

Market power in emission permit markets: Theory and evidence from the EU ETS

Beat Hintermann, University of Basel

March 1, 2015

Abstract

A well-known result about market power in emission permit markets is that efficiency can be achieved by full free allocation to the dominant firm. I show that this result breaks down when taking the interaction between input and output markets into account, even if the dominant firm perceives market power in the permit market alone. I then examine the empirical evidence for price manipulation by the ten largest electricity firms during phase I of the EU ETS. I find that firms' excess allowance holdings are consistent with strategic price manipulation, and that they cannot be explained by price speculation or by precautionary purchases to insure against uncertain future emissions. My results suggest that market power is likely to be an empirically relevant concern during the early years of emission permit markets.

JEL codes H23, H32; Q52, Q53, Q54, Q58.

Keywords Emission permit markets, market power, EU ETS, cost pass-through

1 Introduction

Emission permit markets are a direct application of economics to an environmental context, and an increasingly used policy instrument to control the speed of climate change. Because of the transactions costs associated with monitoring, reporting and verifying emissions, emission permit markets typically cover large firms, whereas emissions from smaller firms or households are usually addressed by other policy instruments such as taxes or direct regulation.

Since some of the included sectors have raised concerns about imperfect competition in the past, it stands to reason that emission permit markets themselves may be susceptible to price manipulation by dominant firms. This topic has been analyzed from a theoretical perspective, starting with the seminal papers by Sinn and Schmolzi (1981) and Hahn (1984), and extended to a dynamic setting by Liski and Montero (2005, 2011). They find that a dominant firm will use its market power to inflate the permit price if it acts as a net seller in the market, and to depress it if it acts as a net buyer. The efficient solution consists in allocating the dominant firm exactly the number of permits it will in fact demand, thereby effectively removing it from the market.

This prescription is unsatisfactory for two reasons. First, the implication that dominant firms should be given a sizable gift in the form of full free allocation is difficult to justify on equity grounds. Second, the "full allocation" result has been derived without an interaction between input and output markets. But firms aim to maximize overall profits, not to minimize the compliance costs associated with an environmental regulation; if permit costs influence product markets, the latter therefore have to be taken into account. Misiolek and Elder (1989) show that if a dominant firm is able to manipulate both the input and output markets, Hahn's prescription of full allocation no longer applies. The dominant firm will generally buy more (or sell fewer) permits than in the case where it focuses on compliance costs alone, as this allows it to increase its market share and overall profits relative to the fringe. Similarly, Sartzetakis (1997) shows that a dominant firm may use emission permits to keep rivals out of the product market.

Disegni Eshel (2005) and Hintermann (2011) shift the focus away from exclusionary manipulation and show that a dominant firm with market power in both the permit and the product market will manipulate both prices so as to increase its profits at the expense of

consumers and taxpayers. A dominant net permit buyer may want to increase the permit price, provided that the increase in compliance costs is more than offset by the sum of the revenue increase in the output market and the increase in rents embedded in free allocation. The latter has been discussed under the label of "windfall profits" in the EU ETS (Neuhoff et al., 2006; Sijm et al., 2006). Firms receive a large share of their allocation for free, but nevertheless pass through the opportunity costs of emissions to consumers. This is not a market failure, but rather a sign of a well-functioning market: If emissions are capped, permits become a costly input and should therefore be reflected in the marginal cost of production. The inefficiency arising from price manipulation lies in the fact that the carbon cost that is being passed through does not reflect the economy's opportunity cost to reduce emissions.

In this paper, I show that it is not the presence of "double" market power per se, but simply the transmission of input costs into output prices that leads to a failure of Hahn's prescription of full free allocation. This market interaction is sufficient to generate an incentive even for net permit buyers to increase the permit price, provided that they receive a sufficiently large free allocation. The relevant allocation threshold is below the firm's efficient permit demand and necessarily makes it a net permit buyer. If the firm is allocated more (fewer) free permits than this threshold, it will use its market power to increase (decrease) the permit price. The threshold is firm-specific and can be negative, implying that firms may have an incentive for upward price manipulation even under full permit auctioning.

The second contribution of the paper is empirical. Although there is a sizable theoretical literature about market power in emission permit markets, only a few empirical studies exist to date. A series of papers have addressed the effect of market power by means of controlled laboratory experiments. The magnitude of the inefficiency arising from market power depends on the experimental context, but the common finding is that including the output market in the experimental setting and ascribing market power to participants in both markets significantly affects the outcome, relative to the case where the goal is strictly to minimize compliance costs (Godby, 2002; Muller and Mestelman, 1998).

Some studies have examined the evidence for market power in the context of the regional permit markets in the USA. Ellerman and Montero (2007) and Liski and Montero (2011) find no evidence for market manipulation in the US Acid Rain Program for SO_x emissions.

However, they do not consider cost pass-through, perhaps because the US electricity markets are not deregulated to the same extent as in the EU. Holland and Moore (2012) examine intertemporal permit holdings in California's Regional Clean Air Incentives Market and find that some firms allowed surplus permits to expire unused, which could potentially be due to the over-purchasing of permits. Kolstad and Wolak (2008) report that electricity firms selling their output into the California market paid too high prices for NO_x permits to "cost justify" their electricity bids.

In the empirical part of this paper, I investigate the possibility that the ten electricity firms with the highest emissions inflated the permit price during phase I of the EU ETS. I start by showing that all of these firms received a free allocation that is greater than the threshold derived in the theoretical section, such that their profits were an increasing function of the permit price. Conditional on having market power, it would therefore have been in their economic interests to manipulate the permit price upwards, despite the fact that most of them were net permit buyers.

Because profiting from a high permit price and manipulating the permit price upwards are two different things, I then examine the evidence for the exercise of market power. The EU ETS is concentrated in the sense that a few firms cause most of the emissions, but traditional measures of market power such as the Herfindahl-Hirschman index (HHI) suggest relatively low scope for price manipulation. However, the HHI may not be a good measure for a market where firms' incentives are aligned, because they all profited from a high price. Also, because many smaller firms did not participate in the market due to transactions costs (Jaraitė-Kažukauskė and Kažukauskas, 2014), the relevant size of the market may have been much smaller, and thus the market share of the dominant firms may have been larger than what would be suggested by using the entire cap as the market base. Likewise, price-based tests of market power (such as the Lerner index) are not possible in the current context, because neither firms' marginal abatement costs nor the "right" permit price (i.e., the one that corresponds to market fundamentals) are known, and reasons other than market power could lead to a deviation between the permit price and marginal abatement costs.

As a measure of market power, I instead rely on the notion of excess allowance holdings, which I define as a situation where a firm holds more allowances than it needs to cover its en-

tire emissions during the phase. This is a sufficient condition for firms' setting their marginal abatement cost below the permit price. Six firms in the sample fulfilled this condition and held excess allowances at least during some months of the first phase. Note that due to the no-banking provision between phases I and II, the surplus allowances held at the very end of the phase expired unused.

Holding excess allowances is consistent with firms' attempting to drive up the permit price, but it may also be explained by other reasons. For example, firms may have expected the permit price to increase towards the end of the phase. If so, they should have held short positions in forward contracts expiring at the end of the phase. However, using permit transfers in December 2007, March 2008 and April 2008 as a proxy for futures contracts, most firms in fact held long positions (or acted as net buyers in the spot market during these months). Price speculation is therefore an unlikely explanation for firms' excess allowance holdings.

Last, I examine the possibility that firms engaged in precautionary allowance purchases in order to account for the stochastic nature of future emissions. If firms believed that they were not going to be able to access the permit market later on, they may have wanted to make sure that they held a sufficient number of allowances even in case of a positive demand shock. Using a simple prediction model that is arguably less precise than firms' own forecasts, I compute the remaining emissions until the end of the phase from various vantage points in time and thus arrive at a distribution for total emissions during the phase; due to the availability of output data, I can do this for four firms only. I find that RWE and Vattenfall held more allowances than needed to cover the 95th percentile of this distribution, suggesting that their excess allowance holdings cannot be explained by hedging against carbon risk either. In contrast, the excess allowance holdings by E.ON and EdF are not inconsistent with precautionary purchases. Although my results deliver no conclusive proof for price manipulation, because firms' decision processes are not observable, they do suggest that some firms are likely to have manipulated the allowance price in their favor during phase I of the EU ETS by purchasing more allowances than they needed, or selling fewer than the could have sold.

The next section presents the theory. Section 3 determines firms' incentives for price

manipulation, and section 4 contains the empirical application. Section 5 concludes.

2 Theory

I start by presenting my theoretical model and compare it to that by Hahn (1984). I then discuss the implications of this theory for empirical tests of market manipulation.

