



Information Aggregation in Emissions Markets with Abatement

Estelle Cantillon
SBS-EM, ECARES, Université libre de Bruxelles,
FNRS, Toulouse School of Economics, CEPR

Aurélie Slechten
Lancaster University Management School

December 2018

ECARES working paper 2018-37

Information aggregation in emissions markets with abatement*

Estelle Cantillon[‡] and Aurélie Slechten[†]

November 2018

Abstract

A key policy argument in favor of emissions markets (relative to command-and-control types of regulation) is their ability to aggregate dispersed information and generate price signals to guide firms' trading and abatement decisions. We investigate this argument in a multi-period model where firms receive noisy private signals about their current period emissions and privately observe their previous period emissions before this information is made public to the rest of the market. Firms respond to information by trading and abating emissions. We show that there exists a rational expectations equilibrium that fully aggregates firms' private information, justifying the policy argument in favor of emissions markets, in the absence of other frictions. We also derive predictions about how prices should be reacting to new private or public information and show that the possibility of abatement dampens the impact of shocks on prices. Finally, we show that the information aggregation result breaks down if firms' abatement costs are also private information.

Keywords: Information aggregation, efficient market hypothesis, price formation, emissions trading

JEL Codes: G14, D83, D84, D85, Q58

*We thank Stefan Ambec, Thomas Gall, Hans Gerbach, Pär Holmberg, Ilan Kremer, Thomas-Olivier Léautier, Laurent Linnemer, Matti Liski, Jingfeng Lu, David Martimort, Andy Skrypacz, Ali Tahbaz-Salehi, Thomas Tangerås, Olivier Scaillet, Jean Tirole, Ulrich Wagner and several conference and seminar audiences for comments and suggestions. Financial support from ERC grant 208535 is gratefully acknowledged.

[‡]FNRS, Université Libre de Bruxelles (ECARES), Toulouse School of Economics and CEPR, email : Estelle.Cantillon@ulb.ac.be

[†]Lancaster University Management School, e-mail: a.slechten@lancaster.ac.uk.

1 Introduction

A central argument in favor of market-based policy instruments (as opposed to command-and-control policies) is their ability to decentralize decisions to those economic agents who have the information. Translated in the context of emissions markets, this argument contends that profit maximization by firms subject to the scheme, combined with competitive markets for allowances, will generate the “right” price signal to guide abatement and trading decisions, ultimately minimizing the costs of reducing emissions.

Montgomery (1972), Cronshaw and Kruse (1996), and Rubin (1996) have provided theoretical foundations for this argument - a variant of the First Welfare Theorem - for emissions markets without uncertainty. Yet, uncertainty is paramount in most emissions trading schemes. Business-as-usual emissions are intrinsically uncertain. For example, emissions from the electricity sector are highly dependent on the relative price of gas and coal, the business cycle, as well as on the weather. Firms learn about their emissions over time as this uncertainty unfolds. In addition, different firms have access to different abatement opportunities and these are typically privately known.

The presence of these sources of uncertainty - private information and exogenous risk - raises two questions. First, can competitive emissions markets adequately aggregate private information, a necessary condition for efficiency? Second, what impact does uncertainty have on the welfare properties of competitive emissions markets? The existing literature has partially answered the second question in the context of exogenous risk (Schennach, 2000, Seifert et al., 2008, Hitzemann, 2018) but has not tackled the issue of information aggregation.

To answer these questions, we analyze a multi-period emissions trading model with both private information and aggregate shocks. Our model builds on the canonical model in the literature (Rubin, 1996, Schennach, 2000) with risk neutral price-taking firms in the allowances market, no pass-through of costs, and short-term (instantaneous) abatement. Our main departure point is to explicitly model firms’ sequential arrival of private information and public information releases, and how these get aggregated into the equilibrium price. We assume the most favorable conditions on the market to perform this role by using the dynamic rational expectations equilibrium (Tirole, 1982) as our equilibrium concept. In a dynamic rational expectations equilibrium (REE), firms take the price as given and use the information contained in the price, together with their own signals, to decide on the best strategy to maximize their expected discounted profits, at every point in time. There is no trading friction and a single price prevails.

We model a market with a finite horizon, by which firms need to surrender enough allowances to cover their emissions to date, and where surplus allowances lose their values entirely at the end.

The first phase of the EU emissions trading scheme (2005-2007) is a case in point, but the model also captures markets with restrictions on the banking of unused allowances. To model the gradual arrival of information to the market, we assume that the market is made of several compliance periods. At the time when they decide on their abatement, firms receive a noisy signal about their emissions during the current compliance period. They learn about their actual level of emissions at the end of the compliance period and this information becomes public with a delay, at the time emissions are verified.

We first consider an artificial economy where all private information is public and solve for the equilibrium price and trading at all times when new information reaches the market. Because the horizon is finite, we can solve the model by backward induction using standard dynamic programming techniques. At all times, the equilibrium price is equal to the probability that the market is short at the end of the horizon, times the penalty that firms must pay for failing to deliver enough allowances to cover their emissions. The equilibrium price follows a martingale and converges to zero or the penalty as the horizon gets closer and uncertainty about whether the market is long or short decreases. The gradual arrival of information implies that the impact of exogenous shocks to emissions spreads over two periods: at the beginning of the compliance period, when firms receive a noisy signal about their current period emissions, and at the beginning of the next period, when they learn the realized emissions. Abatement endogenously dampens shocks to emissions as they provide an alternative for firms to cover their emissions: when emissions are high, abatement increases in reaction.

To prove that this full information equilibrium is an equilibrium when information is private, we need to show that the equilibrium price fully aggregates private information, in the sense that the equilibrium price is a sufficient statistic for the information dispersed in the economy and it results in the same allocation as if each trader had access to all private information (Grossman, 1976). The challenge here is that the equilibrium price, a one dimensional variable, is driven by both the expectation and the variance of future emissions, net of abatement, i.e. a two dimensional variable. We show that, when the only source of asymmetric information is private information about emissions, prices can nevertheless be inverted to recover these parameters, i.e. prices are a sufficient statistic. This implies that the full information equilibrium outcome is also the outcome under rational expectations and dispersed information. This is no longer true when abatement costs are also private information.

Relationship with the literature: Understanding the ability of markets to aggregate dispersed information has been a central question in finance since the seminal papers by Grossman

(1976) and Grossman and Stiglitz (1980). This literature has shown that many factors can result in prices not fully aggregating the underlying dispersed information even when traders are price-taking (see Brunnermeier, 2000, for a survey). Two features distinguish emissions markets from other asset markets with dispersed information. First, the value of emissions allowances depends on the probability that the market is going to be short or long and it is the same for all traders, whereas traders in asset markets may value the same piece of information differently based on their current positions or risk attitudes. Second, the presence of abatement in emissions markets implies that prices today endogenously affect the balance between the supply and demand of allowances tomorrow, whereas shocks to the supply of assets are typically exogenous in finance. These two features - on top of less fundamental modeling differences - set our paper apart from recent explorations about information aggregation in dynamic markets (see e.g. He and Wang, 1995, Vives, 2008, and Cespa and Vives, 2012). We will see that the first feature provides favorable conditions for information aggregation, while the second complicates the inference problem because abatement endogenously affects both the expected level of emissions and its variance, conditional on traders' information.

In contrast to finance, the environmental economics literature has paid little attention to issues of information aggregation. In fact, most theoretical analyses of emissions trading since Rubin (1996) take the social planner's perspective as their starting point, arguing e.g. as Schennach (2000, p. 191) that “[Rubin (1996) has shown that] when allowed to trade with one another, units will collectively behave like a central planner who efficiently allocates emission permits to each unit to minimize total costs.” Other recent examples of this social planner's approach include Seifert et al. (2008) and Hintermann (2010). Relatedly, Chesney and Taschini (2012) solve for the informationally efficient equilibrium emissions price in a model with private information but no abatement. Such approaches implicitly assume that the equilibrium price will fully aggregate private information. Our results show that this assumption is warranted when the only source of asymmetric information is private information about emissions but that it is inconsistent with equilibrium behavior with dispersed information when there is also private information about abatement costs.

Our analysis also generates new insights into the determinants of the dynamics of prices in emissions markets. In particular, we go beyond the well-known result according to which emissions prices are martingales and derive explicit forms for the way they get updated, following new information arrival. In particular, when full information aggregation obtains, we show that prices should not respond to the publication of verified emissions since that information was already (privately) known to market participants. This is a special case of the strong form of the Efficient Market Hypothesis.

2 A benchmark model of price formation

We explore price formation in a model designed to capture the salient features of many emissions markets, namely gradual arrival of private information about own emissions, abatement, and unlimited banking and borrowing of allowances within a phase. Figure 1 illustrates the timing of allocation and surrender of emissions allowances in the EU emissions trading scheme (EU ETS). Phases are made of several compliance periods. Allocations for the whole phase are announced at the beginning of the phase but allowances are distributed in yearly installments at the end of February of each year. Emissions for year t are verified by March of year $t + 1$. Firms have then until April 30 to surrender the allowances corresponding to the emissions of the previous year. This timing makes it possible for firms to use some of the current year's allowances to cover the previous year's emissions. It is arguably equivalent to full borrowing as one year of allowances is likely to provide the necessary inter-temporal flexibility in most cases. Failure to surrender sufficient allowances results in a fine, together with the requirement to buy any missing allowances in the market (from the next phase if necessary). Unused allowances can be banked for future years. During Phase I of the EU ETS (2005-07), allowances lost their value at the end of the phase. Banking across phases was allowed starting with phase II (2008-13).




Figure 1: Timing of allocation and surrender of allowances in the EU ETS

Figure 1 also clarifies the timing of information available to the market. While firms could arguably track their individual emissions as they happened, verified emissions for all installations are only made public in April of the year after.¹

¹According to the Report on the functioning of the European carbon market (COM(2015) 576 final), most of the

We model the finite horizon version of this market. The market plays out over $T + 1$ periods. In periods 1 to T , N firms emit carbon according to an i.i.d. exogenous stochastic process, $e_{it} \sim N(\mu_t, \sigma^2)$, where μ_t and σ^2 capture, respectively, the business cycle and firm heterogeneity.² For now, we assume that both μ_t and σ^2 are common knowledge but relax this assumption for μ_t in section 4. The number of firms N is finite (otherwise there would not be any aggregate uncertainty about future emissions) but sufficiently large to make the price-taking assumption reasonable and some of our proofs will rely of N being large enough. We interpret e_{it} as the emissions of firm i during period t , *net* of their received allowances. Let e_t denote period t aggregate emissions, $\sum_i e_{it}$, and $E_t = \sum_{\tau=1}^t e_\tau$ denote aggregate emissions to date. Similarly, $E_{it} = \sum_{\tau=1}^t e_{i\tau}$ denote firm i 's cumulative emissions to date.

To reflect the gradual arrival of information, we further divide each period $t \leq T$ into two subperiods, one beginning at time t and the other beginning at time $t + 0.5$. Subperiods differ in the information available to firms. At integer times $t = 1, 2, \dots$, firms privately learn about their emissions in the previous period, e_{it-1} , and receive a noisy but unbiased signal x_{it} about their current period emissions, equal to $x_{it} = e_{it} + \varepsilon_{it}$, where ε_{it} is i.i.d. distributed according to $N(0, 1)$. At mid-period points, $t = 2.5, 3.5, \dots$, emissions about the previous period are made public. Firms do not receive any additional private information.

