Offset Credits in the EU Emissions Trading System : A Firm-Level Evaluation of Transaction Costs

Helene Naegele*
DIW Berlin

June 15, 2015

Abstract

International offset certificates trade at lower prices than European Union Allowances (EUAs), although they are substitutes within the EU Emissions Trading System (EU ETS) for CO2. Firms therefore had a strong incentive to use the cheaper certificates up to the maximum quantity fixed by the regulator. However, a considerable number of firms did not use their offset credit entitlement and, by doing so, seemingly forwent profits. While most of the literature on emission trading evaluates the efficiency of regulation in a frictionless world, firms in reality incur managerial costs when complying with regulation. This study examines the use of international offset credits in the EU ETS, in order to assess the relevance of such managerial and information-related transaction costs. It establishes a model of firm decision under fixed entry costs and estimates the size of transaction costs rationalizing firm behavior using semi-parametric binary quantile regression methods. These costs appear to be sizable, as they prevent a fifth of all firms, especially small emitters, from using offset certificates. It appears that for most firms the bulk of these transaction costs stems from participation in the EU ETS in general, rather than additional participation in the offset trade.

JEL : C25, D23, H23, Q58.

Keywords : Environmental policy, EU ETS, emission trading, transaction costs, binary quantile estimation.

*Contact under hnaegele@diw.de. The author thanks Christian von Hirschhausen, Pio Baake and especially Aleksandar Zaklan for guidance and comments. I also thank participants at DIW Berlin PhD Workshop and 27th Workshop on International Climate Policy at ESCP Berlin.

1 Introduction

The EU Emissions Trading System (EU ETS) has the objective to achieve the EU's carbon emission goals at minimum cost. Instead of imposing a fixed tax, the policy determines an emission level and lets the market determine the equilibrium price, which economists have been advocating for a long time. Ideally, as this system ensures that all firms incur the same price for emissions, abatement should be realized where it is cheapest, so that the aggregate abatement cost is minimized. However, abatement costs are not the only costs arising from an emission trading scheme: just like any other regulation, this policy has to be implemented and managed by firms, causing a wide range of administrative, managerial and information-related transaction costs. Typically, such transaction costs are unobserved by the econometrician. Presumably, many firms themselves do not account explicitly for the value of their employees' time and resources spent in the course of EU ETS compliance and optimization. As a result, transaction costs often play a minor role both in academic and in policy discussions about emission trading. This study uses firm-level data to estimate these transaction costs and argues that their magnitude is relevant for some of the regulated firms and should thus be taken into account when assessing the efficiency of the EU ETS.

To identify transaction costs, I exploit an important aspect of EU ETS regulation: the possibility to use international offset credits. In order to favor the development of a global emission trading system, the EU ETS has been linked to the international certificate market created through the Kyoto framework. On aggregate, these additional foreign certificates increase the cap for European polluters and decrease their compliance cost. Offset credits have been cheaper than the EU Allowances (EAUs) all through Phase II of the EU ETS (2008-2012). The quantity of offset credits used in the EU is limited by a firm-specific offset entitlement fixed by the regulator. For the firms, offset usage was an unambiguous way to reduce emission cost. In this context, it is puzzling that a considerable share of regulated firms did not use any offset credits.

This study provides both a descriptive and an analytical contribution to the literature. First, it describes the observed offset usage behavior and gives an overview of how much firms could have saved in compliance cost if they had all used their offset entitlements. Among firms that failed to participate in the offset market, there are mostly small firms and more particularly those firms with relatively generous free allocations of European certificates. Across all firms, forgone profits add up to around €1.37 billion. In a second step, I argue that firms' reluctance to participate can be explained by transaction costs. Without such unobserved transaction costs, the offset entitlement would be an unequivocal "free lunch" opportunity. The large share of firms forgoing these profits can only be rationalized by the interference of some unobserved frictions: transaction costs, as defined in this study, can include employees' time/salaries, training and consultancy costs. They are assumed fixed and payable whenever a firm first decides to engage in offset trading or emission trading in general, as such they might also be called "entry costs".

The theoretical section lays out how transaction costs change the firms' optimization problem. It builds on Stavins (1995), however I introduce a second type of certificates and simplify by accounting for fixed transaction costs only. It establishes that such costs can make the firms' free allocation of permits non-neutral, as firms with allocations larger than their emission do not *need* to engage in emission trading: they can avoid transaction costs of active trading, so that they are less likely to use their offset entitlement. The model establishes a link between, on one hand, the decision to participate in the offset market and, on the other hand, the initial allocation status and the potential benefit from offset usage. The fundamental assumption is that a firm renounces the potential benefits from offset trading only if the incurred transaction costs are higher than these benefits.

The empirical section uses this insight to estimate the unobserved transaction costs necessary to rationalize firms' decision not to participate in the offset market. I use a binary quantile regression to estimate a series of quantiles of the transaction cost distribution. The empirical results show that participation in the offset market has relatively low transaction costs for firms which are already actively trading European certificates, with a median of 922. The remaining firms – those which can comply with their EU ETS obligations without buying any certificates – have a median cost of 7,725. For the large majority of firms, the additional cost for offset trading appears small compared to the cost of generally engaging in active emission trading. However, the distribution of these costs appears highly skewed, so that the means are much higher than the medians (29,176 for average general participation plus 52,132 for offset market participation), resulting from some large outliers. A simple regression at the mean would thus be misleading about the costs faced by the majority of firms. Although these transaction costs are often small compared to other production factors, they make active participation unprofitable for 21% of the firms. For bigger firms, investment in offset certificates remains profitable.

This study brings together elements, on the one hand, from literature on the use of offset certificates in the EU ETS (Trotignon, 2012a; Ellerman et al., 2014) and, on the other hand, theoretical literature on the impact of transaction costs on emissions trading (Stavins, 1995; Montero, 1998). Moreover, this research relates to contingent valuation theory and uses binary quantile methodology (Kordas, 2006; Belluzzo Jr, 2004).

While the abatement incentives of cap-and-trade schemes have been amply discussed, most of the literature makes implicit assumptions on the absence of frictions arising from practical management of compliance. However, emission trading – just like any other market transaction – is unlikely to be completely free of transaction costs. In his seminal paper, Coase (1960) underlines that the irrelevancy of initial property allocation for final resource allocation holds only if frictions are negligible. The theoretical importance for cap-and-trade regulation of such "costs to use the price mechanism" has been modeled by Stavins (1995) and Montero (1998). Stavins (1995) shows that a major problem arising from transaction costs is that they make initial allocation non-neutral, so that free allocation (like in Phase II of the EU ETS) has an impact on the resulting market equilibrium. Montero (1998) moreover adds the impact of uncertainty and technology constraints.

In practice, cost-efficiency gains from emission trading schemes are often claimed to be below expectations. Atkinson and Tietenberg (1991) argue that trading has been too scarce to reach a cost-efficient outcome; they claim that this inefficiency stems from the bilateral, sequential nature of trades (instead of a simultaneous centralized market).

Empirical evidence on transaction costs in environmental policy is relatively scarce, as McCann et al. (2005) regret in their literature review on this topic: transaction costs are rarely evaluated, maybe simply because of their latent (unobservable) nature. Literature suggests that transaction costs and other market imperfections have hampered the impact of US environmental trading programs (Tietenberg, 2006; Hahn and Hester, 1989). Concerning the EU ETS, it is found in general that small firms trade more "passively" and that many firms seem to lack institutional capacity for optimal trading (Sandoff and Schaad, 2009). A common strategy among German SMEs is to trade only at the end of the year and *only if* the grandfathered allocation is not sufficient (Loeschel et al., 2011). Surveys show that large emitters set up more sophisticated structures to optimize their compliance and face smaller per-tonne transaction costs (Heindl, 2012; Jaraite and Kazukauskas, 2012; Loeschel et al., 2010, 2011). Similarly supporting the idea of fixed costs, Jaraite et al. (2010) estimate that average participation costs of the largest firms were €0.05 per tonne of emissions, while they were up to €2 per tonne for small firms. Schleich and Betz (2004) underline that allocations are so generous that the average need for additional permits for SMEs is only about 1,250 tCO2e per year, an amount at which participation costs are likely to be higher than the actual certificate cost.

Virtually all empirical work on transaction costs in the EU ETS is based on survey-data, except Jaraite and Kazukauskas (2012): they use transaction data from Phase I (2005-2007), which was the test phase of the policy. They claim that transaction costs were a substantial factor stopping firms from actively trading EUAs, but they do not estimate their magnitude. The observed trading pattern gives another hint that entry costs are potentially relevant: it appears that firms trade rarely and most transactions take place between plants belonging to the same firm (Zaklan, 2012; Jaraite and Kazukauskas, 2012).[1]

The methodology used in this study approaches the problem differently than surveys. Anderson and Sallee (2011) identify marginal costs of regulating fuel-standards by observing to what extent automakers use a regulatory loophole of known costs to avoid the fuel-efficiency standards. Conceptually, this is close to the present study which identifies fixed costs by observing on the opposite what benefits firms forwent in order to avoid trading. Using binary choice to identify a latent variable, this study relates to the revealed-preference methodology used in nonmarket (contingent) valuation of environmental goods (Bennett and Blamey, 2001). However, the firm size distribution is strongly skewed and the fit of parametric models is poor, so that I estimate semi-parametric binary quantile model. Quantile models have been developed by Koenker and Bassett (1978), but have only recently been applied to binary choice (Kordas, 2006). Belluzzo Jr (2004) uses them to estimate the distribution of willingness-to-pay for some public good, which is analogous to the present study: transaction costs are measured here from the observed "unwillingness-to-benefit" of firms.

While the previously cited literature examines trading schemes with only one type of certificates, Trotignon (2012b) describes how offsets have been used in the EU ETS and shows that firms initially used few offsets until a sharp increase in 2011. He estimates the cumulated savings of firms at €1.5 billion. An aggregate view going up to the end of Phase II in 2012 is provided by Ellerman et al. (2014).

The remainder of this article is organized as follows. After explaining the institutional and legal framework of international offset certificates (section 2.1), I briefly explain the aggregate impact of offset trading in the EU ETS (section 2.2) and the definition of transaction costs in this context (section 2.3). I then set up a model of firm-behavior in the reference case, i.e. without any transaction/entry costs (section 3.1), which is extended by adding entry costs (section 3.2). Finally, I present the data and some stylized facts (section 4.1, 4.2 and 4.3), explain the econometric methodology (section 4.4) and present the estimated distribution of transaction costs (section 5).

[1]However, transaction data needed for such analysis is only available for the Phase I of the EU ETS.

2 Background

2.1 Institutional framework

The EU ETS follows a yearly rhythm: each year, the EU gives out EU emission allowances (EUAs) summing up to the overall emission cap for that year. In Phase II – which is the period under study here – virtually all these certificates were distributed free of charge to the regulated firms, according to their historic emission levels (called *grandfathering*). At the end of each year, firms have to report their emissions and hand in an according number of certificates: one for each ton of CO2, hence the use of "tons of CO2 equivalent" (tCO2e) to measure quantities of certificates. Used certificates expire, unused certificates can be banked across years.

In order to coordinate international efforts of emission reduction and to lower abatement cost for EU-based companies, the EU has linked the EU ETS to the international framework established by the United Nations Framework Convention on Climate Change (UNFCCC, 1992) and the Kyoto Protocol. According to these international conventions, firms can generate Certified Emission Reductions (CERs, from Clean Development Mechanism (CDM)) or Emission Reduction Units (ERUs, from Joint Implementation projects (JI)) by reducing emissions in unregulated regions of the world. CERs and ERUs are commonly called *international offset certificates*.[2]

Within their obligations from the EU ETS, firms could substitute a limited amount of European certificates with such offset certificates. Such a substitution was attractive because offset certificates are generally cheaper than EUAs. However, to ensure that the bulk of emission reduction was achieved domestically, the EU has restricted the quantity of offsets usable by each firm. The exact definition of this quantity limit depends on the national government, but most countries used a percentage share of 10 to 20% of the grandfathered allocation as a benchmark, cf. Table 3 on page 24 in the Appendix, yielding a firm-specific offset entitlement.