2.1 Model

In the following, I use a simplified version of the model by Hintermann (2011), which consists of a dominant firm (or several colluding firms) and a price-taking fringe. Fringe firms $i \in (2, \dots, N)$ take output prices p and allowance prices σ as given and choose output q_i , emissions e_i and permit holdings x_i to maximize profits subject to an emissions constraint:

$$\max_{q_i, e_i, x_i} \Pi_i = pq_i - C^i(q_i, e_i) - (x_i - \bar{x}_i)\sigma \quad \text{s. t.} \quad e_i \leq x_i \quad (1)$$

Firm i 's cost function C^i depends on both output and emissions and satisfies $C_q^i \equiv \partial C^i / \partial q_i > 0$, $C_{qq}^i > 0$, $C_e^i < 0$, $C_{ee}^i > 0$, $C_{qe}^i < 0$ and $C_{qq}^i C_{ee}^i - (C_{qe}^i)^2 > 0$. The optimality conditions of this problem are independent of the amount of free allocation \bar{x}_i and state that firms equate price with marginal cost in both markets. This means that the fringe firms' optimal output, emissions and permit decisions are a function of market prices:

$$p = C_q^i; \quad \sigma = -C_e^i \quad (2)$$

$$\Rightarrow q_i^* = q_i^*(p, \sigma); \quad e_i^* = x_i^* = x_i^*(p, \sigma) \quad (3)$$

Permit purchase decisions by the dominant firm, which we call firm 1, have an impact on both the permit and the output price:

$$\frac{\partial \sigma}{\partial x_1} > 0 \quad (4)$$

$$\frac{\partial p}{\partial x_1} = \frac{\partial p}{\partial Q^f} \cdot \frac{\partial Q^f}{\partial \sigma} \cdot \frac{\partial \sigma}{\partial x_1} > 0; \quad Q^f \equiv \sum_{i=2}^N q_i \quad (5)$$

Intuitively, the dominant firm's permit purchases affect the permit price directly in (4),

and the output price indirectly via the fringe's response: An increase in the permit price lowers the fringe's output, which in turn increases the output price.

The dominant firm considers (2-5) when maximizing its own profits. An interior solution with $q_1, x_1, e_1 > 0$ is characterized by the following optimality conditions:¹

$$p = C_q^1 \tag{6}$$

$$\sigma + (x_1 - \bar{x}_1) \frac{\partial \sigma}{\partial x_1} - \frac{\partial p}{\partial x_1} q_1 = \lambda \tag{7}$$

$$-C_e^1 = \lambda \tag{8}$$

$$x_1 \geq e_1; \quad \lambda \geq 0; \quad (x_1 - e_1) \cdot \lambda = 0 \tag{9}$$

Condition (7) is the core result and shows that the dominant firm will generally not equate the permit price with marginal abatement costs (using eq. 8) if it receives an allocation equal to its efficient permit demand. The reason is that buying another permit drives up both prices, and both affect the firm's profits. Depending on the free allocation \bar{x} it receives, its influence on the permit price and the effect of the permit price on the output price, the firm sets its marginal abatement costs above or below the permit price. Conditions (8) and (9) further imply that if the dominant firm purchases more permits than needed for compliance, its emissions constraint and thus its marginal abatement costs in an economic sense as captured by $-C_e^1$ are zero, even if the technical costs of reducing emissions by one ton may be positive.² This is the situation that fits the empirical part of the paper, as (some) firms bought more permits than they needed for compliance.

Substituting the emissions constraint, using the implicit function theorem and denoting J as the Jacobian matrix of the maximization problem (which has a positive determinant by assumption), it follows that the dominant firm's output and its demand for permits are an

¹In order for the second-order conditions to hold, the curvature of the permit demand curve must not be too great. Specifically, it must be the case that $\frac{\partial \partial \sigma}{(\partial x_1)^2} \left(\frac{\partial p}{\partial \sigma} - (x_1 - \bar{x}_1) \right) < C_{ee} + 2 \frac{\partial \sigma}{\partial x}$. Since the RHS is positive, this is trivially the case with a linear permit demand function.

²Usually, $-C_e$ is referred to as marginal abatement costs, because abatement equals negative emissions. However, strictly speaking, $-C_e$ reflects the marginal change in the firms' costs when its emissions increase by one unit. If the firm's emission constraint is not binding, then the firms' costs do not change with a one-unit increase in emissions. This is the interpretation for $-C_e = 0$, even though the model is not contingent on the assumption of zero initial marginal abatement costs. In other words, the engineering cost associated with abating the first unit of emissions may be positive even if $-C_e = 0$. Naturally, for interior solutions this relatively subtle distinction between marginal abatement costs and negative marginal emission costs vanishes.

increasing function of the level of free allocation:

$$\frac{\partial q_1}{\partial \bar{x}_1} = \frac{\left(\frac{\partial p}{\partial x_1} - C_{qe}\right) \cdot \frac{\partial \sigma}{\partial x_1}}{|J|} > 0 \quad (10)$$

$$\frac{\partial x_1}{\partial \bar{x}_1} = \frac{C_{qq} \cdot \frac{\partial \sigma}{\partial x_1}}{|J|} > 0 \quad (11)$$

Under perfect competition it is well-known that firms' output, emissions and abatement decisions are independent of the amount of free allocation. This no longer holds if market power is present: the higher the dominant firm's free allocation, the higher its preferred levels of output and permit demand.

We can solve (7) for the level of free permit allocation that makes the firm equate its marginal abatement costs to the permit price:

$$\bar{x}_1^0 = x_1 - \frac{\partial p / \partial x_1}{\partial \sigma / \partial x_1} \cdot q_1 \quad (12)$$

Here, x_1 is the firm's optimal permit demand, conditional on it receiving a free allocation of \bar{x}_1^0 . Because x_1 is continuous in \bar{x}_1 according to (11), (12) defines a fixed point for which \bar{x}_1^0 is the unique solution.

Since both prices are a positive function of the dominant firm's permit demand, the efficient amount of free allocation \bar{x}_1^0 is below the firm's optimal permit demand. In other words, efficiency requires the firm to be a net permit buyer.³ If the dominant firm receives free permits in excess of \bar{x}_1^0 , it will use its market power to inflate the permit price by under-abating (i.e., the LHS in 7 exceeds the RHS), and vice versa.

Because marginal production costs decrease with emissions ($C_{qe} < 0$), under-abatement further implies that the dominant firm increases its market share at the expense of the fringe (its marginal production costs decrease relatively to the fringe). This is consistent with the core result from the literature about raising rivals' costs (Krattenmaker and Salop, 1986;

³The model by Hintermann (2011) differs from the current one in that he allows the firm to directly manipulate the output price, which leads to additional terms associated with $\partial p / \partial q_1$ and $\partial \sigma / \partial q_1$ in (6). He also computes a threshold of free allocation that makes the firm act as a price taker in the output market, and shows that this is necessarily smaller than \bar{x}_1^0 . The existence of two different thresholds implies that efficiency cannot be obtained by means of free permit allocation alone, which is the familiar result that two market failures require more than one policy instrument for efficiency. Since the assumption here is that the dominant firm sets its marginal production cost equal to the permit price, no such threshold exists, and free allocation as defined by (12) is efficient.

Salop and Scheffman, 1983; Sartzetakis, 1997): A dominant firm has an incentive to increase input costs in order to gain market share at the expense of its rivals, provided that the cost increase hurts it relatively less. If fringe firms equate marginal abatement costs to the permit price but the dominant firm sets $-C_e^1 < \sigma$, an increase in the permit price is relatively less costly for the dominant firm than for the fringe.

I now compare these results with those derived by Hahn (1984). If the output market is neglected, the dominant firm's maximization problem can be written as one of compliance cost minimization:

$$\min_{e_1, x_1} C^1(e_1) + (x_1 - \bar{x}_1)\sigma \quad s.t. \quad e_1 \leq x_1 \quad (13)$$

with $C_e < 0$ and $C_{ee} > 0$. After substituting the constraint and minor rearrangements, the resulting optimality condition becomes

$$\sigma + (x_1 - \bar{x}_1) \frac{\partial \sigma}{\partial x_1} = -C_e \quad (14)$$

from which follows directly that the relevant threshold of free allocation is given by

$$\bar{x}_1^{0, \text{Hahn}} = x_1 \quad (15)$$

If the dominant firm receives a free allocation that is greater than (15) such that it is a net seller of permits, it will inflate the permit price by under-abating, and vice versa for a free allocation below x_1 . The efficient solution consists in taking the problematic firm out of the market by allocating it exactly (15).⁴

2.2 Implications for empirical analyses of market power

Hahn's result has made a lasting impression in the literature on market power in emission permit markets. Relatively recently, it was extended to the dynamic case by Liski and Montero (2011), who show that in the absence of borrowing, the effect of upward price manipulation

⁴Again, the firm's optimal permit demand is a continuous function of its free allocation such that $\bar{x}_1^{0, \text{Hahn}}$ solves the fixed point defined by $\bar{x}_1^{0, \text{Hahn}} = x_1(\bar{x}_1^{0, \text{Hahn}})$. Using Brouwer's fixed point theorem, it is straightforward to show that $\partial x_1 / \partial \bar{x}_1 > 0$.

by a net seller is much more serious than that of downward price manipulation by a net buyer.⁵ Based on the finding that the largest firms in the US Acid Rain Program were net buyers, the authors conclude that strategic price manipulation cannot be expected to be empirically relevant in that market.