Trading takes place at the beginning of each subperiod. Let y_{it} denote the number of allowances traded by firm i at time t (where $y_{it} > 0$ corresponds to a purchase and $y_{it} < 0$ to a sale) and let p_t denote the price of allowances at time t . In addition, at the beginning of each (full) period, firms decide on their abatement a_{it} for the period. Abatement incurs a quadratic cost $c(a_i, \theta_i) = \frac{1}{2}\theta_i a_i^2$, with $\theta_i > 0$, that differs across firms. For now, we assume that these costs are common knowledge and do not vary over time. We relax this assumption in Section 5. For future reference we also let $\frac{1}{\theta} = \sum_i \frac{1}{\theta_i}$. As for emissions, we introduce notations a_t and A_t to denote period t and cumulative aggregate abatement respectively, and A_{it} to denote firm i 's cumulative abatement to date.

Period $T + 1$ is a pure compliance period: there are no emissions nor abatement. Firms learn about their emissions in period T , e_{iT} , decide on their sales or purchases of allowances, y_{iT+1} , and surrender their allowances for compliance. The role of this last period is to provide an opportunity for firms to buy the needed allowances to cover their net emissions from period 1 to period T . With

installations covered by the EU ETS use the calculation-based methodology to monitor their emissions throughout the year. This method weighs all emissions-producing inputs used by the installation by an emission factor to generate an estimate of its emissions. Continuous emissions measurement systems are only used in a very small number of installations (140 installations in 2014).

²Alberola et al. (2008) analyze the empirical drivers of emissions in the EU Emissions Trading Scheme.

full borrowing and banking, there is no loss in generality in assuming that compliance only occurs at time $T + 1$ (the end of the phase). If market participants fail to surrender enough allowances to cover their emissions, they are fined K per missing allowance.³

Figure 2 summarizes the timing of actions and information in the model.




Figure 2: Timing of information and action

Firms are rational but, in the tradition of rational expectations models, they take prices as given.⁴ Their objective is to minimize their expected total cost of compliance which is the sum of incurred abatement costs, trading costs, and penalty if applicable, $\sum_{t=1}^T c(a_{it}, \theta_i) + \sum_{t=1}^{T+1} p_t y_{it} + \max\{0, K(E_{iT} - A_{iT} - Y_{iT+1})\}$ where the last term corresponds to the total penalty.⁵

Since there are only $2T + 1$ moments of time when new information arrives or actions are taken,

³In practice, the penalty may depend on t because a firm that fails to surrender a number of allowances corresponding to its emissions has to pay a fixed penalty and buy any missing allowance on the market. If there are enough allowances in the market in period $T + 1$, cost-minimizing firms will always comply and have enough allowances to cover their emissions (otherwise they pay the fixed penalty K plus the current price of allowances, p_{T+1}). If the market is short at time $T + 1$, firms that do not have enough allowances will need to pay the fixed penalty K plus the expected allowance price in the next phase. As there is no banking or borrowing between phases, the prices in the different phases are statistically disconnected and therefore the fact that the penalty may depend on t (because of the expected price in the next phase) does not affect our results.

⁴The Rational Expectations Equilibrium concept has well-known limitations (Grossman, 1976, Dubey et al., 1987, Jackson, 1991) but has also shown to be a useful empirical benchmark (Plott and Sunder, 1988, Forsythe and Lundholm, 1990). We use it as an idealized form of frictionless market (see also Kovalenkov and Vives, 2014).

⁵We ignore discounting for simplicity. We also ignore any indirect effect of the market for allowances on the markets that firms serve (no pass-through).

we solve for the market equilibrium at these $2T + 1$ moments. Let s_{it} denote the vector of private signals received at time t by trader i . At integer times, $s_{it} = (x_{it}, e_{it-1})$, while at mid-period points, $s_{it+0.5} = (e_{t-1})$. A **fully dynamic Rational Expectations Equilibrium (REE)** (Tirole, 1982) is a sequence of price functions mapping firms' signals $s_t = (s_{1t}, s_{2t}, \dots, s_{Nt})$ (and, implicitly, past history of prices and own actions) into a price $s_t \rightarrow p_t = g_t(s_t)$ and a sequence of trades $y_{it}(s_{it}, p_t)$ and abatement decisions $a_{it}(s_{it}, p_t)$ such that (i) each trader i minimizes his expected total cost of compliance at all t and for all information (s_{it}, p_t) trader i may have, and (ii) markets clear $\sum_i y_{it}(s_{it}, p_t) = 0$ for all t , s_{it} and p_t .

The existence of a fully dynamic REE does not necessarily imply that prices fully aggregate private information. This would be the case only if the price is a sufficient statistic for the private and public information dispersed in the market (Grossman, 1976). In other words, when the market equilibrium fully aggregates private information, observing the equilibrium price leads to the same market outcomes as if all participants pooled their private signals.

3 Equilibrium

3.1 Full information equilibrium

We start by considering the equivalent artificial economy where all private information is public.⁶ Let I_t denote all the information available at time t in the economy, including all past prices, emissions, and signals. Formally:

$$I_t = \{p_1, \dots, p_t, e_1, \dots, e_{t-1}, x_1, \dots, x_t\}$$

where we have innocuously summarized individual past emissions, abatement and signals through their aggregates. Because, in that economy, no new information is revealed at mid-period times (past emissions are already known at the beginning of the period), it is sufficient to solve for the equilibrium at integer times. We do this using backward induction.

At time $T+1$, emissions and abatements are fixed and known so there are only two possibilities: the market is short, i.e. cumulative aggregate emissions are higher than cumulative abatement $E_T > A_T$, in which case $p_{T+1} = K$, the penalty, or $E_T \leq A_T$, in which case $p_{T+1} = 0$. Firm i will need to buy allowances at that price if its cumulative abatement and purchases do not cover its cumulative emissions, $E_{iT} - A_{iT} - Y_{iT} > 0$ (where we have implicitly assumed that no trading takes place at $T + 0.5$ since no new information reaches the market) and will sell allowances otherwise.

⁶This corresponds to Radner (1979)'s notion of full communication equilibrium.

At time T , period $T-1$ aggregate emissions are revealed but only a noisy signal, x_T , is available for period T emissions. Based on this information, firms decide on their trading and abatement to minimize their current and future expected costs of compliance:

$$\min_{a_{iT}, y_{iT}} p_T y_{iT} + c(a_{iT}, \theta_i) + V_{T+1}^i(a_{iT}, y_{iT} | I_T)$$

where $V_{T+1}^i(a_{iT}, y_{iT} | I_T)$ is firm i 's expected continuation cost of compliance at time $T+1$ given choices a_{iT} and y_{iT} , and information I_T . Substituting for $V_{T+1}^i(a_{iT}, y_{iT} | I_T)$ into this expression leads to:

$$\min_{a_{iT}, y_{iT}} p_T y_{iT} + c(a_{iT}, \theta_i) + E[p_{T+1}(E_{iT} - A_{iT} - Y_{iT-1} - y_{iT}) | I_T]$$

Together with market clearing, $\sum_i y_{iT} = 0$, and the fact that firms share the same beliefs about p_{T+1} since they have the same information, the first order conditions for firms' individual optimization problem are given by

$$\begin{aligned} c_a(a_{iT}, \theta_i) &= \theta_i a_{iT} = E[p_{T+1} | I_T] \\ p_T &= E[p_{T+1} | I_T] \end{aligned} \tag{1}$$

At the full information equilibrium at time T , we find the classic result that firms choose to abate up to the point where marginal abatement costs are equal to the current price.

Define $Z = E_T - A_T$, the excess emissions by the end of the phase. From the perspective of time t , Z is a random variable, with:

$$Z = E_{t-1} - A_t + \sum_{k=t}^T \tilde{e}_k - \sum_{k=t+1}^T \tilde{a}_k \tag{2}$$

where a tilda has been added on the values of emissions and abatement that are not known with certainty at time t . With this definition, the expectation of p_{T+1} conditional on the information available at time T can be written as:

$$\begin{aligned} E[p_{T+1} | I_T] &= K \Pr(Z > 0 | I_T) \\ &= K \Pr(\tilde{e}_T > A_{T-1} + a_T - E_{T-1} | I_T) \end{aligned} \tag{3}$$

where the terms on the righthand side of the inequality sign in (3) are either known at time T (A_{T-1} and E_{T-1}) or are contemporaneously determined (a_T), whereas the term on the lefthand side is a random variable, based on I_T . Using the projection theorem for normally distributed variables leads to $\tilde{e}_T | I_T \sim \tilde{e}_T | x_T \sim N(\frac{1}{1+\sigma^2} N\mu_T + \frac{\sigma^2}{1+\sigma^2} x_T, \frac{N\sigma^2}{1+\sigma^2})$. Therefore, combining (1) and (3) implies that the equilibrium price at time T is defined implicitly by the following equality:⁷

⁷This expression uniquely pins down p_T : the LHS is increasing in p_T , whereas the RHS is decreasing in p_T . Moreover, both sides are continuous and the LHS is lower than the RHS for low values of p_T and higher for high values of p_T .

$$p_T = K \left[1 - \Phi \left(\frac{A_{T-1} + \frac{p_T}{\theta} - E_{T-1} - \frac{1}{1+\sigma^2} N\mu_T - \frac{\sigma^2}{1+\sigma^2} x_T}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} \right) \right] \quad (4)$$

where Φ stands for the cumulative distribution function of the standardized normal or, equivalently, using the definition of Z ,

$$p_T = K \left[1 - \Phi \left(\frac{-E(Z|I_T)}{\sqrt{V(Z|I_T)}} \right) \right] \quad (5)$$

Equation (4) shows that the equilibrium price increases with cumulative past emissions (E_{T-1}) and the current aggregate emissions forecast ($\frac{1}{1+\sigma^2} N\mu_T + \frac{\sigma^2}{1+\sigma^2} x_T$) and it decreases in past and current abatement ($A_{T-1} + \frac{p_T}{\theta}$). Uncertainty about emissions, σ^2 , has an ambiguous effect on prices. On the one hand, the variance of Z increases in σ^2 . The effect on price depends on the sign of expected excess emissions, $E(Z|I_T)$. If $E(Z|I_T) < 0$, i.e. the market is expected to be long at time $T + 1$, an increase in σ^2 will increase the price because it increases the probability that the market becomes short at $T + 1$. The effect is opposite if $E(Z|I_T) > 0$. On the other hand, a higher level of uncertainty also means that firms put more weight on signal x_T when updating their prior about period T emissions. This can either contribute towards increasing prices (if x_T is larger than $N\mu_T$) or decreasing prices (otherwise) by changing $E(Z|I_T)$. These effects reinforce one another if the market is expected to be long but emissions in the current period are expected to be higher than expected ($x_T > N\mu_T$) or, alternatively, if the market is expected to be short and current period emissions are expected to be lower than $N\mu_T$.

Finally, note that the numerator in (4) scales proportionally with the number of firms N , whereas its denominator only scales at the rate of \sqrt{N} . This implies that the price converges to K or 0 (depending on the sign of the numerator) as the number of firms increases. Intuitively, as the number of firms increases, aggregate uncertainty about net emissions decreases and the market is either long or short.