These quantity limits for offset use were set for Phase II. In the middle of Phase II (April 2009), an EU directive[3] announced that the usage limits of certain offsets should be transferable (*bankability*) into Phase III (2013-2020), however it was unclear what amounts and which types of certificates.[4] Due to institutional obstacles, the final regulation ensuring the bankability and its conditions, finally only appeared in November 2013,[5] i.e. *after* the original claim on Phase II expired. The present study assumes that firms could not predict that regulation in March 2013.

The new legislation states that firms can use in total over 2008-2020 the highest of the following: (a.) the international credit entitlement specified in the national allocation plan in Phase II; or (b.) 11% of the free allocation of EU allowances granted to them in Phase II; or (c.) 4.5% of their verified emissions in Phase III. The first two points simply aim at ensuring equity for companies from different countries. The last point is aiming at new entrants and firms considerably extending their activity; it barely impacts firms already established in Phase II.

[2]As CERs and ERUs can be used interchangeably under this legislation, I will from now on only use the term "offsets" while everything applies equally to CERs and ERUs.

[3]Directive 2009/29/EC

[4]Offset certificates have been criticized because they rely on the fundamental criterion of *additionality*, which is virtually impossible to ensure completely. Some types of certificates, in particular those from "industrial gas" projects, seem to easily manipulated so that they are not further accepted in Phase III of the EU ETS.

[5]Commission Regulation (EU) No 1123/2013

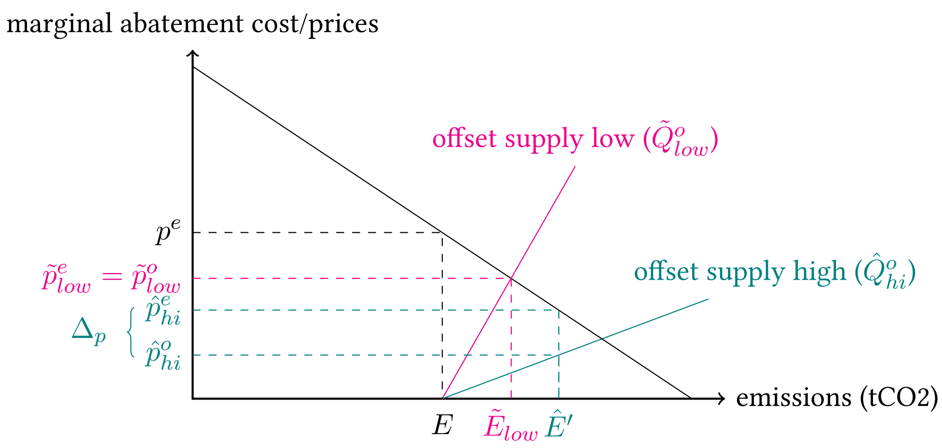

Figure 1: Aggregate market equilibrium, with two alternative offset supply levels

2.2 Why are offset certificates cheaper? – Theory of the aggregate impact of offset credits

International credits cover emissions from geographic regions which are not previously included in the scope of EU ETS. As such, they are a "spatial flexibility mechanism" (Stevens and Rose, 2002) allowing firms to abate where it is cheapest (other countries, esp. developing countries) and have the abatement credited *via* the use of offset credits for their EU-located plants. The creation of offset certificates increases the overall cap imposed by the EU ETS. Potentially, the cap could increase by an amount equal to the sum of all firms' offset entitlements. In practice, it depends on prices whether the regulatory offset quantity limit or the supply of offset certificates will determine the overall amount of certificates available.

The resulting market equilibrium is illustrated in Figure 1: without offset credits, the standard result holds, so that the market clears at the regulated maximum emission level E at price p^e, equal to the marginal abatement cost at E (Trotignon, 2012b). Offset certificates are perfect substitutes to EUAs up to the regulatory quantity limit. When offsets are expensive to produce (supply Q^o_{low}), their availability increases the overall cap, lowers the price and moves the equilibrium to $\tilde{E}_{low}$, where prices equalize at the level for which offset supply clears. This equalizes EUA and offset prices $\tilde{p}^e_{low} = \tilde{p}^o_{low}$. If offset creation is cheap (supply Q^o_{hi}), firms would like to buy more offset certificates than allowed. The aggregate offset quantity limit E' is binding in that case. The resulting constrained equilibrium at E', does not ensure equal prices anymore: EUAs trade at marginal abatement cost $\hat{p}^e_{hi}$ of the new emission level E'. The over-supply of offset certificates drives their price down to $\hat{p}^o_{hi}$ such that $Q^o_{hi} + E \leq E'$.

The price differential $\Delta_p = \hat{p}^e - \hat{p}^o$ is thus always positive ore zero; its magnitude depends on the difficulty to generate offsets of offsets and on the stringency of the offset usage quantity limit. Note that even if EUA and offset prices are not equalized, the introduction of offset credits nevertheless reduces EUA prices from p^e to $\hat{p}^e_{hi}$.

2.3 Definition of transaction costs

Beyond the direct cost of the emission certificates, the EU ETS causes a number of information-related and management frictions which I will summarize under the term transaction costs. As Heindl (2012) explains, the EU ETS produces costs through different channels:

- monitoring costs, as firms have to observe and report their emissions (which before were largely unknown),
- service charges of the EU registry ("formal" administrative costs),
- salaries of people employed by the firm for trading and information gathering,
- bargaining, contracting and transaction fees of the actual trading transactions.

The first two sources of participation costs are unavoidable and should thus not explain firms' non-participation in the offset market. The latter two sources of participation costs are directly related to firms' trading behavior and might explain why firms do not venture into the offset market. In the following, the term transaction cost will be defined as costs arising from trade (direct transaction costs) and from information gathering about market structure and management (indirect costs); it is likely to include personnel salaries, recruiting cost, consulting fees, etc. It *does not* include monitoring and reporting of emissions, administrative cost for EU/national agencies, and generally any other "unavoidable" cost. This is a more narrow definition than in some other work which considers the overall cost of establishing, managing, monitoring and enforcing a policy (Krutilla and Krause, 2010; Joas and Flachsland, 2014).[6]

These transaction costs are *a priori* unobservable. However, a multitude of news and data providers (Point Carbon), consulting firms (ICIS/Tschach) and financial transaction services (brokerage like TFS Green, exchange platforms like ICE) has emerged. Firms using this sort of costly services indicates that there must be some information problem. As the final transaction is virtually costfree, transaction costs are in this context largely due to upfront costs of information procurement. Just as an example, setting up a trading account at the ICE – the biggest exchange, clearing about 90% of emission certificate trade in Europe – costs €2,500 in direct fees,[7] while an individual transaction thereafter costs only cents.[8]

3 Model

First, this static model describes firm's optimization problem in presence of two types of emission certificates without transaction costs. In a second step, I examine how optimal behavior changes in presence of transaction costs. Simply put, firms always want to use offset credits, unless transaction costs are prohibitively high compared to potential profits from using cheap offset credits.

3.1 Emission trading with offset credits: least-cost scenario

For the purpose of this study, it is useful to look at firms' optimization problem aggregated over Phase II, which is qualitatively equivalent to looking just at the last year of offset validity.[9]As a reference case, this subsection examines emission trading with two types of certificates *without* participation

[6]In particular, this study concentrates on firms' costs and does not take into account what Joas and Flachsland (2014) call "public-sector costs" incurred by the regulatory authority.

[7]As indicated on https://www.theice.com/fees (01.03.2015)

[8]Convery and Redmond (2007) establish a list of direct transaction fees: brokers have relatively large minimum trade sizes and take between 1 and 5 cent fee per certificate (tCO2e). Exchanges take smaller trades and charge between 0.5 and 3 cent per certificate.

[9]Note that allocations and offset entitlements for the whole period were known to the firms from the beginning (from the NAPs).

costs. I show that firms can separate the decision of optimal emission levels (and produced quantities) from the partitioning between European and offset certificates. This point will considerably simplify analysis in the subsequent section 3.2, where we will concentrate on the latter decision.

In absence of offsets, it is straightforward and has been shown (e.g. by Montgomery, 1972) that there is a market equilibrium ensuring that marginal abatement cost is constant across firms and equal to the EUA price p^e. The present model adds a second type of certificates, so that each firm i solves the following optimization problem:

$$\max_{Y_i, E_i, Q_i^e} \pi = pY_i - C(Y_i, E_i) - T(Q_i^o, Q_i^e) + p^e A_i, \tag{1}$$

$$\text{subject to } E_i = Q_i^o + Q_i^e, \tag{2}$$

$$T(Q_i^o, Q_i^e) = p^o Q_i^o + p^e Q_i^e, \tag{3}$$

$$Q_i^o \leq K_i, \tag{4}$$

where equation (1) is the profit maximization with $C(Y_i, E_i)$ the production cost, which depends on emissions E_i and output Y_i.[10] As usual, I assume that increasing production Y_i (at a fixed emission level) increases cost $C_Y(Y_i, E_i) > 0$ and reducing emissions (at a given production level) increases cost $C_E(Y_i, E_i) < 0$.[11] Q_i^o are the offsets surrendered in Phase II and Q_i^e the surrendered European certificates. $T(Q_i^o, Q_i^e)$ is the cost of complying to ETS rules, i.e. the cost of buying the certificate quantities Q_i^e and Q_i^o necessary to justify the emission level E_i (equations (2) and (3)). Firms are given free certificate allocation A_i at the beginning of Phase II; they can sell superfluous certificates at market price p^e. The firm-specific constant K_i in equation (4) is the regulatory limit for offset usage. The overall amount of EUAs in the market is fixed by total allocations over all firms.

The firm has to solve three problems simultaneously: decide on the optimal produced quantity Y_i, determine the optimal emission level E_i and split compliance (i.e. an amount of certificates equal to E_i) between international offset and European certificates. The first-order conditions require quantity to be chosen optimally given production cost $C(Y_i, E_i)$ and prices. Let us assume that the production function C and prices p are such that there exists a function $Y_i^*(E_i)$ giving the optimal quantity produced for any given emission level at given prices.[12] Compliance cost $T(Q^o, Q^e)$ results from the cost incurred for both types of certificates. To satisfy the first-order condition for emissions, marginal abatement cost has to be equal to marginal compliance cost:

$$p\frac{\partial Y_i^*(E_i)}{\partial E_i} - \frac{\partial C(Y_i^*(E_i), E_i)}{\partial E_i} = \frac{\partial T}{\partial E_i} \tag{5}$$

$$= \frac{\partial T}{\partial Q^o}\frac{\partial Q^o}{\partial E_i} + \frac{\partial T}{\partial Q^e}\frac{\partial Q^e}{\partial E_i}$$

$$= p^o\frac{\partial Q^o}{\partial E_i} + p^e\frac{\partial Q^e}{\partial E_i} \tag{6}$$

The compliance cost arises from an optimal partitioning of certificates between EUAs and offsets,

[10] Emissions E_i and output Y_i as well as all other variables are here summed over the whole Phase II (2008-2012)

[11] C_Y and C_E denote the partial derivatives.

[12] A competitive market hypothesis simplifies this part, but is not essential to the subsequent argument, as long as there is a single equilibrium quantity $Y^*(E^*)$.

