Journal of Environmental Economics and Management 80 (2016) 20–36

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Emissions trading systems with cap adjustments ☆

Sascha Kollenberg [a], Luca Taschini [b,*]

[a] University of Duisburg-Essen, Universitätsstraße 12, 45141 Essen, Germany
[b] The Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

ARTICLE INFO

Article history:
Received 25 May 2015
Available online 15 September 2016

JEL classification:
H23
E63
Q52
Q58

Keywords:
EU ETS reform
Dynamic allocation
Policy design
Responsiveness
Resilience
Supply management mechanism
Risk-aversion

ABSTRACT

Emissions Trading Systems (ETSs) with fixed caps lack provisions to address systematic imbalances in the supply and demand of permits due to changes in the state of the regulated economy. We propose a mechanism which adjusts the allocation of permits based on the current bank of permits. The mechanism spans the spectrum between a pure quantity instrument and a pure price instrument. We solve the firms' emissions control problem and obtain an explicit dependency between the key policy stringency parameter —the adjustment rate—and the firms' abatement and trading strategies. We present an analytical tool for selecting the optimal adjustment rate under both risk-neutrality and risk-aversion, which provides an analytical basis for the regulator's choice of a responsive ETS policy.

Introduction

The economic slowdown following the financial crisis has substantially increased the academic interest in how economic shocks interact with climate change policies. Of particular interest has been the design of two carbon pricing instruments aimed at reducing carbon dioxide emissions: emissions trading systems (ETSs) and carbon taxes. Shocks like business cycles, technological progress or the introduction of new overlapping policies can influence the efficacy of existing instruments, as has been evidenced in the case of the European Union Emissions Trading System (EU ETS). EU ETS observers have suggested that the collapse and the continuing low level of the permits price since 2008 has been the consequence of two effects. On the one hand, the economic recession and renewables-promoting policies have led to a significant drop in permit demand; on the other, the system has been unable to respond to changes in economic circumstances (Grosjean et al., 2014 and

☆ Part of Kollenberg's research was supported by the University of Duisburg-Essen. Part of Taschini's research was supported by the Centre for Climate Change Economics and Policy, which is funded by the UK Economic and Social Research Council (ESRC). This project has received funding from the European Union's Seventh Programme for research, technological development and demonstration under grant agreement No 308481 and the German Federal Ministry of Education and Research under promotional reference No 01LA1115A. The authors would like to thank Dallas Burtraw, Corina Comendant, Harrison Fell, Carolyn Fischer, Rüdiger Kiesel, Andrei Marcu, Steve Salant, Audun Sætherø, Peter Vis, Peter Zapfel, and other participants of the Workshop on EU ETS Market Stability Reserve policies held at DIW and CEPS. The usual disclaimers apply.
* Corresponding author.
E-mail addresses: sascha.kollenberg@uni-due.de (S. Kollenberg), l.taschini1@lse.ac.uk (L. Taschini).

http://dx.doi.org/10.1016/j.jeem.2016.09.003
0095-0696/

Ellerman et al., 2015a).

To address the allowance supply imbalance in the EU ETS, at least temporarily, EU regulators have proposed two main plans. Under the first one, known as 'back-loading', EU regulators have reduced the number of permits in the market via near-term auctions, reintroducing the quantity removed later on. Under the second one, a Market Stability Reserve (MSR) system will be implemented in order to make the ETS "more resilient to supply–demand imbalances so as to enable the ETS to function in an orderly fashion" (European Parliament and Council, 2015). The MSR can be viewed as an extended 'back-loading'; it will shift permit allocation into the future but within the bounds of the pre-determined cap. In contrast to 'back-loading', however, the MSR will adjust auction permits *in response* to changes in the inventories of unused permits. Notwithstanding, because the cap stringency will remain unaltered, by design the MSR will not be able to make the ETS fully responsive to external shocks.

It has long been known that responsive policy instruments have the potential to improve upon fixed policies. In his seminal paper, Weitzman (1974) recognizes that an ideal instrument of central control would be a contingency message that gives instructions dependent on the state of the world that is revealed. However, because implementing such contingent instruments can be complicated in practice, he concentrates his analysis on single-order policies, so-called pure instruments.

More recently, new models have made substantial contributions to our understanding of climate policy responsiveness, suggesting various proposals for improving the design of carbon pricing instruments. Doda (2016) provides a comprehensive review of the literature on how to build responsiveness into such design. Some of these improved instruments, for instance, presume intensity targets or indexed regulation in order to allow the regulator to condition policy stringency on observable economic indicators (Jotzo and Pezzey, 2007; Ellerman and Wing, 2003 and Newell and Pizer, 2008, among others, study indexing rules; Heutel, 2012 and Golosov et al., 2014 study climate policy cyclicality). Carbon pricing instruments can also be made responsive to economic shocks by mixing elements of a carbon tax into an ETS, a possibility first recognized by Roberts and Spence (1976).

Such instruments, known as hybrid, work by imposing a price ceiling and/or floor to an ETS and thus adjust the stringency of the cap in response to changes in price levels: permits are added to the market when the price crosses the ceiling and are removed when the price crosses the floor. Adapting Robert and Spence's original insight to pollution control, several recent studies find that hybrid instruments can improve welfare relative to a pure quantity or price instrument (e.g. Unold and Requate, 2001; Pizer, 2002; Hepburn, 2006; Fell and Morgenstern, 2010; Grüll and Taschini, 2011, and Fell et al., 2012). In our paper, we focus on quantity-based hybrid mechanisms and show that their performance is superior to fixed-cap instruments.

Our work ties together the literature on responsive policy instruments by deriving a responsive policy stringency mechanism that spans the spectrum between a standard emissions trading system with a fixed cap—pure quantity instruments[1]—and a carbon tax with a fixed price—pure price instrument. Because our work focuses on quantity-based mechanisms, the spectrum reflects the transition between a fixed cap and a fully floating cap (i.e. the cap is responsive in the latter case). At the pure-price extreme, the cap is *fully* floating because the current allocation of permits is perfectly adjusted in response to shocks. As the policy stringency moves from the pure-quantity to the pure-price extreme, it goes through the following generalised transformations: pure fixed cap (i.e. pure-quantity extreme) > partially-responsive fixed cap (e.g. EU ETS MSR) > floating cap (the current allocation of permits is adjusted to compensate shocks, but not fully) > fully floating cap (i.e. pure-price extreme). Any point on the policy stringency spectrum except the extremes describes a version of a hybrid policy.

For a floating-cap mechanism, the degree of policy stringency is determined by the mechanism's permit–allocation adjustment rate. The rate controls the quantity of available permits (i.e. the bank) within a confidence interval, by adjusting the current allocation of permits. In our model, therefore, the adjustment rate serves as a lever to control system responsiveness within the setting of a hybrid mechanism indexed to a quantity indicator (i.e. the bank).

The analysis abstracts away from the comparison of individual pure instruments—price vs. quantity—and focuses instead on a responsive ETS policy. It is based on the understanding that a constant ETS cap may be too lenient in economic recessions, and too stringent in economic expansions. We show that the proposed responsive mechanism can improve market efficiency. More importantly, we also implement the mechanism and answer the question: how should the government pick the optimal adjustment rate?

Central to our model is the premise that compliance cost minimization in an ETS is achieved by exploiting inter-temporal differences in abatement costs via banking and borrowing. With this in mind, we model the responsiveness mechanism such that it adjusts the stringency of the cap in response to changes in the aggregate bank of permits. We also analytically describe how firms respond to changes in expectations about the future cap and permit net demand. We show that the optimal policy stringency can be achieved via an adjustment rate, which we quantify. The adjustment rate is a key parameter of the policy stringency rules and is indexed to the aggregate bank.

The proposed mechanism allows us to identify a trade-off between the two extremes of the policy stringency spectrum, i.e. a pure fixed cap (a pure quantity instrument) and a fully floating cap (a pure price instrument).[2] The trade-off arises from the fact that in an ETS with banking and borrowing provisions, permits are inter-temporally tradable, which provides firms

[1] Note that although the MSR maintains a fixed cap, we do not consider it as the quantity extreme of the spectrum, because the MSR does provide for some flexibility, albeit with only temporal effects. In other words, the MSR is a (partially) responsive system with a fixed policy stringency, and lies in the proximity of the quantity extreme. To avoid ambiguity, we refer to this extreme as a *standard* fixed cap.

[2] We describe later on why a fully floating cap corresponds to a pure price instrument. As a summary, when market shocks are perfectly offset by changes in the cap, expectations about future required abatement do not change. In other words, firms remain on their pre-shock abatement paths, which could not occur if the equilibrium permit price were to change—thus, the price is fixed.

with improved cost minimization opportunities (Pizer and Prest, 2016). When the cap stringency is relaxed (the policy stringency moves away from the quantity extreme), the cost arbitrage opportunities are reduced. As policy stringency nears the price extreme, inter-temporal trading thins out, and firms are no longer able to benefit from differences in marginal costs across time.

Finally, our model is sufficiently general to allow us to qualitatively describe the impact of the EU ETS MSR on market behavior. As mentioned previously, because the EU ETS MSR maintains a fixed cap, it shifts permit allocation only temporarily and thus does not alter the system's stringency. We show that a mechanism allowing cap stringency to be adjusted can reduce overall compliance costs.

The remainder of the paper is organized as follows. In the next two sections we introduce our main assumptions and define the key decision-making variables; this includes a qualitative description of how the EU ETS MSR impacts abatement decisions. In the subsequent section we present the decision problem and discuss the impact on market behavior of policy instruments with fixed and responsive stringencies. In 'The responsive policy stringency mechanism', we propose a responsive mechanism that covers the spectrum between standard fixed-cap and fully floating cap instruments. In the subsequent section, we show how to select an optimal adjustment rate. In 'The effect of the adjustment rate on the risk premium', we relax the assumption of risk neutrality and explore the mechanism´s impact on the risk premium corresponding to investments in abatement. Finally we have the Conclusions.

