)'s findings that core dealers (top 5% traders in terms of eigenvector centrality) in securitization markets deliver price improvements between 0.40 and 0.64 p.p. to their clients. Second, our results indicate that small compliance traders suffered a price disadvantage of the order of 2 p.p. Third, compliance traders from the energy sector and non-compliance traders benefit from a 1.6-1.8 p.p. advantage premium relative to compliance traders from non-energy sectors (the omitted category in the regressions).

5 Discussion

The EU carbon market was very fragmented during its first phase and our results show that this had consequences: prices systematically differed across trading venues and traders, reflecting both local exchange conditions and traders' characteristics. These findings shed light on our understanding of financial market fragmentation, on the one hand, and on the design of emissions markets, on the other hand.

The literature on market fragmentation typically distinguishes between over-the-counter trading, where prices depend on traders' identity, and situations where trading is split across multiple trading venues, each characterized with centralized, anonymous, pricing. In practice, these two modes of trading coexist in many markets and our results indicate that the boundary between the two is not as clearcut as previously thought: trading on exchanges displays some of the patterns typically associated with over-the-counter trading, namely, better connected traders getting better terms. The picture that emerges, therefore, is more one of a continuum of trading mechanisms, where exchanges provide vehicles to pool information and connectivity from many traders and reduce - but don't eliminate - idyosyncratic advantage. The centrality premium that traders were able to obtain on exchanges during the first phase of the EU ETS (around 0.5 p.p. for the better connected traders) is small relative to the exchange-specific advantage they got, but is not negligible, and it is aligned with centrality premia found in OTC markets (see e.g. Hollifield et al., 2017).

Our results also bear lessons for the design of emissions markets. The central objective of emissions markets is to encourage the efficient allocation of abatement efforts across the firms subject to the regulation through the generation of an informative price signal. Firms with cheaper abatement opportunities than the going price will prefer to abate. Firms with higher abatement costs will prefer to buy emission allowances. Market frictions increase price volatility and hinder the efficient allocation of abatement, reducing the cost effectiveness of emissions trading as a regulatory instrument.

Our findings show that the laissez-faire approach to market development that the EU took for its emissions trading scheme hampered the ability of market participants to get a full picture of the prevailing balance between supply and demand in the market, and failed to ensure an equal playing field among traders, and singularly, compliance traders. The vast majority of compliance traders used the over-the-counter market where prices tended to be higher, on average, than on exchanges. But, even on exchanges, prices differed systematically, in a way that penalized smaller compliance traders and compliance traders from the non-energy sectors.

Emissions trading schemes are *designed* markets. Different jurisdictions have made other choices regarding who has access to their markets and how trading is organized. In the Korea emissions trading scheme (ETS), spot transactions take place over-the-counter or on the Korea Exchange (KRX), where designated market makers ensure a level-playing field for all traders. In the Chinese ETS, allowances are exclusively traded on the Shanghai Environment and Energy Exchange and non-compliance firms are excluded. In California, spot allowances are traded over-the-counter but they coexist with quarterly auctions run by the California Air Resources Board (CARB) that serve as the primary market. It is an open question to what extent these different designs facilitated participation and price discovery by compliance traders.[13]

Today's EU carbon spot market has consolidated somewhat. There are three exchanges left

[13] Joskow et al. (1998) provide an early study of how market support mechanisms can help market participants in an emissions market discover the equilibrium price.

serving the market (ICE Endex, EEX and Nasdaq Oslo), each offering daily futures, a close substitute to spot allowances. Allowances are also auctioned daily by the EEX as part of the primary market and the OTC market, which represented close to 70% of trading volumes in phase I, now only represents around 15%. Concerns remain, however, regarding the market's ability to provide a level-playing field (see e.g. ESMA (2022)'s review of the market).

References

[1] Abrell, J., Cludius, J., Lehmann, S., Schleich, J. and Betz, R. (2022). Corporate emissions-trading behaviour during the first decade of the EU ETS. *Environmental and Resource Economics*, 83(1), 47-83.

[2] Adamic, L., Brunetti, C., Harris, J. H., & Kirilenko, A. (2017). Trading networks. The Econometrics Journal, 20(3), S126-S149

[3] Babus, A. and Kondor, P. (2018). Trading and information diffusion in over-the-counter markets. *Econometrica*, 86(5), 1727-1769.

[4] Balietti, A.C. (2016). Trader types and volatility of emission allowance prices. Evidence from EU ETS Phase I. *Energy Policy*, 98, pp.607-620.

[5] Barclay, M.J., Hendershott, T. and Jones, C.M. (2008). Order consolidation, price efficiency, and extreme liquidity shocks. *Journal of Financial and Quantitative Analysis*, 43(1), 93-121.

[6] Betz, R.A. and Schmidt, T.S. (2016). Transfer patterns in Phase I of the EU Emissions Trading System: a first reality check based on cluster analysis. *Climate Policy*, 16(4), 474-495.

[7] Borghesi, S. and Flori, A. (2018). EU ETS Facets in the Net: Structure and Evolution of the EU ETS Network. *Energy economics*, 75, 602-635.

[8] Brancaccio, G., Li, D., & Schürhoff, N. (2020). Learning by trading: The case of the US market for municipal bonds. Mimeo.

[9] Brogaard, J., Hendershott, T., & Riordan, R. (2014). High-frequency trading and price discovery. *The Review of Financial Studies*, 27(8), 2267-2306.