For future reference, it is useful to examine how the price at time T responds to the new information reaching the market at time T , namely x_T and e_{T-1} . Let $f_T(\cdot)$ describe the function implicitly defined by (4) that maps x_T and e_{T-1} into p_T . Linearizing f_T around the expectation of its arguments conditional on I_{T-1} , $E[x_T | I_{T-1}] = N\mu_T$ and $E[e_{T-1} | I_{T-1}] = \frac{1}{1+\sigma^2} N\mu_{T-1} + \frac{\sigma^2}{1+\sigma^2} x_{T-1}$ yields:

$$f_T(N\mu_T, E[e_{T-1} | I_{T-1}]) + \nabla f_T(N\mu_T, E[e_{T-1} | I_{T-1}]) \begin{bmatrix} x_T - N\mu_T \\ e_{T-1} - E[e_{T-1} | I_{T-1}] \end{bmatrix} \quad (6)$$

In words: relative to its expectation from the perspective of period $T - 1$, the price at time T reacts to higher than expected signals about current emissions, $x_T - N\mu_T$, and to higher past emissions

than previously foreseen, $e_{T-1} - E[e_{T-1} | I_{T-1}]$. In Lemma 1, we show that the multiplier of these surprises is positive and that it is smaller for surprises about current emissions because firms account for the fact that x_T is a noisy signal and therefore only partially update their beliefs about current period emissions. Lemma 1 also shows that the presence of abatement, which manifests it-self in the presence of $\frac{p_T}{\theta}$ in the righthand side of (4) lowers the magnitude of these price adjustments: abatement dampens shocks. All proofs can be found in the Appendix.

Lemma 1 (Price adjustment to new information at time T). *Let $f_T(\cdot)$ describe the function implicitly defined by (4) that maps x_T and e_{T-1} into p_T . Then:*

- (i) Prices respond positively to shocks in emissions and the more so to realized emissions than to noisy signals about current emissions: $\frac{\partial}{\partial x_T} f_T(x_T, e_{T-1}) = \frac{\sigma^2}{1+\sigma^2} \frac{\partial}{\partial e_{T-1}} f_T(x_T, e_{T-1}) > 0$.
- (ii) The possibility of abatement dampens the price response to emissions shocks: $\frac{\partial}{\partial x_T} f_T(x_T, e_{T-1}) < \frac{\partial}{\partial x_T} \tilde{f}_T(x_T, e_{T-1})$ and $\frac{\partial}{\partial e_{T-1}} f_T(x_T, e_{T-1}) < \frac{\partial}{\partial e_{T-1}} \tilde{f}_T(x_T, e_{T-1})$ where $\tilde{f}(x_T, e_{T-1})$ is the analogous price function in an environment without abatement, evaluated at the same value of excess emissions Z .
- (iii) Abatement only partially adjusts to emissions shocks so that the net effect of a positive (negative) emissions shock remains positive (negative): $\frac{\partial}{\partial x_T} E[e_T - a_T | I_T], \frac{\partial}{\partial e_{T-1}} E[e_{T-1} - a_T | I_T] \in (0, 1)$

Moving to time $T - 1$, firm i 's optimization problem is given by:

$$\min_{a_{iT-1}, y_{iT-1}} p_{T-1} y_{iT-1} + c(a_{iT-1}, \theta_i) + V_T^i(a_{iT-1}, y_{iT-1} | I_{T-1})$$

or, equivalently, substituting for the expected continuation cost and leveraging the law of iterated expectations and the equilibrium at time T which imply that $E[p_T | I_{T-1}] = E[p_{T+1} | I_{T-1}]$:

$$\min_{a_{iT-1}, y_{iT-1}} p_{T-1} y_{iT-1} + c(a_{iT-1}, \theta_i) + E[c(a_{iT}, \theta_i) + p_{T+1}(E_{iT} - A_{iT} - Y_{iT-1}) | I_{T-1}]$$

Market clearing and common information (and therefore beliefs) lead to the familiar first order conditions for firms' individual optimization problem :

$$\begin{aligned} c_a(a_{iT-1}, \theta_i) &= \theta_i a_{iT-1} = E[p_{T+1} | I_{T-1}] \\ p_{T-1} &= E[p_{T+1} | I_{T-1}] \end{aligned} \tag{7}$$

As before, abatements equalize marginal costs with the current price of allowances. The difference is that firms now need to form expectations both about future emissions *and* about future abatements. Indeed, $E[p_{T+1} | I_{T-1}]$ is now given by:

$$K \Pr(Z > 0 | I_{T-1}) = K \Pr(E_{T-2} - A_{T-1} + \tilde{e}_{T-1} + \tilde{e}_T - \tilde{a}_T > 0 | I_{T-1})$$

where we have again used a tilde to highlight unknown variables from the perspective of $T - 1$. From (1), we know that $a_T = \sum_i a_{iT} = \frac{p_T}{\theta}$. However, p_T is not known at time $T - 1$ but a function of random variables x_T and e_{T-1} , $f_T(x_T, e_{T-1})$. Using the linearization (6) of $f_T(\cdot)$ around the expectation of its arguments given I_{T-1} , $Z|I_{T-1}$ can be approximated by the sum of independent normal variables, and thus a normal variable it-self, with mean and variance given in Lemma 2.

Lemma 2 (Distribution of excess emissions as of time T-1). $Z|I_{T-1}$ is distributed approximately normally with the following mean and variance

$$\begin{aligned} E[Z|I_{T-1}] &= E_{T-2} + \frac{1}{1+\sigma^2} N\mu_{T-1} + \frac{\sigma^2}{1+\sigma^2} x_{T-1} + N\mu_T - A_{T-2} - \frac{2p_{T-1}}{\theta} \\ V(Z|I_{T-1}) &= \frac{N\sigma^2}{(1+\sigma^2)} + \delta_T N\sigma^2 \end{aligned}$$

$$\text{with } \delta_T = \left(1 - \frac{1}{\theta} \frac{\partial}{\partial e_{T-1}} f_T(N\mu_T, E[e_{T-1}|x_{T-1}])\right)^2 = \left(\frac{K\phi_T}{\theta\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} + 1\right)^{-2} \in (0, 1) \text{ and } \phi_T = \phi\left(\frac{-E(Z|I_{T-1})}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}}\right).$$

Here again, we see that abatement changes the dynamics of excess emissions, Z , as information received at time $T - 1$ impacts both the expected excess emissions and its variance (through ϕ_T), a property absent in models without abatement (Chesney and Taschini, 2012) and in finance models that rely on normally distributed noises. The reason is that abatement is an endogenous response to shocks.

The expression for the variance in Lemma 2 captures the net contribution of emission shocks at time $T - 1$, $\frac{N\sigma^2}{(1+\sigma^2)}$ (since a_{T-1} is contemporaneously determined it does not contribute noise to $Z|I_{T-1}$), and the contribution from future shocks at time T , $\delta_T N\sigma^2$.

Remember that $N\sigma^2$ corresponds to the variance of emissions in period T , absent abatement. Therefore, δ_T can be viewed as the dampening due to abatement. Looking at the expression for δ_T in Lemma 2, it is easy to see that it is decreasing in K and increasing in θ . A higher penalty increases abatement by increasing the expected price everything else equal. This reduces the net effect of emission shocks, reducing the variance. On the other hand, an increase in θ means that abatement is getting more expensive, which reduces firms' adjustment to emissions shocks, everything else equal.

The level of dampening also depends on the current value of excess net emissions (through ϕ_T). It is maximal when $E(Z|I_{T-1}) = 0$ and minimal (i.e. δ_T converges to 1) when $E[Z|I_{T-1}]$ is either very large or very low. The reason is that the level of excess net emissions influences the returns to abatement. When excess net emissions are very large, the price is close to its ceiling, K . In that case, the marginal benefit of abatement is limited because it will not change the probability that the market is short very much. Likewise, when net emissions are very small (large in absolute value but negative), the equilibrium price is close to its floor, zero. Again the marginal benefit of

abatement is small. Instead, the marginal returns from abatement is high when $E(Z|I_{T-1})$ is close to zero, which explains that these are also circumstances where dampening is at its highest level.

Given Lemma 2, the price at time $T - 1$ is implicitly defined by:

$$p_{T-1} = K \left[1 - \Phi \left(\frac{-E(Z|I_{T-1})}{\sqrt{V(Z|I_{T-1})}} \right) \right] \quad (8)$$

where p_{T-1} appears both on the left of this equation but also on the right, in $E[Z|I_{T-1}]$. We show in the Appendix that (8) uniquely pins down the equilibrium price at $T - 1$. The intuition is similar to period T , namely that the righthand side, i.e. the probability of being short at the end of the phase, is decreasing with p_{T-1} because abatement increases with p_{T-1} . The difference with period T is that $V(Z|I_{T-1})$ also changes with p_{T-1} . The proof shows that the direct effect through abatement is larger than the indirect effect through a change in variance.

Lemma 3 (Uniqueness of the full information price equilibrium at time $T - 1$). *The full information equilibrium price implicitly defined by (8) is unique for N large enough.*

We can replicate this approach for $T - 2$ and before, leading to Theorem 1:

Theorem 1 (Full information equilibrium). *For N large enough, the full information equilibrium price is unique and implicitly defined by the following equation:*

$$p_t = K \left[1 - \Phi \left(\frac{-E(Z|I_t)}{\sqrt{V(Z|I_t)}} \right) \right] \text{ for all } t = 1, \dots, T \quad (9)$$

with

$$E(Z|I_t) = E_{t-1} + \frac{1}{1+\sigma^2} N\mu_t + \frac{\sigma^2}{1+\sigma^2} x_t + \sum_{k=t+1}^T N\mu_k - A_{t-1} - \frac{(T-t+1)}{\theta} p_t \quad (10)$$

$$V(Z|I_t) = \begin{cases} \frac{N\sigma^2}{1+\sigma^2} + \sum_{k=t+1}^T E \left[\left(1 - \frac{(T-k+1)}{\theta} \frac{\partial f_k}{\partial e_{k-1}} \right)^2 | I_t \right] N\sigma^2 & \text{for } t \leq T - 1 \\ \frac{N\sigma^2}{1+\sigma^2} & \text{for } t = T \end{cases} \quad (11)$$

where $\frac{\partial f_k}{\partial e_{k-1}}$ is shorthand notation for $\frac{\partial}{\partial e_{k-1}} f_k(N\mu_k, E[e_{k-1}|I_{k-1}])$.

At all times, in equilibrium, firms equate marginal abatement cost to the price and are indifferent between trading today or later. Equilibrium prices follow a martingale $E[p_t|I_{t-1}] = p_{t-1}$.

Theorem 1 shows that the equilibrium price takes the same form at all times. It is equal to the expected penalty at time $T + 1$, which depends on the level of penalty, K , and the probability that emissions exceed the target. The expression for expected excess emissions, $E(Z|I_t)$, consists of the

expectation of past, current and future emissions, based on information at time t , $E_{t-1} + \frac{1}{1+\sigma^2}N\mu_t + \frac{\sigma^2}{1+\sigma^2}x_t + \sum_{k=t+1}^T N\mu_k$, minus past, current and future abatement, where the term $\frac{(T-t+1)}{\theta}p_t$ captures the fact that firms expect to abate exactly the same amount per period going forward since p_t is the best forecast of future prices (intertemporal smoothing given that abatement costs are convex). The expression for the variance accounts for the dampening of future shocks through abatement.

Corollary 1 derives the law of motion for expected excess emissions and formalizes the extent to which uncertainty about the eventual balance between emissions and abatement decreases as time goes by.