The model

Firms face an inter-temporal optimization problem where, at each point in time, they have to decide how much they want to offset their emissions (by abating and by trading permits), considering the current and future costs of reducing emissions. Central to each firm's decision-making is its current bank of permits and the number of permits it expects to receive in the future. Given this, the required abatement—the difference between future permit allocation and future emissions before abatement (future counterfactual emissions)—is the key quantity that every firm must assess at each point in time up to the end of the regulated period. During this period, changes in firms' expectations about their required abatement affect how much abatement and banking occur. If the cap is fixed and if the permit allocations schedule remains unaltered, shocks to counterfactual emissions are equally transferred to the required abatement. If, on the other hand, a responsive mechanism is in place, the current allocation of permits is adjusted, accounting for changes to expected emissions. In the extreme case, adjustments of the allocation could perfectly offset demand shocks. That is, a shock to the demand-side is completely offset by an adjustment of the supply-side. Between these two extremes, no response and full response, there is a spectrum of quantity-based policies, each characterized by its cap stringency. We model a responsive mechanism, defined by rules for allocation adjustment, that describes the policy spectrum and show how the regulator can identify the level of the policy stringency that minimizes expected compliance costs.

Allowance supply and demand

Let t denote the current time, where $0 \leq t \leq T$ and T is the end of the regulated period. Firms are assumed to be atomistic in a perfectly competitive market for emission allowances: let firms $i \in I$ be continuously distributed in a set I, under a probability measure m. Each firm is characterized by her initial bank of allowances B_0^i, her cumulated (counterfactual) emissions process $E^i(0, t)$, and her cumulated permit allocation process $A^i(0, t)$. Because $A^i(0, t)$ spans the regulated period up to time t, it incorporates the post-shock cap adjustment if the policy stringency is responsive. Although we adjust the cap along the policy spectrum, we assume that the long-term cap is binding, i.e. period-wide counterfactual emissions are higher than the period-wide pre-adjustment cap. We will see later that the resulting uncertainty in the required abatement is fundamental to the regulated firms' inter-temporal optimization.

We now analytically describe the main expressions of the firms' individual and aggregate permit positions. We first consider the bank of permits, which we use in the proposed mechanism as a quantity indicator for permit allocation adjustments. The bank of allowances held by an individual firm at time t is:

$$B_t^i = B_0^i + A^i(0, t) - E^i(0, t) + \int_0^t \alpha_s^i \, ds - \int_0^t \beta_s^i \, ds,$$

where α_t^i denotes instantaneous abatement and $|\beta_t^i|$ is the number of permits sold ($\beta_t^i > 0$) or bought ($\beta_t^i < 0$). By time T, the sum of permits received (the initial bank plus permits allocated from time 0 to T) should equal the sum of permits demanded by a firm over the entire period (permit demand here is period-wide emissions less net permit purchase). In other words, by the end of the regulated horizon firms are in compliance and the bank is completely deployed, i.e. $B_T = 0$.[3] We assume that compliance is only required at time T and allow for unrestricted banking and borrowing during the regulated

[3] Accordingly, after T, permits have no (compliance) value. Later we consider the case where the final bank is non-zero to illustrate some important concepts in the context of responsive mechanisms.

period.[4] This condition allows us to transparently analyze the policy stringency spectrum from a standard fixed-cap to a fully floating cap. When borrowing constraints are binding, i.e. when the bank equals zero, the firms' inter-temporal optimization breaks down. The analysis of the effect of different policy stringencies would only be relevant before this breakdown. In a companion paper Kollenberg and Taschini (2016) , we show how to transfer our model from a setup without borrowing limitations but with a fixed time horizon T to a no-borrowing framework over an infinite time horizon, where the inter-temporal breakdown is endogenous. In the context of this paper, however, we choose a setup with a fixed time horizon and unlimited borrowing in order to make our results analytically more transparent.

In our model, we can observe the impact of the responsive mechanism (which alters the level of policy stringency) on a key state variable of the system: the time-t expectation of the firm's required abatement. We define required abatement as future counterfactual emissions over $[t, T]$ minus the total number of permits to be received over the same period (which, importantly, incorporates cap adjustments), net of the existing bank of permits. We can express the time-t expectation of the required abatement as

$$R_t^i := \mathbb{E}_t\Big[E^i(t, T) - A^i(t, T)\Big] - B_t^i.$$

The expected required abatement consists of two components. The first component corresponds to the firm's expectation about its future non-covered emissions $E^i(t, T) - A^i(t, T)$. Note that because $A^i(t, T)$ incorporates future permit allocation adjustments, a responsive mechanism may alter the required abatement, ultimately changing the inter-temporal problem that firms face. The second component is the existing bank of permits. Firms may decide to use the bank to cover future permits or to sell the bank and use the proceeds to finance abatement. The firm's expected required abatement is thus the residual permit demand (positive or negative) that it expects to have before it takes any abatement measures or trades any permits at time t. Each firm uses its required abatement to decide on its optimal abatement and trading strategies. By definition, and recalling that firms have to be in compliance at time T, the final residual permit demand is zero, i.e. $R_T^i = B_T^i = 0$.

The expression for a firm's expected required abatement is key to understanding how firms react to changes in policy stringency or to newly available information on realized emissions. In fact, the expression allows us to readily describe the effect of various events that are of particular relevance to the debate on the EU ETS structural reform, and specifically, the market impact of a mechanism that adjusts permit allocation within the bounds of a given cap—the EU ETS MSR.

We start by noting that under unlimited banking and borrowing, firms adjust period-by-period abatement and trading in order to spread the shock effect over time. The expected required abatement changes, which reflects the firms' period-by-period adjustments. Under the conditions of unrestricted banking and borrowing, therefore, a mechanism that only changes the timing of permit allocation, but retains the pre-shock policy stringency, has no effect on the system. Firms will tap into the market's banking and borrowing potential with no constraints, neutralizing the effects of a partially responsive fixed-cap mechanism. This means that the EU ETS MSR—a partially responsive fixed-cap mechanism—can affect the EU ETS only if borrowing constrains are binding. If this is the case, a shift in the permit allocation program over time can generate temporary permit scarcity, resulting in a short-term permit price increase. Subsequently, when permits are reintroduced into the market, prices decrease (Salant 2016 and Perino and Willner, 2015). If the mechanism allows policy stringency to partially offset market shocks, changes in the expected required abatement are smaller and, consequently, the firms' period-by-period adjustments are smaller too. Although the firms' abatement strategies change as a result of the shocks, they do not change as much as they would if the cap were kept fixed. If, on the other hand, the mechanism allows policy stringency to perfectly offset market shocks, $\mathbb{E}_t\Big[A^i(0, T) - E^i(0, T)\Big]$ is a constant and the firms' expectations about their future required abatement remain the same. In other words, a fully responsive policy with a fully floating cap allows firms to retain their existing abatement strategies when shocks occur, whereas a partially responsive policy with a pure fixed cap does not.

The inter-temporal decision problem

The firm's dynamic cost minimization problem is

$$\min_{\alpha^i, \beta^i} \quad \mathbb{E}\left[\int_0^T e^{-rt}\left(\Pi^t \alpha_t^i + \varrho(\alpha_t^i)^2 - P_t\beta_t^i + \nu(\beta_t^i)^2\right)dt\right],$$
$$\text{s. t.} \quad B_T^i = 0. \tag{1}$$

where r is the risk-free interest rate and we assume the usual quadratic functional form for the abatement cost curve with

[4] Under the current EU ETS Directive, borrowing (to some extent) is implicitly possible within a trading phase due to next-year free allocation preceding the submission of current-year emissions. The shift from free allocations to auctioning limits this opportunity and requires the introduction of borrowing constraints.

Π_t and ϱ representing the intercept and the slope of the marginal cost curve, respectively.[5] Firms can sell and buy allowances $|\beta_t^i|$ at a price P_t. In addition to the positive (negative) cost βP_t when buying (when selling) $|\beta|$ permits, we assume that firms face transaction costs per trade. Among others, Frino et al. (2010) and Medina et al. (2014) document non-negligible transaction costs in the EU ETS. In our framework, we assume linear marginal trading costs of $P_t - 2\nu\beta$. This expression ensures uniqueness of the equilibrium and allows us to derive the equilibrium in closed form. The introduction of transaction costs does not limit the generality of our results. On the contrary, they are valid for any assumption about market transaction costs, including negligible transaction costs.

In general, the difference between the value of ν and the value of ϱ can be interpreted as a firm's propensity for abatement relative to trading. Both quantities parametrize the firms' costs associated with adjusting their strategies in response to shocks on required abatement. For our analysis of a quantity-based mechanism, the relative cost difference between trading and abatement is irrelevant; for this reason, we can focus on the effect of their sum $\nu + \varrho$, which has a quantifiable impact on aggregate compliance costs.