Following this logic, market power would also not be expected to play a role in the EU ETS, because the largest included firms are power producers, which were on average under-allocated and were, in fact, the only net buyers in the the first phase (Hintermann, 2010). Furthermore, with a decrease in the share of free permit allocation and a corresponding increase in auctioning in later market phases, all firms will eventually become net buyers, such that the theory developed by Hahn (1984) and Liski and Montero (2011) implies that market power is simply not a relevant concern in the context of the EU ETS, and presumably the same applies to other emission permit markets. This, along with the scarcity of data, may be the reason why so few empirical tests for market power in permit markets exist, and none at all in the context of the EU ETS. For a review of the existing empirical literature, see Montero (2009).

The situation changes when the output market is taken into account. The relevant question is no longer whether a large firm is a net seller, but whether its free allocation exceeds the threshold (12), which is much more likely; in fact, in this paper I show that the threshold is exceeded for all large electricity firms covered by the EU ETS. For some firms, the threshold is negative, implying that they would profit from increasing the permit price even under full auctioning, which is the situation for many power generators since the start of phase III. Most firms in my sample are net permit buyers. The theory developed by Hahn (1984) and Liski and Montero (2011) correctly states that an increase in the permit price increases these firms' compliance costs. However, my results show that the profit increase in the output market more than offsets the compliance cost increase, such that firms' overall profits are a positive function of the permit price. Since firms presumably care more about the bottom line

⁵The authors assume that the dominant firm has an initial stock which it can bank, but that eventually it will become a permit buyer, because its marginal abatement costs are above those of the fringe. If the dominant firm is a net seller over the time frame of the program, it will restrict permit supply in the beginning, which raises the market price above the efficient level and shifts the moment at which the firm's stock is exhausted further into the future. Conversely, if it is a net buyer, it will increase its sales, thus depressing the price and shifting the moment of exhaustion forward in time. If firms are not allowed to borrow permits, the fringe firms know that the dominant firm cannot keep its own demand low forever, because eventually it will face prohibitively high abatement costs. Knowing this, they will restrict their own permit demand and thus keep the price near the efficient level.

than about compliance costs, this changes the situation qualitatively, and price manipulation in the EU ETS becomes not so unthinkable after all.

Importantly, this conclusion does *not* depend on explicit market power in both markets, as could perhaps be inferred from the analyses by Misiolek and Elder (1989), Disegni Eshel (2005) and Hintermann (2011).⁶ The functional form in (12) and thus the underlying intuition is in fact identical to equation (6) in Hintermann (2011). In that paper, the dominant firm also quantity-controls the output price, such that its optimal choices of output, emissions and permits holdings will generally differ from the solution of the present model. Furthermore, allocating the dominant firm the number of free permits implied by condition (12) will lead to an efficient outcome under single but not under double market power due to the additional market failure in the output market. However, the main insight of the two models is the same: The dominant firm's permit purchases affect costs and profits in both markets due to the *interaction* between the permit and the output market, and this reduces the relevant allocation threshold for upward price manipulation to below full allocation.

My empirical strategy consists of the following steps. In the next section, I investigate whether the largest firms' free allocation exceeded (12) and find this to be the case. On its own, this states that firms profited from a high permit price, but it provides no evidence for price manipulation. In order to examine the possibility that they actively inflated the permit price, I focus on firms' permit holdings over time in section 4. Under-abatement and over-purchasing of permits (or under-selling for net buyers) would be an indication for the exercise of market power. To exclude other reasons for over-purchasing, I test whether this could be explained by precautionary purchases or by price speculation, and find this to be the case for some but not all firms in the sample.

⁶By "explicit" market power in both markets, I mean that the dominant firm also considers the effect of its output quantity on the output and the permit price, not just the effect of its permit purchases. Of course, one might argue that the dominant firm manipulates both prices in the present model as well, because it considers the fact that the permit price affects the output price. However, taking into account a market interaction is qualitatively different from attempting to increase the output price by strategically withholding output. In the current model, the dominant firm sets price equal to marginal cost in the output market. The market interaction means the marginal production cost is a function of the permit price.

3 Free allocation for large power producers in the EU ETS

In this section, I examine whether the ten largest electricity firms covered by the EU ETS would have found it profitable to inflate the permit price according to condition (12). I start with a brief description of the relevant background.

3.1 Market background

The European Union Emissions Trading Scheme (EU ETS) is a cap-and-trade system that covers CO₂ emissions from about 11,000 installations that belong to energy-intensive sectors in all EU countries. It is the largest such system to date and includes around 42 percent of the EU's emissions. The EU ETS served as one of the EU's primary instruments to achieve their Kyoto targets, and it continues to play an important role in order to reach the EU's reduction target by 2020 and beyond.

The first market "phase" covered the years 2005–2007 and was intended as a trial run for phase II, which coincided with the Kyoto compliance period 2008–2012. The market is currently in its third phase, which covers the years 2013–2020. Apart from their association with different compliance periods defined by international agreements, the phases differ with respect to rules about allocation and the banking of emission permits. While permits from phase I could not be banked into phase II as this would have effected the EU's Kyoto targets, there are no limits on banking starting in phase II. The no-banking provision between the first two phases is important in the context of this paper, because it removes banking as a potential reason for purchasing more permits than needed for compliance. Permit borrowing is allowed within, but not across phases.

Permits were allocated mostly for free during phases I and II, but auctioning is the main method of allocation since phase III. Importantly, firms in the Power & Heat sector no longer receive any free allocation. For a more detailed discussion of the institutional setup of the EU ETS and a review of the related literature, please refer to the symposium by Ellerman et al. (2015), Hintermann et al. (2015) and Martin et al. (2015).

Table 1 contains total allocation and emissions by year for the first two market phases, and Figure 1 shows the permit price in the EU ETS during phases I–II. The price for permits (called EU allowances, or EUAs, giving the holder the right to emit one ton of CO₂) was


Table 1: Allocation and emissions in EU ETS, 2005–2012

Year	Allocation (Mt)	Emissions (Mt)	Surplus (Mt)	Surplus (%)
2005	2,313	2,014	299	12.93
2006	2,277	2,036	241	10.58
2007	2,278	2,165	113	4.96
Phase I	6,868	6,215	653	9.51
2008	2,136	2,120	16	0.77
2009	2,186	1,880	306	14.01
2010	2,220	1,939	281	12.67
2011	2,222	1,904	318	14.33
2012	2,222	1,867	355	15.99
Phase II	10,988	9,710	1,278	57.78

Source: European Union Transaction Log (EUTL).

above €20 for much of phase I, even surpassing €30 before dropping to about half its value when verified emissions were published for the first time in April 2006. The allowance price stabilized around €10–15 for another half year before approaching zero towards the end of the phase. Similarly, the price rose at the beginning of phase II, but then declined at the onset of the financial crisis of 2008.

Figure 1: Allowance spot prices in the EU ETS, 2005–2012



Source: European Energy Exchange (EEX).

Given the no-banking provision between the first two phases, the apparent over-allocation in phase I should have led to a price close to zero in hindsight. A number of papers have attempted to explain the price behavior. Some have focused on fundamentals such as fuel prices, economic activity and the weather (Hintermann, 2010; Mansanet-Bataller et al., 2007).

Others have tried to explain the price path by means of financial theory (Ben; Cetin and Verschuere, 2009; Seifert et al., 2008), or options formulae (Chesney and Taschini, 2012; Hintermann, 2012). In this paper, I examine the empirical evidence for the possibility that large firms manipulated the allowance price upwards during the first two years of the phase.

The EU ETS is quite concentrated. The bottom 80 % of firms account for about 3 % of emissions, whereas the top 1 % of firms are responsible for 59 %.⁷ The largest firms in the EU ETS in terms of emissions are electricity producers, so if market power is an empirically relevant issue, it would arguably arise from the power sector.

3.2 Allocation, emissions and permit holdings by firm

In order to determine whether these firms received an allocation that exceeds the threshold (12), we have to find an empirical substitute, because the price responses in the definition of \tilde{x}_1^0 are not observed. Totally differentiating the output and permit price and dividing the former by the latter gives

$$\frac{dp}{d\sigma} = \frac{\partial p/\partial x_1 \cdot dx_1 + \partial p/\partial q_1 \cdot dq_1}{\partial \sigma/\partial x_1 \cdot dx_1 + \partial \sigma/\partial q_1 \cdot dq_1} \quad (16)$$

If both the output and the permit price only depend on the dominant firm's permit purchases as assumed in the theory section, such that $\partial p/\partial q = \partial x/\partial q = 0$, the allocation threshold simplifies to⁸

$$\tilde{x}_1^0 = x_1 - \frac{dp}{d\sigma} q_1 \quad (17)$$

This threshold can also be obtained by examining how (all) firms' profits change in response to an exogenous variation in the permit price. Applying the envelope theorem to (1)

⁷For the 10 firms in my sample, I performed the aggregation of the installation data in EUTL to the firm level using information from firms' 2007 business reports. For the remaining firms, I used a dataset produced by Jaraite et al. (2013), which aggregates the installation-level information in the EUTL to the firm level (global ultimate owner) based on 2006 ownership information in Bureau van Dijk's Amadeus database.