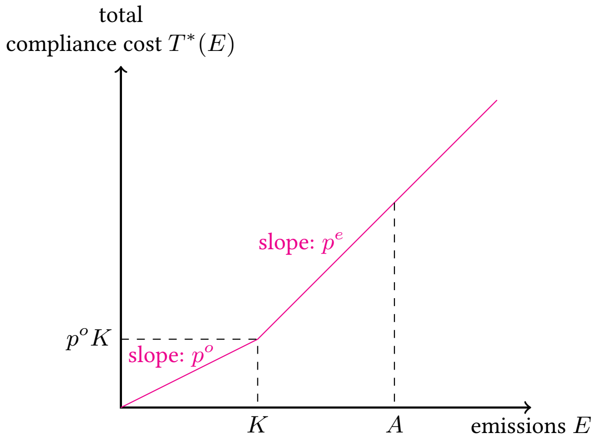

Figure 2: Deciding on the optimal quantity of EUAs and offsets (no entry cost, reference case)

given the price differential and the quantity restriction on offsets. The marginal cost is either p^e or p^o depending on which sort of certificate is used to cover the last (marginal) emission. As previously seen, offsets are perfect substitutes for EUAs up to a certain quantity limit; their price is thus at most as high as an EUA's price, but never higher: $p^e - p^o =: \Delta_p \geq 0$. For the purpose of this study, I will only consider situations in which offset certificates are strictly cheaper than EUAs, as the alternative where both prices are equal is qualitatively not different from a system without offsets.[13]

The result is straightforward and illustrated in Figure 2: as a perfect substitute at lower price, offset credits are clearly preferable to EUAs, up to the regulatory quantity limit. Only if emissions are above K_i, the firm will comply for remaining emissions by using the more expensive EUAs. Compared to a system with only EUAs, the firm saves an amount equal to $\Delta_p K_i$. The compliance cost can be simplified to $T^*(E_i)$ giving for all emission levels E_i the compliance cost resulting from an optimal split between European and offset certificates. The equation for $T^*(E_i)$ then enters the optimization problem as a constraint:

$$\max_{E_i} \pi(Y^*(E_i), E_i) = pY^*(E_i) - C(Y^*(E_i), E_i) - T^*(E_i) \tag{7}$$

$$\text{such that } T^*(E_i) = \begin{cases} p^o E_i, & \text{if } 0 < E_i \leq K_i \\ p^e(E_i - K_i) + p^o K_i, & \text{if } K_i < E_i \end{cases} \tag{8}$$

3.2 Entry costs for both certificate markets

We will now see how fixed participation costs on the offset market change the firm's problem. I assume that firms face some entry cost to generally participate in the certificate market, i.e. the cost of setting up a "trading department" no matter the type of certificates, and an additional cost to participate in the offset market. They can avoid both costs if they use only their freely allocated certificates. Obviously, firms with optimal emissions bigger than their allocation will be forced to buy certificates and cannot avoid the general participation cost. Compliance cost from equation (3) has now two additional terms:

[13] Moreover, the data reveals that in practice there has always been a clear price discount for offset certificates.

$$T(Q_i^o, Q_i^e, E_i) = p^o Q_i^o + p^e Q_i^e + \mathbb{1}^o T^o + \mathbb{1}^e T^e, \tag{9}$$
$$= p^e E_i + \mathbb{1}^e (T^e + \mathbb{1}^o (T^o - \Delta_p Q_i^o)), \tag{10}$$
$$\text{where } \mathbb{1}^o = 1 \text{ iff } Q_i^o > 0 \tag{11}$$
$$\mathbb{1}^e = 1 \text{ iff } Q_i^o > 0 \text{ and/or } Q_i^e > A_i \tag{12}$$

where a firm incurs general information entry costs T^e if it buys any certificates, but needs to pay an additional information cost T^o to participate in the less well-known offset market. This specification also implements the idea that firms which are "long" in equilibrium, i.e. which got more free allocations than needed for their optimal emissions ($A_i > E_i$), are not obliged to actively trade certificates. "Short" firms need to enter the market to buy some certificates anyways and should thus consider the general participation cost T^e sunk when deciding about offset usage.[14] The impact of transaction costs on offset usage and incurred total cost depends on the relative magnitudes of T^o, $T^o + T^e$ and $K_i \Delta_p$.

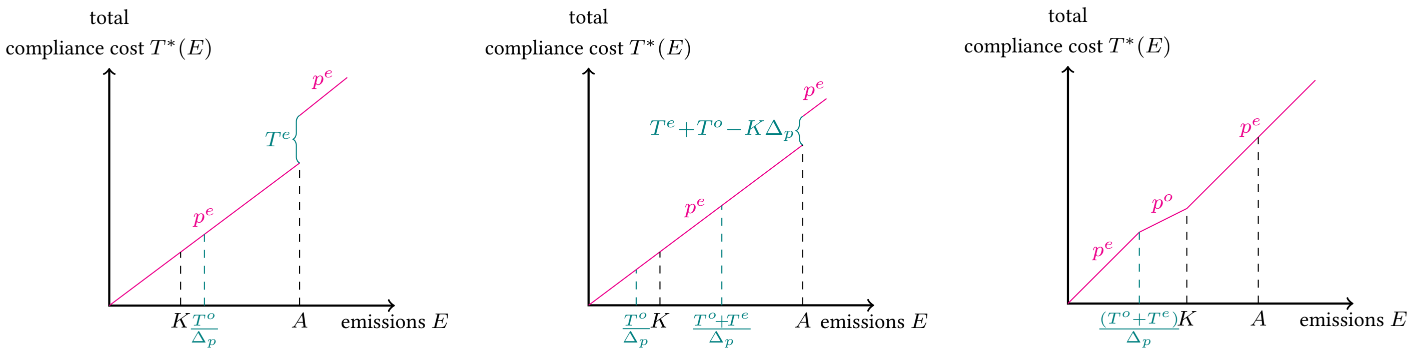

(a) High entry costs (b) Intermediary entry costs (c) Low entry costs

Figure 3: Deciding on the optimal quantity of European and offset certificates (with entry cost on both markets)

Figure 3 illustrates the three possible situations. In Figure 3a, offset entry costs are high so that $T^o > K_i \Delta_p$. In this case, entering the offset market is not useful at any emission level. Firms still have to incur entry cost T^e to enter the EUA market if their emissions are higher than their free allocation, which results in a discontinuity at $E_i = A_i$. In Figure 3b, T^o is relatively low, but $T^o + T^e$ is so high that offsets alone are unprofitable. As $T^o < K_i \Delta_p$, firms which already incur entry price T^e (because $E_i > A_i$) will also buy offset certificates. There is thus a similar discontinuity as in 3a, but the jump is reduced from T^e to $T^e + T^o - K_i \Delta_p$, because the firm cashes in some gains from offset usage. Finally, Figure 3c shows the situation if both entry costs are relatively low so that a firm uses offsets as soon as its emissions are above the threshold.

Cases (a) and (b) illustrate situations in which entry costs may make initial allocation non-neutral, as they produce a jump in the cost curve. The direct effect of participation costs on total compliance cost $T^*(E)$ does not impact the marginal cost-benefit analysis: both above and below A_i, firms face a marginal price of p^e. Even in case (c), the slope of the cost curve is p^e for most firms and only the very

[14] This definition is not ideal: rather than conditioning on firms which are forced to trade, one would like to condition on firms actually trading. However, the data does not allow this distinction. As a consequence, the estimate for transaction costs may be downward biased.

smallest firms face a marginal certificate price of p^o.

Let "allocation status" $\mathbb{1}_i^{long}$ be a dummy variable indicating that allocation A_i is larger than emissions E_i^*, such that optimizing compliance simplifies to the decision whether to use offset certificates:[15]

$$\max_{\{\mathbb{1}_i^o\}} \mathbb{1}_i^o(\Delta_p K_i - T^o - \mathbb{1}_i^{long} T^e) \tag{13}$$
$$\text{where } \mathbb{1}_i^o = 1 \text{ iff } Q_i^o > 0$$

A firm participates if it is worth incurring the entry costs, which depend on the allocation status – long or short – of the firm. The empirical section will use the prediction that a short firm not participating must imply that $\Delta_p K_i < T^o$ and a long firm not participating shows that $\Delta_p K_i < T^o + T^e$, while the same inequalities are inversed for participating firms. Note that this solution still implies an "all-or-nothing" decision as long as entry cost is fixed. Observe as well that in spite of these frictions, marginal abatement cost is still equalized across the large majority of firms at the level of p^e.

$$\mathbb{1}_i^o = \begin{cases} 1 & \text{if } \Delta_p K_i > T^o - \mathbb{1}_i^{long} T^e, \\ 0 & \text{otherwise.} \end{cases}$$

An important assumption is that firms take prices as given here: every individual firm is too small to consider its own impact on the price level. On the aggregate, p^e depends on the number of firms using offset certificates. To the extent that transaction costs reduce access to the offset market, they are not neutral for p^e and thus for Y^* and E^*. There will be a second-order effect, as participation costs impact the demand for offsets: this decreases the offset price p^o and increases the EUA price p^e. While these price effects are essential for a general equilibrium and welfare assessment, they are not informative on transaction costs and lay beyond the scope of this study.

4 Data and research design

A preliminary descriptive data analysis reveals some stylized facts, that my empirical analysis will rely on: (a.) offset certificates are indeed cheaper than European certifiactes, (b.) a large majority of firms has emissions superior to their offset entitlement and inferior to their free allocation, (c.) a non-negligible number of firms does not use their offset entitlements (22%) and (d.) the size distribution of firms is very unequal.

4.1 Data sources

This study mainly relies on the data of the European ETS Registry (European Union Transaction Log, EUTL) which is a compilation of member states' national registries of Phase I and II (2005-2012). This comprehensive administrative data comprises the allocated EUAs, verified emissions and surrendered certificates (EUAs, CERs and ERUs) for *all* 13,590 plants subject to ETS compliance obligations. Moreover, a matching with Bureau van Dijk's Orbis company database reveals ownership structures that

[15] See Appendix on page 27 for more details on the potential interaction of transaction costs and allocation status.

relate many of these individual plants.[16] This matching is important as the relevant decision is likely to happen at an aggregated firm level, even though regulation, allocation and offset entitlements are defined at plant level. After some data cleaning,[17] there remain around 9,000 plants belonging to 4,578firms. Almost half of the plants belong to firms owning just one plant.

The plant-specific offset quantity limit K_i is the product of a nation-specific quota of offsets multiplied by the sum A_i of a plant's free allocations over Phase II. The magnitude of this quota has been chosen by national governments, but the EU has restricted the maximum to 22%, as in Germany or Spain. For the purpose of this study, the limits have been computed by this rule and verified using the International Credit Entitlement tables published by the EUTL in 2014. Allocations have been relatively generous so that 80% of the firms could cover all of their emissions using only grandfathered allocations (they will further be called "long" firms). However, offset entitlement K_i has been so small that only a meager 2.8% of firms is able to comply by using offsets only. Table 1 shows that free allocation has on average been just above emissions. Firms have a wide variety of sizes, with some firms owning up to 156 plants and being active in 11 sectors and 16 countries.

	Mean	Median	SD	Min	Max
Number of countries active	1.13	1	.728	1	17
Number of plants	1.88	1	5.03	1	158
Number of sectors active (NACE definition)	1.12	1	.566	1	11
Free allocated EUAs (ktCO2e)	1,975	112	13,831	.015	380,586
Emissions (ktCO2)	1,919	78.5	16,148	.003	563,608
International credit entitlement (ktCO2e)	272	12	2,335	.001	91,537
Used offset credits (ktCO2e)	208	8.34	1,494	0	55,536
Savings from offset use (k €)	799	31.2	5,836	0	217,412
Unexploited profits from CERs (k €)	627	22	7,370	.00465	200,316
Firms using all offset entitlement (in %)	50.5	100	50	0	100
Firms using no offsets (in %)	22	0	41.4	0	100
Observations	4578				

Table 1: Descriptive firm statistics

4.2 Price spread and realized savings

As argued before, offsets are expected to trade at an inferior price to EUAs or at best at equal price if the offset supply is relatively scarce. Indeed, offsets have always traded at a positive discount from EUAs. Figure 4 shows that the price differential was rather small in the beginning. After few months, the spread became clearer and offsets have been up to €7 cheaper than EUAs. The spread increased with time and was rather volatile. On average the price difference was €3.60.