In Appendix A we solve the optimization problem in (1) and obtain the market equilibrium as a set $\{\alpha_t^i, \beta_t^i, P_t;\ t \in [0, T],\ i \in I\}$, where the price process P_t satisfies the market-clearing condition $\int_I \beta_t^i dm(i) = 0$ for all t.[6] In equilibrium, the abatement and trading strategies are:

$$\alpha_t^i = \frac{P_t - \Pi_t}{2(\nu + \varrho)} + \frac{\nu r\, R_t^i}{(e^{r(T-t)} - 1)(\nu + \varrho)} \quad \text{and} \quad \beta_t^i = \alpha_t^i - \frac{r\, R_t^i}{e^{r(T-t)} - 1},$$

and the price process is

$$P_t = \Pi_t + R_0 \frac{2\varrho r e^{rt}}{e^{rT} - 1} + 2\varrho r e^{rt} \int_0^t \frac{d\xi_s}{e^{rT} - e^{rs}},$$

where

$$d\xi_s = d\mathbb{E}_s\big[E(0, T) - A(0, T)\big].$$

The process ξ reflects the changes in the firms' expectations about their required abatement as a result of shocks to counterfactual emissions $E(t, T)$. These changes may be partially or fully offset by changes to the total number of permits $A(t, T)$. Suppose that the regulator implements a responsive mechanism that allows it to control the stringency of the cap. In other words, the regulator can choose the degree by which the total number of permits is adjusted to meet changes in expected total emissions. The degree by which the cap stringency changes is determined by an adjustment rate δ, which quantifies the extent to which emissions shocks are offset by allocation adjustments. When the policy stringency is at maximum (pure fixed cap), $\delta=0$, regulated firms adjust their period-by-period abatement in order to unrestrictedly spread the effect of emission shocks over time. By increasing δ, the policy stringency is relaxed and emission shocks are increasingly being absorbed by changes in the cap. At the other extreme, when policy stringency is at its minimum, emission shocks are fully offset by changes to the total number of permits. Note that via ξ_s, emission shocks and cap adjustments are incorporated into the firms' inter-temporal optimization problem. In this way, our model allows us to capture the market's reaction to changes in policy stringency. In order to study this market reaction, we first make some general observations on how aggregate abatement is affected by the adjustment rate δ. Later we propose a responsive mechanism and show how the regulator can choose an optimal adjustment rate to minimize expected total compliance costs.

For a given adjustment level δ, the aggregate abatement is

$$\alpha_t = r e^{rt} \frac{R_0(\delta)}{e^{rT} - 1} + r e^{rt} \int_0^t \frac{d\xi_s(\delta)}{e^{rT} - e^{rs}}. \tag{2}$$

The process ξ reflects the changes in a firms' expectations about their required abatement and determines how firms adjust their abatement strategies following a change in policy stringency. Abatement is now a function of the adjustment rate. The first term on the right-hand side of Eq. (2) is the expected required abatement given the information available at time 0. The term $R_0(\delta)$ represents the total permit demand before firms make any abatement decisions. Abatement is then spread over the remaining timeframe to minimize discounted compliance costs.

At each time t, new information about changes in future required abatement becomes available and adjustments to the equilibrium abatement may occur. This is represented by the second term on the right-hand side of Eq. (2). When the time s expectation about the expected required abatement changes by $d\xi_s$, the necessary abatement is spread over the remainder of the regulated period. For example, a negative shock to expected emissions decreases permit demand and decreases the firms' expected required abatement.

In reviewing their strategies, firms decide how to respond to these changes in their expectations of future total emissions

[5] The intercept Π_t is assumed to increase at the risk-free rate r. A detailed discussion of the calibration of marginal abatement costs can be found in Landis (2015).

[6] The equilibrium price process P_t does not allow for an individual deviation from the equilibrium strategies to safe costs. Hence, the notion of the equilibrium employed here is that of a Nash equilibrium. We refer to the appendix for a derivation of the equilibrium.

and total permit allocation. For a given change in expectations, each firm adjusts its abatement and trading strategies such that it remains on its cost-minimizing path. If the cap adjustment perfectly compensates the shock on expectations of future emissions, there is no change in expected required abatement. In this case, the mechanism is perfectly responsive: the term $d\xi_s$ is zero, and abatement increases at the risk-free rate r. In equilibrium, the permit price equals the marginal abatement cost, and therefore the price also increases at rate r. Hence, permit prices (in present terms) remain constant. In other words, the mechanism lies in the pure price extreme of the spectrum. On the one end of the spectrum, where the cap is fully floating (pure price instrument), the permit price is thus constant in present terms. By contrast, at the fixed-cap end of the spectrum, $A(t, T)$ remains unchanged throughout the regulated period and the permit price varies in response to emissions shocks.

We now propose a responsive mechanism that spans the spectrum of policy stringency and show how the regulator can identify an optimal adjustment rate. In particular, we find that there is a trade-off between the level of responsiveness which lowers the firms' costs of adjusting to shocks in emissions, and the potential for firms to benefit from inter-temporal cost-saving opportunities.

The responsive policy stringency mechanism

Our proposed responsive mechanism is indexed to the aggregate bank of permits such that at each time $t < T$, if the current aggregate bank is above a level $c \geq 0$, a fraction δdt of the difference $|B_t - c|$ is permanently removed from the scheduled permit allocation. Conversely, $\delta \cdot |B_t - c|\, dt$ permits are permanently added to scheduled allocations if the aggregate bank is lower than c. In a discretized setting, dt would, for example, correspond to a year and δ would be the percentage of $|B_t - c|$ that is added to or removed from next year's allocation. In a continuous-time setting, δ is the adjustment rate of the allocation of permits.

Note that when $\delta \cdot |B_t - c|\, dt$ permits are removed from or added to the scheduled permit allocation, the cap is instantaneously decreased or increased. However, future allocation adjustments may again change the cap. Thus, over the entire length of the regulated period, cumulated changes to the cap may cancel out, and the time–T realized cap may equal the cap before the implementation of the mechanism. The expected change to period-wide allocations $A(0, T)$ (the realized cap at time T) may differ from $\delta \cdot |B_t - c|\, dt$ and is subject to the equilibrium dynamics. We explore this interdependency in Appendix A in order to solve for the equilibrium under the adjustment mechanism. This step of the analysis is fundamental to the quantification of aggregate compliance costs under the responsive mechanism, which we discuss in the next section.

To provide the intuition behind the adjustment rate, consider for the moment the case where $c > 0$. For illustration purposes, we also require firms, in aggregate terms, to hold a positive number of c permits at time T and the expected abatement requirement becomes

$$R_t = \mathbb{E}_t[E(t, T) - A(t, T)] - B_t + c.$$

Later on, we apply the mechanism using the more natural assumption of $c=0$.

We start with the case of an extremely high adjustment rate δ, i.e. close to 100% per unit of time. This rate corresponds to an (almost) fully floating cap.[7] Any deviation of the bank from the target bank c is continuously, and almost perfectly, offset by adjustments to the cap. The bank is kept in a very tight band around c and emission shocks are almost fully offset by changes to the allocation of permits.

Consider next the case of a low adjustment rate δ. This corresponds to a floating cap, where permit allocation is adjusted to partially compensate the shocks. The bank fluctuates around the target level c. The lower the adjustment rate, the larger the fluctuations. It is possible that, instead of selecting an absolute level for the target bank, the regulator may prefer controlling the bank within an interval. The regulator's objective would be to adjust the permit allocation such that the desired interval is respected throughout the regulated period with some level of confidence. Within the setting of our mechanism, this could be achieved using the adjustment rate δ. In this case, the level $c > 0$ does not necessarily represent a specific desired bank level, but simply an appropriate level within the desired interval such that the interval bounds are maintained.[8] We show that this interval can be represented as a confidence interval for the bank, given a confidence level. In the following we provide an analytically tractable relation between the adjustment rate δ and the confidence intervals of B_t. This will serve to illustrate the considerations described above. Additionally, we use the concepts introduced here to discuss the trade-off along the policy stringency spectrum discussed in the next section.

To illustrate the concept of a bank confidence interval, consider a pre-adjustment allocation schedule f_t. In the presence of the responsive mechanism, permit allocation during each time interval $[t, t + dt]$ is $f_t dt + \delta(c - B_t)dt$. Accordingly, the change in the permits bank can be expressed as

$$dB_t = f_t dt + \delta(c - B_t)\, dt - E(t, t + dt) + \alpha_t dt,$$

where $E(t, t + dt)$ denotes the cumulated emissions between t and $t + dt$, which may be subject to uncertainty.

[7] An adjustment rate equal to 100% corresponds to a fully floating cap, the pure price extreme, where the bank is maintained at exactly c.

[8] In fact, the regulator may choose to cease allocation adjustment as long as the bank is within the desired interval; equivalently, the regulator could formulate his mechanism using the interval bounds as threshold levels for the aggregate bank.

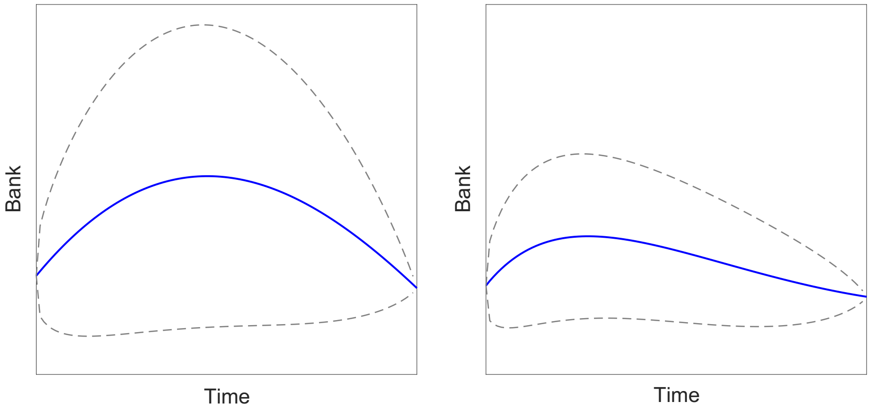

Fig. 1. Aggregate bank when the responsive mechanism is deactivated, $\delta = 0$ (left diagram) and when the adjustment rate is $\delta = 5\%$ per year (right diagram). The solid lines show the expected banking curves obtained for an initial bank of $B_0 = 2$ Bn permits, expected yearly emissions of 2 Bn tonnes and a constant allocation of 250 Mn tonnes per year; c is set to 500 Mn permits. The dashed lines show the 5% and 95% quantiles obtained for a standard deviation in yearly emissions of 200 Mn tonnes. Abatement costs are parametrized by $\Pi_0 = 5$ Euros/tonne and $\varrho = 0.25$ Euros per squared tonne. The risk-free interest rate is $r = 3\%$.

In Appendix A.1 we explicitly derive the distribution of the bank B_t, depending on the adjustment rate, when cumulative emissions $E(t, t + dt)$ are normally distributed with a given mean and a given volatility. The distribution of the bank allows us to compute the probability that the time-t bank stays within a given interval, or, equivalently, to choose an adjustment rate δ such that a selected interval can be maintained with a given probability.