Corollary 1 (Dynamics of excess emissions). *Expected excess emissions take the following recursive form:*

$$E[Z|I_t] = E[Z|I_{t-1}] + \begin{bmatrix} \frac{\sigma^2}{1+\sigma^2} & 1 & -\frac{(T-t+1)}{\theta} \end{bmatrix} \begin{bmatrix} x_t - N\mu_t \\ e_{t-1} - E[e_{t-1}|I_{t-1}] \\ p_t - p_{t-1} \end{bmatrix} \quad (12)$$

Their expected variance declines over time:

$$E[V(Z|I_{t+1})|I_t] = V(Z|I_t) - N\sigma^2 \left(1 - \frac{(T-t)}{\theta} \frac{\partial f_{t+1}}{\partial e_t}\right)^2$$

Expression (12) shows that beliefs about expected excess emissions are updated following arrival of new information about current period emissions ($x_t - N\mu_t$) and previous period emissions ($e_{t-1} - E[e_{t-1}|I_{t-1}]$), as well as changes in prices which lead market participants to update their beliefs about future abatements.

Corollary 1 shows that the variance of excess emissions only decreases in expectation. It eventually does decrease as expected given that uncertainty reduces over time.⁸ Equilibrium prices therefore eventually converge to 0 or K .

Finally, though equilibrium prices are not linear in signals, we can derive an approximate linear recursive form for them that highlights their martingale property and the nature of updating following new information arrival. Corollary 2 generalizes (6).

Corollary 2 (Approximate dynamics of equilibrium prices). *Equilibrium prices follow approximately the following recursive form:*

$$p_t = p_{t-1} + \nabla f_t(N\mu_t, E[e_{t-1}|x_{t-1}]) \begin{bmatrix} x_t - N\mu_t \\ e_{t-1} - E[e_{t-1}|x_{t-1}] \end{bmatrix}$$

with $\frac{\partial}{\partial x_t} f_t(N\mu_t, E[e_{t-1}|x_{t-1}]) = \frac{\sigma^2}{1+\sigma^2} \frac{\partial}{\partial e_{t-1}} f_t(N\mu_t, E[e_{t-1}|x_{t-1}])$.

⁸To see this, compare the expressions for $V(Z|I_T)$ in (4) and $V(Z|I_{T-1})$ in Lemma 2

3.2 Information aggregation

We now turn to the equilibrium in the economy with private information. Let I_{it} denote firm i 's information set at time t . Formally,

$$I_{it} = \{p_1, \dots, p_t, e_1, \dots, e_{t-2}, e_{t-1}, x_{i1}, \dots, x_{it}\} \text{ for integer values of } t$$

The next theorem shows that the full information equilibrium price is a sufficient statistic for market participants' private information. Therefore, the full information equilibrium is also an equilibrium in the economy with private information. Because the full information equilibrium is unique, so is the REE that fully aggregates private information.⁹

Theorem 2 (Full information aggregation). *There exists a unique dynamic Rational Expectations Equilibrium that fully aggregates market participants' private information. This equilibrium is characterized at integer times by the same equations (9) - (11) as the full information equilibrium. Moreover, $p_t = p_{t+0.5}$: prices are not affected by the publication of aggregate emissions at mid-period points because this information has already been integrated into the price.*

Theorem 2 shows that the strong form of the Efficient Market Hypothesis holds in our setting: all information - private and public - is integrated into the price as soon as it arises. This result is remarkable because, *a priori*, two parameters are needed to identify the distribution of excess emissions ($E[Z|I_t]$ and $V(Z|I_t)$) and there is only one price. The technical reason why the price is nevertheless a sufficient statistic for the information dispersed in the economy is the combination of two factors: first, the price reveals $\frac{E[Z|I_t]}{\sqrt{V(Z|I_t)}}$ (this is seen in (9)), second, $V(Z|I_t)$ is a known function of $E[Z|I_t]$ and the ratio $\frac{E[Z|I_t]}{\sqrt{V(Z|I_t)}}$ is strictly increasing in $E[Z|I_t]$. Together, these elements imply that $E[Z|I_t]$ and $V(Z|I_t)$ can be identified separately.

Economically, this comes from the fact that all incentives are eventually driven by $E[Z|I_t]$ only. In particular, the marginal returns from abatement depend on the market expectation about excess emissions, $E[Z|I_t]$, which in turn impacts the level of dampening of emissions shocks, and therefore $V(Z|I_t)$. This reduces the information to infer to a single dimensional variable, and the equilibrium price is therefore able to perform its information aggregation role.

We now examine the welfare properties of the rational expectations equilibrium. Montgomery (1972), Cronshaw and Kruse (1996) and Rubin (1996) have shown theoretically that competitive emissions markets with no uncertainty minimize aggregate compliance costs: individual optimization by firms results in a total compliance cost that is equivalent to the aggregate least-cost solution,

⁹This does not rule out other - non full information aggregating - REE (see e.g. Pálvölgyi and Venter, 2015, for a recent contribution to these questions).

a corollary of the First Welfare Theorem. Theorem 3 shows that, when full information aggregation takes place at equilibrium, this result generalizes to a market with private information and uncertainty, with the caveat that the welfare criterion now is expected least-cost solution conditional on all the information in the economy (constrained efficiency).

Theorem 3 (Constrained efficiency of the Rational Expectations Equilibrium). *When full information aggregation occurs, equilibrium decentralized abatement decisions minimize the expected sum of current and future abatement costs at all times, conditional on the information available in the economy.*

As Kling and Rubin (1997) have noted before, the least cost solution may not be the solution that maximizes social welfare if firms discount future payoffs differently from Society, if social costs of emissions differ over time, or if the number of allowances is set suboptimally. These caveats also apply here. In particular, K may differ from the marginal social cost of emissions beyond the initial target but the market minimizes compliance cost, including this penalty.

4 Aggregate uncertainty about emissions

So far, the only source of uncertainty in the model was private information about past and current emissions. In particular, the average difference between business-as-usual emissions and allowances, μ_t , was common knowledge. This is a strong assumption for at least two reasons. First, in the context of the launch of the EU ETS, business-as-usual emissions were hastily computed before phase I and allowances allocations decided on this basis. Ex-post, it turned out that some of these business-as-usual emissions were over-estimated. Second, μ_t can be interpreted as capturing the business cycle, with high values of μ_t during booms and low values during recessions. While it may be reasonable that the current business cycle is observable and common knowledge, business cycles are intrinsically uncertain and therefore future μ_t 's are unlikely to be known with certainty today.

In this section, we explore two polar cases that relax the assumption that μ_t is common knowledge and capture these sources of uncertainty:

Case 1: Unknown business-as-usual emissions Suppose that individual emissions are still drawn from a normal distribution $N(\mu_t, \sigma^2)$ as earlier, but now $\mu_t = \mu$ for all t and unknown to firms which hold a common prior $N(\mu_0, \tau_0^2)$ about μ .

Case 2: Uncertain future business cycle Suppose that individual emissions are still drawn from a normal distribution $N(\mu_t, \sigma^2)$ as earlier, but now μ_t is only observed at time

t . Future average emissions are not known but correlated over time. They follow the process $\mu_{t+1} = \mu_t + \rho\eta_t$ where η_t is i.i.d. $N(0, 1)$ and ρ is common knowledge.

We first discuss how these alternative assumptions affect the full information equilibrium. Under case 1, firms now not only learn about their past and current emissions at every period (and this information is shared), they also update their beliefs about μ . The following lemma describes how beliefs about μ and therefore beliefs about e_{it} are updated.

Lemma 4 (Updated beliefs about μ and future emissions under case 1). *Suppose that case 1 holds. In the full information economy, firms' posterior beliefs at time t about average emissions μ are normally distributed according to $N(\mu_t, \tau_t^2)$ with*

$$\begin{aligned}\mu_t &= \left(\frac{\sigma^2(1+\sigma^2)}{N\Delta}\right)\mu_{t-1} + \frac{(1+\sigma^2)\tau_{t-1}^2}{\Delta} \left[\frac{e_{t-1}}{N} + \frac{\sigma^2}{\sigma^2+1} \frac{x_t}{N}\right] \\ \tau_t^2 &= \frac{\sigma^2(1+\sigma^2)}{\Delta N} \tau_{t-1}^2\end{aligned}$$

where $\Delta = \tau_{t-1}^2(1 + 2\sigma^2) + \frac{\sigma^2(1+\sigma^2)}{N}$. Posterior beliefs at time t about current and future individual emissions are normally distributed according to $N(\mu_k, \sigma_k^2)$ with $\mu_k = \mu_t$ given above and $\sigma_k^2 = \sigma^2 + \tau_t^2$ for $k > t$.

Lemma 4 shows that, every period, firms update their beliefs about current and future emissions, following their learning about past emissions $\{e_{1t-1}, \dots, e_{Nt-1}\}$ and current emissions $\{x_{1t}, \dots, x_{Nt}\}$. They put more weight on their prior when realized emissions are noisy (σ^2 large) and their prior precise (τ_{t-1}^2 small). As time passes, their posterior about μ becomes increasingly precise (τ_t^2 decreases with t). Firms have the same prior about all future emissions.

Under case 2, firms observe the existing business cycle (and therefore learn about the current value of μ_t). They update their beliefs about future business cycles accordingly:

Lemma 5 (Updated beliefs about future emissions under case 2). *Suppose that case 2 holds. Posterior beliefs at time t about future individual emissions are normally distributed according to $N(\mu_k, \sigma_k^2)$ with $\mu_k = \mu_t$ and $\sigma_k^2 = (k - t)\rho^2 + \sigma^2$ for $k > t$.*

Lemma 5 shows that firms' best prediction about average individual emissions is the same for all future periods but become less precise as the time horizon increases. Future emissions are positively correlated.

Looking back at how we derived the full information equilibrium in the benchmark model, it is easy to see that the full information equilibrium in these richer informational environments is characterized by the same equations as in Theorem 1 except that, at time t , the values for $\{\mu_k\}_{k \geq t}$ and σ^2 used to compute $E[Z|I_t]$ and $V(Z|I_t)$ are now replaced by the values $\{\mu_k\}_{k \geq t}$ and $\{\sigma_k^2\}_{k \geq t}$

described in Lemma 4 or Lemma 5, whichever applies. Importantly, in both cases, $V(Z|I_t)$ is still a known function of $E(Z|I_t)$ and common knowledge parameters $K, \theta, N, T, \sigma^2$ and now μ_0, τ_0^2 and ρ (cf. intermediate step in the proof of Theorem 1). Thus, the argument for why prices separately identify $E[Z|I_t]$ and $V(Z|I_t)$ in Theorem 2 still goes through and full information aggregation obtains.

Theorem 4 (Full information aggregation when μ_t is not common knowledge). *Consider the extension of the benchmark model where either business-as-usual emissions are unknown (case 1) or future business cycles are uncertain (case 2). There exists a unique dynamic Rational Expectation Equilibrium that fully aggregates market participants' private information. This equilibrium is characterized at integer times by the same equations (9) - (11) as the full information equilibrium, with the difference that, at time t , the values for $\{\mu_k\}_{k \geq t}$ and σ^2 used to compute $E[Z|I_t]$ and $V(Z|I_t)$ are now replaced by the values $\{\mu_k\}_{k \geq t}$ and $\{\sigma_k^2\}_{k \geq t}$ described in Lemma 4 (case 1) or Lemma 5 (case 2). Moreover, $p_t = p_{t+0.5}$: price are not affected by the publication of aggregate emissions at mid-period points because this information has already been integrated into the price.*

Theorem 4 shows that the information aggregation result derived in the context of the benchmark model is robust to extensions to richer informational environments about the determinants of emissions.