This price spread has allowed firms to achieve considerable savings, culminating around 217 million for the largest firm.[18] Altogether, firms have saved an overall amount of €3.6 billion. However, the

[16]For more information on this matching, see Jaraite et al. (2013) or their website http://fsr.eui.eu/CPRU/EUTLTransactionData.aspx

[17]Plants from countries which do not participate in the standard way described in section 2.1 (Norway, Estonia, Liechtenstein; 135 plants) and some which have offset-use beyond the legal limit (most likely because of merger and acquisition transactions which are unobserved in this data set; 94 plants) are excluded. Around 3,000 plants never register any emission or cease existing in 2011 and 2012, so they are excluded as well.

[18]This number is approximately computed by multiplying the annual average price spread with the amount of offset

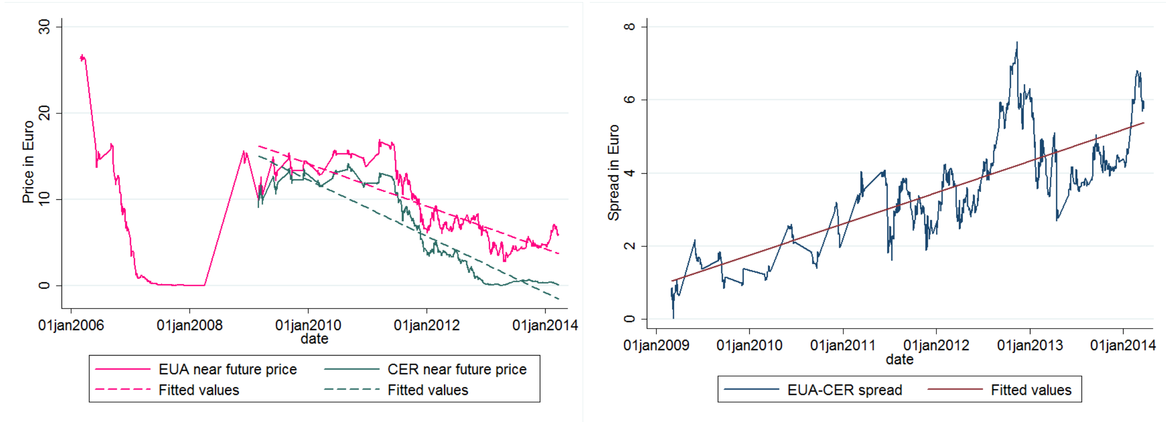

(a) Prices of EUAs and CERs on the secondary market

(b) EUA-CER price spread

Figure 4: Prices of EU and offset certificates (source: www.theice.com)

additional unused 288 million tCO2e certificates could have generated another 1.37 billion at 2012 prices. On average a firm has saved €806 thousand while the median is only €31.7 thousand, which results from some firms saving large amounts.

These numbers take prices as given, so they cannot be interpreted as the "overall saving" for firms from offset usage: as seen in section 2.2, the counterfactual EUA price in absence of offset credits would have been higher than the observed prices. The *de facto* achieved savings from offset usage are probably higher than my estimates used in Table 1. Stephan et al. (2014) estimate demand elasticity relatively high, so that actual firms' savings are as high as €20 billion, as offset availability decreased the overall stringency of the cap. Moreover, it does not include the incurred transaction costs.

4.3 Evidence for transaction costs

As mentioned before, many firms did not use their offset entitlements. Given the large supply of offset certificates (see Appendix on page 25 for details) and their relatively low price, this is surprising. One potential explanation for firms not participating in the market is that their expected pay-off was not high enough to cover transaction costs of information procurement, such as the cost of hiring additional personal or devoting existing resources to compliance optimization.

The stylized facts supporting this idea are (a.) a largely binary behavior between using either the maximum allowed or no offsets at all, (b.) the non-neutrality of EUA allocation status for participating in the offset market and (c.) an increasing likelihood of participating in the offset market as firm size and offset entitlement increase.

Supporting the idea of fixed participation cost, firms have mostly followed an "all or nothing" strategy in their offset usage: Figure 5 shows the surrendered offsets as a percentage of the total offset entitlement. One can see two frequency spikes: over half of the firms use *all* their offset entitlements and almost a quarter of the firms use *none*. Finally the last quarter of firms use some but not all of their offset entitlement. While per-unit costs would lead to a marginal trade-off and intermediary usage rates, a fixed cost incurred for market entry could explain such a binary behavior. It is interesting to note that many multi-plant firms with intermediary usage are composed of plants that expose an all

certificates surrendered in that year, because the actual transaction details (date/price) are not observed.

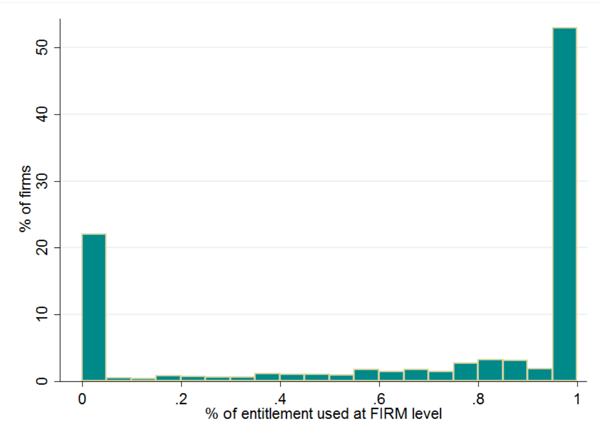

Figure 5: Ratio of used offset credits over overall offset entitlement (source: EUTL and own computations)

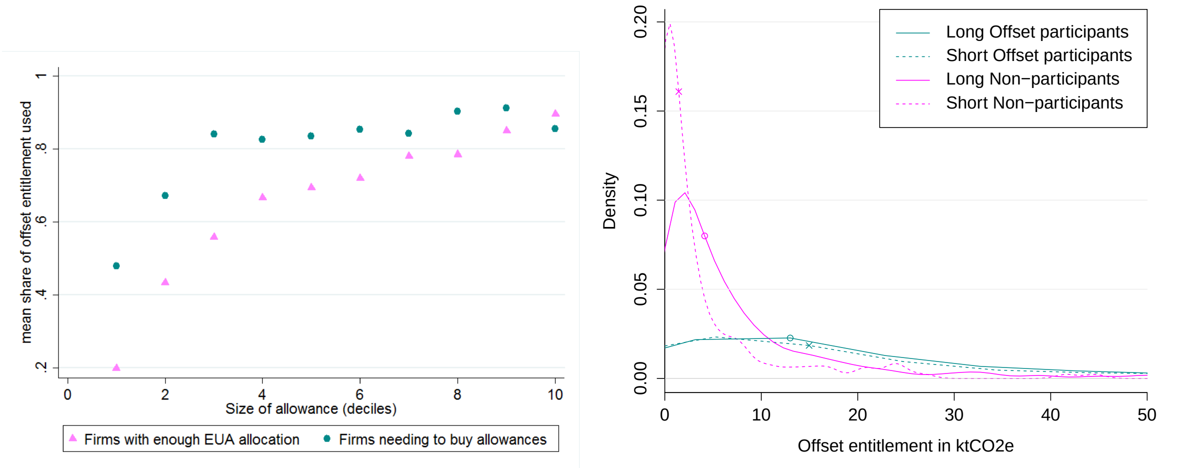

(a) Offset usage by allocation and according to size deciles of offset entitlement, or potential profit

(b) Distribution of offset entitlement in the different groups of firms

Figure 6: Relationship between offset participation, offset entitlement and allocation status (source: EUTL and own computations)[19]

or nothing behavior: it seems likely that this results from coordination problems within firms.

The main consequence of transaction costs is that they make initial allocation non neutral (Stavins, 1995). With fixed costs, firms with large credit entitlements should participate more as entry costs become small compared to the potential gain. Figure 6a shows the interaction between the size and the allocation effect: at lower size deciles, firms use relatively rarely offsets, with a large difference between long and short firms. As size increases, firms become more likely to use more offsets, while at the same time the difference between long and short firms becomes less marked. At the tenth size decile, virtually all firms participate and there is no significant difference between long and short firms' behavior.

Assuming that firms take rational decisions, plants that do not participate must estimate their

[19]Density estimation using gaussian kernel from density() in R, with smoothing bandwidths calculated by Silverman's rule of thumb; for readability, the graph is cut at 50 ktCO2e, although both densities continue beyond. Crosses and circles indicate median values.

participation cost to be *higher* than their potential profit, so that the mean offset entitlement multiplied by the mean price spread should give us a *lower bound* of the magnitude of these transaction costs (similar to the reasoning in Attanasio and Paiella, 2011). At the same time, the opposite is true for participating firms. These two distributions largely overlap, but Figure 6b shows that the means and medians are strongly different. In general non-participating firms tend to be smaller, with half of firms below 3,600 tCO2e of offset entitlements (while the median is 16,600 tCO2e for participating firms). Nevertheless, the distributions both stretch out until above 50,000 tCO2e, showing that the separation is not clear cut. The largest non-participant firm has 262,000 tCO2e entitlement, and the 9 percentiles of the potential profit distribution above this value *all* participate. Among participating firms, the size distribution of long and short firms is similar. On the opposite, small short firms are overrepresented in the non-participating group. This gives us an order of magnitude of avoided transaction costs.

Figure 6b shows that the size distribution of firms' offset entitlements is highly unequal, and similar inequality is true for emissions, number of plants and grandfathered allocations. The empirical methods used will need to be chosen such that they are robust to these rare and extremely large outlier firms.

4.4 Econometric methodology

The model gives us an indication about the link between firm behavior (using any offset credits or not) and the magnitudes of the unknown (to be estimated) entry costs T^o and T^e relative to the known quantities A_i, E_i and $\Delta_p K_i$. We want to measure the unobserved (latent) transaction cost TC_i^*, while observing only the binary outcome $\mathbb{1}_i^o$ equal to 1 if TC_i is smaller than $\Delta_p K_i$:[20]

$$\mathbb{1}_i^o = \mathbb{1}\{\Delta_p K_i > TC_i^*\} \tag{14}$$

$$= \mathbb{1}\{\Delta_p K_i > T^o + T^e \mathbb{1}_i^{long} + \epsilon_i\} \tag{15}$$

In this binary choice setup, $\Delta_p K_i$ is the firm-specific cut-off value relevant for the decision to participate. Other than in most binary choice settings, e.g. standard probit, a firm-specific cut-off allows us to identify an intercept as it fixes a scale for the two estimated parameters T^o and T^e in terms of units of $\Delta_p K_i$ (Euros).[21] This method to use preference revelation is similar to the methodology of contingent valuation, often used to analyze "willingness-to-pay" (WTP). Here, rather than estimating WTP, I interpret the foregone profits as "unwillingness-to-benefit" to identify transaction costs. Note that unlike most of the contingent valuation literature this study does not use survey methods, i.e. I am working with revealed preferences rather than stated preferences.

If the error term was assumed to be iid following a normal distribution, equation (15) would describe a probit model in which coefficients are normalized so that the coefficient of the potential profit equals 1. The other coefficients then measure transaction costs in Euros. This relates to the contingent valuation literature, where willingness-to-pay for certain attributes is usually estimated by normalizing the utility of income to 1.[22] However, the stylized facts presented in section 4 strongly suggest that this homoskedastic normality assumption does not hold. Firms have a highly unequal distribution of size

[20] We observe a transformation of the latent variable by an indicator function, which is a monotone transformation. See Koenker and Hallock (2001) on the equivariance of quantile estimates to monotone transformations.

[21] K_i is measured in tCO2e of offset entitlement and Δ_p is the average price spread measured in €per tCO2e.