Fig. 1 shows the relation between the adjustment rate and the aggregate bank and Appendix A details the expressions for the bank distribution.

The figure shows the aggregate bank quantiles for a 90% confidence level when the responsive mechanism is inactive (left diagram) and when it is active with a positive adjustment rate (right diagram). In the latter case, market changes force the mechanism to adjust the cap up or down so that the resulting aggregate bank is contained within a tighter band. The right diagram suggests that this is possible depending on the selected adjustment rate. By contrast, the left diagram shows that when the adjustment rate is zero, firms adjust their period-by-period abatement in order to spread the effect of emission shocks, ultimately leading to a higher dispersion of the bank.

The optimal adjustment rate

We now express total aggregate compliance costs as a function of the adjustment rate δ and analytically derive the optimal adjustment rate that minimizes total costs. Because carbon dioxide—the focus of our study—is a stock pollutant, the marginal benefit curve is flat relative to the marginal abatement curve. As such, the problem of minimizing expected costs is the same as that of maximizing expected benefits minus costs; total costs and total net benefits differ by a constant (Newell and Pizer, 2008). Hence, we can abstract from the quantification of the damage caused (or avoided) by the adjustment of the cap and consider only total compliance costs.

For the sake of expositional clarity, we assume here that firms are identical. In particular, all firms have the same initial bank and the same emissions process.[9] This assumption does not qualitatively change our results concerning total aggregate costs under different adjustment rates δ. We therefore consider the regulator's problem of choosing an adjustment rate δ that minimizes aggregate compliance costs as follows:

$$\min_{\delta} \mathbb{E}\left[\int_0^T e^{-rt}(\Pi_t\alpha_t(\delta) + \varrho\alpha_t^2(\delta))\,dt\right], \tag{3}$$

where $\alpha(\delta)$ denotes the aggregate abatement quantity, which depends on the adjustment rate δ.

In Appendix A.2 we explicitly derive aggregate costs as

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\left(\Pi_t\alpha_t + \varrho\alpha_t^2\right)dt\right] = \Pi_0 R_0 + \varrho r\frac{R_0^2}{e^{rT} - 1} + \varrho r\int_0^T \frac{d\langle\xi\rangle_t}{\left(e^{rT} - e^{rt}\right)}. \tag{4}$$

This equation allows us to decompose the trade-off between the level of policy responsiveness (which lowers the costs of adjusting to changes in expectations of required abatement due to shocks in demand), and the potential for firms to benefit

[9] We emphasize that the model is general enough to allow the impact on aggregate costs to be analyzed based on the distribution of the firms' characteristics (initial bank, emissions distributions). However, this affects the results only nominally and is therefore ignored for simplicity.

from inter-temporal cost-saving opportunities. The trade-off arises from the fact that in an emissions trading system, permits are inter-temporally tradeable. The opportunity to save (or borrow) permits for (from) the next trading period implies an arbitrage condition between current and future permit prices, within the constraints of a firm's cost minimization (Newell and Pizer, 2008).

Analytically, the trade-off arises from the interaction of the three terms on the right hand-side of Eq. (4). The first term $\Pi_0 R_0$, represents the total permit demand before firms make any abatement decisions, where $R_0 = \mathbb{E}_0\left[E(0,T) - A(0,T)\right] - B_0 + c$. By increasing δ, the bank is constrained around c: permits are removed (added) from (to) the bank, the expected abatement requirement R_0 increases (decreases), and the abatement costs—parametrized by the intercept Π_0 of marginal abatement costs—increases (decreases). In particular, if the initial bank is higher than c (e.g. $B_0 > 0$ and $c=0$), the mechanism is expected to decrease the time-T realized cap. Thus, the larger the adjustment rate δ, the larger the total permit demand (pre-abatement decisions) and the larger the R_0. Notice that the first component of the aggregate costs does not depend on the coefficient ϱ, which is the slope of the marginal abatement cost function. In this sense, the first component of the aggregate costs represents the costs that are imposed simply by the expected adjustment of the cap.

The second term $\varrho r (R_0)^2/(e^{rT} - 1)$ quantifies the cost of bank 'tightness' around the level c, as described in the previous section. The higher the adjustment rate, the tighter the interval is around the bank and the lower the benefit is from inter-temporal trading. This component emerges from the fact that, under the usual assumption that marginal abatement costs are increasing in emissions reduction, firms first accumulate permits and then draw them down (Rubin, 1996 and Schennach, 2000). This is observable in Fig. 1. The steeper the marginal abatement cost (i.e. the larger the ϱ), the greater the incentive to spread abatement costs over time. This is possible via banking. When permit allocation is lowered, the incentive to bank in the short term is reduced because a higher bank would lead to a proportionally higher reduction of the cap. Abatement is thus postponed—the abatement curve is skewed, which leads to higher overall costs due to the firms' convex cost structure.

The third term $\varrho r \int_0^T \frac{d\langle \xi \rangle_t}{\left(e^{rT} - e^{rt}\right)}$ captures the cost saving due to lower frequency of abatement and trading. The higher the adjustment rate, the more responsive the system and the larger the cap adjustment in response to market shocks. When changes in the permit demand are offset by changes in permit supply (allocation), the variability of the expected required abatement is reduced. Firms can adjust their abatement and trading strategies less frequently in order to stay on their cost-minimizing paths.

We graphically illustrate the trade-off between the effects of these components in Fig. 2. We plot the aggregate compliance costs for yearly adjustment rates between zero and 100% for a set of selected parameters of the problem.

When the yearly adjustment rate is zero, the system's original cap is unaffected—this represents the fixed cap extreme of the spectrum. Regulated firms adjust their behavior, spreading the effect of market shocks over time and within the bounds of the pre-determined cap. Permit prices vary reflecting period-by-period abatement adjustments. When the adjustment rate is high, the system's cap stringency is reduced; the regulator adjusts the allocation in response to shocks, dynamically changing the cap. This represents the fully floating cap extreme of the spectrum (effectively, a pure carbon tax).

Between the extremes, we find that when the yearly adjustment rate is gradually increased, the variability of supply and demand is reduced and firms face lower costs of adjusting to shocks. Once the adjustment rate overtakes a certain value, however, the opportunity costs of foregone inter-temporal trading override the benefits associated with cap adjustments. In this case, the firms' loss of benefits from exploiting differences in marginal abatement costs across time exceed the firms' cost savings caused by the shock-mitigating effect of a responsive policy.

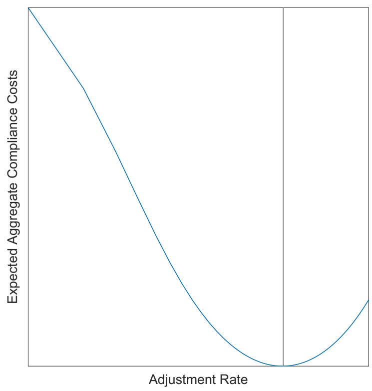

Fig. 2. Expected aggregate compliance costs as a function of the yearly adjustment rate, obtained for an initial bank of $B_0 = 2$ Bn permits, expected yearly emissions of 2 Bn tonnes and a constant allocation of 250 Mn tonnes per year; and $c=0$. Abatement costs are parametrized by $\Pi_0 = 5$ Euros/tonne and $\varrho = 0.25$ Euros per squared tonne. The risk-free interest rate is $r=3\%$. The figure shows the curve on a logarithmic scale on both axes.

The effect of the adjustment rate on the risk premium

In the previous section we showed that by choosing an adjustment rate δ, the regulator can decide on the balance between a low policy responsiveness where the burden of adjustment to shocks is mostly borne by firms (high cap stringency), and a high policy responsiveness where the burden of adjustment is shared between the firms and the regulator (low cap stringency). In responding to shocks, the regulator can curb the impact of emission shocks to the system, ultimately limiting price variability.

When emissions are uncertain, firms cannot perfectly predict the number of permits they will require in the future. Unforeseen changes to future counterfactual emissions will impact the required abatement and ultimately the equilibrium price, thus increasing the risk of abatement investments.

The extent to which shocks to counterfactual emissions are reflected in the required abatement depends on the adjustment rate δ, which determines how permit allocation responds to changes in expected emissions. The adjustment rate δ and, accordingly, the level of cap stringency, has a profound impact on the perceived riskiness of investments in abatement or permits. Depending on their risk appetite, firms would demand a positive risk-premium q_t on top of the risk-free rate r.

To calculate the present value of their investments, firms discount their future cash flows from trading permits by a discount rate that includes the risk-premium q_t. If alternative investments promise higher returns (discounted by their corresponding risk-adjusted rates), firms prefer to postpone abatement and use their current bank of permits to offset emissions. In turn, lower abatement levels result in lower permit prices. Intuitively, a larger discount rate due to a positive risk premium should thus imply a lower level of aggregate abatement and, consequently, a lower aggregate bank compared to the case of a null risk premium. This has been observed by Fell (2015) and in their sensitivity analyses.

To illustrate the impact of the responsive mechanism on investment risk premia, we first consider the case of a fully flexible cap with a yearly adjustment rate of 100%. Shocks to emissions are perfectly compensated by the cap adjustment. Accordingly, firms stay on their pre-adjustment abatement paths, and discounted permit prices remain constant. Because the future required abatement—and therefore also the abatement investments—are certain, the rate of return from trading permits equals the risk-free rate r. By contrast, when the adjustment rate is lowered, the uncertainty (variability) about the future required abatement increases and, consequently, permit prices become volatile.

