

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

Estelle Cantillon[†] and Aurélie Slechten[‡]

July 2024

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 into its exchange, counterparty and trader drivers and show that *where* traders traded and *how connected* they and their counterparties were with the rest of the market covary with the terms they were able to obtain. Such features are expected to characterize OTC transactions but not, typically, anonymous 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.

*We thank George Deltas, Luca Merlino, Marco Pagano, Michael Pahle, Markus Trunschke and seminar and conference audiences 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

1 Introduction

The European Union Emissions Trading Scheme (EU ETS) is the largest carbon emissions market in the world by traded value. 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 - Directorate-General for Environment 2005, p. 14). In practice, trading picked up slowly and remained highly fragmented for a long time.

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 the unusually rich information they contain about trader characteristics. We find that the law of one price fails generically. Prices on different trading venues reflect local supply and demand conditions, as well as the position of the exchange in the network formed by the transactions between market participants. Our results suggest that the fragmentation of the EU carbon market hampered price aggregation and put less connected, small and 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 - when price discovery and transactions are not taking place in a single venue - 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. Markets in their early stages are typically less liquid, magnifying the impact of market microstructure. But the EU carbon market is especially interesting because trading during that crucial period was spread across 11 exchanges and over-the-counter. The market brought together compliance firms, that 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 contains the universe of transactions during the first phase. During that period, about 56% of transactions were carried out over-the-counter, the rest on exchanges. Compliance traders made up 91% of market participants and mostly traded over-the-counter: Only 8% of them used exchanges. Each trading venue attracted a different pool of 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-vary with exchange, counterparty and trader-specific factors. Such relationship captures the characteristics of the market equilibrium that arose. 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 market, price advantage is solely driven by intraday variation in prices. In a fragmented market, deviations from market-wide prices arise from exchange-specific factors, such as local supply and demand conditions, or counterparty and trader-specific characteristics.

We regress transaction-level price advantage on trader and exchange characteristics. Following the recent literature on over-the-counter markets, we use the network formed by the transactions between market participants as a 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 traders in the network formed by traders 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 degree to which the profile of surplus allowances of exchange participants differs from the market-wide profile.

We find that the price that traders get on exchanges depends on the local 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 ratio on an exchange significantly reduce the price advantage that a seller can get (the same result holds, *mutatis mutandis*, for buyers). More 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 that obtained from exchange and counterparty characteristics. Still, it is in the range of effects found in other over-the-counter markets. Finally, we find that small compliance traders and compliance traders from non-energy sectors received significantly worse terms than other traders.

Relationship with the literature. Market fragmentation can arise because transactions are spread across several exchanges, because transactions are over-the-counter, or both. There is a long and rich literature in finance on market fragmentation, starting with Demsetz (1968) and Smidt (1971). One take-away is that market fragmentation is typically associated with trader heterogeneity and self-sorting of traders into trading venues according to their trading needs and preferences, and trading venue characteristics (Pagano, 1989; Madhavan, 1995; Easley et al., 1996). A second take-away is that market fragmentation does not necessarily result in inferior price discovery, if price information is sufficiently well distributed (Jensen, 2007; Barclay et al., 2008) or if some traders multi-home (Foucault and Menkveld, 2008; Chen and Duffie, 2021). Our empirical analysis confirms these findings and, in particular, shows that realized prices in better-connected exchanges respond less to local market conditions than in lesser-connected exchanges.


Our paper builds on the more recent literature on over-the-counter markets that seeks to understand price formation through the lens of traders' connections. One approach, proposed by

Duffie et al. (2005), views price formation in OTC markets as the result of random encounters between traders. An alternative approach, formalized by Babus and Kondor (2018), views price formation as the result of persistent relations between traders. A seminal contribution along this line is Baker (1984), who documented how trading relationships among floor traders at a US stock option exchange changed with market size, and how these were systematically related to price volatility. More recently, Adamic et al. (2017), Di Maggio et al. (2017), Hollifield et al. (2017) and Li and Schürhoff (2019), among others, provided evidence from different OTC markets that the structure of trading relationships indeed helps explain observed financial outcomes.