⁸The output price may depend on the dominant firm's output, even if the firm does not take this relationship into account when deciding on its optimal output and permit purchase decisions. Hintermann (2011) shows that if the dominant firm's output is nondecreasing in σ (which will be the case if the demand response is sufficiently small and the dominant firm expands its market share at the cost of the fringe when σ increases), it follows that $\frac{dp}{d\sigma} < \frac{\partial p/\partial x_1}{\partial \sigma/\partial x_1}$, such that (17) > (12). This means that if a firm's free allocation exceeds the "empirical" threshold, it will also exceed the theoretical one.

with respect to the permit price and solving for the quantity of free allocation that leaves profits unchanged leads to (17). Thus, if the firm receives a greater allocation than \tilde{x}_1^0 , its profits are an increasing function of the permit price, and vice versa, independent of its degree of market power.

The ratio $dp/d\sigma$ refers to the average pass-through of carbon costs to electricity prices (in euro/Mwh per euro/tCO₂). Estimates vary and depend on the method, the market under consideration, assumptions about endogeneity of input and output prices, the load period and also the time frame; for a review and the empirical challenges of estimating cost pass-through, see Fell et al. (2015). Using hourly data, Fabra and Reguant (2014) and Hintermann (2014) show for the Spanish and German wholesale market, respectively, that carbon costs are passed through fully in the hourly spot prices. I employ a range for average cost pass-through of 0.4–1.0 tCO₂/MWh with a central value of 0.7, which means that an increase in the allowance price of 1 €/tCO₂ leads to an average increase in electricity prices of 0.7 €/MWh.⁹

Table 2 shows generation, emissions, allocation, and the corresponding allocation threshold \tilde{x}^0 for the largest ten electricity firms in Europe in terms of fossil generation and emissions. All of them received a free allocation in excess of the threshold even for $dp/d\sigma = 0.4$, indicating that their profits are an increasing function of the carbon price, even though most of them were net permit buyers. This is consistent with the results by Bushnell et al. (2013), who find a positive correlation between the allowance price and electricity firms' stock performance.

Depending on the value of cost pass-through, the threshold of efficient allocation is negative for some firms, suggesting that they profit from an increase in the permit price even under full auctioning. The underlying reason is that most firms generate electricity using a range of sources including nuclear and renewable energy, and thus have an average emission intensity that is below that of the marginal generator, which is fossil-based during most hours (Hintermann, 2014). With full pass-through of carbon costs, firms with an average emission intensity below $dp/d\sigma$ will profit from a positive carbon price even in the absence of free allo-

⁹To put these numbers into perspective, a combined-cycle gas turbine produces around 0.42 tCO₂/MWh of electricity, whereas the emission intensity of a hard coal plant is around 0.96 tCO₂/MWh. Under the assumption of complete pass-through of marginal carbon costs to electricity prices, the range of 0.4–1.0 tCO₂/MWh spans the range of emission intensities between gas and coal generation. Alternatively, this range is also consistent with a situation where coal is marginal all the time, but costs are passed through between 40% and 100%. Based on the empirical evidence in the literature as cited in the text, it is very unlikely that the average cost pass-through is less than 0.4 tCO₂/MWh in the German, French or Spanish electricity markets.

cation, because the average increase in the electricity price more than compensates them for their average increase in compliance costs.¹⁰ The firm with the highest allocation threshold is RWE, which produces a large share of its generation using lignite and coal.

Table 2: Output, permit holdings and emissions of firms in sample, 2005–2007

	Emissions		Allocation		Generation		$\bar{x}^0 = x - d\sigma / dp$ (Mt)			Surplus (Mt)
	(Mt)	Share	(Mt)	Share	(TWh)	Share	$d\sigma/dp=0.4$	$d\sigma/dp=0.7$	$d\sigma/dp=1$	
RWE	443	7.1%	414	6.6%	656	6.8%	191	-6	-203	11.0
E.ON	292	4.7%	257	4.1%	607	6.3%	55	-127	-309	5.5
Vattenfall	263	4.2%	272	4.3%	512	5.3%	62	-92	-246	3.4
EdF	174	2.8%	166	2.6%	1,728	17.8%	-510	-1,029	-1,547	6.8
Suez	161	2.6%	146	2.3%	469	4.8%	-26	-166	-307	1.4
ENEL	154	2.5%	129	2.0%	310	3.2%	30	-63	-156	0.0
Endesa	131	2.1%	103	1.6%	393	4.1%	-26	-144	-261	0.1
CEZ	120	1.9%	127	1.3%	120	1.9%	72	36	0	0.4
EnBW	43	0.7%	44	0.7%	222	2.3%	-44	-111	-177	1.4
Iberdrola	36	0.6%	25	0.3%	36	0.6%	22	11	0	0.2
Sum	1,817	29.2%	1,683	26.6%	5,053	52.2%				30.2

Source: Own computations based on data from EUTL and firms' business reports

The main result of this section is that all firms in the sample received a higher free allocation than the threshold defined by (17), such that their profits are an increasing function of the permit price. Whether this led them to actively manipulate the permit price upwards is another matter, to which I turn in the following section.

4 Empirical evidence for price manipulation

I start discussing some traditional measures of market power and why they are not appropriate in the current context. I then turn to a test of market power that relies on an over-purchasing of allowances.

4.1 Traditional measures of market power

One of the most frequently used measures for market power is the Hirschman-Herfindahl Index (HHI), which is based on firms' market shares. The ten firms in the sample generate about 52 % of the electricity demanded in the EU and were responsible for 29% of all emissions

¹⁰To see this, set $x_1 = e_1$ and $\bar{x}_1^0 = 0$ in (17) and denote the firm's emission intensity by $\rho \equiv e_1/q_1$. If the permit price increases by $d\sigma$, its cost will increase by $\rho \cdot d\sigma \cdot q_1$, whereas its revenue increases by $dp/d\sigma \cdot d\sigma \cdot q_1$. If $dp/d\sigma > \rho$, the latter exceeds the former and the firm's profits increase with the permit price.

in the EU ETS (see Table 2). The HHI for the permit market, using the cap as the market base, is 0.02. This would generally be considered a low degree of concentration from an antitrust perspective. For example, the US Department of Justice considers markets with an HHI of over 0.15 to be moderately concentrated, and above 0.25 to be highly concentrated (US Department of Justice, 2010), whereas the EU states that there is no concern with firm mergers in a market with a HHI of below 0.1 (European Commission, 2004).

However, the emissions cap may be the wrong basis for calculating the HHI. Market participation in the EU ETS was quite low initially, because many firms first had to become acquainted with the new market. Especially smaller firms, which tended to be over-allocated, appeared to be cautious about the new market and were reluctant to sell their surplus allowances (Jaraite et al., 2010). This lack of liquidity suggests that electricity firms may have had the power to influence the permit price by sustained purchases, especially during the initial phase, despite owning a minority of all permits. Instead, one might define the relevant basis as the share of the total cap allocated to installations that traded at least once by a certain date. By April 2008, the last month to surrender allowances for phase I, this "marketed" share is 95% of the cap, whereas the corresponding number by April 2007 and April 2006 are 70% and 51%, respectively. Using the latter as the basis for the HHI would lead to a value of 0.08.¹¹ Note that this may still severely underestimate the number of permits that were held by firms that actively observed the price and readjusted their permit holdings and abatement strategies on a continuous basis, as is implicitly assumed in a market.

Furthermore, the results from the previous section show that all firms profited from a higher permit price. This profit increase does not come at the expense of industry rivals, but of consumers (who pay higher prices) and taxpayers (who forgo a reduction in their tax rates from recycling the value of auctioned allowances). It is possible that although a firm may not use resources to drive up the allowance price itself but free-ride on the efforts of others, it will also not spend resources to counteract these efforts. In contrast, a firm may well invest resources when competing for output market share. For these reasons, it is not obvious that the HHI is a meaningful measure in a new market where firms are in the process of learning.

¹¹These numbers are based on a subset of countries for which it is possible to match the installation ID between EUTL's "Operator Holding Accounts" and "Transactions" subsections. This subset includes 16 EU countries that together account for 86 % of total allocation: CZ, DE, EE, ES, FI, FR, GB, HU, IT, LT, LU, LV, NL, PL, PT and SE. Note also that for the ten firms in the sample, I use actual allowance holdings; for the remaining firms, I did not compute holdings information and instead assume that holdings are equal to emissions.

Another well-known measure of market power is the Lerner index, defined as the proportional markup (the difference between price and marginal cost, divided by the price), or in its adjusted version, as the absolute value of the inverse of a firm's price elasticity of demand. The Lerner index describes a situation where a firm uses its market power to make a profit by selling something above marginal cost. But as I show above, most firms are net permit buyers, not sellers. Pindyck (1985) proposes an equivalent measure for the case of monopsony, where the dominant firm uses its market power to buy below marginal value (and thus decrease the price). But again, this does not fit the present situation where firms are net buyers, but due to the permit-electricity-market interaction find it optimal to increase the permit price. Perhaps yet an alternative Lerner-type index could be devised that takes this market interaction into account. Besides the permit price, such an index would have to incorporate either firms' marginal abatement costs, or the elasticity of the permit supply that they face when making their permit purchase decisions. Unfortunately, neither of these two is known.