It is natural to expect that the risk-premium q_t is a monotonically increasing function of the volatility of permit prices. Because the variability of the future required abatement depends on the adjustment rate, we can express the risk-premium q_t as a function of δ. To study the impact of the adjustment rate on risk-premia, we first consider the permit price return when firms are risk-neutral:

$$\frac{dP_t}{P_t} = rdt + \frac{2\varrho V_t \kappa_t}{P_t} dW_t$$

The price return (per time-unit) equals the risk-free rate r plus a stochastic component determined by the uncertainty around the future required abatement. When firms are risk-neutral, the expected rate of return is $\mathbb{E}[dP_t/P_t] = rdt$. When firms are risk-averse, the expected rate of return includes the risk-premium q_t and equals $\mathbb{E}[dP_t/P_t] = (r + q_t)dt$.

We then obtain the expression for the permit price return under risk-aversion as follows:

$$\frac{dP_t}{P_t} = (r + q_t)dt + \frac{2\varrho V_t(\delta, r)\kappa_t}{P_t} dW_t. \tag{5}$$

For simplicity, let us assume that q_t increases linearly with the volatility of permit prices, i.e. q_t is proportional to the volatility coefficient in (5) as follows:

$$q_t = k \cdot \frac{2\varrho V_t(\delta, r)\kappa_t}{P_t}, \tag{6}$$

where the constant k represents the (overall) level of the firms' risk-aversion. Note that in equilibrium we obtain

$$V_t(\delta, r) = \frac{\delta + r}{e^{(\delta+r)(T-t)} - 1}.$$

The variable k can be used to calibrate the model to historical prices. However, in the context of analyzing the policy spectrum, we use it to normalize q_t to values between 0 and 1.

Eq. (6) models the impact of the adjustment rate δ on the risk-premium. When the adjustment rate is relatively high and the policy is very responsive, $V_t(\delta, r)$ is very small (the denominator increases exponentially in δ, while the numerator increases linearly). Thus, the risk-premium in Eq. (6) is very small and the permit price increases by a value equal or just above the risk-free rate r. A decreasing adjustment rate increases $V_t(\delta, r)$, and therefore also the risk premium.

Fig. 3 shows the relation between the adjustment rate δ, the risk premium q_t and the aggregate compliance costs. Appendix B provides the detailed solution to the problem under risk-aversion.[10]

We simulate 500 random paths of counterfactual emissions to illustrate the impact of δ on risk-premia and costs; each dot represents the outcome of one simulation. For illustration purposes, we choose seven different adjustment rates

[10] In Appendix B we show that the model under risk-neutrality can be transferred to the case of risk-aversion since the necessary change of measure is viable under Eq. (6). In particular, we are able to calculate the risk-premia and expected aggregate compliance costs for different choices of δ.

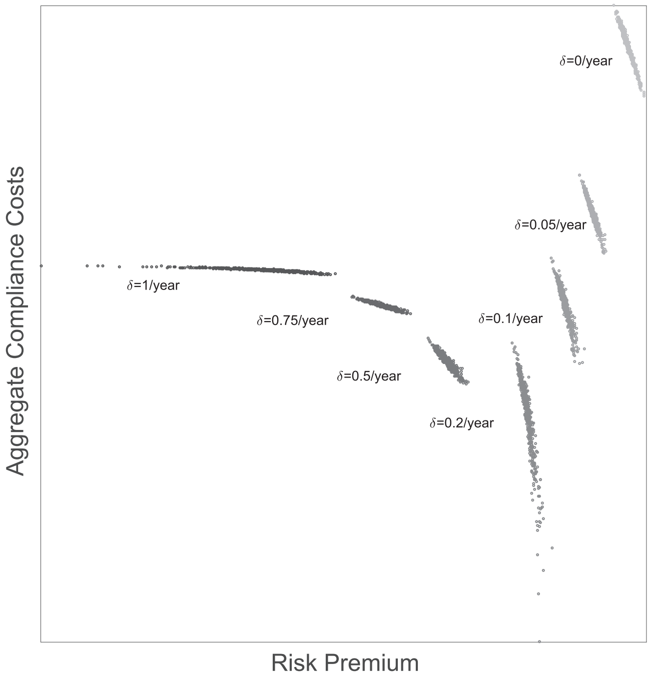

Fig. 3. Realized aggregate compliance costs and risk-premia for 500 simulated emissions paths under different adjustment rates; obtained for an initial bank of $B_0 = 2$ Bn permits, expected yearly emissions of 2 Bn tonnes and a constant allocation of 250 Mn tonnes per year; the standard deviation in yearly emissions is 200 Mn tonnes and $c=0$. Abatement costs are parametrized by $\Pi_0 = 5$ Euros/tonne and $\varrho=0.25$ Euros per squared tonne. The risk-free interest rate is $r=3\%$. The figure shows a scatterplot on a logarithmic scale on both axes.

representing different points on the policy spectrum. The adjustment rate increases from sparse clouds for $\delta = 0$ per year (in light grey) to concentrated clouds for $\delta = 1$ per year. Both aggregate compliance costs (y-axis) and risk-premia (x-axis) are normalized to values between 0 and 1. Note the logarithmic scale on both axes of Fig. 3. The null adjustment rate of a fixed-stringency policy corresponds to the sparse clouds in the top right corner of the diagram. This is the case of the fixed-cap extreme, where permit prices are volatile; the future required abatement is uncertain and the risk premium is at maximum. The burden of adjustment to shocks is completely borne by firms and, accordingly, aggregate costs are expected to be high. As the adjustment rate increases (this corresponds to the descending clouds), total compliance costs decrease. The future required abatement is less uncertain and permit price volatility decreases; so does the associated risk premium. As the risk premium continues to decrease, total compliance costs first decrease and then start to increase again. Similar to Fig. 2, this cost inversion reflects the trade-off between the two extremes of the policy stringency spectrum. In particular, as the adjustment rate nears the fixed price extreme, firms face decreasing inter-temporal trading opportunities, which undermines their ability to benefit from differences in marginal costs across time.

Conclusions

Most existing emissions trading systems have implemented fixed caps. By design, therefore, such systems lack provisions to address permit demand imbalances resulting from economic shocks. A currently proposed reform of the EU ETS aims to introduce a Market Stability Reserve that, within the bounds of a fixed cap, intends to increase the system responsiveness by temporarily adjusting permit allocation to the state of the aggregate bank of permits.

We study a responsive mechanism that can adjust policy stringency along a spectrum connecting a pure quantity instrument (i.e. an emissions trading system with a fixed cap) and a pure price instrument (i.e. a carbon tax with a fully floating cap). The mechanism changes the allocation of permits in response to shocks, adjusting the cap stringency to a level desired by the regulator. We solve the inter-temporal problem of regulated firms in the presence of such a mechanism and study how the level of policy stringency affects equilibrium abatement and trading strategies. Finally we show that by adjusting the stringency of the cap, the regulator can improve the overall cost effectiveness of the system.

Within the setting of the proposed mechanism, we identify a trade-off between the two extremes of the policy stringency spectrum. The trade-off arises from the fact that in an emissions trading system with banking and borrowing provisions, firms can benefit from inter-temporal abatement cost differences depending on where the policy stringency lies on the spectrum. When the cap stringency is relaxed (the policy stringency moves away from the pure fixed cap extreme), the cost arbitrage opportunities are reduced. As policy stringency nears the fully floating cap (or fixed price) extreme, inter-temporal trading thins out, and firms are no longer able to benefit from differences in marginal costs across time. In exchange, because the mechanism becomes more responsive to demand–supply imbalances, firms benefit from lower costs of having to adjust their strategies in response to shocks. We show that by selecting the right level of cap stringency, expected compliance costs can be minimized.

We solve the firms´ cost-minimization problem under both risk-neutrality and risk-aversion. Using a risk-averse setting allows us to describe the more realistic case of a risky investment in abatement when firms are unable to perfectly predict their future required number of permits. We show how the adjustment rate can impact the investment risk premium and through that, how it can impact the equilibrium dynamics and the expected total compliance costs. The adjustment rate becomes a significant determinant of the permit price dynamics. When the adjustment of the current permit allocation offsets the shock impact perfectly, the discounted permit prices are constant and the rate of return approaches the risk-free discount rate. Conversely, when the cap is fixed, permit prices vary reflecting period-by-period abatement adjustments, and the risk premium is at its maximum. Similar to risk-neutrality, we observe a policy stringency trade-off under risk aversion too: as the adjustment rate increases, the risk premium continues to decrease, with total compliance costs first decreasing and then increasing again.

Appendix A. The model under risk-neutrality

In the following sections we provide the derivations of the key results.

We consider an inter-temporal optimization problem where, at each point in time, firms have to decide how much they want to offset their emissions (by abating and by trading permits), considering the current and future costs of reducing emissions. We model a single finite compliance period $[0, T]$ of a secondary emissions permits market in a partial equilibrium framework under perfect competition. Firms $i \in I$ are assumed to be atomistic, that is, individual quantities x^i are continuously distributed under a measure m^x such that aggregate quantities can be obtained by integration, $x = x^I = \int_I x^i dm^x(i)$. Each firm continuously minimizes expected abatement and trading costs at each point in time $t \in [0, T]$, where each firm's instantaneous cost function is given by

$$v^i(\alpha_t^i, \beta_t^i) = \Pi_t \alpha_t^i + \varrho \cdot (\alpha_t^i)^2 - P_t \beta_t^i + \nu \cdot (\beta_t^i)^2.$$

That is, each firm has the same marginal abatement cost curve $\Pi_t + 2\varrho\alpha_t^i$, where we assume that the intercept Π_t increases by the risk-free rate r and $\varrho > 0$ is constant. Firms face transaction costs per trade. More specifically, we assume marginal trading costs to be linear in the number of permits sold ($\beta_t^i > 0$) or bought ($\beta_t^i < 0$). The parameter ν represents the magnitude of transaction costs and P_t denotes the time-t permit price. For convenience, we write $R_t^i = \mathbb{E}_t[Y^i(t, T)]$, where $Y^i(t, T)$ is the time-t residual required abatement:

$$Y^i(t, T) = Y^i(0, T) - \int_0^t \alpha_s^i \, ds + \int_0^t \beta_s^i \, ds,$$

where

$$Y^i(0, T) = E^i(0, T) - A^i(0, T) - B_0^i + c.$$

We note that

$$d\mathbb{E}_t[Y^i(t, T)] = (\beta_t^i - \alpha_t^i)dt + d\mathbb{E}_t[Y^i(0, T)].$$

We assume that instantaneous emissions $E^i(t, t + dt)$ are normally distributed with deterministic and bounded variance $(\kappa_t^i)^2 dt$, distributed around an average of g_t^i dt . In other words, we assume (recorded) cumulated emissions to be given by

$$E^i(0, t) = \int_0^t g_s^i ds + \int_0^t \kappa_s^i dW_s^{\mathbb{Q}},$$

where $W^{\mathbb{Q}}$ is a Brownian motion with respect to $\mathbb{Q}$. We assume that recording of emissions ends an (arbitrarily short) time before the compliance date, such that compliance is always possible. Accordingly, we assume for some $\tilde{t} < T$, that $\kappa_s^i = g_s^i = 0$ for $\tilde{t} \le s \le T$, meaning that each firm is given the interval $[\tilde{t}, T]$ to abate any remaining non-covered emissions. Furthermore, for $t=T$ we have the compliance constraint $B_T^i = B_T = c$; i.e. $Y^i(T, T) = Y(T, T) = 0$.