[22] The standard normalization of a probit sets the standard deviation σ to 1; in contrast, the standard deviation is a free parameter here.

and activity range. Moreover, if the distribution of transaction costs is also skewed, an estimation of the mean cost is not the most representative summary statistic as they might be driven by some large outliers. Following empirical work from Kordas (2006) and Belluzzo Jr (2004), I estimate a range of binary quantile regressions to analyze the whole conditional distribution of transaction costs rather than the conditional mean. This semi-parametric method is more robust to non-symmetric error distributions and outliers. For all quantiles $\tau \in [0,1]$, I define the conditional quantile $Q(\tau)$ of the probability to participate in the offset market which depends on the distribution of the latent transaction cost TC_i^*:

$$Q_{\{\mathbb{1}_i^o\}}(\tau|\mathbb{1}_i^{long}, \Delta_p K_i) = \mathbb{1}\{\Delta_p K_i \geq T_\tau^o + T_\tau^e \mathbb{1}_i^{long} + \epsilon_i\} \tag{16}$$

The parametric probit draws its identification from the conditional mean assumption $E(\epsilon_i|x) = 0$, while the following methodology estimates the median and thus is identified over the assumption that the conditional median error is zero. The earliest estimator using this semiparametric assumption is the maximum score estimator from Manski (1975). At the median (with $\tau = .5$) this estimator maximizes the number of "correct predictions" using an indicator function:

$$\max_{T_\tau^o, T_\tau^e} S_{n\tau}(T_\tau^o, T_\tau^e; \Delta_p K_i) = n^{-1} \sum_{i=1}^{n} [\mathbb{1}_i^o - (1-\tau)] \mathbb{1}\{\Delta_p K_i - T_\tau^o - T_\tau^e \mathbb{1}_i^{long} \geq 0\} \tag{17}$$

Similarly to the median, we can estimate other conditional quantiles. While this estimator is relatively intuitive, it is not continuous which makes it difficult to optimize and determine standard errors. To solve this problem, Horowitz (1992) has formulated a smoothed maximum score estimator using a kernel function to get a continuous function of the estimated parameters. He provides comprehensive asymptotic theory for this estimator, which has been extended to other quantiles than the median by Kordas (2006). The smoothed binary quantile estimator at quantile $\tau \in (0,1)$ is the solution to the problem:

$$\max_{T_\tau^o, T_\tau^e} S_{n\tau}^*(T_\tau^o, T_\tau^e; h_n, \Delta_p K_i) = n^{-1} \sum_{i=1}^{n} [\mathbb{1}_i^o - (1-\tau)] \Phi\left((\Delta_p K_i - T_\tau^o - T_\tau^e \mathbb{1}_i^{long})/h_n\right) \tag{18}$$

where $\Phi(\cdot)$ is a continuous, differentiable kernel function and h_n an appropriate bandwidth which tends to zero as sample size increases.

As the estimation of this model involves the optimization over a complex function, I use R to implement Kordas' S-Plus/Fortran code to perform simulated annealing following the algorithm of Goffe et al. (1994). Simulated annealing has the advantage to be more robust to starting values, local optima and discrete parts of the objective function. With a large sample such as the one used in this study, results of Manski's discrete quantile maximum estimator and Horowitz' smoothed estimator turn out to be virtually identical. Standard errors for both estimators are calculated by bootstrap methods.

5 Estimation results

The binary quantile regression gives an overview of the distribution of transaction costs. Note that I do not estimate transaction costs for different sizes of firms, but the distribution of transaction costs from which each firm draws its realized transaction cost. As this distribution does not follow a known functional form, it is described here by estimating 19 quantiles, from the 5th to the 95th quantile in 5 point steps.[23]

The results are shown in the first two columns of Table 2 and Figure 7. The transaction cost components are measured in units of potential profit, i.e. in euros. For short firms, the median cost T^c is estimated around €922, which means that a firm with €922 potential offset benefit has a 50% chance of participating. While transaction costs are relatively low around €500 for the lower half of the transaction cost distribution, their values are high at the upper end. The distribution for T^e indicates that long firms (with generous initial allocations) are much more reluctant to participate. At the median, their behavior is consistent with an *additional* cost equivalent to €7,725. This goes up to the higher quantile estimates around €41,900 for $\tau = 0.95$.

The right-hand side of Table 2 shows the result if we separate manufacturing and electricity-generating firms.[24] The results are similar; again, the general transaction cost T^e is higher than T^o at most quantiles, even though it is unprecisely estimated for the electricity firm's cost distribution.

	All firms		Manufacturing		Electricity	
τ	$\hat{T}^C$	$\hat{T}^E$	$\hat{T}^C$	$\hat{T}^E$	$\hat{T}^C$	$\hat{T}^E$
0.05	34.1***	-8.9	943.8**	-21.4	29.5***	-7.1
	[24; 142]	[-98; 895]	[335; 1,314]	[-90; 2,691]	[22; 158]	[-98; 1,175]
0.1	30.8***	16.8	959.0**	932.2	30.0***	5.9
	[27; 304]	[-96; 3,256]	[353; 1,307]	[-67; 2,983]	[21; 256]	[-99; 2,102]
0.25	338.0***	2,954.9***	672.0***	2,924.9***	331.1***	2,848.7***
	[34; 810]	[2,350; 4,576]	[333; 1,375]	[952; 4,567]	[31; 877]	[833; 7,651]
0.5	922.8***	7,725.4***	967.0**	5,319.8***	879.9***	7,421.0***
	[343; 2,779]	[3,910; 10,606]	[277; 1,479]	[3,812; 10,305]	[331; 5,039]	[3,067; 14,172]
0.75	9,090.3***	18,133.2***	1,027.2***	21,032.0***	12,278.0***	14,847.3**
	[2,755; 12,944]	[8,601; 32,162]	[428; 12,221]	[9,830; 34,854]	[3,122; 25,653]	[1,613; 31,461]
0.9	28,537.3***	57,053.5**	20,890.4***	61,491.3**	85,807.4***	32,899.3*
	[17,785;100,673]	[1,069; 172,396]	[11,080; 96,965]	[6,385; 161,611]	[31,176;169,200]	[248; 153,951]
0.95	167,473.4***	41,932.0***	96,681.2***	106,720.6*	166,164.3***	25,773.5*
	[69,623;304,577]	[13,398;477,539]	[22,811;300,785]	[689;513,383]	[60,919;827,249]	[790;1,036,612]
Mean	52,132	29,176	82,566	25,018	60,825	63,388
Probit	109,557***	44,302***	169,389***	94,359***	-45,689	1,889
N	4,578		3,018		1,537	

Function optimized by simulated annealing, significance and 95% confidence intervals are determined by bootstrap (500 replications). Columns 1 and 2 show the result of the binary quantile regression, dependent variable is the offset participation dummy, regressors are offset entitlement (normalized to one), allocation dummy and a constant. Columns 3 to 6 show the result of the same regression with moreover dummies for sector affiliation (and their interaction with the allocation dummy).

Table 2: Quantile estimates from the smoothed maximum score estimator

The quantile analysis reveals that the transaction cost distribution spans a large range and is strongly skewed : while the difference between the median quantile and lower quantiles is quite small, there are large outliers driving the estimates of the highest quantiles. Consequently, the means (Table 2)[25] can be misleading about the transaction cost distribution. The probit estimates largely overstate

[23]For better readability, Table 2 shows only the deciles, while Figure 7 shows the full estimates for quantile steps of 0.05 (19 estimates).

[24]Sector-specific graphical representations of these results are in the Appendix on page 31

[25]Means from the quantile regression are computed with the following steps: (a.) estimate quantile parameters in 5% steps

these numbers in the general case and for the manufacturing sector, while they do not find significant transaction costs in the electricity and heat generation sector.[26]

Note that for virtually all quantiles, the impact of allocation status is stronger than the offset-specific transaction cost: the bulk of transaction costs stems from the general cost component T^e. This means that it is not the offset trading *per se*, but rather the cost of emission trading in general that stopped firms from using their offset entitlement. However this finding is completely hidden if we look only at the means, both from probit and from quantile regression (Table 2), which show that transaction cost for offset are on average larger than the ones for general trading. There are some large outliers in the distribution of T^o.

Figure 8 which compares the estimated probability of participating in the offset market from the probit and quantile model to the observed frequencies at different entitlement magnitudes. For the quantile estimation, this graph is simply the mirror image of Figure 7. Particularly for smaller emitters, the quantile method predicts participation probability much better than the probit. Analoguously, the fit of the quantile estimation is also quite good if evaluated with the method outlined by Kordas (2006) (cf. Appendix on page 32).

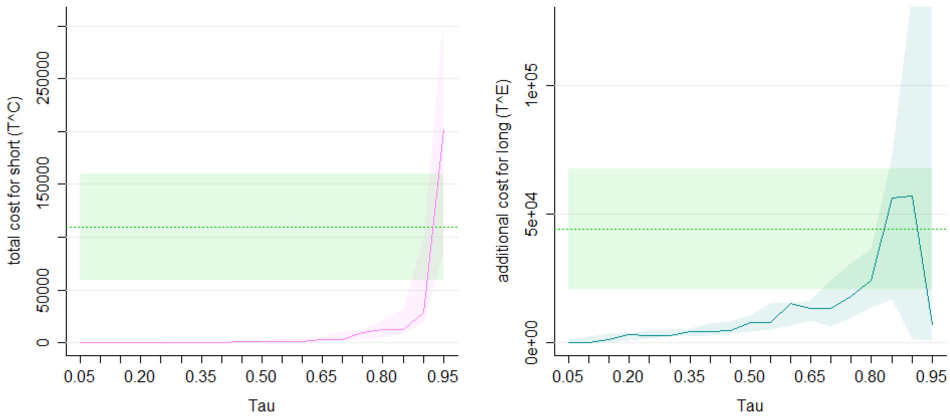

Figure 7: Estimated transaction cost (in €) - quantile plot
Quantile estimates for all 5th percentiles from 5% to 95%. The dotted green line is the mean estimate from probit, the shaded bands represent the 95% confidence intervals.

from the 5th percentile to the 95th, (b.) predict participation probability depending on firm characteristics (see Appendix G), (c.) impute transaction cost from τ equal to predicted probability, (d.) take average across all observed firms.

[26]More detail on these parametric estimations can be found in the Appendix on page 29.

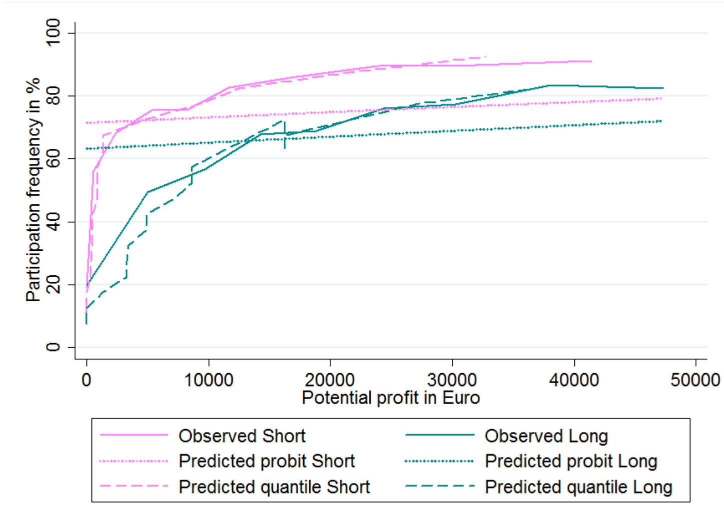

Figure 8: Observed frequencies and predicted probabilities of quantile method and probit (cut at 40,000 tCO2e for better readability)

6 Concluding remarks

Within their obligations from the EU ETS, firms had the opportunity to reduce expenses by using their right to substitute European certificates with international offset certificates. *A priori* it is thus always profitable to use offset certificates. However, many firms do not make full use of their offset entitlement. After briefly explaining the aggregate mechanics of offsets in the EU ETS, this study explains the theoretical impact of fixed transaction costs on offset usage and estimates the distribution of transaction costs necessary to rationalize firms' participation in the offset market.