Even if marginal abatement costs could somehow be estimated, there are at least two reasons other than market power that could explain why firm under-abate by setting their marginal abatement costs below the permit price. First, firms had reasons to expect that free allocation in phase II was going to be based on emissions during phase I (which indeed turned out to be the case), creating a strong disincentive to abate. The reason is that in addition to the actual cost associated with abatement, reducing another unit of emissions leads to a decrease in the expected future free allocation (Bohringer and Lange, 2005; Harstad and Eskeland, 2010); this has become known under the term of allocation "updating". And second, if abatement requires irreversible investment, the uncertainty inherent in future abatement costs leads to a real option value of waiting until more information becomes available (Dixit and Pindyck, 1994; Zhao, 2003). Firms will only engage in irreversible investment if they are compensated for this loss of flexibility, which means that the allowance price is equal to the technical cost of reducing emissions by a ton plus the option value.

4.2 Evidence based on allowance holdings

Because neither the HHI nor price-based measures are applicable in the present context, I assess the probability of firms exerting market power by focusing on their allowance holdings. More specifically, I investigate whether firms over-purchased (or under-sold) allowances to the extent that their expected emissions constraint was no longer binding. This strategy is robust to the lack of information about firms' marginal abatement costs. Conditions (8)–(9) show that for $x_1 > e_1$, it follows that $-C_e^1 = \lambda = 0$. Therefore, a positive permit price combined with $x_1 > e_1$ directly implies that $\sigma > -C_e^1 = 0$, which is the condition for the firm increasing the permit price by setting its marginal abatement costs below it. Last, note that $x_1 > e_1$ is a sufficient but not a necessary condition for $\sigma > -C_e^1$, as under-abatement is possible for an interior solution as well.

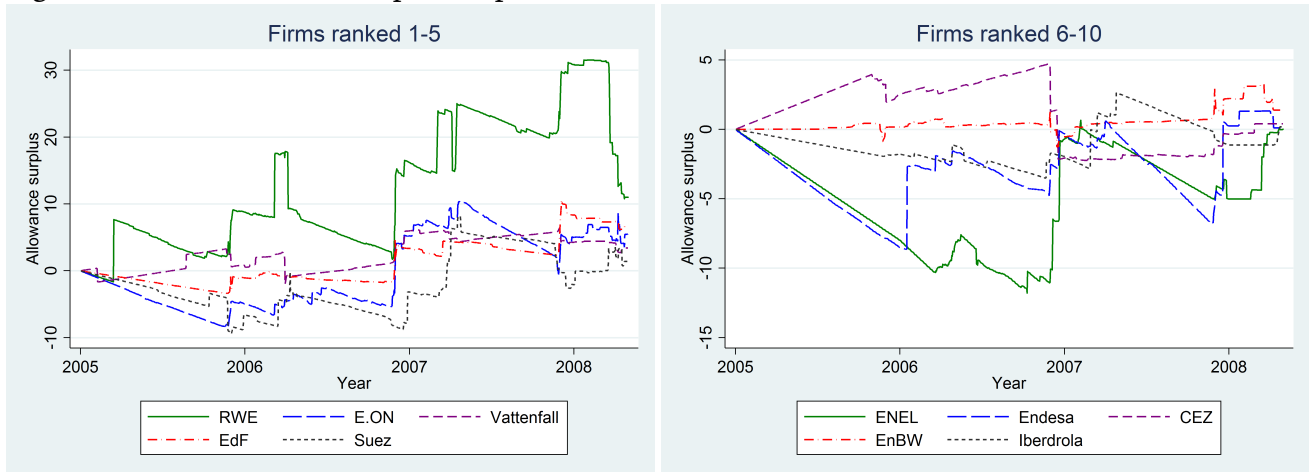
Focusing on excess allowance holdings also removes the two alternative reasons for deviation between the allowance price and marginal abatement costs discussed in the preceding subsection: While allocation "updating" can be expected to reduce abatement, it should not lead to over-purchases, as future allocation is based on past emissions, not permit holdings. Likewise, it is not clear how irreversible investment under uncertainty should lead firms to hold more allowances than necessary for compliance.

Figure 2 shows a time series of firms' cumulative allowance surplus over time, defined as cumulative allowance allocation plus trades minus cumulative emissions, under the assumption that annual allocation and emissions are spread evenly across the year. A positive value means that a firm holds more permits than needed to cover its cumulative emissions to date, and vice versa. Firms appear to have chosen different strategies: Some built up a permit surplus early in the phase, whereas other firms held a deficit in the beginning, building up a surplus only later on. Most held an allowance surplus at the end of the phase.

Holding a temporary allowance deficit is not necessarily inconsistent with price-taking behavior. Firms only have to surrender allowances once a year (end of April) to cover the emissions from the previous calendar year. Because they receive the next year's free allocation before April, their emissions constraint is effectively not binding until the end of the phase.¹²

¹²This is not visible in Figure 2, because I distributed free allocation throughout the year to show the actual net position. In practice, firms received their free allocation in February of each year. Also note that the final month relevant for emissions (January 2005 through December 2007) differs from the period during which firms could reach compliance (through April 2008).

Figure 2: Net allowance surplus in phase I



Likewise, a temporary allowance surplus could arise naturally if a firm is slightly short of allowances on average, but rather than buying a small number every week it purchases a larger number in the beginning of each year. Finally, the EUTL records actual allowance transfers only, whereas a significant share of allowances are traded via future contracts, meaning that the day of the trade and the physical transfer of allowances may not coincide.¹³ The fact that many futures contracts expire at the end of quarters, and especially at the end of the year, most likely explains the discrete changes in the cumulative allowance surplus. All of this implies that focusing on a temporary allowance surplus or deficit is not useful.

As a more meaningful indication for excess allowance holdings, and one that is not simply explained by firms' timing of allowance purchases and sales or by the divergence of futures trade and delivery dates, I propose the following definition:

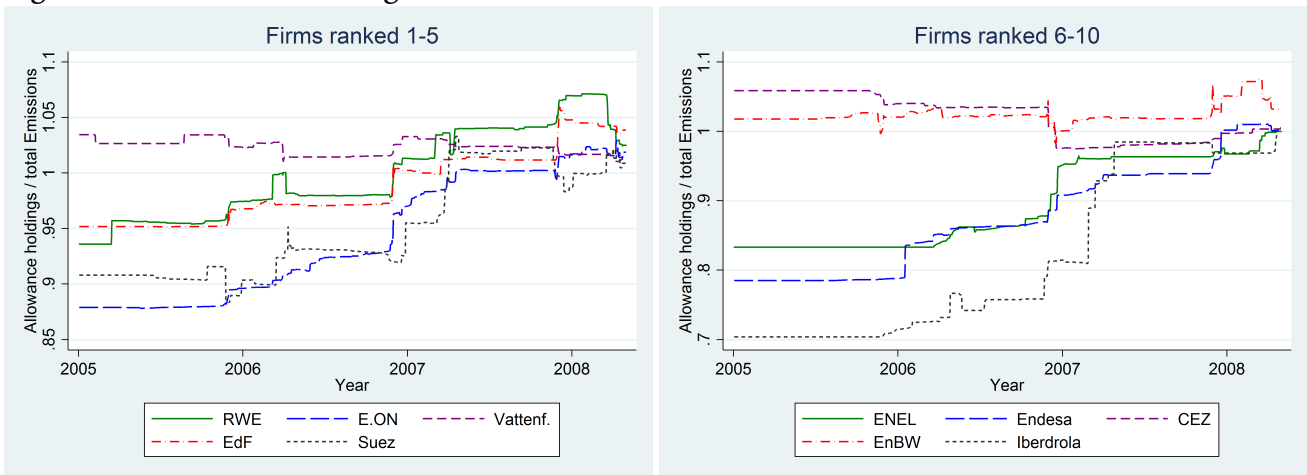
Definition: A firm that is a net allowance buyer holds an excess of allowances if these exceed the firm's *entire* emissions during phase I at any one point in time. For a net seller, an additional requirement is that it holds an allowance surplus at the end of the phase.

The additional requirement for the net seller is necessary, because a price-taking net seller may have sold excess allowances in a forward trade that expired at the end of the phase. Figure 3 shows total allowance holdings (cumulative trades plus the entire free allocation for the phase) relative to total emissions for the firms in the sample; at the very beginning of the phase, this ratio is equivalent to firms' net positions in the market. A holdings-to-emissions

¹³EUA futures contracts are traded on EEX, ICE, Nasdaq OMX, and until 2012 also on BlueNext. All EUA futures contracts traded on these exchanges were settled physically in the relevant time period (this continues to be the case for ICE and Nasdaq OMX, whereas EEX has switched to financial settlement).

ratio above unity is a sign of excess holdings.

Figure 3: Allowance holdings relative to total emissions



The figure shows the main result of this subsection: The five largest firms in terms of emissions plus EnBW (ranked 9th) engaged in excess allowance holdings according to the definition given above, either by over-purchasing or under-selling. In contrast, the three Spanish firms and CEZ do not fulfill this definition.

Excess allowance holdings are consistent with the theory developed in Section 2: Because firms' profits increase with the permit price (due to their allocation exceeding the threshold, as shown in Section 2), manipulating the allowance price upwards by over-purchasing or under-selling allowances is a profit-maximizing strategy. However, there are two alternative explanations that are consistent with excess allowance holdings: Price speculation, and hedging against stochastic carbon emissions. I examine these alternative explanations in the next two subsections.