Our paper’s contribution to this literature is threefold. First, we document price formation in a market that differs starkly from the dealer-dominated markets that have been studied so far. The EU carbon market is extremely diverse, with transactions going through exchanges, financial intermediaries, or directly between compliance firms. Second, we show that trader connectivity also matters for exchange transactions, and not only for OTC transactions. Since exchanges and OTC markets often coexist for the same assets, this suggests a more fungible boundary between the two market arrangements than previously thought, and networks as a unifying framework to look at price formation in both. Third, and despite documenting that trader connectivity matters, we show that exchange characteristics - including the mix of traders they attract - remain the primary determinant of price advantage: *where* traders trade is more important than *how* connected they are. We can do this because our data cover several exchanges and we observe traders’ identities. This allows us to account for both exchange and trader characteristics in a unified setting.

We are not the first to exploit the transaction log of the EU ETS. The existing literature has documented several specificities of this 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 compliance firms is highly heterogeneous (Martino and Trotignon, 2013; Zaklan, 2013; Betz and Schmidt, 2016; Jaraitė-Kažukauskė and Kažukauskas, 2015; Abrell et al., 2022; Zaklan, 2023). 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; Karpf et al., 2018). Fourth, large compliance traders tend to use exchanges 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 price, we can go one step further and *quantify* the impact of this fragmentation on market participants.

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



Notes: This graph represents the network of transactions between May 2005 and 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 (exchange-based transactions have the exchange on one of the sides of the transaction). 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.

2 The EU carbon market

The setting for our analysis is the first phase of the EU ETS, which was established by the European Union as part of its commitment under the Kyoto Protocol. The first phase of the EU ETS covered emissions in 2005-07. During that period, close to 11,000 installations from the most energy-intensive sectors in the economy (electricity generation, basic chemistry, cement, steel, glass and ceramics, ...) 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.¹ 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 (CMCEK) and Slovak (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 allowance futures.

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 that are part 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 market participants with many connections 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 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 the 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 Baliatti (2016), we restrict attention to transactions before May 2007 to avoid the period of very low prices at the end of Phase I.²

We use three sources of data. The first source is the Community Independent Transaction Log 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 timestamp 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

¹Electricity producers represented approximately 60% of the emissions covered by the ETS at the time.

²Given 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.

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 surrendered for compliance. Our third source of data is 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.³ We also remove transactions 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.

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.⁴ 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.⁵

³To do this, we first use fuzzy matching (Levenshtein distance) based on the name, 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 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.

⁴At 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.

⁵We include futures positions in the computation of accumulated net surplus because they represent future commitments to buy or sell.

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	27.01	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	19.17	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	0.82	0.14	0.51
ECX	14.3	7.40	46.8	63.3	29.34	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	92.69	-	0.32
Total	1,019.6	37.49	1,242.0	91.3	94.07	-	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).

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.⁶ 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 connectivity of a trader or an exchange in a network by accounting both for the number of traders with which they traded and the connectivity of these trading partners.⁷ 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

⁶The 12-month window is motivated by the low frequency of trades in this market and the annual compliance cycle.

⁷Eigenvector centrality as a proxy for trader connectivity has been used e.g. by Adamic et al. (2017), Di Maggio et al. (2017) and Hollifield et al. (2017). An alternative measure of centrality is the number of connections, used e.g. by Adamic et al. (2017) and Kondor and Pintér (2022). Li and Schürhoff (2019) finds that most centrality measures are correlated.

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 our analysis). Trading was fragmented: 56% of transactions and 69% of transaction volume took place over-the-counter. The remainder was split among 11 exchanges. The table documents large differences across exchanges 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 volume. 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 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 brevity 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 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 attracts 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 several covariates of realized prices in fragmented markets. We explore these relationships in our sample of exchange-based spot transactions (panel A of Table 2).