Prior work has used survey data to show that compliance to the EU ETS generates managerial costs. To the best of my knowledge, this is the first study which establishes a framework to assess transaction costs empirically through the use of comprehensive administrative data. These entry costs are estimated to be at the median €7,725 (average €29,176) for general participation in the certificate market (be it EUA or offset certificates) plus €922 (average €52,132) for offset participation. Overall, the empirical results underline that the behavior on the offset market is significantly impacted by initial allocation: it appears that for most firms transaction costs are largely due to general participation in emission trading, rather than the offset market specific setup costs. However, this average hides a large heterogeneity that is best captured by a quantile estimation that suggests that these means are driven by some large outliers.

Environmental policy aims at reducing ecological harm at minimum cost to society. Most academic and policy-related work accounts only for direct compliance or abatement cost of the EU ETS. However – just like any regulation – the EU ETS causes administrative and management-related transaction costs. My estimates suggest that these costs are relevant in practice: firms significantly deviate from the least-cost scenario. Indeed, designing policy is "an empirical matter" as Montero (1998) puts it. Usually, optimal regulation aims at giving the optimal incentive structure, while this study argues that regulatory complexity also creates costs. As the objective of a regulation becomes more complicated, there appears to be a trade-off between incentive perfection and a need to keep complexity for the regulated firms at bay – incentives only work as intended if they are understood and implemented at low cost. In this perspective, this paper aims at contributing to the practical debate about the shape of environmental policy. Empirical evidence for transaction costs calls for more simple permit designs, rather than more sophisticated (but complicated) policy designs. The problem is even more stringent if the costs impact firms differently, such as the ones estimated in this study, where only large firms benefit from the cost reduction of offset certificates.

Note that I address only part of the actually arising transaction costs: all other costs that are unavoidable – such as monitoring and administration costs – do not affect behavior and can thus not be captured with my methodology. In a way, my estimates are thus the lower bound of the costs that should be included in the policy discussion. More importantly, these transaction costs are in no way synonymous to the overall efficiency loss: while effort spent in information gathering is certainly not welfare-improving, a real welfare effect analysis would need to look at the bigger picture of the general equilibrium. It would be interesting to estimate the impact of offset certificates on EUA prices, as well as to further dig into the price distortions (both on EUAs and offsets) caused by transaction costs.

The estimated transaction cost is a "black box" measuring *all* the frictions stopping firms from investing in offsets. It remains to be analyzed, what exactly these costs include and how they could consequently be reduced to implement a less distortionary policy. In fact, this study cannot differentiate between financial costs and more "behavioral" reasons, such as inattention, salience, risk aversion, misperceptions, etc. However, we are talking about the behavior of firms, so that psychological factors should play much less a role than they do for consumer decisions.

References

Aldrich, E. L and Koerner, C. L. Unveiling Assigned Amount Unit (AAU) trades: current market impacts and prospects for the future. *Atmosphere*, 3(1):229–245, 2012.

Anderson, S. T and Sallee, J. M. Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards. *The American Economic Review*, 101(4):1375–1409, 2011.

Atkinson, S and Tietenberg, T. Market failure in incentive-based regulation: the case of emissions trading. *Journal of Environmental Economics and management*, 21(1):17–31, 1991.

Attanasio, O. P and Paiella, M. Intertemporal consumption choices, transaction costs and limited participation in financial markets: reconciling data and theory. *Journal of Applied Econometrics*, 26(2): 322–343, 2011.

Belluzzo Jr, W. Semiparametric approaches to welfare evaluations in binary response models. *Journal of Business & Economic Statistics*, 22(3):322–330, 2004.

Bennett, J and Blamey, R. *The choice modelling approach to environmental valuation*. Edward Elgar Publishing, 2001. ISBN 178195660X.

Coase, R. H. Problem of Social Cost, The. *Journal of Law & Economics*, 3:1, 1960.

Convery, F. J and Redmond, L. Market and price developments in the European Union emissions trading scheme. *Review of Environmental Economics and Policy*, 1(1):88–111, 2007.

Ellerman, A. D, Marcantonini, C, and Zaklan, A. The EU ETS: Eight Years and Counting. *Robert Schuman Centre for Advanced Studies Research Paper*, (2014/04), 2014.

Elsworth, R, Worthington, B, and Morris, D. Help or Hindrance? Offsetting in the EU ETS. Report, Sandbag, 2012.

Goffe, W. L, Ferrier, G. D, and Rogers, J. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 60(1):65–99, 1994.

Hahn, R. W and Hester, G. L. Where Did All the Markets Go-An Analysis of EPA's Emissions Trading Program. *Yale J. on Reg.*, 6:109, 1989.

Heindl, P. Transaction costs and tradable permits: Empirical evidence from the EU emissions trading scheme. *ZEW-Centre for European Economic Research Discussion Paper*, (12-021), 2012.

Horowitz, J. L. A Smoothed Maximum Score Estimator for the Binary Response Model. *Econometrica*, 60(3):505–531, 1992.

Jaraite, J and Kazukauskas, A. Firm Trading Behaviour and Transaction Costs in the European Union's Emission Trading System: An Empirical Assessment. *CERE Working Paper*, (2012:9), 2012.

Jaraite, J, Convery, F, and Di Maria, C. Transaction costs for firms in the EU ETS: lessons from Ireland. *Climate Policy*, 10(2):190–215, 2010.

Jaraite, J, Jong, T, Kazukauskas, A, Zaklan, A, and Zeitlberger, A. Matching EU ETS Accounts to Historical Parent Companies. A Technical Note. *A Technical Note*, 2013.

Joas, F and Flachsland, C. The (ir)relevance of transaction costs in climate policy instrument choice: an analysis of the EU and the US. *Climate Policy*, pages 1–24, 2014.

Koenker, R and Bassett, J, Gilbert. Regression Quantiles. *Econometrica*, 46(1):33–50, 1978.

Koenker, R and Hallock, K. Quantile regression: An introduction. *Journal of Economic Perspectives*, 15 (4):43–56, 2001.

Kordas, G. Smoothed binary regression quantiles. *Journal of Applied Econometrics*, 21(3):387–407, 2006.

Krutilla, K and Krause, R. Transaction Costs and Environmental Policy: An Assessment Framework and Literature Review. *International Review of Environmental and Resource Economics*, 4(3-4):261–354, 2010.

Loeschel, A, Kiehl, K, Heindl, P, Lo, V, Koschel, H, and Koesler, S. Effizienzpotenziale des Emissionshandels noch nicht ausgeschoepft-Strategien und Management deutscher Unternehmen. Report, 2010.

Loeschel, A, Brockmann, K. L, Heindl, P, Lutz, B, and Schumacher, J. Hoher Anpassungsbedarf im EU-Emissionshandel ab 2013-deutliche Defizite bei der Vorbereitung in den Unternehmen. Report, 2011.

Manski, C. F. Maximum score estimation of the stochastic utility model of choice. *Journal of Econometrics*, 3(3):205–228, 1975.

McCann, L, Colby, B, Easter, K. W, Kasterine, A, and Kuperan, K. Transaction cost measurement for evaluating environmental policies. *Ecological economics*, 52(4):527–542, 2005.

Montero, J.-P. Marketable pollution permits with uncertainty and transaction costs. *Resource and Energy Economics*, 20(1):27–50, 1998.

Montgomery, W. D. Markets in licenses and efficient pollution control programs. *Journal of economic theory*, 5(3):395–418, 1972.

Sandoff, A and Schaad, G. Does EU ETS lead to emission reductions through trade? The case of the Swedish emissions trading sector participants. *Energy Policy*, 37(10):3967–3977, 2009.

Schleich, J and Betz, R. EU emissions trading and transaction costs for small and medium sized companies. *Intereconomics*, 39(3):121–123, 2004.

Stavins, R. N. Transaction costs and tradeable permits. *Journal of environmental economics and management*, 29(2):133–148, 1995.

Stephan, N, Bellassen, V, and Alberola, E. Use of Kyoto credits by European industrial installations: from an efficient market to a burst bubble. *Climate report*, 43, 2014.

Stevens, B and Rose, A. A dynamic analysis of the marketable permits approach to global warming policy: A comparison of spatial and temporal flexibility. *Journal of environmental economics and management*, 44(1):45–69, 2002.

Tietenberg, T. H. *Emissions trading: principles and practice.* Resources for the Future, 2006.

Trotignon, R. Combining cap-and-trade with offsets: lessons from the EU-ETS. *Climate Policy*, 12(3): 273–287, 2012a.

Trotignon, R. *In search of the carbon price: The european CO2 emission trading scheme: from ex ante and ex post analysis to the protection in 2020.* Thesis, 2012b.

Williams, R. Fitting heterogeneous choice models with oglm. *Stata Journal*, 10(4):540–567, 2010.

Zaklan, A. *Econometric Analyses of Carbon Resource Markets.* Thesis, 2012.

A National offset entitlement rules

	Annual Cap Ph.II (MMt CO2e)	Offset limit (%)	Annual offset limit (MMt CO2e)	Banking/ Borrowing	Industry	Energy	Other sector differentiation	Included in this study
Austria	30.7	10	3.1	Yes/yes				
Belgium	58.5	8.4	4.9	-	Flanders 24% Walloon 4%	Flanders 7% Walloon 8%		
Bulgaria	42.3	12.6	5.3	Yes/yes				
Cyprus	5.48	10	0.5	Yes/yes				
Czech Rep.	86.8	10	8.7	Yes/yes				
Denmark	24.5	17	4.2	Yes/yes	6.50%	28.70%		
Estonia	12.72	10	1.3	No/no	(started only in 2011)			No
Finland	37.6	10	3.8	Yes/Yes	8 / 8.5%	8.5 /9.5 /23.9%		
France	132.8	13.5	17.9	Yes/Yes				
Germany	453.1	22	99.7	Yes/Yes				
Greece	69.1	9	6.2	Yes/Yes				
Hungary	26.9	10	2.7	-				
Ireland	22.3	10	2.2	Yes/Yes	5%	11%	Cement 11%	
Italy	195.8	15	29.4	Yes/no	7.2%	Electricity 19.3%	"Other" combustion 7.2%	
					Ferrous metal 16.7%	Refineries 13.2%		
Latvia	3.43	10	0.3	Yes/Yes				
Lithuania	8.8	20	1.8	No/no				No
Luxembourg	2.5	10	0.3	Yes/Yes				
Malta	2.1	10	0.2	Yes/Yes				No
Netherlands	85.8	10	8.6	Yes/Yes				
Norway		13		Yes/No	13% of actual emissions (rather than allocation)			No
Poland	208.5	10	20.9	Yes/No				
Portugal	34.8	10	3.5	Yes/Yes				
Romania	75.9	10	7.6	Yes/Yes				
Slovakia	30.9	7	2.2	Yes/Yes				
Slovenia	8.3	15.8	1.3	Yes/Yes				
Spain	152.3	20.6	31.4	Yes/No	7.90%	42%		
Sweden	22.8	10	2.3	Yes/Yes				
UK	246.2	8	19.7	Yes/No	8%	9.30%		

Table 3: Offset limits collected from National Allocation Plans by Elsworth et al. (2012)

B Availability of offset certificates

The present study only examines the demand side of offset certificates. One possible alternative explanation for our finding would be that the constraint came actually from the supply side. However, this section argues that offset certificates have been available in larger quantities than demanded at all times throughout Phase II.