4.3 Price speculation

Firms may have held more allowances than they needed to cover their emissions if they were expecting the allowance price to increase towards the end of the phase. If the sectors covered by the EU ETS had been unable to cut their emissions to below the cap, the allowance price would have been equal to the sum of the penalty for noncompliance (40 euro per ton of emissions that was not covered by an allowance) plus the forward price for phase II, because in addition to paying the penalty firms had to surrender the missing allowances in

the following year (Hintermann, 2012). Considering that the allowance price never exceeded 33 euro, large-scale noncompliance would indeed have led to a price increase towards the end of the phase.

However, a firm that expects a price increase would over-purchase allowances in the beginning of the phase and sell at the end, by taking a short position in a future contract and/or selling allowances on the spot market. The most likely expiration dates for end-of-phase futures are December 2007 (last end-of-year contract), March 2008 (last quarterly contract) and April 2008 (last month to sell any allowances). Figure 4 shows the total and individual transfers during these months for the six firms that held excess allowances. RWE, E.ON, EdF recorded positive transfers, which is inconsistent with price speculation as a reason for their excess allowance holdings. Vattenfall, Suez and EnBW indeed sold more permits than they purchased during these months, but neither of these firms sold their entire surplus, such that all of them ended up with too many allowances at the end of the phase.

Figure 4: Transfers at the end of the phase



Price speculation can therefore be ruled out as an explanation for excess allowance holdings for RWE, E.ON and EdF, and it can only serve as a partial explanation for Vattenfall, Suez and EnBW. I will return to this issue in the context of a robustness check.

4.4 Hedging against stochastic emissions

Future emissions are stochastic from the firm's point of view because of uncertain demand for electricity as well as stochastic fuel prices. Holding an excess of allowances can then be

viewed as a method of hedging against the risk of having to pay the penalty for noncompliance in case the cap turns out to be binding (Chesney and Taschini, 2012; Hintermann, 2012), rather than a sign of price manipulation. In addition to the financial risk, firms may also want to avoid a situation of noncompliance for reasons of public perception. Note also that the trading data shows that firms partly balance out their futures positions over time, because on many days firms' registry accounts report both an inflow and an outflow of allowances; this is consistent with firms hedging against uncertain emissions and adjusting their portfolios as new information becomes available.¹⁴

If a risk-averse firm is uncertain that it will be able to purchase more allowances at the end of the phase, it may choose to hold enough allowances to cover its entire emissions for the phase with a sufficiently high probability. Formally, let a firms' allowance surplus at time t , and its remaining cumulative emissions until the end of the phase T be defined by

$$S_t \equiv \bar{x} + \sum_{k=0}^t x_k - \sum_{k=0}^t e_k \quad (18)$$

$$V_t \equiv \sum_{k=t+1}^T e_k \quad (19)$$

Because firms' future emissions are uncertain, V_t is a stochastic variable, whereas S_t is known and continuously adjustable at any point in time. If a firm wants to limit the probability of not being able to cover its emissions by the end of the phase to some value α , it will choose S_t^* such that

$$\text{Prob}(S_t^* \leq V_t) = \alpha \quad \forall t \quad (20)$$

Because emissions are only available on an annual frequency, I estimate the probability that a firms' total allowance holdings are insufficient to cover its total emissions for the phase, or

$$\text{Prob}\left(S_t^* + \sum_{k=0}^t e_k \leq V_t + \sum_{k=0}^t e_k\right) \leq \alpha \quad (21)$$

¹⁴Presumably, the long and the short positions for the same delivery date are taken on different trading dates in response to changing expectations on total emissions.

Since adding a known constant to both sides of the equation has no effect on the probability, this is equivalent to (20).

In the following I compute firms' allowance holdings that are consistent with $\alpha = 0.05$ for the beginning of each quarter between 2006q3 and 2007q4 and compare them with firms' actual holdings. I do this by engaging in an exercise of "uninformed prediction" and using the estimates to compute emission forecasts and the associated prediction error. Arguably, firms' own forecasts of their generation and emissions are arguably more precise than one based on publicly available information.¹⁵ A conclusion that a firm's excess allowance holdings cannot be explained by hedging against carbon risk, because it holds more allowances than S_t^* is therefore conservative in the sense that the firm's true prediction error, and thus its own solution for S_t^* , is smaller than the one presented here.

Because information about emissions is available only since 2005 and even then only on an annual level, I base my regressions on electric generation by fossil sources and argue that to the extent that firms are able to predict fossil generation, they are also able to predict the associated CO₂ emissions. For E.ON and Vattenfall, fossil generation is available on a quarterly level since 2004, and for RWE since 2005 (data from business reports). EdF publishes output by energy source on an annual level since 2000. For the other two firms for which Figure 3 shows excess allowance holdings, Suez-Electrabel and EnBW, not enough generation information is available to estimate a meaningful prediction model. For this reason, I restrict the analysis in this subsection to RWE, E.ON, Vattenfall and EdF.

I estimate separate sets of regression depending on the relevant prediction horizon. For the earliest vantage point (July 1, 2006), this implies estimating six different regression with the minimum lag order for the explanatory variables ranging between one and six. For the latest vantage point (October 1, 2007), only a single regression is required, with a lag order of one. Using the individual quarterly predictions I compute the amount of expected fossil generation during the next 1-6 quarters (for the German firms) or the next 1-2 years (for EdF) as well as the forecast error. Finally, I convert the output prediction and the forecast error into CO₂ emissions using the respective firms' emissions intensity.¹⁶

¹⁵Firms have superior information in the form of more precise and frequent generation data in the past and planned outages. More importantly, their output is influenced (although not completely determined) by a series of strategic decisions, which are proprietary information.

¹⁶For EdF, annual data is available for the years 2000-2013 for generation within France only. To compute total emissions by EdF, I multiplied generation in France by emissions from total generation, divided by fossil

As explanatory variables, I use lags of the dependent variable, country-level fossil generation, electricity prices and stock market indices (I use the FTSE index for Germany and France), all of them lagged between once and six times. I further include quarterly dummies, a trend, and a dummy marking the nuclear accident in Fukushima that is equal to zero before 2011q1, and to one afterwards (regressions involving German firms only; note that nuclear generation in Germany was significantly reduced after the Fukushima accident and has not reached previous levels since).¹⁷

Figure 5: Actual and predicted fossil generation (3 quarter-horizon)