The equilibrium consists of abatement- and trading strategies α_t^i and β_t^i for each firm i and the market clearing price process P_t. In equilibrium, individual deviations from the equilibrium do not yield expected additional cost savings for any firm. The market is assumed to be free of arbitrage and complete. We can therefore postulate the existence of a martingale measure $\mathbb{Q}$ that is equivalent to the real-world measure $\mathbb{P}$. We first assume that firms are risk-neutral. Accordingly, all expectations in this section are taken under the measure $\mathbb{Q}$. In Appendix B, we transfer our results to risk-averse firms by deriving the change of measure from $\mathbb{Q}$ to $\mathbb{P}$.

We begin by assuming Markovian strategies $\alpha^j = \alpha(t, Z_t^j)$, $\beta^j = \beta(t, Z_t^j)$ for every firm $j \in I \setminus \{i\}$ except for i. These strategies are given as functions of each firm´s individual state processes Z_t^j, which will be specified later. We show that it is optimal for firm i to replicate the other firms' strategies given below, as a function of her own state process Z_t^i. For convenience, we define

$$h_t = \frac{re^{rt}}{e^{rT} - e^{rt}}.$$

For each firm $j \in I \backslash \{i\}$, let her abatement and trading strategies be given by

$$\alpha_t{}^j = \frac{P_t - \Pi_t}{2(\nu + \varrho)} + \frac{\nu}{\nu + \varrho} h_t R_t^j \quad \text{and} \quad \beta_t{}^j = \frac{P_t - \Pi_t}{2(\nu + \varrho)} - \frac{\varrho}{\nu + \varrho} h_t R_t^j.$$

The market clearing condition $\beta^I = 0$ yields

$$P_t = \Pi_t + 2\varrho h_t R_t^I. \tag{7}$$

Substituting for the strategies α_t^j, β_t^j above, we obtain the dynamics for the process R_t^j:

$$dR_t{}^j = (\beta_t{}^j - \alpha_t{}^j)\, dt + d\mathbb{E}_t[Y^j(0, T)] = -\frac{re^{rt}}{e^{rT} - e^{rt}} R_t^j\, dt + d\mathbb{E}_t[Y^j(0, T)].$$

Solving the above, we obtain:

$$R_t^j = R_0^j \frac{e^{rT} - e^{rt}}{e^{rT} - 1} + (e^{rT} - e^{rt}) \int_0^t \frac{d\mathbb{E}_s[Y^j(0, T)]}{e^{rT} - e^{rs}}.$$

Integrating over I yields, together with Eq. (7) that

$$P_t = \Pi_t + 2\varrho h_t R_t^I = \Pi_t + 2\varrho \frac{re^{rt}}{e^{rT} - 1} R_0^I + 2\varrho r e^{rt} \int_0^t \frac{d\mathbb{E}_s[Y^I(0, T)]}{e^{rT} - e^{rs}}.$$

In particular, we observe that P has the following dynamics

$$dP_t = rP_t dt + 2\varrho h_t d\mathbb{E}_t[Y^I(0, T)].$$

Let the random shocks to $\mathbb{E}_t[Y^I(0, T)]$ be governed by a driftless diffusion

$$d\mathbb{E}_t[Y^I(0, T)] = \sigma_t^I dW_t^{\mathbb{Q}},$$

where σ_t^I is deterministic and $W^{\mathbb{Q}}$ is a Brownian motion under the measure $\mathbb{Q}$. In the next section, we analyze how the variance $(\sigma_t^I)^2$ is affected by a mechanism that adjusts the permits allocation based on the aggregate bank. The process σ_t^i describes how changes in the expected total required abatement of permits are distributed across the set of firms I. We abstract from specific assumptions about the form of σ_t^i. However, we note that it is reasonable to assume different σ_t^i for different firms, since pre-abatement emissions levels and permits allocations can vary depending on the type of industry in consideration.

We consider changes in pre-abatement permits demand and in the (possibly contingent) permits allocation. Their degree of impact on firms can vary. However, all firms are subject to systemic shocks. Hence, we consider the same Brownian motion $W^{\mathbb{Q}}$ for each $i \in I$, whereas differences in size, technology, etc. are represented by the distribution of σ_t^i across I. Accordingly, shocks to $\mathbb{E}_t[Y^i(0, T)]$ are represented by

$$d\mathbb{E}_t[Y^i(0, T)] = \sigma_t^i dW_t^{\mathbb{Q}}.$$

We now consider the problem of optimal pollution control and permits trading for firm i. Let p denote an observed permit price and let $P^{t,p}$ denote the price process with time-t value $P_t^{t,p} = p$. Analogously, let $\Pi_t^{t,\pi} = \pi$. At time $s \geq t$, firm i has to bear costs v_s^i given by

$$v_s^i(\alpha_s^i, \beta_s^i) = \Pi_s^{t,\pi} \alpha_s^i + \varrho \cdot (\alpha_s^i)^2 - P_s^{t,p} \beta_s^i + \nu \cdot (\beta_s^i)^2.$$

Firm i's problem is to find (Markovian) abatement- and trading strategies α^i and β^i respectively, such that, for all $t \in [0, T)$, the cost function J, given by

$$J(t, R^i, p, \pi, \alpha^i, \beta^i) = \mathbb{E}_t\left[\int_t^T e^{-rs} v_s^i(\alpha_s^i, \beta_s^i)\, ds\right],$$

is minimized by α^i, β^i for all $\pi > 0$, $p \geq 0$, and such that the constraint $\mathbb{E}_t[R_T^i] = 0$ is satisfied for all $t \in [0, T)$. Let $w(t, R^i, p, \pi) = \inf_{(\alpha^i, \beta^i)} J(t, R^i, p, \pi, \alpha^i, \beta^i)$ denote the value function for firm i.

The firm observes the state process $Z_t^i = (R_t^i, P_t, \Pi_t)$, where

$$dR_t^i = (\beta_t^i - \alpha_t^i)\, dt + d\mathbb{E}_t[Y^i(0, T)] = (\beta_t^i - \alpha_t^i)dt + \sigma_t^i\, dW_t^{\mathbb{Q}},\ dP_t = rP_t dt + 2\varrho h_t d\mathbb{E}_t[Y^I(0, T)] = rP_t dt + 2\varrho h_t \sigma_t^I\, dW_t^{\mathbb{Q}},\ d\Pi_t = r\Pi_t dt.$$

Let the firm's filtration $(\mathcal{F}_t^i)_{t \geq 0}$ be generated by the process Z^i and accordingly, let $(\mathcal{F}_t^I)$, generated by Z^I, denote the aggregate filtration. The Hamilton–Jacobi–Bellman (HJB) Equation associated to the minimization problem above is given by

$$0 = D_t w + \inf_{a,b}\Big[(b-a)D_{R^i}w + rpD_p w + r\pi D_\pi w + \frac{1}{2}\mathrm{tr}(\Sigma\Sigma' D_z^2 w) + e^{-rt}(\pi a + \varrho a^2 - pb + \nu b^2)\Big]$$
$$= D_t w + rpD_p w + r\pi D_\pi w + \frac{1}{2}\mathrm{tr}(\Sigma\Sigma' D_z^2 w) + \inf_{a,b}\Big[(b-a)D_{R^i}w + e^{-rt}(\pi a + \varrho a^2 - pb + \nu b^2)\Big],$$

where Σ is the vector

$$\Sigma = \begin{pmatrix} \sigma_t^i \\ 2\varrho h_t \sigma_t^I \\ 0 \end{pmatrix}$$

which implies that

$$\mathrm{tr}(\Sigma\Sigma' D_z^2 w) = (\sigma_t^i)^2 D_{R^i}^2 w + 2\varrho h_t \sigma_t^i \sigma_t^I D_p D_{R^i} w + 2\varrho h_t \sigma_t^i \sigma_t^I D_{R^i} D_p w + 4\varrho^2 h_t^2 (\sigma_t^I)^2 D_p^2 w.$$

We notice that the minimizers a, b in the above equation have to satisfy

$$a = \frac{1}{2\varrho}\Big(e^{rt}D_{R^i}w - \pi\Big) \quad \text{and} \quad b = \frac{1}{2\nu}\Big(p - e^{rt}D_{R^i}w\Big). \tag{8}$$

Furthermore, we notice that the second-order condition is satisfied for all a, b. This yields the following

Lemma 1. *The HJB equation can be rewritten as*

$$0 = e^{rt}\Big(D_t w + rpD_p w + r\pi D_\pi w\Big) + \frac{e^{rt}}{2}\mathrm{tr}(\Sigma\Sigma' D_z^2 w) - \frac{1}{4\varrho}\big(e^{rt}D_{R^i}w - \pi\big)^2 - \frac{1}{4\nu}\big(p - e^{rt}D_{R^i}w\big)^2. \tag{9}$$