5 When information aggregation fails

Another relevant source of private information in practice is abatement costs. These were assumed to be common knowledge in the benchmark model but are unlikely to be common knowledge in practice, especially in markets covering firms from multiple industries as is the case of the EU ETS.

To explore the impact of information asymmetries about abatement costs, we return to the benchmark model but now let firms' abatement cost parameters $(\theta_1, \dots, \theta_N)$ be private information, with some common prior over the joint distribution. This does not change the full information equilibrium which is still given by Theorem 1. However, now $V(Z|I_t)$ is no longer a known function of $E[Z|I_t]$. This is already seen in period $T - 1$ where $V(Z|I_{T-1}) = \frac{N\sigma^2}{(1+\sigma^2)} + \delta_T N\sigma^2$ where δ_T is defined in Lemma 2 and depends on both θ and $E[Z|I_{T-1}]$. Therefore, whereas p_t may still reveal $\frac{E[Z|I_t]}{\sqrt{V(Z|I_t)}}$, market participants are unable to separately identify $E[Z|I_t]$ and $V(Z|I_t)$. As a result, full information aggregation at integer times fails. For the same reason, the publication of past emissions at mid-period does not help re-establish information aggregation, even if prices may react to such publication. This is formalized in the next Theorem whose proof is omitted.

Theorem 5 (Failure of information aggregation when abatement costs are private information). *When both emissions and abatement costs are private information, equilibrium prices can no longer fully aggregate private information. The strong form of the Efficient Market Hypothesis fails.*

Theorem 5 establishes that the full information equilibrium, or equivalently, the informed social planner approach to emissions markets is not appropriate when there is asymmetric information about abatement costs. Intuitively, a low price can be the result of either lower emissions than anticipated or low abatement costs. Low emissions and low abatement costs should also impact future beliefs differently but firms have no way to tell from the price which event applies.

6 Concluding comments

A central argument in favor of market-based policy instruments such as emissions markets is their ability to decentralize decisions to those who have the information, thereby arguably fostering greater efficiency. Our paper examined this argument, using tools from finance to model explicitly the nature of information asymmetry and uncertainty in such markets.

We have found that markets can fully aggregate private information and generate informative prices if the source of informational asymmetry is due to firms' emissions. For the optimist, these results - especially Theorems 2-4 - can be seen as providing microfoundations - under some conditions - for the social planner approach to emissions markets.

In our view, however, a more suitable take-away from the analysis is that information aggregation is difficult to obtain. Despite the extremely favorable modeling assumptions that we have adopted (no friction beyond asymmetric information, a single price, ... etc), information aggregation breaks down as soon as there is some information asymmetry regarding abatement costs. Even when abatement costs are common knowledge, the inference problem - from prices to a sufficient statistic about market participants' information - is highly complex and non linear, raising doubts that real markets could solve it. The fact that market participants in the EU ETS closely watch the yearly publication of past realized emissions and that prices react to this is evidence that information aggregation probably fails in this market.

An interesting venue for future research is to understand what competitive emissions markets can realistically achieve when there is private information about abatement costs and, in particular, how market design can support the information aggregation process. Another important open question is to understand the efficiency implications for emissions prices as an investment signal in these markets, when information aggregation fails.

References

- [1] ALBEROLA, Emilie, Julien CHEVALLIER and Benoît CHÈZE (2008): “Price drivers and structural breaks in European carbon prices 2005–2007”, *Energy Policy*, Vol. 36, p. 787–797.
- [2] BRUNNERMEIER, Markus K. (2000): *Asset Pricing under Asymmetric Information*, Oxford University Press.
- [3] 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(5), p. 447-475.
- [4] CESPÀ, Giovanni and Xavier VIVES (2012): “Dynamic trading and asset prices: Keynes vs. Hayek”, *Review of Economic Studies*, Vol. 79, p. 539–580.
- [5] CRONSHAW, Mark and Jamie KRUSE (1996): “Regulated firms in pollution permit markets with banking”, *Journal of Regulatory Economics*, Vol. 9, p. 179-189.
- [6] DUBEY, Pradeep, John GEANAKOPOLOS and Martin SHUBIK (1987): “The revelation of information in strategic market games: A critique of rational expectations equilibrium”, *Journal of Mathematical Economics*, Vol. 16(2), p. 105-137.
- [7] FORSYTHE, Robert and Russell LUNDHOLM (1990): “Information aggregation in an experimental market”, *Econometrica*, Vol. 58(2), p. 309-347
- [8] GREENE, William H. (2011): *Econometric Analysis*, Pearson.
- [9] GROSSMAN, Sanford (1976): “On the efficiency of competitive stock markets where trades have diverse information”, *Journal of Finance*, Vol. 31(2), p. 573-585.
- [10] GROSSMAN, Sanford and Joseph STIGLITZ (1980): “On the impossibility of informationally efficient markets”, *American Economic Review*, Vol. 70(3), p. 393-408.
- [11] HE, Hua and Jiang WANG (1995): “Differential information and dynamic behavior of stock trading”, *Review of Financial Studies*, Vol. 8(4), p. 919-972.
- [12] HINTERMANN, Beat (2010): “Allowance price determinants in the EU ETS”, *Journal of Environmental Economics and Management*, Vol. 59(1), p. 43–56.
- [13] HITZEMANN, Steffen and Marliese UHRIG-HOMBURG (2018): “Equilibrium price dynamics of emission permits”, *Journal of Financial and Quantitative Analysis*, Vol. 54(3), p. 1653-78.

- [14] JACKSON, Matthew (1991): “Equilibrium, price formation and the value of private information”, *Review of Financial Studies*, Vol. 4, p. 1-16.
- [15] KOVALENKOV, Alexander and Xavier VIVES (2014): “Competitive rational expectations equilibria without apology”, *Journal of Economic Theory*, Vol. 149(1), p. 211-235.
- [16] KLING, Catherine and Jonathan RUBIN (1997): “Bankable permits for the control of environmental pollution”, *Journal of Public Economics*, Vol. 64(1), p. 101-115.
- [17] MONTGOMERY, WD. (1972): “Markets in licenses and efficient pollution control programs”, *Journal of Economic Theory*, Vol. 5, p. 395-418.
- [18] PÁLVÖLGYI, Dömötör and Gyuri VENTER (2015): “Multiple equilibria in noisy rational expectations economies”, mimeo.
- [19] PLOTT, Charles R. and Shyam SUNDER (1988): “Rational expectations and the aggregation of diverse information in laboratory security markets”, *Econometrica*, Vol. 56(5), p. 1085-1118.
- [20] RADNER, Roy (1979): “Rational expectations equilibrium: Generic existence and the information revealed by prices”, *Econometrica*, Vol. 47(3), p. 655-678.
- [21] RUBIN, Jonathan (1996): “A model of intertemporal emission trading, banking and borrowing”, *Journal of Environmental Economics and Management*, Vol. 31, p. 269-286.
- [22] SCHENNACH, Susanne (2000): “The economics of pollution permit banking in the context of Title IV of the 1990 Clean Air Act Amendments”, *Journal of Environmental Economics and Management*, Vol. 40, p. 189-210.
- [23] SEIFERT, Jan, Marliese UHRIG-HOMBURG and Michael WAGNER (2008): “Dynamic behavior of CO₂ spot prices”, *Journal of Environmental Economics and Management*, Vol. 56, p. 180-194.
- [24] TIROLE, Jean (1982): “On the possibility of speculation under rational expectations”, *Econometrica*, Vol. 50, p. 1163–82.
- [25] VIVES, Xavier (2008): *Information and Learning in Markets*, Princeton University Press.

Appendix: Proofs

Proof of Lemma 1. (i) Applying the implicit function theorem on (4) yields

$$\frac{\partial}{\partial e_{T-1}} f_T(x_T, e_{T-1}) = \frac{1}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} \frac{K\phi_1}{1 + \frac{K\phi_1}{\theta\sqrt{\frac{N\sigma^2}{1+\sigma^2}}}} = \frac{K\theta\phi_1}{\theta\sqrt{\frac{N\sigma^2}{1+\sigma^2}} + K\phi_1},$$

where $\phi_1 = \phi\left(\frac{A_{T-1} + \frac{p_T}{\theta} - E_{T-1} - \frac{1}{1+\sigma^2}N\mu_T - \frac{\sigma^2}{1+\sigma^2}x_T}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}}\right)$, and likewise,

$$\frac{\partial}{\partial x_T} f_T(x_T, e_{T-1}) = \frac{\sigma^2}{1+\sigma^2} \frac{1}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} \frac{K\phi_1}{1 + \frac{K\phi_1}{\theta\sqrt{\frac{N\sigma^2}{1+\sigma^2}}}} = \frac{\sigma^2}{1+\sigma^2} \frac{\partial}{\partial e_{T-1}} f_T(x_T, e_{T-1}),$$

proving the first part of the claim.

(ii) In the absence of abatement, p_T does not appear on the righthand side of (4), leading to simple

derivatives: $\frac{\partial}{\partial e_{T-1}} \tilde{f}_T(x_T, e_{T-1}) = \frac{1}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} K\phi_2$ and $\frac{\partial}{\partial x_T} \tilde{f}_T(x_T, e_{T-1}) = \frac{\sigma^2}{1+\sigma^2} \frac{1}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} K\phi_2$, where

$$\phi_2 = \phi\left(\frac{-E_{T-1} - \frac{1}{1+\sigma^2}N\mu_T - \frac{\sigma^2}{1+\sigma^2}x_T}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}}\right).$$

The second claim follows directly from the comparison between the expressions of $\frac{\partial}{\partial x_T} f_T(x_T, e_{T-1})$ and $\frac{\partial}{\partial x_T} \tilde{f}_T(x_T, e_{T-1})$, and $\frac{\partial}{\partial e_{T-1}} f_T(x_T, e_{T-1})$ and $\frac{\partial}{\partial e_{T-1}} \tilde{f}_T(x_T, e_{T-1})$ respectively when $\phi_1 = \phi_2$.