All issued CERs have to be authorized and validated by UNEP Risoe, which establishes a central registry. It is difficult to determine the number of CERs available for compliance in the EU ETS, as "Appendix I countries" (mostly Europe, United States and Japan) can use CERs to comply to their Kyoto obligations. The amount shown in Figure 9 is technically thus an upper bound of the amount disposable for compliance in the EU ETS.

On the state level, CERs are a substitute for "Assigned Amount Unit" (AAUs). It has been argued that there has been a surplus of AAUs within in the Kyoto framework. These certificates are traded infrequently and bilaterally, mostly directly between participating states, so that there is no clear market price. However, given the large AAU overallocation of ex-Soviet Union states (so-called "hot air"), the evidence suggests that AAUs are sold usually far below the price of EUAs, CERs and ERUs (Aldrich and Koerner, 2012). In that case, Appendix I countries have little interest to comply to their Kyoto obligations by using CERs, so that the total amount of CERs issued is likely to be a good approximation of the amount actually available to firms for compliance in the EU ETS.

Figure 9 shows that the amount of ERUs issued starts slowly but takes off in the last two years of Phase II. Overall, there were enough offsets available to fully exploit potential profits from substitution of EUAs by offsets by the end of Phase II. Although use rose sharply in the last year, it remained below the produced supply.

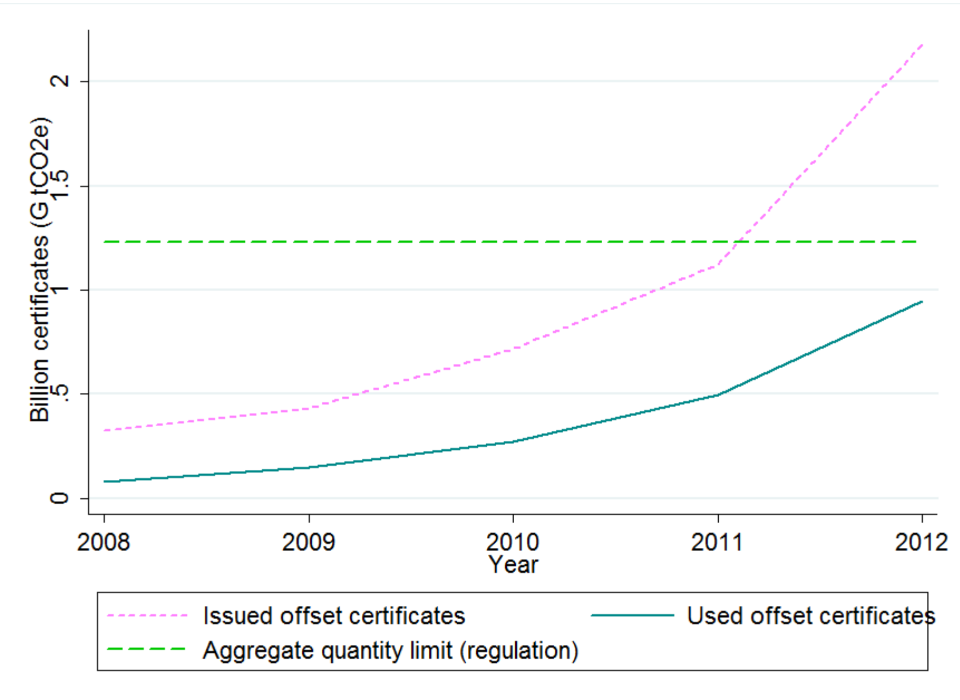

Figure 9: Availability of CERs and ERUs

Some CERs lost validity after compliance deadline of April 2013 because of quality concerns over their environmental integrity. Among canceled CERs were those generated through "industrial gas" methods which had constituted about 60% of the issued CERs. With their convertibility into EUAs these "gray" CERs as they are called now (in opposition to still valid "green" CERs) have lost virtually all their market value and traded at around 1 cent per ton after April 2013. Incentives have thus been particularly strong to submit or sell these certificates before the end of Phase II.

C Level of analysis: firms or plants

The EU registry used for this paper is recorded at the plant-year level. However, Jaraite et al. (2013) have matched plants to the companies which own them and many of these plants belong to the same firms. A "firm" is defined by ownership: the firm is defined an entity of which no other entity holds more than 50.01%. On average a firm owns 16 plants and the biggest firm in the sample holds 156 plants. As firms usually share many services (such as emission certificate trading) across their plants, it would make sense to analyze behavior at this level. Indeed, Table 4 shows that the large majority (71%) of plants belong to a firm that follows the same strategy as the overall company.

	Installation strategies			
	1_Used none	2_Used some	3_Used all	Total
	b	b	b	b
1_Firm used none	200	0	0	200
2_Firm used some	535	1047	2198	3780
3_Firm used all	0	0	862	862
Total	735	1047	3060	4842

Table 4: Comparison of plants' strategy and owning firm's strategy (counts, only among firms with more than 1 plant)

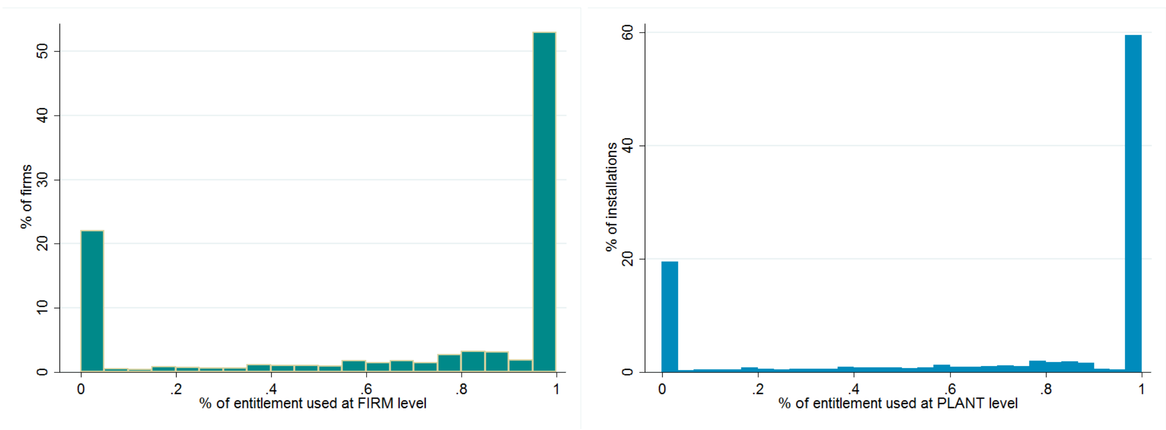

Figure 10: Utilization and timing of plants and of firms (source: EU CITL registry)

D Are emissions constrained by transaction costs?

In section 3.2, I claimed that transaction costs do not constrain emissions. However, firms face a cost curve that "jumps" when emissions increase beyond the level of allocations (fig. 3a). I use the fact that "short" firms, with emissions above allocations, are constrained to trade while "long" firms can choose whether to incur trading entry costs. However, there is the possibility that firms adjust their emissions and production to avoid transaction costs and be "long" even though they would have been short in absence of transaction costs. This section argues, that this case is unlikely to be relevant in practice.Firms choose their production and emissions given production cost and certificate prices; the additional transaction cost is likely to be smaller than the cost of adjusting emissions and production.

First, note that the firm faces the same marginal cost p^e for emissions both below and above the jump of Figure 3a and 3b, so that marginal abatement cost does not play a role. However, overall compliance cost increases; the firm thus compares two situations: one where emissions are reduced to allocation level A_i, so that optimal production is $Y^*(A_i)$ and entry costs are *not* incurred, and another situation where $E_i^* > A_i$ is chosen such that marginal abatement cost equals p^e and entry cost is incurred. The firm reduces its emissions to A_i if the change in profit Δ_π resulting from this reduction is positive:

$$\Delta_\pi = \pi(Y^*(A_i), A_i) - \pi(Y^*(E_i^*), E_i^*) \quad (19)$$

$$= (Y^*(A_i) - Y^*(E_i^*))p - C(Y^*(A_i), A_i) + C(Y^*(E_i^*), E_i^*) - T^*(A_i) + T^*(E_i^*) \quad (20)$$

By assumption, we are looking here at cases where optimal emissions $E_i^* > A_i$ and thus $Y^*(E_i^*) > Y^*(A_i)$; by definition of the optimal emission level E_i^*, Δ_π would always be negative without the transaction cost terms of equation (20) (or, to be more exactly, the left-hand side of equation (21)). As seen on Figure 3, the change in incurred transaction cost is either T^e, as on Figure 3a, or $T^e + T^o - \Delta_p K$, see Figure 3b. Firms would thus reduce their production and emissions only if:

$$(Y^*(A_i) - Y^*(E_i^*))p - C(Y^*(A_i), A_i) + C(Y^*(E_i^*), E_i^*) - p^e(A - E^*) < T^e + min\{T^o - \Delta_p K_i, 0\} \quad (21)$$

Anecdotal and survey evidence (Loeschel et al., 2010, 2011) shows that firms do not have a precise control over their emissions, or rather that there are considerable transaction costs here again to attain such a control. Only large companies track their emissions over the year. The trading scheme's incentives to reduce emissions do not work on a short-term "accurate to the tonne" level, but rather on a long-term technology-inducing level.

Most technologies are such that in the short term the actual technological margin to reduce emissions without a complete corresponding reduction of output is limited; reducing emissions by a certain share is thus equivalent to reducing production by a the same share. After all, emissions are just one production cost factor among many others and flexibility of the cost function is usually low (meaning that emission reductions are to a large extent matched by reductions in the produced outcome). Emission reductions are mostly accomplished in the long term through technical change, whereas this study is looking at short term behavior. Even for a small difference between E_i^* and A_i it is likely that Δ_π is strictly negative.

A notable exception might be emission savings by electricity generating plants, as some firms have a scope for fuel-switching across different plants and emission costs are a more important cost factor in this industry. This sector will thus be examined more in detail.

A priori this case thus seems not so relevant in practice. However, to be clear, it cannot be verified properly, as information on prices p, quantities Y and production costs are not available, neither cost function $C(Y, E)$ nor the profit change Δ_π can be estimated. Instead, one way of gathering (descriptive) evidence on this point comes from checking that we do not observe any crowding of emission levels around $E = A$.

Figure 11 shows the ratio of emissions to allocations; crowding of emissions at allocation would mean that there would be a spike where this ratio equals 1. Figure 11a shows this statistic for all firms, the mode is below one, but distribution is smooth without any sign of discontinuity. Restraining emissions to become short only makes sense for firms which do not use offset certificates (as they incur transaction costs anyways), therefore fig. 11b shows only the firms which do not participate in the offset market: again, there is no indication for crowding at $E = A$. Finally, fig. 11c shows that for electricity firms this is even less relevant. Moreover,virtually all electricity firms participate in the offset market, so that they incur transaction costs anyways.

While theoretically not fully sound, the assumption of exogenous allocation status thus seems empirically valid. Indeed, we see in the data that virtually no firm restrains its emissions to the level of allocations.