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
<u>Transaction characteristics</u>							
Size (10,000tCO ₂)	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
<u>Trader-month characteristics</u>							
Nb transactions per month	2,034	9.28	0.06	0.25	2.40	13.75	37.67
Accumulated net surplus (mtCO ₂)	2,034	0.34	-1.08	-0.00	0.00	0.10	2.83
Trader centrality	2,034	0.32	0.01	0.01	0.05	0.42	1.48
<u>Trader characteristics</u>							
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
<u>Transaction characteristics</u>							
Size (10,000tCO ₂)	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
<u>Trader-month characteristics</u>							
Nb transactions per month	33,228	1.13	0.05	0.10	0.20	0.50	4.00
Accumulated net surplus (mtCO ₂)	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
<u>Trader characteristics</u>							
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 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.⁸ Formally, for each transaction by trader i on exchange k and day t , we define the price advantage of this transaction as:⁹

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

where p_{ikt} denotes the transaction price and \bar{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 (\mathcal{S}) or on the buy side (\mathcal{B}).¹⁰ 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).¹¹ 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 typically lies 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 is likely to covary at equilibrium with exchange and trader characteristics. It motivates the following empirical specification:

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

where X_{kt} is a vector of exchange-specific and counterparty-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): $\text{local mkt conditions}_{kt} = (\bar{S}_t - S_{kt})1_{i \in \mathcal{S}} + (S_{kt} - \bar{S}_t)1_{i \in \mathcal{B}}$. An increase in this variable indicates more favorable local market conditions for traders on exchange k , relative to the rest of the market.

⁸This is similar in spirit to the bid-ask spread measure used in the empirical literature on OTC markets (Di Maggio et al., 2017; Hollifield et al., 2017; Li and Schürhoff, 2019). Our measure accounts for the fact that our transactions are exchange transactions and not dealer transactions, and that we only observe transaction prices and not the order book. Our price advantage measure corresponds to the transaction (price impact) component of Kondor and Pintér (2022)’s trading performance measure.

⁹In 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.

¹⁰Exchange-based transactions in our data have the exchange on one 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.

¹¹In practice, this means that, for 90% of the days in our sample, the market-wide price corresponds to the OTC price.

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 across the different venues, prices tend to converge across trading venues and reflect market-wide conditions (Madhavan, 1995; Barclay et al., 2008; Foucault and Menkveld, 2008). We capture these ideas by controlling for exchange eigenvector centrality, a proxy for market connectivity, 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 effects and controlling for the monthly ratio of sellers to buyers active on the exchange, the Herfindahl Hirschmann Index and the average centrality of the counterparty side on the exchange during the day of the transaction.

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. In the context of emissions markets, traders' bargaining power also depends on their commitments (emissions and allowances surrenders or settlements of futures contracts).¹² 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.¹³

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 (Di Maggio et al., 2017; Li and Schürhoff, 2019). 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 and Table 5 in the Appendix provides definitions 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 disappears if the exchange is well connected with the rest

¹²Lehmann et al. (2024) show that firms' trading profits on the EU ETS is related to their net positions.

¹³In our main regressions, we use a cutoff of 1 million tCO₂, which corresponds to the top quintile of initial 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)
<i>Exchange characteristics</i>						
Local Mkt Conditions	17.870*** (5.377)	17.898*** (5.024)	17.678*** (5.068)	14.195** (5.807)	11.109** (5.357)	11.098** (5.328)
Local Mkt Cond. × centrality	-19.033*** (6.114)	-19.062*** (5.692)	-18.759*** (5.736)	-14.186** (6.396)	-10.969* (5.920)	-10.984* (5.884)
Exchange centrality	-9.978*** (2.596)	-8.825*** (2.311)	-8.793*** (2.283)	6.348** (2.706)	5.166** (2.414)	5.011** (2.400)
<i>Counterparty characteristics</i>						
Nb. Sellers / Nb. Buyers	8.248*** (2.226)	6.325*** (2.077)	6.002*** (2.025)	-3.491* (1.859)	-2.219 (1.708)	-2.047 (1.677)
Counterparty average centrality	-1.158* (0.594)	-0.929 (0.601)	-1.014* (0.522)	-1.603*** (0.423)	-1.344*** (0.399)	-1.303*** (0.405)
Counterparty HHI			1.004 (1.374)			-1.047 (1.054)
Powernext	5.898* (3.203)	5.412* (3.176)	5.172 (3.377)	-2.743 (2.245)	-5.654** (2.450)	-5.066** (2.460)
EXAA	-0.839 (2.324)	0.878 (2.551)	-0.028 (3.431)	1.860 (1.385)	-1.432 (1.795)	-0.418 (1.950)
Nord Pool	3.319 (3.554)	3.647 (3.542)	2.890 (4.217)	-1.127 (2.328)	-4.413* (2.605)	-3.401 (2.754)
SENDECO2	-10.294*** (3.395)	-7.573** (3.395)	-8.396** (4.000)	6.328** (2.929)	4.333 (2.923)	5.303* (3.038)
<i>Trader characteristics</i>						
Trader centrality		0.422*** (0.099)	0.423*** (0.099)		0.435*** (0.158)	0.442*** (0.160)
Trader surplus (mtCO2)		0.015 (0.027)	0.015 (0.027)		0.071* (0.039)	0.074* (0.039)
Small compliance trader		-2.290*** (0.820)	-2.255*** (0.818)		-1.859** (0.730)	-1.824** (0.724)
Energy sector		0.961 (0.594)	0.992* (0.600)		1.636* (0.878)	1.592* (0.870)
Non-compliance trader		0.746 (0.599)	0.784 (0.602)		1.817** (0.895)	1.769** (0.886)
<i>Transaction characteristics</i>						
Transaction Vol. (log)	-0.075 (0.129)	-0.167 (0.123)	-0.170 (0.121)	0.092 (0.127)	0.082 (0.127)	0.073 (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	0.119	0.131	0.132