In order to enforce the constraint $\mathbb{E}_t[R_T^i] = 0$ for all t, we impose the singular terminal condition

$$\lim_{t \nearrow T} w\big(t, R^i, p, \pi\big) = \begin{cases} 0: R^i = 0, \\ \infty: R^i \neq 0. \end{cases} \tag{10}$$

Theorem 1. *The HJB Eq. (9), together with the terminal condition (10) is solved by*

$$w(t, R^i, p, \pi) = \frac{r\nu\varrho(R^i)^2}{(e^{rT} - e^{rt})(\nu + \varrho)} + e^{-rt}\left(\pi + \frac{\varrho(p - \pi)}{\nu + \varrho}\right)R^i + \frac{(1 - e^{r(T-t)})(p - \pi)^2}{4re^{rt}(\nu + \varrho)} + \int_t^T C_s\, ds$$

where

$$C_s = \frac{r\nu\varrho(\sigma_s^i)^2}{(e^{rT} - e^{rs})(\nu + \varrho)} + \frac{2\varrho^2 h_s \sigma_s^i \sigma_s^I e^{-rs}}{\nu + \varrho} + \frac{\varrho^2 h_s^2 (\sigma_s^I)^2 (1 - e^{r(T-s)})}{re^{rs}(\nu + \varrho)} \quad \textit{for } t \le s < T.$$

The above theorem can be proven by simple differentiation. The verification argument for w is straightforward but lengthy. Thus, we omit the full proof. We note that standard arguments of verification confirm α^i, β^i as the firm's optimal strategies. Substituting $D_{R^i}w$ in Eq. (8) yields

$$\alpha_t^i = \frac{P_t - \Pi_t}{2(\nu + \varrho)} + \frac{\nu}{\nu + \varrho} h_t R_t^i \quad \text{and} \quad \beta_t^i = \frac{P_t - \Pi_t}{2(\nu + \varrho)} - \frac{\varrho}{\nu + \varrho} h_t R_t^i.$$

This proves the equilibrium strategies α^i, β^i to be given as above for all $i \in I$. Furthermore, the aggregate abatement is given by

$$\alpha_t = re^{rt}\frac{R_0^I}{e^{rT} - 1} + re^{rt}\int_0^t \frac{d\mathbb{E}_s[Y^I(0, T)]}{e^{rT} - e^{rs}}$$

and, accordingly, the market-clearing price process is given by

$$P_t = \Pi_t + 2\varrho\frac{re^{rt}}{e^{rT} - 1} R_0^I + 2\varrho re^{rt}\int_0^t \frac{d\mathbb{E}_s[Y^I(0, T)]}{e^{rT} - e^{rs}}.$$

A.1. The bank under the responsive mechanism

We give a brief derivation of the closed-form expression for the time-t bank of permits. In the following, we omit the superscript I for aggregate quantities. For convenience we define $d\varepsilon_t = \kappa_t dW_t^{\mathbb{Q}}$. The dynamics of the aggregate bank is given by

$$dB_t = f_t dt + \delta(c - B_t)dt - g_t dt - d\varepsilon_t + \alpha_t dt. \tag{11}$$

Notice that Eq. (11) yields an expression for B_t in terms of the process α_t:

$$B_t = B_0 e^{-\delta t} + \int_0^t e^{\delta(s-t)}(\alpha_s + f_s + \delta c - g_s)\, ds - \int_0^t e^{\delta(s-t)}\, d\varepsilon_s. \tag{12}$$

Recall that aggregate abatement is given by

$$\alpha_t = re^{rt}\frac{R_0}{e^{rT} - 1} + re^{rt}\int_0^t \frac{d\mathbb{E}_s[Y(0, T)]}{e^{rT} - e^{rs}}$$

and notice that the dynamics of $d\mathbb{E}_s[Y(0, T)]$ are given by

$$d\mathbb{E}_t[Y(0, T)] = d\mathbb{E}_t\left[\int_0^T d\varepsilon_s - \int_0^T \delta B_s\, ds\right] = d\varepsilon_t + \delta d\mathbb{E}_t\left[\int_0^T B_s\, ds\right].$$

In order to obtain α_t and B_t distributions, we solve the above expression. To simplify notation, recall that

$$h_t = \frac{re^{rt}}{e^{rT} - e^{rt}}.$$

This yields

$$d\alpha_t = r\alpha_t dt + h_t d\mathbb{E}_t[Y(0, T)]. \tag{13}$$

Since the constraint $\mathbb{E}_t[B_T] = c$ is satisfied for all t, we have

$$d\mathbb{E}_t\left[\int_0^T e^{\delta s}\alpha_s\, ds\right] = d\mathbb{E}_t\left[\int_0^T e^{\delta s}(g_s - f_s)\, ds - B_0 + \int_0^T e^{\delta s}\, d\varepsilon_s\right] = e^{\delta t} d\varepsilon_t. \tag{14}$$

We use the dynamics of α_t in Eq. (13) to establish:

$$\int_0^t \alpha_s e^{\delta s}\, ds = \frac{e^{\delta t}}{\delta + r}\left(\alpha_t - \alpha_0 e^{-\delta t} - \int_0^t e^{\delta(s-t)} h_s\, d\mathbb{E}_s[Y(0, T)]\right),$$

which, together with Eq. (14), yields:

$$d\mathbb{E}_t[Y(0, T)] = \frac{V_t(\delta, r)}{h_t}\, d\varepsilon_t, \tag{15}$$

where $V_t(\delta, r) = (\delta + r)/(e^{(\delta+r)(T-t)} - 1)$. From this we finally obtain

$$B_t = B_0 e^{-\delta t} + \frac{r(e^{rt} - e^{-\delta t})}{(\delta + r)(e^{rT} - 1)}R_0 - \frac{e^{rt}}{V_t(\delta, r)}\int_0^t e^{-rs}\, V_s(\delta, r)\, d\varepsilon_s + \int_0^t e^{\delta(s-t)}(f_s - g_s + \delta c)\, ds. \tag{16}$$

Note that $R_0 = \mathbb{E}_0[Y(0, T)]$. Using the compliance condition $B_T = c$, we can derive the expression for $\mathbb{E}_0[Y(0, T)]$:

$$\mathbb{E}_0[Y(0, T)] = -\frac{(\delta + r)(e^{rT} - 1)}{r(e^{rT} - e^{-\delta T})}\left(B_0 e^{-\delta T} + \int_0^T e^{\delta(s-T)}(f_s - g_s + \delta c)\, ds - c\right).$$

Thus, the time-t bank B_t is indeed determined in closed-form by Eq. (16).

We obtain that $B_t \sim \mathcal{N}(a_t, b_t^2)$ where

$$a_t = B_0 e^{-\delta t} + \frac{r(e^{rt} - e^{-\delta t})}{(\delta + r)(e^{rT} - 1)}R_0 + \int_0^t e^{\delta(s-t)}(f_s - g_s + \delta c)\, ds$$

is the mean, and

$$b_t^2 = \frac{e^{2rt}}{V_t^2(\delta, r)}\int_0^t e^{-2rs}\, V_s^2(\delta, r)\, \kappa_s^2\, ds$$

is the variance. Now, let λ denote the probability that the bank stays within the band $[l_t, u_t]$. We can then compute the following

$$\lambda = \Phi\left(d_t^{(1)}\right) - \Phi\left(d_t^{(2)}\right),$$

where $\Phi(\cdot)$ represents the cumulative distribution function of the standard normal distribution and

$$d_t^{(1)} = \frac{u_t - a_t}{b_t} \quad \text{and} \quad d_t^{(2)} = \frac{l_t - a_t}{b_t}.$$

Thus, λ can be expressed as a function $C(\delta) = \lambda$. Any bank interval can therefore be maintained with a confidence level λ when the adjustment rate δ is set to $C^{-1}(\lambda)$ and vice-versa.

A.2. Expected aggregate compliance costs

Corollary 1. *We find that expected aggregate compliance costs for identical firms are given by*

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\left(\Pi_t\alpha_t + \varrho\alpha_t^2\right)dt\right] = \Pi_0 R_0 + \varrho r\frac{R_0^2}{e^{rT}-1} + \varrho r\int_0^T \frac{d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)}.$$

Proof.

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\left(\Pi_t\alpha_t + \varrho\alpha_t^2\right)dt\right] = \mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\Pi_t\,\alpha_t\,dt\right] + \varrho\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\alpha_t^2\,dt\right] = \Pi_0\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T \alpha_t\,dt\right] + \varrho\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\alpha_t^2\,dt\right]. \tag{17}$$

Recall that we have for all δ

$$\int_0^t e^{\delta s}\alpha_s ds = \frac{1}{\delta + r}\left(e^{\delta t}\alpha_t - \alpha_0 - \int_0^t e^{\delta s}h_s d\mathbb{E}_s[Y(0,T)]\right).$$

We can use the above with $\delta = 0$ and $t{=}T$ to resolve the left-hand integral in Eq. (17):

$$\Pi_0\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T \alpha_t\,dt\right] = \frac{\Pi_0}{r}\left(\mathbb{E}^{\mathbb{Q}}\left[\alpha_T\right] - \alpha_0 - \mathbb{E}^{\mathbb{Q}}\left[\int_0^T h_t d\mathbb{E}_t[Y(0,T)]\right]\right).$$

By the expression for aggregate abatement we obtain

$$\mathbb{E}^{\mathbb{Q}}[\alpha_T] - \alpha_0 = re^{rT}\frac{R_0}{e^{rT}-1} - r\frac{R_0}{e^{rT}-1} = rR_0.$$

We thus arrive at

$$\Pi_0\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T \alpha_t\,dt\right] = \Pi_0 R_0. \tag{18}$$