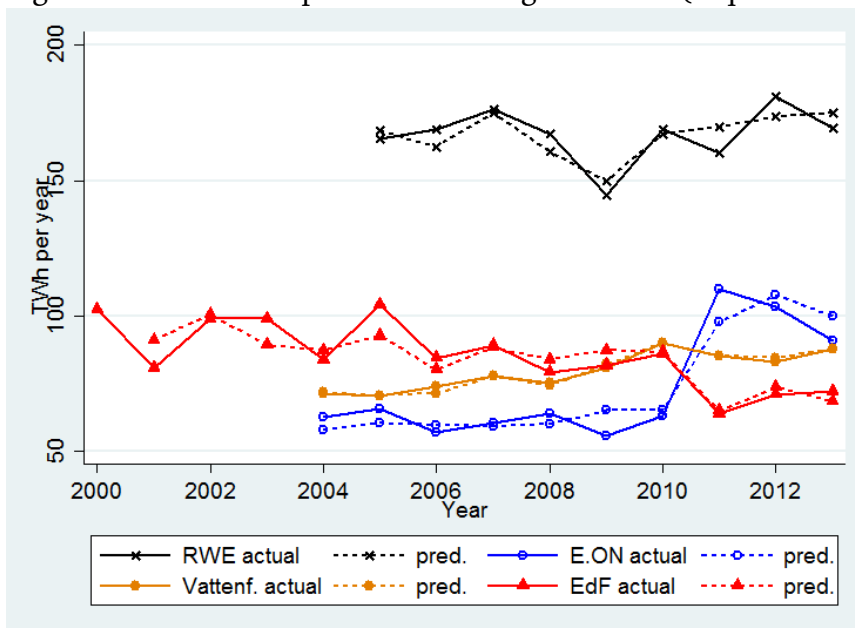


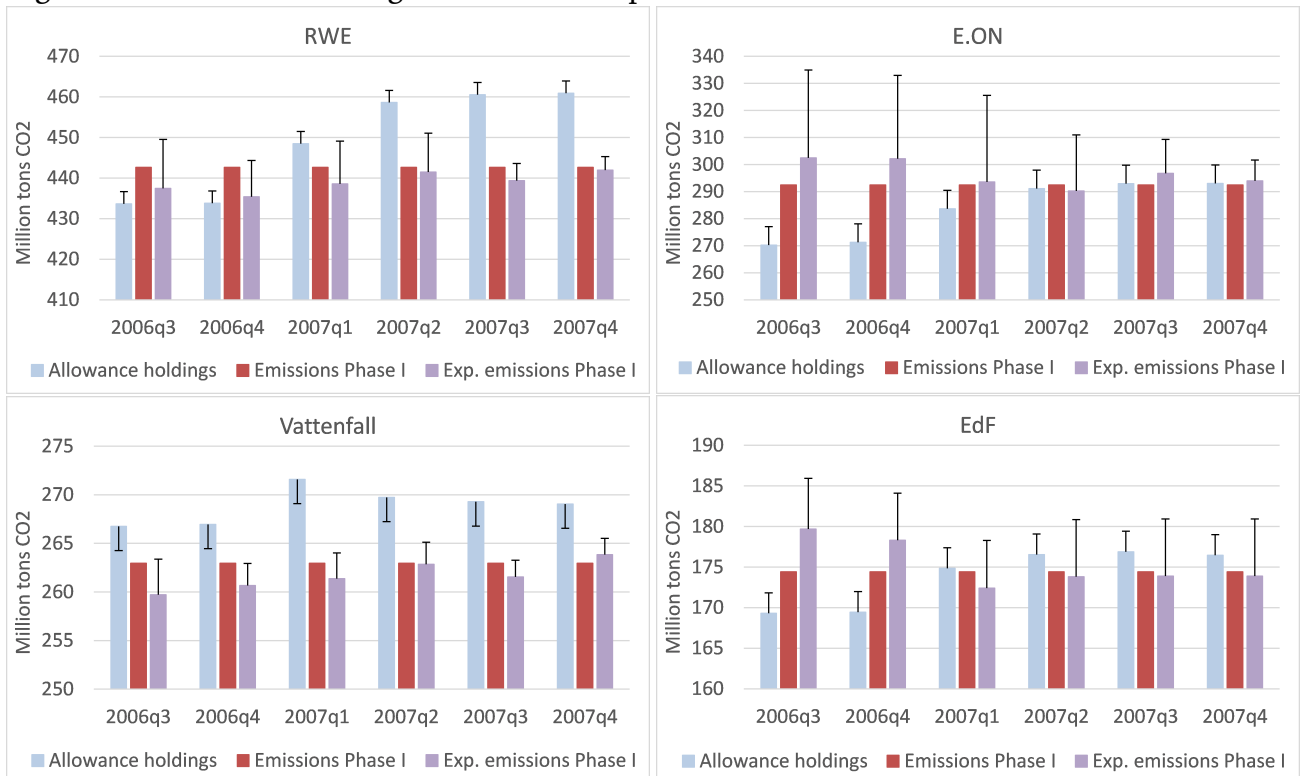
Figure 5 shows annual fossil generation along with a prediction made at the start of the second quarter. This prediction is used to compute the remaining emissions in the phase from the vantage point of April 1, 2007. The quality of the prediction varies dramatically between firms. This has to do with firms' fossil generation portfolio (RWE and Vattenfall produce a significant share of their generation using lignite, which varies much less than generation by hard coal or natural gas), but also with the available data. The least precise prediction is possible for EdF, which publishes only annual output data.

Figure 6 shows firms' total allowance holdings along with their actual and predicted emissions during phase I, with the prediction made at the beginning of the quarter noted on the generation in France only. The underlying assumption is that EdF is able to predict its total fossil generation, including generation in other countries, as well as it can predict generation in France.

¹⁷Note that the inclusion of data through 2013 in the production forecasts introduces another reason for my forecast to be less precise than that made by firms at the time, because the significant increase in solar and wind generation capacity in recent years has made it more difficult to predict fossil generation due to the intermittency of these energy sources.

horizontal axis. While the actual emissions during the phase are constant, their prediction varies with the vantage point in time. The error bar on top of the prediction is the 95%-upper confidence interval, based on the standard deviation of the forecast error and the degrees of freedom given by the number of predictions minus one. As would be expected, this error decreases as the forecast horizon decreases; in other words, the prediction of phase-I emissions in July 2006 is less precise than the prediction made in October 2007.

Figure 6: Allowance holdings vs. actual and predicted emissions



Note: The positive error bars on expected emissions represent the 95th percentile of the prediction error. The error bars on allowance holdings represent the net transfer position during the months December 2007, March 2008 and April 2008.

The figure contains the main results of this subsection and shows that RWE and Vattenfall held excess allowances even when accounting for the uncertainty of future emissions. These firms held more allowances than needed to cover the 95th percentile of their emission prediction at several vantage points in time. In contrast, the excess allowance holdings by E.ON and EdF never exceed the 95th percentile such that they could be explained by these firms engaging in precautionary allowance purchases to hedge against their carbon risk.

As a robustness check, I treat the allowance transfers in December 2007, March 2008 and April 2008 as future contracts that were made well before these expiry months. By doing

this I assume that firms knew that they were going to buy or sell these additional allowances at the different vantage points, although the transfers were not yet visible in the registry. The addition of this potential futures balance is represented by the error bar (positive or negative) on allowance holdings. Qualitatively, not much changes: Both RWE and Vattenfall continue to hold more allowances than can be explained by hedging against carbon risk. For RWE, the 95th percentile is now exceeded from an additional vantage point. The fact that Vattenfall's excess allowance is beyond the upper confidence interval of expected emissions even when accounting for its allowance sales towards the end of the phase suggests that price speculation does not explain its excess permit holdings. For E.ON and EdF, the conclusion remains unchanged: Their excess allowance holdings could be a sign of price manipulation, but also be due to precautionary purchases in the presence of carbon risk.

5 Conclusions

Because firms aim to maximize overall profits rather than minimize emission compliance costs, the interaction between the emission and output markets is crucial for understanding the effect of imperfect competition on permit market efficiency. I show that even net permit buyers can have an incentive to inflate the permit price by over-purchasing permits and under-abating emissions. I derive these results under the assumption of "single" market power in the sense that the dominant firm ignores the effect of its output quantity decision on the product price. These results are a special case of the results in Hintermann (2011), but they highlight the insight that it is not the simultaneous manipulation of two markets that leads to this qualitative departure from Hahn's (1984) seminal prescription of full allocation, but the simple fact that emission costs are reflected in the output price.

Using data about permit holdings and generation during phase I of the EU ETS, I show that the largest ten electricity firms would indeed have profited from increasing the allowance price, even though most of these firms were net allowance buyers. Focusing on excess allowance holdings as a potential sign of price manipulation, I find that six firms held more allowances than they needed to cover their entire emissions during phase I. This result is consistent with strategic price manipulation. I then examine two alternative explanations for why firms may have wanted to hold excess allowances.

First, excess allowance holdings may be explained by price speculation if firms believed the permit price to increase towards the end of the phase. However, the pattern and amount of allowance transfers at the very end of the phase suggests that price speculation is an unconvincing explanation for firms' excess allowance holdings earlier in the phase. Second, because firms did not know with certainty what their emissions were going to be, they may have engaged in precautionary allowance purchases. Based on a set of prediction models I find that this explanation may apply to E.ON and EdF (for which only a very imprecise prediction was possible due to the lack of quarterly output data), but most likely not to RWE and Vattenfall, because these firms held allowances even in excess of the 95th percentile of expected emissions. This conclusion remains unchanged when treating transfers at the end of the phase as the delivery of futures contracts purchased by firms at an earlier time in the phase. I conclude that price manipulation is a likely explanation for the excess allowance holdings by RWE and Vattenfall, and a possibility (along with alternative explanations) for E.ON, EdF, Suez-Electrabel and EnBW. In contrast, the remaining four firms in my sample did not engage in excess allowance holdings. This rules out universal collusion among European electricity generators.

There may be yet other explanations for why firms chose to buy more allowances than they needed (or sell fewer than they could have), and for which I do not test here. It is unlikely that there will ever be a conclusive proof of market power, because firms' choices are partly driven by expectations and beliefs that are not observable. However, considering that these are sophisticated firms with longstanding experience in making strategic decisions and that their profits were an increasing function of the permit price, explaining their excess allowance holdings as resulting from bad luck, limited rationality or incomplete information does not seem very convincing.

On the other hand, the market concentration according to the Herfindahl Index is rather low. This suggests that price manipulation is more likely to have taken place during the first 1–2 years of the market, when many firms had not yet begun to participate so that the large electricity firms held a large share of the marketed allowances and were the only large-scale buyers in the market. In contrast, it is difficult to imagine how price manipulation can be sustained if all firms actively participate in the market and continuously adjust their allowance

holdings according to their marginal costs of emissions abatement. Limiting upward price manipulation to the beginning of the market is also consistent with the very low price during the last few months of phase I. Unfortunately, I cannot test for the strength of market power over time, because I cannot differentiate between spot and future trades in my data. An allowance that enters a firm's registry on a given day could have been purchased on that day in a spot transaction, but also much earlier in a forward contract; also, it could already have been sold in another, not yet expired futures contract.

A natural way to test for the persistence of price manipulation in the EU ETS would be to use data from phases II or III. However, this poses two empirical challenges. For one, data about firms' allowance transfers are available only with a delay of five years, such that the complete data for phase II is not available to date. More importantly, because allowances can be banked between the second and later phases, a finding of excess allowances in phase II is indistinguishable from banking allowances into phase III. Investigating allowance price manipulation in later market phases requires a measure for market power that is different from the one used in the context of phase I, where surplus allowances expired unused.