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 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 of the fact that prices on exchanges tend to be lower than on the OTC market, and exchange connectivity brings prices across trading venues closer to one another: the resulting higher price reduces exchange buyers' advantage (negative coefficient) and increases 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 3.7 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 3.2 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 a smaller price advantage. 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. An increase of one standard deviation in the counterparty average centrality is associated with a 0.8 p.p. increase in the seller'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 10.3 p.p. depending on the specification) and a seller advantage, possibly reflecting the high fraction of non-energy compliance traders on SENDECO2. The total number of allowances bought on SENDECO2 during our sample period was around 1.87 million ton CO₂. Considering the hypothetical market-wide price, a price disadvantage of 8% on SENDECO2 and a price advantage of 5% on Powernext, the total cost of buying those allowances on SENDECO2 was approximately 1.6 million euros higher than on Powernext.

Looking at trader characteristics reinforces the picture that traders' relative position matters for the trading terms they get on exchanges, though the possible endogeneity of some of these variables calls for caution when interpreting the size of these coefficients. First, trader centrality is statistically and economically significant. A trader with an eigenvector centrality one standard deviation above the mean, is associated with a 0.29 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 and with Di Maggio et al. (2017)'s finding of a 0.5 p.p. extra markup on non-dealer clients in the US corporate bond market. What is remarkable about our finding of a centrality premium for traders is that it holds for trading on an exchange - typically

thought to put all traders on the same footing - and not only for over-the-counter trading as previously documented. 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). The total of allowances traded by small compliance traders on exchanges is 1.8 million. With a price disadvantage of 2% (compared to the hypothetical market-wide price), the additional transaction cost paid by small compliance traders (compared to large compliance traders) was around 350,000 euros.

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.¹⁴ 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 - idiosyncratic advantage. The centrality premium that traders were able to obtain on exchanges during the first phase of the EU ETS 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. Di Maggio et al. 2017 and 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 efforts, reducing the cost effectiveness of emissions trading as a regulatory instrument.

But the first order effect of these market frictions lies elsewhere. 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

¹⁴In the exchange context, better connected traders are necessarily traders that multi-home.

singularly, compliance traders for whom the market was primarily designed. 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. Our empirical analysis explored these redistributive effects at some depth. We can also zoom out and quantify them at the aggregate level by summing the financial value of market advantage (the price differential relative to the market-wide frictionless price p_t times the transaction volume) across all trading venues, for compliance and non-compliance traders separately. This provides an indication of the size of the financial transfer between these two types of market participants due to market frictions. We find a net aggregate market disadvantage of -912,494 euros for compliance traders and a net aggregate market advantage of 834,738 euros for non-compliance traders.¹⁵ This is small in absolute terms but high given the small size of the market then.