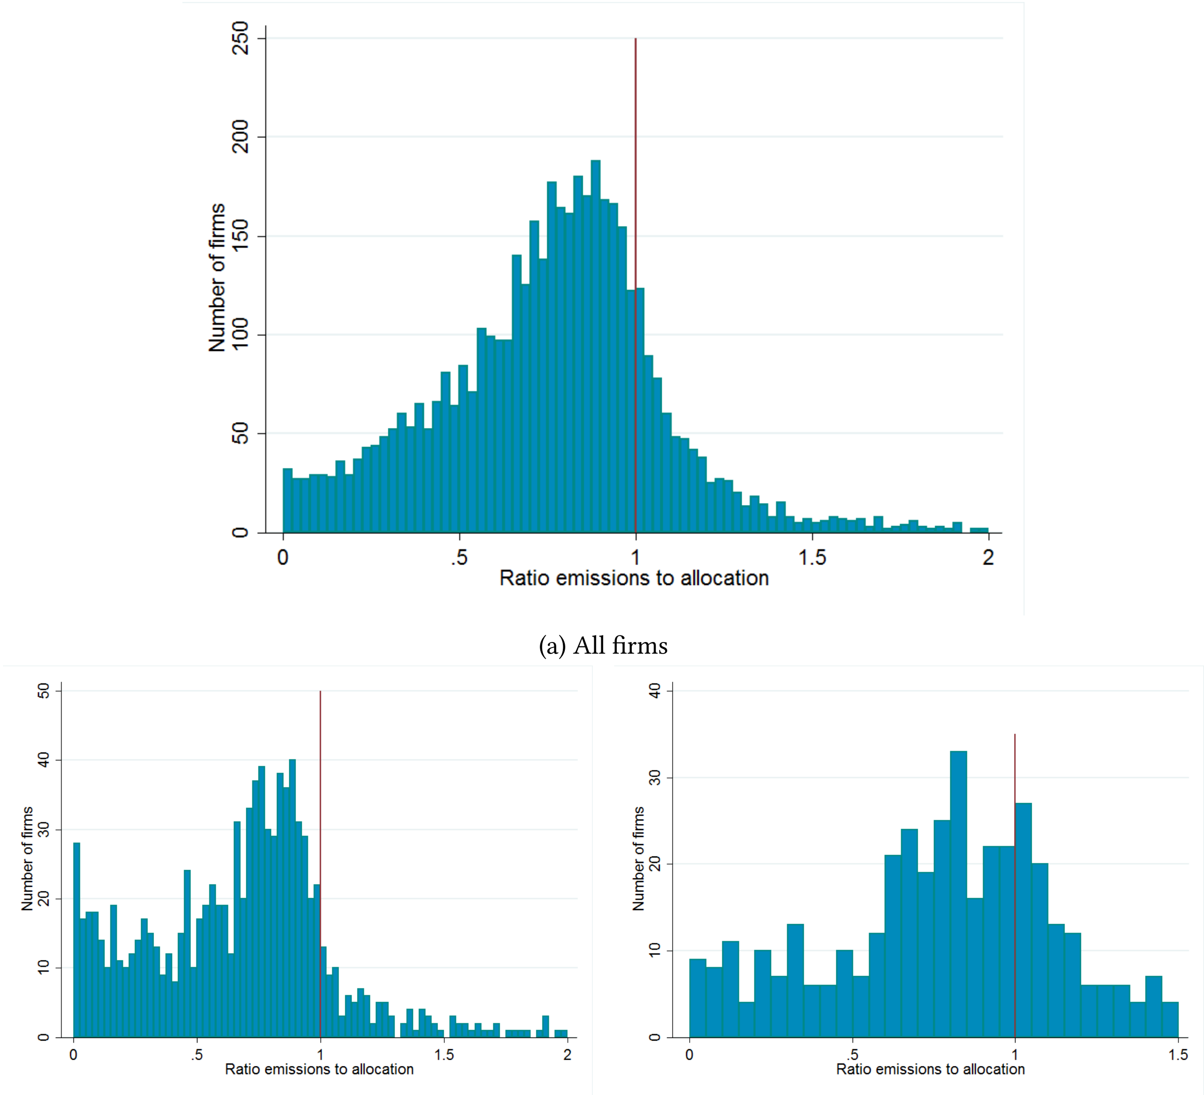

(a) All firms

(b) Only firms using with no offset usage

(c) Only firms in electricity and heat generation

Figure 11: Ratio of emissions to allocations

E Parametric estimation results

In section 4.4, I noted that a standard way to estimate the parameters of equation (15) would be to assume a standard normally distributed error term ϵ_i. The model then becomes similar to a standard probit model: K_i is included as a regressor and coefficients are then normalized such that the coefficient on K_i equals -1.

The stylized facts presented in section 4 strongly suggest that this homoskedastic normality assumption does not hold. As shown before, the distribution of offset entitlements is highly skewed with some firms more than 500 times bigger than the (relatively low) median. Moreover, some firms with high K_i still do not exploit their offset entitlement, so that the distribution of ϵ_i from the transaction cost equation (15) is likely to have some large outliers. The (conditional) mean is a statistic much more sensitive to outliers than the (conditional) median; differently put, the normal distribution assumption has light tails which consequently give large weight to outliers.

A slightly more flexible functional form relaxing the homoskedasticity assumption, would be the mixed probit: error terms are still assumed to have a normal distribution, but the variance scales with he size (here K_i) of the firm. In such a location-scale model, the variance of each ϵ_i depends on some scaling variable and a parameter γ (to be estimated):

$$\epsilon_i \sim \mathcal{N}(0, \sigma_i^2), \text{ where } \sigma_i = \exp(K_i \gamma) \tag{22}$$

This section shows the results for both assumptions, while claiming that they are not an accurate description of the data. The results the probit estimation in both the homoskedastic and (linearly) heteroskedastic versions are shown in Table 5.[27] The costs indicated are measured in euros, as they are normalized by the cut-off value's $\Delta_p K_i$ coefficient. The estimate for T^o, the transaction cost for offset usage, is larger than the estimate for T^e, while both are significant. When I include the sectors, the results are somewhat surprising and the estimates for transaction costs in the electricity and heat generation sector are big but not significant.

	Probit	Heterosk. probit	Probit with sectors	Heterosk. probit with sectors
$\hat{T}^C$ (intercept)	109557***	102660***		
	(4.24)	(4.36)		
$\hat{T}^E$ ($\mathbb{1}^{long}$)	44302***	42798***		
	(3.70)	(3.79)		
$\hat{T}^C$ Manufacturing			169389***	166379***
			(4.48)	(4.52)
$\hat{T}^E$ Manufacturing			94359***	93081***
			(4.39)	(4.43)
$\hat{T}^C$ Electricity			-45689	288125***
			(-2.47)	(4.44)
$\hat{T}^E$ Electricity			1889	2208
			(0.11)	(0.13)
[1em] σ	192950***	182835***	193286***	190246***
	(5.77)	(6.04)	(5.85)	(5.93)
γ		6.96e-08***		2.07e-08***
		(18.15)		(12.07)
R2	.1274	.128	.1371	.1376
Completely determined	371	.	367	.
N	4578	4578	4578	4578

t statistics in parentheses; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Parametric mean estimates for transaction costs

[27] Estimated using Stata oglm command by Williams (2010).

F Sector-specific quantile graphics

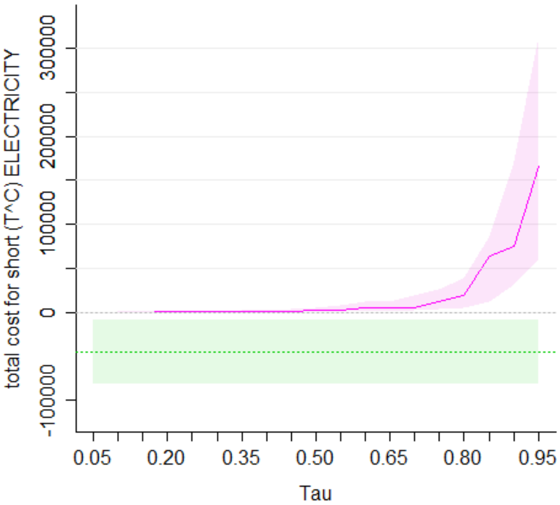

(a) Offset-specific cost

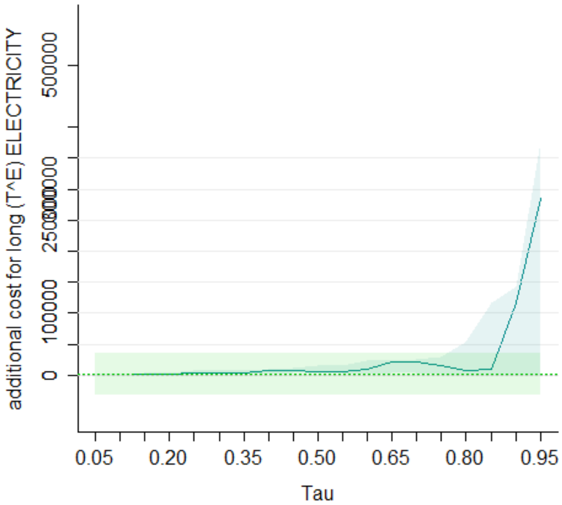

(b) General cost

Figure 12: Quantile estimation results (in €), only electricity- and heat-generation

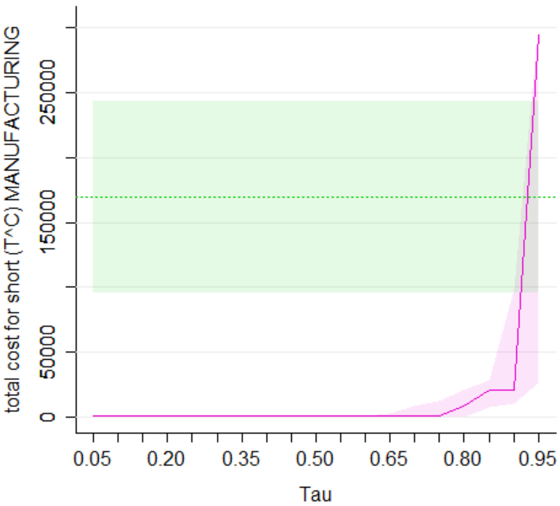

(a) Offset-specific cost

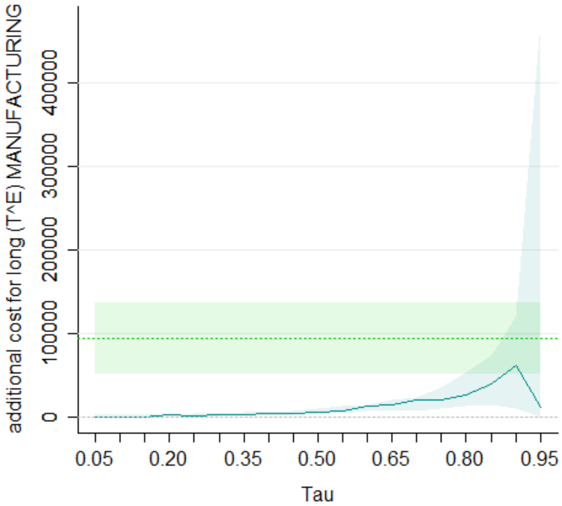

(b) General cost

Figure 13: Quantile estimation results (in €), only manufacturing

G Quantile regression fit

Kordas (2006) suggests to verify the fit of the quantile regressions by predicting probabilities for each observation and verifying that each group has a participation rate close to the predicted probability. Predicting probabilities from the binary quantile regression is simple: one needs to find the smallest quantile $\hat{\tau}_i$ such that the profit-net-of-transaction costs is positive:

$$\hat{\tau}_i = argmin\{\tau : \Delta_p K_i - T^o - \mathbb{1}_i^{long} T^e \geq 0\} \tag{23}$$

Then this gives us an interval for the conditional participation probability:

$$\hat{P}_i \in [1 - \hat{\tau}_i, 1 - \hat{\tau}_{i,-1}] \tag{24}$$

where $\hat{\tau}_{i,-1}$ is the quantile immediately preceding $\hat{\tau}_i$.

For the data used in this study this gives the predicted and observed probabilities displayed in Table 6. Except for the lowest quantile, the models seem to fit the data reasonably well. On the opposite, the probit model predicts for *all* firms a participation probability above 50%: one could says that *all* non-participating firms are unpredicted outliers (false-negatives) with the probit model.

Predicted probability (in 10 point intervals)	10%	20%	30%	40%	50%	60%	70%	80%	90%
Number of observations	72	123	72	118	84	377	650	1047	2022
Observed frequency	13%	16%	32%	44%	48%	58%	71%	81%	94%

Table 6: Specification test of binary regression quantile models