(iii) From the first order conditions for abatement (1), $a_T = \frac{1}{\theta} f_T(x_T, e_{T-1})$. Therefore, using the derivations in (i):

$$\frac{\partial}{\partial x_T} a_T = \frac{\sigma^2}{1+\sigma^2} \frac{K\phi_1}{\theta\sqrt{\frac{N\sigma^2}{1+\sigma^2}} + K\phi_1} < \frac{\sigma^2}{1+\sigma^2} = \frac{\partial}{\partial x_T} E[e_T|I_T]$$

Likewise,

$$\frac{\partial}{\partial e_{T-1}} a_T = \frac{K\phi_1}{\theta\sqrt{\frac{N\sigma^2}{1+\sigma^2}} + K\phi_1} < 1 = \frac{\partial}{\partial e_{T-1}} E[e_{T-1}|I_T]$$

□

Proof of Lemma 2. By definition, $Z = E_{T-2} - A_{T-1} + \tilde{e}_{T-1} + \tilde{e}_T - \tilde{a}_T$, where the variables highlighted with a tilda are unknown as of time $T-1$. Aggregate emissions \tilde{e}_{T-1} and \tilde{e}_T are independently distributed, conditional on I_{T-1} , with

$$\tilde{e}_T | I_{T-1} \sim N(N\mu_T, N\sigma^2)$$

and

$$\tilde{e}_{T-1} | I_{T-1} \sim N \left(\frac{1}{1 + \sigma^2} N \mu_{T-1} + \frac{\sigma^2}{1 + \sigma^2} x_{T-1}, \frac{N \sigma^2}{1 + \sigma^2} \right)$$

Aggregate abatement \tilde{a}_T can be expressed as a function of the unknown variables (as of time $T-1$), e_{T-1} and x_T , $\frac{f_T(x_T, e_{T-1})}{\theta}$, where f_T is the function defined implicitly by (4) that maps e_{T-1} and x_T into p_T . Linearising f_T around the expectations of its arguments, conditional on I_{T-1} , implies that $Z|I_{T-1}$ can be approximated by the sum of three independent normally distributed variables, plus some constant terms:¹⁰

$$\begin{aligned} Z|I_{T-1} &= \left(E_{T-2} - A_{T-1} + e_{T-1} + e_T - \frac{f_T(x_T, e_{T-1})}{\theta} \right) | I_{T-1} \\ &\simeq E_{T-2} - A_{T-1} + \left(1 - \frac{1}{\theta} \frac{\partial f_T}{\partial e_{T-1}} \right) e_{T-1} | I_{T-1} + \left(1 - \frac{1}{\theta} \frac{\partial f_T}{\partial x_T} \right) e_T - \frac{1}{\theta} \frac{\partial f_T}{\partial x_T} \varepsilon_T \\ &\quad - \frac{1}{\theta} f_T(N \mu_T, E[e_{T-1} | I_{T-1}]) + \frac{1}{\theta} \left(\frac{\partial f_T}{\partial x_T} N \mu_T + \frac{\partial f_T}{\partial e_{T-1}} E[e_{T-1} | I_{T-1}] \right) \end{aligned}$$

where $\frac{\partial f_T}{\partial e_{T-1}}$ is shorthand notation for $\frac{\partial}{\partial e_{T-1}} f_T(N \mu_T, E[e_{T-1} | x_{T-1}])$ and $\frac{\partial f_T}{\partial x_T}$ is shorthand notation for $\frac{\partial}{\partial x_T} f_T(N \mu_T, E[e_{T-1} | x_{T-1}])$ derived in Lemma 1. Therefore $Z|I_{T-1}$ is approximately normally distributed with mean:

$$\begin{aligned} E[Z|I_{T-1}] &= E_{T-2} - A_{T-2} - \frac{p_{T-1}}{\theta} + \frac{1}{1 + \sigma^2} N \mu_{T-1} + \frac{\sigma^2}{1 + \sigma^2} x_{T-1} + N \mu_T - \frac{1}{\theta} f_T(N \mu_T, E[e_{T-1} | x_{T-1}]) \\ &= E_{T-2} - A_{T-2} - 2 \frac{p_{T-1}}{\theta} + \frac{1}{1 + \sigma^2} N \mu_{T-1} + \frac{\sigma^2}{1 + \sigma^2} x_{T-1} + N \mu_T \end{aligned}$$

where we have used the fact that $f_T(N \mu_T, E[e_{T-1} | x_{T-1}]) = E[p_T | I_{T-1}] = p_{T-1}$ by the law of iterated expectations. The variance of $Z|I_{T-1}$ can be derived using the law of total variance:

$$\begin{aligned} V[Z|I_{T-1}] &= E[V[Z|I_T] | I_{T-1}] + V[E[Z|I_T] | I_{T-1}] \\ &= \frac{N \sigma^2}{(1 + \sigma^2)} + V \left[\frac{1}{(1 + \sigma^2)} N \mu_T + \frac{\sigma^2}{(1 + \sigma^2)} x_T + E_{T-1} - A_{T-1} - \frac{f_T(x_T, e_{T-1})}{\theta} | I_{T-1} \right] \\ &= \frac{N \sigma^2}{(1 + \sigma^2)} + V \left[\left(\frac{\sigma^2}{(1 + \sigma^2)} - \frac{1}{\theta} \frac{\partial f_T}{\partial x_T} \right) x_T + \left(1 - \frac{1}{\theta} \frac{\partial f_T}{\partial e_{T-1}} \right) e_{T-1} | I_{T-1} \right] \\ &= \frac{N \sigma^2}{(1 + \sigma^2)} + \left(\frac{\sigma^2}{(1 + \sigma^2)} \right)^2 \left(1 - \frac{1}{\theta} \frac{\partial f_T}{\partial e_{T-1}} \right)^2 N (\sigma^2 + 1) + \left(1 - \frac{1}{\theta} \frac{\partial f_T}{\partial e_{T-1}} \right)^2 \frac{N \sigma^2}{(1 + \sigma^2)} \\ &= \frac{N \sigma^2}{(1 + \sigma^2)} + \left(1 - \frac{1}{\theta} \frac{\partial f_T}{\partial e_{T-1}} \right)^2 N \sigma^2 = \frac{N \sigma^2}{(1 + \sigma^2)} + \left(\frac{K \phi_T}{\theta \sqrt{\frac{N \sigma^2}{1 + \sigma^2}}} + 1 \right)^2 N \sigma^2 \end{aligned} \tag{13}$$

where the second and third lines use the expression for $E[Z|I_T]$ and the linearization of p_T as a function of x_T and e_{T-1} derived in the main text, and the fourth line leverages Lemma 1 ($\frac{\partial f_T}{\partial x_T} = \frac{\sigma^2}{1 + \sigma^2} \frac{\partial f_T}{\partial e_{T-1}}$) and the fact that $V(e_{T-1} | I_{T-1}) = V(e_{T-1} | x_{T-1}) = \frac{N \sigma^2}{1 + \sigma^2}$. The last line follows from the expression for $\frac{\partial f_T}{\partial e_{T-1}}$ derived in the proof of Lemma 1. \square

¹⁰This is the delta method in econometrics (see Greene, 2003): by using a first-order Taylor series expansion around the conditional means of its arguments, x_T and e_{T-1} , the non linear function f_T can be approximated by a linear function that takes normally distributed values if x_T and e_{T-1} are normally distributed.

Proof of Lemma 3. Given that the lefthand side of (8) is strictly increasing in p_{T-1} , that both sides are continuous and that the lefthand side is smaller than the righthand side at $p_{T-1} = 0$, a sufficient condition for (8) to admit a unique solution is that $\frac{E(Z|I_{T-1})}{\sqrt{V(Z|I_{T-1})}}$ is decreasing in p_{T-1} . Let $w_{T-1} = E[Z|I_{T-1}]$. Using Lemma 2, we can write $V(Z|I_{T-1})$ as a function of w_{T-1} :

$$h_{T-1}(w_{T-1}) \equiv V(Z|I_{T-1}) = \frac{N\sigma^2}{1 + \sigma^2} + \delta_T(w_{T-1})N\sigma^2 \quad (14)$$

where $\delta_T(w_{T-1}) = \left(\frac{K\sqrt{1+\sigma^2}}{\theta\sqrt{N\sigma^2}} \phi \left(\frac{-w_{T-1}\sqrt{1+\sigma^2}}{\sqrt{N\sigma^2}} \right) + 1 \right)^{-2}$. Since $E[Z|I_{T-1}]$ is decreasing in p_{T-1} , we have

$$\text{sign} \left(\frac{\partial}{\partial p_{T-1}} \left(\frac{E(Z|I_{T-1})}{\sqrt{V(Z|I_{T-1})}} \right) \right) = -\text{sign} \left(\frac{\partial}{\partial w_{T-1}} \left(\frac{w_{T-1}}{\sqrt{h_{T-1}(w_{T-1})}} \right) \right)$$

Expanding this expression,

$$\frac{\partial}{\partial w_{T-1}} \left(\frac{w_{T-1}}{\sqrt{h_{T-1}(w_{T-1})}} \right) = h_{T-1}(w_{T-1})^{-\frac{1}{2}} - \frac{1}{2} h_{T-1}(w_{T-1})^{-\frac{3}{2}} w_{T-1} h'_{T-1}(w_{T-1}),$$

we see that a necessary and sufficient condition for $\frac{E(Z|I_{T-1})}{\sqrt{V(Z|I_{T-1})}}$ to be decreasing in p_{T-1} is that

$$h_{T-1}(w_{T-1}) - \frac{1}{2} h'_{T-1}(w_{T-1}) w_{T-1} > 0$$

Substituting for $h_{T-1}(\cdot)$ and $h'_{T-1}(\cdot)$ using (14) and dividing by $N\sigma^2$, this is equivalent to:

$$\frac{1}{1 + \sigma^2} + \delta_T(w_{T-1}) - \frac{1}{2} \delta'_T(w_{T-1}) w_{T-1} > 0 \quad (15)$$

The strategy of the proof is to show that (15) holds for N large enough. We do this by showing that $\lim_{N \rightarrow \infty} \delta_T(w_{T-1}) = 1$ and $\lim_{N \rightarrow \infty} \delta'_T(w_{T-1}) w_{T-1} = 0$.

Claim 1: $\lim_{N \rightarrow \infty} \delta_T(w_{T-1}) = 1$

This follows directly from the fact that ϕ is bounded:

$$\lim_{N \rightarrow \infty} \left(\frac{K\sqrt{1+\sigma^2}}{\theta\sqrt{N\sigma^2}} \phi \left(\frac{-w_{T-1}\sqrt{1+\sigma^2}}{\sqrt{N\sigma^2}} \right) + 1 \right)^{-2} = 1$$

Claim 2: $\lim_{N \rightarrow \infty} \frac{1}{2} \delta'_T(w_{T-1}) w_{T-1} = 0$

$$\begin{aligned} & \lim_{N \rightarrow \infty} \frac{1}{2} \delta'_T(w_{T-1}) w_{T-1} \\ &= \lim_{N \rightarrow \infty} \left(\frac{K\sqrt{1+\sigma^2}}{\theta\sqrt{N\sigma^2}} \phi \left(\frac{-w_{T-1}\sqrt{1+\sigma^2}}{\sqrt{N\sigma^2}} \right) + 1 \right)^{-3} \frac{K(1+\sigma^2)}{\theta\sigma^2} \phi' \left(\frac{-w_{T-1}\sqrt{1+\sigma^2}}{\sqrt{N\sigma^2}} \right) \frac{w_{T-1}}{N} \\ &= \frac{K(1+\sigma^2)}{\theta\sigma^2} \lim_{N \rightarrow \infty} \phi' \left(\frac{-w_{T-1}\sqrt{1+\sigma^2}}{\sqrt{N\sigma^2}} \right) \frac{w_{T-1}}{N} \\ &= 0 \end{aligned}$$

since $\frac{w_{T-1}}{N}$ converges to a finite number (w_{T-1} scales with N) and $\lim_{x \rightarrow \pm\infty} \phi'(x) = 0$

Putting claim 1 and claim 2 together leads to:

$$\lim_{N \rightarrow \infty} \frac{1}{1 + \sigma^2} + \delta_T(w_{T-1}) - \frac{1}{2} \delta'_T(w_{T-1}) w_{T-1} = \frac{1 + 2\sigma^2}{1 + \sigma^2} > 0$$

Continuity then implies that this inequality also holds for N large enough. \square

Proof of Theorem 1. It is easy to check that the solution for the equilibrium price at time T and $T - 1$ satisfies (9) to (11). To prove that Theorem 1 holds for $t < T - 1$, we first show by induction that $p_t = E[p_{T+1}|I_t]$ for all t . This clearly holds for $t = T, T - 1$. Suppose it holds for $k + 1$ and above, and consider the optimization problem of market participant i at time k :

$$\min_{a_{ik}, y_{ik}} p_k y_{ik} + c(a_{ik}, \theta_i) + V_{k+1}^i(a_{ik}, \theta_i | I_k)$$

Using the fact that $p_t = E[p_{T+1}|I_t]$ for $t > k$ and the law of iterated expectations, this optimization problem can be rewritten as:

$$\min_{a_{ik}, y_{ik}} p_k y_{ik} + c(a_{ik}, \theta_i) + E \left[\sum_{l=k+1}^T c(a_{il}, \theta_i) + p_{T+1}(E_{iT} - A_{iT} - Y_{ik}) | I_k \right]$$

where all transactions after time k cancel out in expectations. Common information and beliefs, and market clearing, then lead to the familiar first order conditions for equilibrium:

$$\begin{aligned} p_k &= E[p_{T+1}|I_k] \\ c_a(a_{ik}, \theta_i) &= \theta_i a_{ik} = E[p_{T+1}|I_k] \end{aligned}$$

Because this was shown for an arbitrary k , we conclude that $p_t = E[p_{T+1}|I_t]$ (and $a_{it} = \frac{p_t}{\theta_i}$) for all t . In particular, this implies that p_t is a martingale: $E[p_t|I_{t-1}] = p_{t-1}$.