Regarding the right-hand term in Eq. (17), we have

$$\int_0^T e^{-rt}\alpha_t^2\,dt = -\frac{1}{r}\int_0^T \alpha_t^2\,de^{-rt} = \frac{1}{r}\left(\alpha_0^2 - e^{-rT}\alpha_T^2 + \int_0^T e^{-rt}d\alpha_t^2\right).$$

We can use the fact that

$$d\alpha_t^2 = 2\alpha_t d\alpha_t + d\langle\alpha\rangle_t = 2\alpha_t(r\alpha_t dt + h_t d\mathbb{E}_t[Y(0,T)]) + h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t = 2r\alpha_t^2 dt + 2\alpha_t h_t d\mathbb{E}_t[Y(0,T)] + h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t$$

to deduce that

$$\begin{aligned}\int_0^T e^{-rt}\alpha_t^2\,dt &= \frac{1}{r}\left(\alpha_0^2 - e^{-rT}\alpha_T^2 + \int_0^T e^{-rt}2r\alpha_t^2 dt + \int_0^T e^{-rt}2\alpha_t h_t d\mathbb{E}_t[Y(0,T)] + \int_0^T e^{-rt}h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t\right)\\ &= \frac{1}{r}\left(\alpha_0^2 - e^{-rT}\alpha_T^2 + \int_0^T e^{-rt}2\alpha_t h_t d\mathbb{E}_t[Y(0,T)] + \int_0^T e^{-rt}h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t\right) + 2\int_0^T e^{-rt}\alpha_t^2 dt.\end{aligned}$$

This implies that

$$\int_0^T e^{-rt}\alpha_t^2\,dt = \frac{1}{r}\left(e^{-rT}\alpha_T^2 - \alpha_0^2 - \int_0^T e^{-rt}2\alpha_t h_t d\mathbb{E}_t[Y(0,T)] - \int_0^T e^{-rt}h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t\right)$$

and consequently

$$\varrho\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\alpha_t^2\,dt\right] = \frac{\varrho}{r}\left(e^{-rT}\left(\mathrm{Var}[\alpha_T] + \left(\mathbb{E}^{\mathbb{Q}}[\alpha_T]\right)^2\right) - \alpha_0^2 - \int_0^T e^{-rt}h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t\right),$$

where we used, in particular, that $d\langle\mathbb{E}[Y(0,T)]\rangle_t/dt$ is deterministic. Since $d\mathbb{E}_t[Y(0,T)]/dW_t^{\mathbb{Q}}$ is deterministic and bounded in $[0,T]$ we obtain

$$e^{-rT}\mathrm{Var}[\alpha_T] = e^{-rT}r^2e^{2rT}\int_0^T \frac{d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)^2} = r^2\int_0^T \frac{e^{rT}d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)^2}.$$

Also notice that

$$\int_0^T e^{-rt}h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t = r^2\int_0^T \frac{e^{rt}d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)^2}$$

and hence

$$e^{-rT}\mathrm{Var}[\alpha_T] - \int_0^T e^{-rt}h_t^2 d\langle\mathbb{E}[Y(0,T)]\rangle_t = r^2\int_0^T \frac{d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)}.$$

Furthermore,

$$e^{-rT}\left(\mathbb{E}^{\mathbb{Q}}[\alpha_T]\right)^2 - \alpha_0^2 = e^{-rT}r^2e^{2rT}\frac{R_0^2}{(e^{rT}-1)^2} - r^2\frac{R_0^2}{(e^{rT}-1)^2} = r^2\frac{R_0^2}{e^{rT}-1}.$$

Together, we find that

$$\varrho\,\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\alpha_t^2\,dt\right] = \varrho r\frac{R_0^2}{e^{rT}-1} + \varrho r\int_0^T \frac{d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)}. \tag{19}$$

(Eqs. (18) and 19) together yield the result:

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt}\left(\Pi_t\alpha_t + \varrho\alpha_t^2\right)dt\right] = \Pi_0R_0 + \varrho r\frac{R_0^2}{e^{rT}-1} + \varrho r\int_0^T \frac{d\langle\mathbb{E}[Y(0,T)]\rangle_t}{\left(e^{rT}-e^{rt}\right)}. \qquad \square$$

Appendix B. The model under risk-aversion

We solved the model above under the assumption that all firms are risk-neutral. We now evaluate the equilibrium dynamics under the objective measure $\mathbb{P}$ by introducing a risk-adjusted discount rate.

Under $\mathbb{P}$, let the price dynamics be given by

$$dP_t = (r + q_t)P_t dt + 2\varrho V_t(\delta, r)\kappa_t dW_t^{\mathbb{P}}, \tag{20}$$

where $W_t^{\mathbb{P}}$ is a standard Brownian motion under $\mathbb{P}$. Recall that under $\mathbb{Q}$, the price process has dynamics

$$dP_t = rP_t dt + 2\varrho V_t(\delta, r)\kappa_t dW_t^{\mathbb{Q}}. \tag{21}$$

From Eqs. (20) and (21) we observe that the $\mathbb{Q}$-Brownian motion $W_t^{\mathbb{Q}}$ has to satisfy

$$dW_t^{\mathbb{Q}} = dW_t^{\mathbb{P}} + \frac{q_tP_t}{2\varrho V_t(\delta, r)\kappa_t}\,dt.$$

We assume that q_t is proportional to the volatility coefficient in permit price returns; i.e. that the expression

$$\frac{q_tP_t}{2\varrho V_t(\delta, r)\kappa_t}$$

is a constant. By Girsanov's Theorem we therefore obtain that for

$$\vartheta_t = rt + \frac{1}{2}\int_0^t\left(\frac{q_sP_s}{2\varrho V_s(\delta, r)\kappa_s}\right)^2 ds + \int_0^t \frac{q_sP_s}{2\varrho V_s(\delta, r)\kappa_s}\,dW_s^{\mathbb{P}},$$

the process $e^{-\vartheta_t + rt}$ is a $\mathbb{P}$-martingale. And hence we obtain the desired change of measure by the Radon–Nikodým derivative $d\mathbb{Q}/d\mathbb{P}|_{\mathcal{F}_t} = e^{-\vartheta_t + rt}$.

References

Doda, B., 2016. How to price carbon in good times… and bad!. WIREs Clim. Change 7 (7), 135–144.

Ellerman, D., Valero, V., Zaklan, A., 2015. An analysis of allowance banking in the EU ETS. EUI Working Paper RSCAS 2015/29.

Ellerman, D., Wing, I., 2003. Absolute versus intensity-based emission caps. Clim. Policy 3 (1), 7–20.

EP, 2015. Market stability reserve for the Union greenhouse gas emission trading scheme. Approved text, European Parliament.

Fell, H., 2015. Comparing Policies to Confront Permit Over-allocation. Unpublished working paper.

Fell, H., Morgenstern, R., 2010. Alternative approaches to cost containment in a cap-and-trade system. Environ. Resour. Econ. 47 (2), 275–297.

Fell, H., Morgenstern, R., Palmer, K., 2012. Soft and hard price collars in a cap-and-trade system: a comparative analysis. J. Environ. Econ. Manag. 64, 183–198.

Frino, A., Kruk, J., Lepone, A., 2010. Liquidity and transaction costs in the European carbon futures market. J. Deriv. Hedge Funds 16 (2), 100–115.

Golosov, M., Hasslet, J., Krussell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. Econometrica 82, 41–88.

Grosjean, G., Acworth, W., Flachsland, C., Marschinski, R., 2014. After monetary policy, climate policy: is delegation the key to EU ETS reform?. Clim. Policy 1, 1–25.

Grüll, G., Taschini, L., 2011. Cap-and-trade properties under different hybrid scheme designs. J. Environ. Econ. Manag. 61 (1), 107–118.

Hepburn, C., 2006. Regulation by prices, quantities, or both: a review of instrument choice. Oxf. Rev. Econ. Policy 22 (2), 226–247.

Heutel, G., 2012. How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks. Rev. Econ. Dyn. 15, 244–264.

Jotzo, F., Pezzey, J., 2007. Optimal intensity targets for greenhouse gas emissions trading under uncertainty. Environ. Resour. Econ. 38 (2), 259–284.

Kollenberg, S. and Taschini, L., Dynamic Supply Adjustment and Banking Under Uncertainty: The Market Stability Reserve. Working Paper.March 2016, Available at SSRN: http://dx.doi.org/10.2139/ssrn.2742256Landis, F., 2015. Interpolation of Abatement Cost According to PACE. Unpublished working paper.

Medina, V., Pardo, A., Pascual, R., 2014. The timeline of trading frictions in the European carbon market. Energy Econ. 42 (3), 378–394.

Newell, R.G., Pizer, W.A., 2008. Regulating stock externalities under uncertainty. J. Environ. Econ. Manag. 56, 221–233.

Perino, G., Willner, M., 2015. The Price and Emission Effects of a Market Stability Reserve in a Competitive Allowance Market. WiSo-HH Working Paper Series, No. 28.

Pizer, W.A., 2002. Combining price and quantity controls to mitigate global climate change. J. Public Econ. 85, 409–434.

Pizer, W.A., Prest, B., 2016. Prices versus Quantities with Policy Updating. NBER Working Paper 22379.

Roberts, M.J., Spence, M., 1976. Effluent charges and licenses under uncertainty. J. Public Econ. 5 (3), 193–208.

Rubin, J.D., 1996. A model of intertemporal emission trading, banking, and borrowing. J. Environ. Econ. Manag. 31, 269–286.

Salant S. W., What ails the European Union's emissions trading system? J. Environ. Econ. Manag. 2016, This issue. Schennach, S.M., 2000. The economics of pollution permit banking in the context of title iv of the 1990 clean air act amendments. J. Environ. Econ. Manag. 40, 189–210.

Unold, W., Requate, T., 2001. Pollution control by options trading. Econ. Lett. 73, 353–358.

Weitzman, M.L., 1974. Prices vs. quantities. Rev. Econ. Stud. 41 (4), 683–691.