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, 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 that serve as the primary market. It is an open question to what extent these different designs facilitate participation and price discovery by compliance traders.¹⁶

Today's EU carbon spot market has consolidated somewhat. There are three exchanges left 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.¹⁷

¹⁵The two terms do not sum to zero given the definition of the market-wide frictionless price based on the volume-weighted median price, rather than the volume-weighted average price.

¹⁶Joskow et al. (1998) provide an early study of how market support mechanisms can help market participants in an emissions market discover the equilibrium price.

¹⁷See e.g. European Securities and Markets Authority 2022's review of the market.

References

- Abrell, J., J. Cludius, S. Lehmann, J. Schleich, and R. Betz (2022). Corporate emissions-trading behaviour during the first decade of the EU ETS. *Environmental and Resource Economics* 83(1), 47–83.
- Adamic, L., C. Brunetti, J. H. Harris, and A. Kirilenko (2017). Trading networks. *The Econometrics Journal* 20(3), S126–S149.
- Babus, A. and P. Kondor (2018). Trading and information diffusion in over-the-counter markets. *Econometrica* 86(5), 1727–1769.
- Baker, W. E. (1984). The social structure of a national securities market. *American Journal of Sociology* 89(4), 775–811.
- Baliotti, A. C. (2016). Trader types and volatility of emission allowance prices. Evidence from EU ETS Phase I. *Energy Policy* 98, 607–620.
- Barclay, M. J., T. Hendershott, and C. M. Jones (2008). Order consolidation, price efficiency, and extreme liquidity shocks. *Journal of Financial and Quantitative Analysis* 43(1), 93–121.
- Betz, R. A. and T. S. Schmidt (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.
- Borghesi, S. and A. Flori (2018). EU ETS facets in the net: Structure and evolution of the EU ETS network. *Energy Economics* 75, 602–635.
- Chen, D. and D. Duffie (2021). Market fragmentation. *American Economic Review* 111(7), 2247–2274.
- Cludius, J. and R. Betz (2020). The role of banks in EU emissions trading. *The Energy Journal* 41(2), 275–300.
- Degryse, H., F. De Jong, and V. v. Kervel (2015). The impact of dark trading and visible fragmentation on market quality. *Review of Finance* 19(4), 1587–1622.
- Demsetz, H. (1968). The cost of transacting. *The Quarterly Journal of Economics* 82(1), 33–53.
- Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics* 124(2), 266–284.
- Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica* 73(6), 1815–1847.
- Easley, D., N. M. Kiefer, and M. O’Hara (1996). Cream-skimming or profit-sharing? The curious role of purchased order flow. *The Journal of Finance* 51(3), 811–833.
- European Commission - Directorate-General for Environment (2005). EU Emissions Trading: An open scheme promoting global innovation to combat climate change. Publications Office.

- European Securities and Markets Authority (2022). Emission allowances and associated derivatives, final report. ESMA70-445-38.
- Foucault, T. and A. J. Menkveld (2008). Competition for order flow and smart order routing systems. *The Journal of Finance* 63(1), 119–158.
- Fruchterman, T. M. J. and E. M. Reingold (1991). Graph drawing by force-directed placement. *Software: Practice and experience* 21(11), 1129–1164.
- Hintermann, B. (2010). Allowance price drivers in the first phase of the EU ETS. *Journal of Environmental Economics and Management* 59(1), 43–56.
- Hollifield, B., A. Neklyudov, and C. Spatt (2017). Bid-ask spreads, trading networks, and the pricing of securitizations. *The Review of Financial Studies* 30(9), 3048–3085.
- Jaraitė-Kažukauskė, J. and A. Kažukauskas (2015). Do transaction costs influence firm trading behaviour in the European emissions trading system? *Environmental and Resource Economics* 62, 583–613.
- Jensen, R. (2007). The digital divide: Information (technology), market performance, and welfare in the South Indian fisheries sector. *The Quarterly Journal of Economics* 122(3), 879–924.
- Joskow, P. L., R. Schmalensee, and E. M. Bailey (1998). The market for sulfur dioxide emissions. *American Economic Review*, 669–685.
- Karpf, A., A. Mandel, and S. Battiston (2018). Price and network dynamics in the European carbon market. *Journal of Economic Behavior & Organization* 153, 103–122.
- Kondor, P. and G. Pintér (2022). Clients’ Connections: Measuring the Role of Private Information in Decentralized Markets. *The Journal of Finance* 77(1), 505–544.
- Lehmann, S., J. Schleich, and J. Pinkse (2024). Allocation or skill? what is driving corporate trading performance in the eu ets. *The Energy Journal*, forthcoming.
- Li, D. and N. Schürhoff (2019). Dealer networks. *The Journal of Finance* 74(1), 91–144.
- Madhavan, A. (1995). Consolidation, fragmentation, and the disclosure of trading information. *The Review of Financial Studies* 8(3), 579–603.
- Martino, V. and R. Trotignon (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, 1–42.
- Pagano, M. (1989). Trading volume and asset liquidity. *The Quarterly Journal of Economics* 104(2), 255–274.
- Smidt, S. (1971). Which road to an efficient stock market: free competition or regulated monopoly? *Financial Analysts Journal* 27(5), 18–20.