My results suggest that market power is not only a theoretical curiosity, but an issue that is likely policy-relevant in the context of emission permit markets, especially (and perhaps exclusively) so in the beginning. Cap-and-trade systems are being set up in various countries and regions, and regulatory authorities should be cautious about large firms' incentives to manipulate prices during the initial years. In particular, allocating firms a large share of their permits for free is most likely the wrong policy in terms of market efficiency, because it creates incentives to manipulate the permit price upwards, and thus also the societal cost of achieving the emissions goal. Since the chief argument in favor of emission permit markets is their cost advantage relative to command-and-control approaches, price manipulation could severely reduce or even reverse the efficiency gains of these markets. On the other hand, the experience of the EU ETS suggests that firms do learn and that market participation increases over time, such that the relevance of market manipulation can be expected to decrease with the maturity of a market. From a policy perspective, keeping the learning phase to a minimum may therefore be the most promising means of fighting market power in emission permit markets.

References

- Bohringer, Christoph and Andreas Lange (2005) “On the Design of Optimal Grandfathering Schemes for Emission Allowances.,” *European Economic Review*, Vol. 49, No. 8, pp. 2041 – 2055.
- Bushnell, James B., Howard Chong, and Erin T. Mansur (2013) “Profiting from regulation: Evidence from the European carbon market,” *American Economic Journal: Economic Policy*, Vol. 5, No. 4, pp. 78–106.
- Cetin, Umut and Michel Verschuere (2009) “Pricing and Hedging in Carbon Emissions Markets.,” *International Journal of Theoretical and Applied Finance*, Vol. 12, No. 7, pp. 949 – 967.
- Chesney, Marc and Luca Taschini (2012) “The Endogenous Price Dynamics of Emission Allowances and an Application to CO2 Option Pricing.,” *Applied Mathematical Finance*, Vol. 19, No. 5-6, pp. 447 – 475.
- Disegni Eshel, Dafna M. (2005) “Optimal Allocation of Tradable Pollution Rights and Market Structures.,” *Journal of Regulatory Economics*, Vol. 28, No. 2, pp. 205 – 223.
- Dixit, Avinash K. and Robert S. Pindyck (1994) *Investment under uncertainty*: Princeton University Press.
- Ellerman, A Denny and Juan-Pablo Montero (2007) “The efficiency and robustness of allowance banking in the US Acid Rain Program,” *The Energy Journal*, pp. 47–71.
- Ellerman, Denny, Claudio Marcantonini, and Aleksandar Zaklan (2015) “The EU ETS: Eight years and counting,” *Review of Environmental Economics and Policy (forthcoming)*.
- European Commission, (2004) *Guidelines on the assessment of horizontal mergers under the Council Regulation on the control of concentrations between undertakings*: Official Journal of the European Union, 2004/C 31/03.

- Fabra, Natalia and Mar Reguant (2014) “Pass-through of Emissions Costs in Electricity Markets,” *American Economic Review*, Vol. 104, No. 9, pp. 2872–2899.
- Fell, Harrison, Beat Hintermann, and Herman Vollebergh (2015) “Carbon content of electricity futures in Phase II of the EU ETS,” *The Energy Journal*, Vol. 36, No. 4, pp. xx–xx.
- Godby, Robert (2002) “Market Power in Laboratory Emission Permit Markets.,” *Environmental and Resource Economics*, Vol. 23, No. 3, pp. 279 – 318.
- Hahn, Robert W. (1984) “Market Power and Transferable Property Rights.,” *Quarterly Journal of Economics*, Vol. 99, No. 4, pp. 753 – 765.
- Harstad, Bård and Gunnar S. Eskeland (2010) “Trading for the future: Signaling in permit markets,” *Journal of Public Economics*, Vol. 94, No. 9, pp. 749–760.
- Hintermann, Beat (2010) “Allowance Price Drivers in the First Phase of the EU ETS.,” *Journal of Environmental Economics and Management*, Vol. 59, No. 1, pp. 43 – 56.
- (2011) “Market Power, Permit Allocation and Efficiency in Emission Permit Markets.,” *Environmental and Resource Economics*, Vol. 49, No. 3, pp. 327 – 349.
- (2012) “Pricing Emission Permits in the Absence of Abatement.,” *Energy Economics*, Vol. 34, No. 5, pp. 1329 – 1340.
- (2014) “Pass-through of CO₂ emission costs to hourly electricity prices in Germany.” CESifo Working Paper No. 4964.
- Hintermann, Beat, Sonja Peterson, and Wilfried Rickels (2015) “Price and market behavior in Phase II of the EU ETS,” *Review of Environmental Economics and Policy* (forthcoming).
- Holland, Stephen P. and Michael R. Moore (2012) “When to pollute, when to abate? Intertemporal permit use in the Los Angeles NO_x market,” *Land Economics*, Vol. 88, No. 2, pp. 275–299.
- Jaraite, Jurate, Frank Convery, and Corrado Di Maria (2010) “Transaction Costs for Firms in the EU ETS: Lessons from Ireland.,” *Climate Policy*, Vol. 10, No. 2, pp. 190 – 215.

- Jaraite, J., T. Jong, A. Kazukauskas, A. Zaklan, and A. Zeitlberger (2013) “Ownership Links and Enhanced EUTL Dataset.” European University Institute, Florence; available at <http://fsr.eui.eu/CPRU/EUTLTransactionData.aspx>.
- Jaraitė-Kažukauskė, Jūratė and Andrius Kažukauskas (2014) “Do Transaction Costs Influence Firm Trading Behaviour in the European Emissions Trading System?” *Environmental and Resource Economics*. Online First.
- Kolstad, Jonathan T. and Frank A. Wolak (2008) “Using Environmental Emissions Permit Prices to Raise Electricity Prices: Evidence from the California Electricity Market.” Working paper, available at http://web.stanford.edu/group/fwolak/cgi-bin/sites/default/files/files/kolstad_wolak_final.pdf.
- Krattenmaker, Thomas G. and Steven C. Salop (1986) “Competition and Cooperation in the Market for Exclusionary Rights,” *American Economic Review*, Vol. 76, No. 2, pp. 109 – 113.
- Liski, Matti and Juan-Pablo Montero (2005) “A Note on Market Power in an Emission Permits Market with Banking,” *Environmental and Resource Economics*, Vol. 31, No. 2, pp. 159 – 173.
- (2011) “Market Power in an Exhaustible Resource Market: The Case of Storable Pollution Permits,” *Economic Journal*, Vol. 121, No. 551, pp. 116 – 144.
- Mansanet-Bataller, Maria, Angel Pardo, and Enric Valor (2007) “CO2 Prices, Energy and Weather,” *Energy Journal*, Vol. 28, No. 3, pp. 73 – 92.
- Martin, Ralf, Mirabelle Muuls, and Ulrich J. Wagner (2015) “The impact of the EU ETS on firms: What is the evidence after nine years?” *Review of Environmental Economics and Policy* (forthcoming).
- Misiolek, Walter S. and Harold W. Elder (1989) “Exclusionary Manipulation of Markets for Pollution Rights,” *Journal of Environmental Economics and Management*, Vol. 16, No. 2, pp. 156 – 166.
- Montero, Juan-Pablo (2009) “Market power in pollution permit markets,” *The Energy Journal*, Vol. 30, No. 2 (special issue), pp. 1–28.

- Muller, R. Andrew and Stuart Mestelman (1998) "What Have We Learned from Emissions Trading Experiments?," *Managerial and Decision Economics*, Vol. 19, No. 4-5, pp. 225 – 238.
- Neuhoff, Karsten, Kim Keats Martinez, and Misato Sato (2006) "Allocation, Incentives and Distortions: The Impact of EU ETS Emissions Allowance Allocations to the Electricity Sector.," *Climate Policy*, Vol. 6, No. 1, pp. 73 – 91.
- Pindyck, Robert S. (1985) "The measurement of monopoly power in dynamic markets," *Journal of Law and Economics*, Vol. 28, No. 1, pp. 193–222.
- Salop, Steven C. and David T. Scheffman (1983) "Raising Rivals' Costs.," *American Economic Review*, Vol. 73, No. 2, pp. 267 – 271.
- Sartzetakis, Eftichios Sophocles (1997) "Raising Rivals' Costs Strategies via Emission Permits Markets.," *Review of Industrial Organization*, Vol. 12, No. 5-6, pp. 751 – 765.
- Seifert, Jan, Marliese Uhrig-Homburg, and Michael Wagner (2008) "Dynamic Behavior of CO₂ Spot Prices.," *Journal of Environmental Economics and Management*, Vol. 56, No. 2, pp. 180 – 194.
- Sijm, Jos, Karsten Neuhoff, and Yihsu Chen (2006) "CO₂ Cost Pass-Through and Windfall Profits in the Power Sector.," *Climate Policy*, Vol. 6, No. 1, pp. 49 – 72.
- Sinn, Hans-Werner and Ulrich Schmoltzi (1981) "Eigentumsrechte, Kompensationsregeln und Marktmacht-Anmerkungen zum 'Coase Theorem'. (Property Rights, Compensation Rules and Market Power-Comments on the 'Coase Theorem.' With English summary.)," *Jahrbucher fur Nationalokonomie und Statistik*, Vol. 196, No. 2, pp. 97 – 117.
- US Department of Justice, (2010) *Horizontal merger guidelines*: Issued August 19.
- Zhao, Jinhua (2003) "Irreversible abatement investment under cost uncertainties: tradable emission permits and emissions charges," *Journal of Public Economics*, Vol. 87, No. 12, pp. 2765–2789.