From the equilibrium at $T + 1$,

$$p_t = E[p_{T+1}|I_t] = K Pr(Z > 0 | I_t) \tag{16}$$

where, from the perspective of time t , Z is given by:

$$Z = E_{t-1} - A_t + \sum_{k=t}^T e_k - \sum_{k=t+1}^T \frac{p_k}{\theta} \tag{17}$$

In expression (17), e_k is normally distributed and p_k is the result of normally distributed shocks and can be approximated, from the perspective of time t , by the sum of normally distributed random variables. Therefore $Z|I_t$ is approximately normally distributed and thus pinned down by its mean and variance, which we now derive.

Expression for $E[Z|I_t]$:

Taking expectations of (17) given I_t and applying the law of iterated expectations, we get:

$$E[Z|I_t] = E_{t-1} - A_{t-1} + \frac{1}{1+\sigma^2}N\mu_t + \frac{\sigma^2}{1+\sigma^2}x_t + \sum_{k=t+1}^T N\mu_k - (T-t+1)\frac{p_t}{\theta}$$

i.e. (10). For future reference, we note that this implies that $E[Z|I_t]$ takes the following recursive form:

$$E[Z|I_t] = E[Z|I_{t-1}] + \frac{\sigma^2}{1+\sigma^2}(x_t - N\mu_t) + (e_{t-1} - \frac{1}{1+\sigma^2}N\mu_{t-1} - \frac{\sigma^2}{1+\sigma^2}x_{t-1}) - \frac{(T-t+1)}{\theta}(p_t - p_{t-1}) \quad (18)$$

Expression for $V[Z|I_t]$:

To derive the expression for the variance, we again use an induction argument and the expression derived for $E[Z|I_t]$. Suppose that (11) holds for all $t > l$. From the law of total variance,

$$\begin{aligned} V(Z|I_l) &= E[V(Z|I_{l+1})|I_l] + V(E[Z|I_{l+1}]|I_l) \\ &= \frac{N\sigma^2}{1+\sigma^2} + \sum_{k=l+2}^T E\left[\left(1 - \frac{(T-k+1)}{\theta} \frac{\partial f_k}{\partial e_{k-1}}\right)^2 |I_l\right] N\sigma^2 \\ &\quad + V\left(E_l - A_l + \frac{1}{1+\sigma^2}N\mu_{l+1} + \frac{\sigma^2}{1+\sigma^2}x_{l+1} + \sum_{k=l+2}^T N\mu_k - (T-l)\frac{p_{l+1}}{\theta} |I_l\right) \\ &= \frac{N\sigma^2}{1+\sigma^2} + \sum_{k=l+2}^T E\left[\left(1 - \frac{(T-k+1)}{\theta} \frac{\partial f_k}{\partial e_{k-1}}\right)^2 |I_l\right] N\sigma^2 + V\left(e_l + \frac{\sigma^2}{1+\sigma^2}x_{l+1} - (T-l)\frac{p_{l+1}}{\theta} |I_l\right) \end{aligned}$$

From the perspective of time l , p_{l+1} is a function of x_{l+1} and e_l which we denote $f_{l+1}(x_{l+1}, e_l)$.

Linearizing this function around the expectations of its two arguments and accounting for the fact that $\frac{\partial}{\partial x_{l+1}}f_{l+1}(N\mu_{l+1}, E[e_l|I_l]) = \frac{\sigma^2}{1+\sigma^2} \frac{\partial}{\partial e_l}f_{l+1}(N\mu_{l+1}, E[e_l|I_l])$ implies that:

$$\begin{aligned} V\left(e_l + \frac{\sigma^2}{1+\sigma^2}x_{l+1} - (T-l)\frac{p_{l+1}}{\theta} |I_l\right) &= V\left(e_l \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right) + \frac{\sigma^2}{1+\sigma^2}x_{l+1} \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right) |I_l\right) \\ &= \left(\frac{N\sigma^2}{1+\sigma^2} + \frac{N\sigma^4}{1+\sigma^2}\right) \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right)^2 = N\sigma^2 \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right)^2 \end{aligned}$$

Therefore:

$$V(Z|I_l) = \frac{N\sigma^2}{1+\sigma^2} + \sum_{k=l+1}^T E\left[\left(1 - \frac{(T-k+1)}{\theta} \frac{\partial f_k}{\partial e_{k-1}}\right)^2 |I_l\right] N\sigma^2 \quad (19)$$

as claimed, or equivalently and for future reference:

$$V(Z|I_l) = E[V(Z|I_{l+1})|I_l] + N\sigma^2 \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right)^2 \quad (20)$$

Putting the expression for $E[Z|I_l]$ and $V[Z|I_l]$ together with (16) yields (9).

The last step in the proof is to check for uniqueness (existence follows directly from Brouwer's fixed point theorem on the domain $[0, K]$). We first show that $V(Z|I_t)$ is fully pinned down by $E[Z|I_t]$.

Intermediary step: $V(Z|I_t)$ is fully pinned down by $E[Z|I_t]$.

The proof proceeds by induction. The claim holds for $V(Z|I_{T-1})$ as shown in the proof of Lemma 3 (and trivially so for $V(Z|I_T)$). Suppose it also holds for all $t > l$. Let $h_t(\cdot)$ describe the function that maps $E[Z|I_t]$ into $V(Z|I_t)$ for $t > l$ i.e. $V(Z|I_t) = h_t(E[Z|I_t])$. Then (20) can be rewritten as:

$$V(Z|I_l) = E[h_{l+1}(E[Z|I_{l+1}]|I_l)] + N\sigma^2 \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right)^2 \quad (21)$$

We first consider the second term of the righthand side. By definition, $\frac{\partial f_{l+1}}{\partial e_l}$ describes how p_{l+1} changes with e_l , evaluated at $x_{l+1} = N\mu_{l+1}$ and $e_l = E[e_l|I_l]$. Applying the implicit function theorem to (9) to obtain an expression for $\frac{\partial}{\partial e_l} f_{l+1}(N\mu_{l+1}, E[e_l|I_l])$ and rearranging, we get that :

$$\left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right)^2 = \left(1 + \frac{K(T-l)}{\theta} \phi(R) \frac{\partial R}{\partial e_l}\right)^{-2} \quad (22)$$

where $R = \frac{E[Z|I_{l+1}]}{\sqrt{V(Z|I_{l+1})}}$ evaluated at $x_{l+1} = N\mu_{l+1}$ and $e_l = E[e_l|I_l]$ and the corresponding equilibrium price. Since, by assumption, $V(Z|I_{l+1})$ is fully pinned down by $E[Z|I_{l+1}]$ and $E[Z|I_{l+1}] = E[Z|I_l]$ at $x_{l+1} = N\mu_{l+1}$ and $e_l = E[e_l|I_l]$ (no new information), then (22) is fully pinned down by $E[Z|I_l]$ (and the knowledge of the model parameters, K, θ, N, T and σ^2).

We now turn to the first term of the righthand side of (21). The martingale property of expected excess emissions, (18), implies that $E[h_{l+1}(E[Z|I_{l+1}]|I_l)]$ can be rewritten as

$$E[h_{l+1}(E[Z|I_l] + \Delta_{l+1})|I_l] \quad (23)$$

where

$$\Delta_{l+1} = \frac{\sigma^2}{1 + \sigma^2}(x_{l+1} - N\mu_{l+1}) + (e_l - \frac{1}{1 + \sigma^2}N\mu_l - \frac{\sigma^2}{1 + \sigma^2}x_l) - \frac{(T-l)}{\theta}(p_{l+1} - p_l)$$

is approximately normally distributed from the perspective of time l , with $E[\Delta_{l+1}|I_l] = 0$ and $V(\Delta_{l+1}|I_l) = N\sigma^2 \left(1 - \frac{(T-l)}{\theta} \frac{\partial f_{l+1}}{\partial e_l}\right)^2$, a function of $E[Z|I_l]$ (following the same arguments as above). Therefore, $E[h_{l+1}(E[Z|I_{l+1}]|I_l)]$ depends both directly, as an argument in h_{l+1} , and indirectly, through the distribution of Δ_{l+1} , on $E[Z|I_l]$.

Uniqueness of equilibrium price:

A sufficient condition for the equilibrium price to be unique is that the righthand side of (9) is decreasing in p_t , or equivalently, that $\frac{E(Z|I_t)}{\sqrt{V(Z|I_t)}}$ is decreasing in p_t . Leveraging the intermediary step, let $w_t = E[Z|I_t]$ and $h_t(w_t) = V(Z|I_t)$. Since $E[Z|I_t]$ is decreasing in p_t , we have

$$\text{sign} \left(\frac{\partial}{\partial p_t} \left(\frac{E(Z|I_t)}{\sqrt{V(Z|I_t)}} \right) \right) = -\text{sign} \left(\frac{\partial}{\partial w_t} \left(\frac{w_t}{\sqrt{h_t(w_t)}} \right) \right)$$

so we need to show that:

$$\frac{\partial}{\partial w_t} \left(\frac{w_t}{\sqrt{h_t(w_t)}} \right) > 0 \text{ for all } t \quad (24)$$

Let $\delta_{t+1}(w_t) = \left(1 - \frac{(T-t)}{\theta} \frac{\partial f_{t+1}}{\partial e_t}\right)^2$, the dampening factor for period $t+1$. From (19), $h_t(w_t)$ can be expressed as

$$h_t(w_t) = \frac{N\sigma^2}{1 + \sigma^2} + N\sigma^2 \delta_{t+1}(w_t) + N\sigma^2 \sum_{k=t+2}^T E[\delta_k(w_{k-1})|I_t] \quad (25)$$

for $t < T$, where, using (22),

$$\delta_{t+1}(w_t) = \left(1 + \frac{(T-t)}{\theta} K \phi \left(\frac{-w_t}{\sqrt{h_{t+1}(w_t)}} \right) \left(\frac{\partial}{\partial w_t} \frac{w_t}{\sqrt{h_{t+1}(w_t)}} \right) \right)^{-2} \quad (26)$$

From Lemma 2, we know that $\delta_T(w)$ is strictly bounded between 0 and 1, and likewise, we established in Lemma 3 that $\frac{\partial}{\partial w} \frac{w}{\sqrt{h_{T-1}(w)}} > 0$ for N large enough, which implies (from (26)) that $\delta_{T-1}(w)$ is also strictly bounded between 0 and 1. Referring back to (25), this implies that $h_{T-2}(w_{T-2})$ scales up with N (denoted $h_{T-2}(w_{T-2}) = O(N)$).