[10] Chen, Daniel and Darrell Duffie (2021), Market fragmentation, American Economic Review, 111(7), 2247-74

[11] Cludius, J. and Betz, R. (2020). The role of banks in EU emissions trading. *The Energy Journal*, 41(2), 275-300.

[12] Degryse, H., De Jong, F. and Kervel, V.V. (2015). The impact of dark trading and visible fragmentation on market quality. *Review of Finance*, 19(4), 1587-1622.

[13] Di Maggio, M., Kermani, A. and Song, Z. (2017). The value of trading relations in turbulent times. *Journal of Financial Economics*, 124(2), 266-284.

[14] Donna, Javier, Pablo Schenone, Gregory Veramendi (2020), Networks, frictions, and price dispersion, Games and Economic Behavior, 124, 406-31

[15] Duffie, D., Gârleanu, N. and Pedersen, L.H. (2005). Over-the-counter markets. *Econometrica*, 73(6), 1815-1847.

[16] European Securities and Markets Authority (ESMA). (2022). Emission allowances and associated derivatives, final report, ESMA70-445-38.

[17] European Commission - Directorate-General for Environment (2005), EU Emissions Trading: An open scheme promoting global innovation to combat climate change, Publications Office.

[18] Fruchterman, T. M. J., and Reingold, E. M. (1991). Graph Drawing by Force-directed Placement, *Software-Practice and Experiences*, 21 (11), 1129-1164.

[19] Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental Economics and Management*, 59(1), 43-56.

[20] Hollifield, B., Neklyudov, A., and Spatt, C. (2017). Bid-ask spreads, trading networks, and the pricing of securitizations. *The Review of Financial Studies*, 30(9), 3048-3085.

[21] Jaraité-Kažukauské, J. and Kažukauskas, A. (2015). Do transaction costs influence firm trading behavior in the European Emissions Trading System. *Environment and Resource Economics*, 62, 583-613.

[22] Jensen, R. (2007). The digital provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector. *Quarterly Journal of Economics*, 122(3), 879-924.

[23] Joskow, P. L., Schmalensee, R., and Bailey, E. M. (1998). The market for sulfur dioxide emissions. *American Economic Review*, 88(4), 669-685.

[24] Kondor, P. and Pinter, G. (2022). Clients' Connections: Measuring the Role of Private Information in Decentralised Markets. *The Journal of Finance,* 77(1), 505-544.

[25] Li, D. and Schürhoff, N. (2019). Dealer Networks. *The Journal of Finance*, 74(1), 91-144.

[26] Martino, V., and Trotignon, R. (2013). Back to the future: A comprehensive analysis of carbon transactions in Phase 1 of the EU ETS. *Cahiers de la Chaire Economie Climat*, 27, Université Paris Dauphine.

[27] Zaklan, A. (2013). Why Do Emitters Trade Carbon Permits? Firm-Level Evidence from the European Emission Trading Scheme, DIW working paper 1275.

[28] Zhu, H. (2014). Do dark pools harm price discovery?. *The Review of Financial Studies*, 27(3), 747-789.

Appendix

Table 4: Descriptive Statistics for regression variables

Panel A: Trader is on the buy-side	unit of obs.	N	mean	SD	0.50	min	max
Local Mkt Cond. (mtCO2)	exchange-month	82	-0.27	0.50	-0.17	-1.96	0.29
Exchange centrality	exchange-month	82	0.55	0.41	0.68	0.00	1.54
Nb. Sellers / Nb. Buyers	exchange-month	82	1.08	0.34	1.06	0.20	2.00
Counterparty average centrality	exchange-day	816	0.85	0.71	0.72	0.00	4.45
Counterparty HHI	exchange-day	816	0.63	0.33	0.56	0.10	1.00
Trader centrality	trader-month	638	0.59	0.69	0.29	0.00	4.45
Trader Surplus (mtCO2)	trader-month	638	0.64	2.46	0.03	-4.84	16.97
Small compliance traders	trader	143	0.50	0.50	0.00	0.00	1.00
Energy Sector	trader	143	0.25	0.44	0.00	0.00	1.00
Non-compliance traders	trader	143	0.30	0.46	0.00	0.00	1.00
Transaction Vol. (10,000 tCO2)	transaction	5,319	1.15	1.38	1.00	0.00	32.15
Panel B: Trader is on the sell-side							
	unit of obs.	N	mean	SD	0.50	min	max
Local Mkt Cond. (mtCO2)	exchange-month	81	0.28	0.50	0.18	-0.29	1.96
Exchange centrality	exchange-month	81	0.55	0.41	0.68	0.02	1.54
Nb. Sellers / Nb. Buyers	exchange-month	81	1.07	0.33	1.06	0.20	2.00
Counterparty average centrality	exchange-day	769	0.86	0.61	0.82	0.00	4.45
Counterparty HHI	exchange-day	769	0.59	0.31	0.51	0.14	1.00
Trader centrality	trader-month	655	0.61	0.69	0.33	0.00	4.45
Trader Surplus (mtCO2)	trader-month	655	0.54	1.92	0.05	-5.11	12.68
Small compliance traders	trader	124	0.42	0.50	0.00	0.00	1.00
Energy Sector	trader	124	0.38	0.49	0.00	0.00	1.00
Non-compliance traders	trader	124	0.34	0.48	0.00	0.00	1.00
Transaction Vol. (10,000 tCO2)	transaction	5,184	1.16	1.32	1.00	0.00	30.00

Notes: Our unit of analysis in Section 4 is at the transaction level. However, most of our variables are measured/collected at a different level reported in Column 2 (unit of obsevration). N is the number of unique observations for each variable.