- Zaklan, A. (2013). Why Do Emitters Trade Carbon Permits? Firm-Level Evidence from the European Emission Trading Scheme. DIW working paper 1275.
- Zaklan, A. (2023). Coase and cap-and-trade: evidence on the independence property from the european carbon market. *American Economic Journal: Economic Policy* 15(2), 526–558.
- 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	-0.16	-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	0.62	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	0.17	-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	0.53	1.94	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 observation). N is the number of unique observations for each variable.

Table 5: Definitions of regression variables

Variable	Unit of obs.	Definition
Trader surplus	trader-month	Trader accumulated net surplus of allowances up to that month (it includes net spot transactions, intrapolated settled transactions, and for compliance traders, intrapolated free allocations and surrender)
Exchange average surplus S_{kt}	exchange-month	Average accumulated net surplus of traders active on a specific exchange during the month of day t
Market-wide average surplus \bar{S}_t	month	Average accumulated net surplus of traders active on the market (OTC and exchanges) during the month of day t
Local Mkt Cond.	exchange-month	For an exchange-based transaction i on exchange k in month of day t , it is defined as $(\bar{S}_t - S_{kt})1_{i \in \mathcal{S}} + (S_{kt} - \bar{S}_t)1_{i \in \mathcal{B}}$ where \mathcal{B} is the set of buy-side transactions on exchanges and \mathcal{S} is the set of sell-side transactions on exchanges
Exchange centrality	exchange-month	month-level eigenvector centrality of an exchange in the network formed by all the transactions (OTC + exchanges) that took place over the past 12 months
Trader centrality	trader-month	month-level eigenvector centrality of a trader in the network formed by all the transactions (OTC + exchanges) that took place over the past 12 months
Nb. Sellers / Nb. Buyers	exchange-month	Monthly number of active sellers on the exchange divided by the monthly number of active buyers on that exchange
Counterparty average centrality	exchange-day	For buy-side transactions, this is defined as the average eigenvector centrality of the trader selling allowances on that exchange during a specific day. The definition for sell-side transactions is similar.
Counterparty HHI	exchange-day	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 a value between $1/n$, where n is the number of active buyers that day (least concentrated) and 1 (most concentrated). The definition for the sell-side is similar.
Small compliance traders	trader	Dummy variable equal to 1 if the trader is a compliance trader with a total annual initial allocation below 1 million tCO ₂
Energy sector	trader	Dummy variable equals 1 if the trader is a compliance trader from the energy sector. The energy sector comprises NACE 35 (Electricity, gas, steam and air-conditioning supply) and NACE 19.20 (Manufacture of refined petroleum products)
Non-compliance traders	trader	Dummy variable equal to 1 if the trader does not own any regulated installations (i.e., that do not receive or surrender allowances for compliance)
Transaction vol. (10,000 tCO ₂)	transaction	volume of the transaction

Note: The number of traders active on a trading venue (or the market) is computed on the basis of traders who have traded on that trading venue (or on the market) in the past 12 months. Exchange-based transactions in our data have the exchange on one of the sides 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.