Define the function $g_t(w_t, N) = \frac{w_t}{\sqrt{h_t(w_t, N)}}$ (where with a slight abuse of notation we have made explicit the dependence of h_t on N). When $t = T-2$, this function scales with $N^{\frac{1}{2}}$ since its numerator scales with N and we have just established that its denominator scales with $N^{\frac{1}{2}}$. Let's rewrite $w_t = N\bar{w}_t$ to highlight the dependence of the numerator on N . Taking the total derivative of $g_{T-2}(w_{T-2}, N)$ with respect to N , this means that, for $\bar{w}_t > 0$ and N large enough:

$$\frac{d}{dN} g_{T-2}(w_{T-2}, N) = \bar{w}_{T-2} \frac{\partial}{\partial w_{T-2}} g_{T-2}(w_{T-2}, N) + \frac{\partial}{\partial N} g_{T-2}(w_{T-2}, N) > 0$$

Because the second term is negative, we can conclude that

$$\frac{\partial}{\partial w_{T-2}} g_{T-2}(w_{T-2}, N) \equiv \frac{\partial}{\partial w_{T-2}} \left(\frac{w_{T-2}}{\sqrt{h_{T-2}(w_{T-2})}} \right) > 0$$

for N large enough. We can apply this argument recursively to establish that (24) for all t . \square

Proof of Corollary 1. The corollary follows directly from (18) and (20). \square

Proof of Corollary 2. Corollary 2 follows directly from a linearization of (9) around $x_t = N\mu_t$ and $e_{t-1} = E[e_{t-1}|I_{t-1}]$. \square

Proof of Theorem 2. From equation (9), market participants can infer $\frac{E(Z|I_t)}{\sqrt{V(Z|I_t)}}$ from the observation of the price. Moreover, the argument establishing the uniqueness of the equilibrium price in the proof of Theorem 1 established that $\frac{E[Z|I_t]}{\sqrt{V(Z|I_t)}}$ is strictly increasing in $E[Z|I_t]$. Therefore, $E[Z|I_t]$ and $V(Z|I_t)$ are separately identified from the observation of equilibrium prices. These are also sufficient statistics for Z conditional on I_t . Therefore, $E[p_{T+1}|I_{it}] = E[p_{T+1}|I_t]$ for all i and market participants minimize their expected compliance costs at the full information equilibrium actions. \square

Proof of Theorem 3. We show that the solution $\{a_{1t}^{FB}, \dots, a_{it}^{FB}, \dots, a_{Nt}^{FB}\}_{t=1}^T$ to the dynamic aggregate cost minimization problem is identical to the abatement decisions arising in the market equilibrium.

At time $T + 1$, emissions and abatements are fixed so there is no decision to take. Either cumulative aggregate emissions are higher than cumulative abatement, in which case a penalty of K per missing allowance is imposed, or $E_T \leq A_T$, in which case no further cost is incurred. Therefore, at time $T + 1$, aggregate abatement costs are equal to $\max \{K (E_T - A_T), 0\}$.

The cost minimization problem from time T 's perspective is given by:

$$\min_{a_{1T}, \dots, a_{NT}} \sum_{i=1}^N c(a_{iT}, \theta_i) + E [\max \{K (e_T + E_{T-1} - A_T - A_{T-1}), 0\} | I_T] \quad (27)$$

Let $\varphi(\cdot | I_T)$ denote the pdf of e_T , the only variable that is not known with certainty in (27), conditional on information I_T . Expression (27) can be rewritten as:

$$\min_{a_{1T}, \dots, a_{NT}} \sum_{i=1}^N c(a_{iT}, \theta_i) + K \int_{a_T + A_{T-1} - E_{T-1}}^{\infty} (e_T + E_{T-1} - A_T) \varphi(e_T | I_T) de_T$$

with first order conditions:

$$\theta_i a_{iT}^{FB} - K \int_{a_T + A_{T-1} - E_{T-1}}^{\infty} \varphi(e_T | I_T) de_T = 0$$

Using the fact that $e_T | I_T \sim N(\frac{1}{(1+\sigma^2)} N\mu_T + \frac{\sigma^2}{(1+\sigma^2)} x_T, \frac{N\sigma^2}{1+\sigma^2})$, this expression can be rewritten as:

$$\theta_i a_{iT}^{FB} = K \left[1 - \Phi \left(\frac{A_{T-1} + a_T^{FB} - E_{T-1} - \frac{1}{(1+\sigma^2)} N\mu_T - \frac{\sigma^2}{(1+\sigma^2)} x_T}{\sqrt{\frac{N\sigma^2}{1+\sigma^2}}} \right) \right]$$

Comparing this expression with the equation characterizing the competition market equilibrium (1) and (4), we conclude that a_{iT}^{FB} is the same as the abatement decision that arises in the market equilibrium, conditional on being in the same situation, i.e. past abatement and past realized emissions being the same.

The rest of the proof proceeds by induction. Suppose that the first best levels of abatement for period $t + 1$ and above are identical to those arising in the market equilibrium and consider the aggregate cost minimization problem at time t :

$$\min_{a_{1t}, \dots, a_{Nt}} \sum_{i=1}^N c(a_{it}, \theta_i) + E \left[\sum_{k=t+1}^T \sum_{i=1}^N c(a_{ik}, \theta_i) + K \int_0^{\infty} Z \varphi(Z | I_t) dZ \right]$$

where Z is defined by (2) as before, and $\varphi(\cdot | I_t)$ denotes its pdf, conditional on I_t . The first-order conditions for $i = 1, \dots, N$ are given by:

$$\theta_i a_{it}^{FB} = K \int_0^{\infty} \varphi(Z | I_t) dZ \quad (28)$$

Because, by assumption, a_{ik}^{FB} is the same as in the market equilibrium for all $k > t$, then $Z - A_{t-1} - E_{t-1}|I_t$ has the same distribution as in the market equilibrium, $Z - A_{t-1} - E_{t-1}|I_t \sim N(E[Z|I_t] - A_{t-1} - E_{t-1}, V(Z|I_t))$, with $E[Z|I_t]$ and $V(Z|I_t)$ described by (10) and (11). Therefore, (28) can be rewritten as:

$$\theta_i a_{it}^{FB} = K \left[1 - \Phi \left(\frac{-E(Z|I_t)}{\sqrt{V(Z|I_t)}} \right) \right]$$

and the claim follows from comparison with (9). \square

Proof of Lemma 4. At every integer time t , two sets of new information, $\{e_{1t-1}, \dots, e_{Nt-1}\}$ and $\{x_{1t}, \dots, x_{Nt}\}$, can be used to update firms' prior beliefs about μ . Because the prior about μ and the e_{it-1} 's and x_{it} 's are normally distributed, the posterior belief about μ is also normally distributed. Let $\mu \sim N(\mu_{t-1}, \tau_{t-1}^2)$, the period $t-1$ posterior about μ . This means that beliefs about e_{it-1} and x_{it} at the end of period $t-1$ are given by $e_{it-1} = \mu_{t-1} + \tau_{t-1}\eta_t + \sigma\nu_{it-1}$ and $x_{it} = e_{it} + \varepsilon_{it} = \mu_{t-1} + \tau_{t-1}\eta_t + \sigma\nu_{it} + \varepsilon_{it}$ with $\eta_t, \nu_{it}, \nu_{it-1}$, and ε_{it} being i.i.d. $N(0, 1)$. Given this, beliefs about μ , $\frac{e_{t-1}}{N}$ and $\frac{x_t}{N}$ are jointly normally distributed:

$$\begin{bmatrix} \mu \\ \frac{e_{t-1}}{N} \\ \frac{x_t}{N} \end{bmatrix} = N \left(\begin{bmatrix} \mu_{t-1} \\ \mu_{t-1} \\ \mu_{t-1} \end{bmatrix}, \begin{bmatrix} \tau_{t-1}^2 & \tau_{t-1}^2 & \tau_{t-1}^2 \\ \tau_{t-1}^2 & \tau_{t-1}^2 + \frac{\sigma^2}{N} & \tau_{t-1}^2 \\ \tau_{t-1}^2 & \tau_{t-1}^2 & \tau_{t-1}^2 + \frac{1+\sigma^2}{N} \end{bmatrix} \right)$$

This means that the updated posterior belief about μ conditional on observing $\frac{e_{t-1}}{N}$ and $\frac{x_t}{N}$ is normally distributed with mean μ_t and variance τ_t^2 given by (application of the projection theorem):

$$\mu_t = \mu_{t-1} + \begin{bmatrix} \tau_{t-1}^2 & \tau_{t-1}^2 \end{bmatrix} \begin{bmatrix} \tau_{t-1}^2 + \frac{\sigma^2}{N} & \tau_{t-1}^2 \\ \tau_{t-1}^2 & \tau_{t-1}^2 + \frac{1+\sigma^2}{N} \end{bmatrix}^{-1} \begin{bmatrix} \frac{e_{t-1}}{N} - \mu_{t-1} \\ \frac{x_t}{N} - \mu_{t-1} \end{bmatrix}$$

where $\Delta = N(\tau_{t-1}^2 + \frac{\sigma^2}{N})(\tau_{t-1}^2 + \frac{1+\sigma^2}{N}) - N\tau_{t-1}^4 = \tau_{t-1}^2(1 + 2\sigma^2) + \frac{\sigma^2(1+\sigma^2)}{N}$, i.e.

$$\mu_t = \left(\frac{\sigma^2(1 + \sigma^2)}{N\Delta} \right) \mu_{t-1} + \frac{(1 + \sigma^2)\tau_{t-1}^2}{\Delta} \left[\frac{e_{t-1}}{N} + \frac{\sigma^2}{\sigma^2 + 1} \frac{x_t}{N} \right]$$

and

$$\tau_t^2 = \tau_{t-1}^2 - \begin{bmatrix} \tau_{t-1}^2 & \tau_{t-1}^2 \end{bmatrix} \begin{bmatrix} \tau_{t-1}^2 + \frac{\sigma^2}{N} & \tau_{t-1}^2 \\ \tau_{t-1}^2 & \tau_{t-1}^2 + \frac{1+\sigma^2}{N} \end{bmatrix}^{-1} \begin{bmatrix} \tau_{t-1}^2 \\ \tau_{t-1}^2 \end{bmatrix}$$

i.e.

$$\tau_t^2 = \frac{\sigma^2(1 + \sigma^2)}{\Delta N} \tau_{t-1}^2$$

Therefore the updated belief about e_{it} is distributed according to $N(\mu_t, \sigma_t^2)$ with $\sigma_t^2 = \sigma^2 + \tau_t^2$. \square

Proof of Lemma 5. From the perspective of time t (and therefore knowledge of μ_t), individual emissions at time $k > t$ take the following form: $e_{ik} = \mu_t + \rho \sum_{l=t+1}^k \eta_l + \sigma \nu_{ik}$ where η_l and ν_{ik} are i.i.d. $N(0, 1)$. The claim follows directly. \square

Proof of Theorem 4. The claim follows directly from the argument in the text just before the statement. \square