Emissions Trading with Rolling Horizons*

Simon Quemin[1,2,3] Raphaël Trotignon[1]

[1] Climate Economics Chair – *Paris-Dauphine University, PSL Research University*

[2] Potsdam Institute for Climate Impact Research – *Member of the Leibniz Association*

[3] Grantham Research Institute – *London School of Economics and Political Science*

January 2021
Accepted at *Journal of Economic Dynamics & Control*

Abstract

We develop an emission permits trading model where covered firms can (1) utilize rolling planning horizons to deal with uncertainty and (2) exhibit bounded responsiveness to supply-side control policies. We calibrate the model to reproduce annual market outcomes in the EU ETS over 2008–2018 and show that a rolling finite horizon reconciles the banking dynamics with discount rates implied by futures contracts' yield curves. It also replicates the price dynamics well compared to a standard infinite horizon, including the new price regime induced by the 2018 market reform. We then use our calibrated model to decompose the impacts of the 2018 reform's design elements, quantify how they hinge on the firms' horizon and responsiveness, and highlight important implications for policy design. For instance, when firms utilize rolling horizons, the Market Stability Reserve can improve effectiveness by frontloading abatement efforts and induce lower cumulative emissions compared to an infinite horizon.

Keywords Emissions trading, Rolling horizon, Bounded rationality, Supply-side policies, EU ETS.

JEL classification D25, D81, E63, H32, Q58.

*Corresponding author: Simon Quemin. Email address: Simon.Quemin@pik-potsdam.de.

«Agents may be easily assumed to be forward-looking, but do their horizons extend thirty years and even more into the future? And, if the reality is truncated horizons that are updated and moved forward as time progresses, what is the relevant time span? More importantly [...] how would such behavior change the equilibrium paths [...]?»

— Ellerman et al. (2015), discussing intertemporal trading in the EU ETS.

1 Introduction

Emissions trading systems (ETS) have become a widely used climate and energy policy tool. Most allow for some form of flexibility in the covered firms' emission streams through banking and borrowing of emission permits over time (ICAP, 2020; World Bank, 2020). That is, firms must comply with the cumulative cap on emissions over the length of the program, but how they manage their emissions' calendars is left to their discretion, at least to a certain extent. In many respects, therefore, determining the optimal timing of emissions and permit usage is isomorphic to Hotelling's problem of efficiently extracting an exhaustible resource sold in a competitive market under conditions of uncertainty.[1] In principle, as firms cost minimize over time, intertemporal trading and absence of arbitrage entail that the permit price reflects the expected discounted long-term scarcity of permits implied by the cap trajectory and that the cumulative cap is attained at least discounted cost in expectation (e.g. Rubin, 1996; Kling & Rubin, 1997; Schennach, 2000; Ellerman & Montero, 2007).

As the opening quote recognizes, however, the market equilibrium outcomes prevailing under this cost-effective approach to timing emissions crucially hinge on the planning horizon firms effectively employ. In this paper, we formalize and answer the three questions it raises with a specific focus on the EU ETS. That is, we first provide ample anecdotal evidence suggesting that firms use 'truncated horizons that are updated and moved forward as time progresses' in their decision-making process, what we call *rolling horizons*. Second, we develop a model of competitive emissions trading where firms can utilize rolling horizons, which we calibrate to market developments over 2008–2017 to obtain a measure of the 'relevant time span'. Third, we analyze how rolling horizons 'change the equilibrium paths' relative to infinite horizons. This has important implications for policy design and evaluation, some of which we highlight by assessing the impacts of the 2018 EU ETS reform with our calibrated model.

[1] One difference from Hotelling's problem arises when borrowing is restricted (if not prohibited as is often the case in practice) since the feasibility of the cost-effective path then depends on the temporal availability of permits and thus on the allocation stream, see Hasegawa & Salant (2014) for a discussion. Other differences include the existence of supply-side controls or the absence of storage costs and depreciation over time.

Specifically, the equilibrium outcomes when firms utilize rolling horizons with a given discount rate may prima facie be seen as qualitatively equivalent to those when firms have an infinite horizon with a larger discount rate.[2] Crucially, however, this may only hold assuming market design considerations away while in practice permit markets are often equipped with supply-side controls, typically in the form of a price collar (e.g. Roberts & Spence, 1976) or a banking corridor in the post-reform EU ETS (e.g. Perino & Willner, 2016). In the presence of a supply control, short- and long-term equilibrium outcomes will depend on the interaction between the control's design features and the firms' horizon. Additionally, they will also depend on the firms' sophistication in understanding the implications of the control on their intertemporal decision making, what we call firms' *responsiveness*. In this paper, we analyze and quantify the interplay between the post-reform EU ETS design elements and the firms' horizon and degree of responsiveness. This delineates the three main contributions of this paper, which we present below in further detail in relation with the existing literature.

As a first contribution, we develop a model of competitive emissions trading under uncertainty with supply-side control which departs from the existing literature along two key dimensions. First, firms can use rolling horizons in the face of uncertainty and limited information about future permit supply and demand. As a case in point, future demand may vary substantially depending on the economic activity and the achievements of complementary policies (see e.g. Borenstein et al., 2019). This essentially involves an iterative optimization procedure over a sliding truncated horizon given realistic forecasts of the relevant exogenous factors. Only the first period of the plan is implemented, and then a new plan is formulated over an equally long horizon with updated forecasts. Second, firms can exhibit two polar degrees of responsiveness. With zero (resp. full) responsiveness, firms have no (resp. a perfect) understanding of how their intertemporal equilibrium decisions interact with the supply control impacts over time. The full responsiveness case coincides with the rational expectations equilibrium framework of Muth (1961) and obtains as the fixed point between the firms' beliefs about the control impacts and its impacts in equilibrium in the spirit of Lucas & Prescott (1971).[3]

By contrast, the archetypal permit trading model considers fully rational firms with an infinite horizon, or at least extending to the end of the program, see Hasegawa & Salant (2015) for a review. The concept of rolling horizons was first introduced by Goldman (1968) and can be

[2]See Section 3.3 for a discussion of the conditions that are necessary for this equivalence to hold.

[3]That is, under full (resp. zero) responsiveness, firms perfectly (resp. do not) understand the control rules and foresee its consequences, which coincides with the assumption of rational competitive markets (resp. an extreme form of bounded rationality). In this paper, full responsiveness does not entail any deviation from competitive behavior, e.g. manipulation, gaming or any form of strategic behavior (see Section 3.5).

Figure 1: Observed daily EUA spot price and yearly bank (Jan 2008 – Dec 2018)

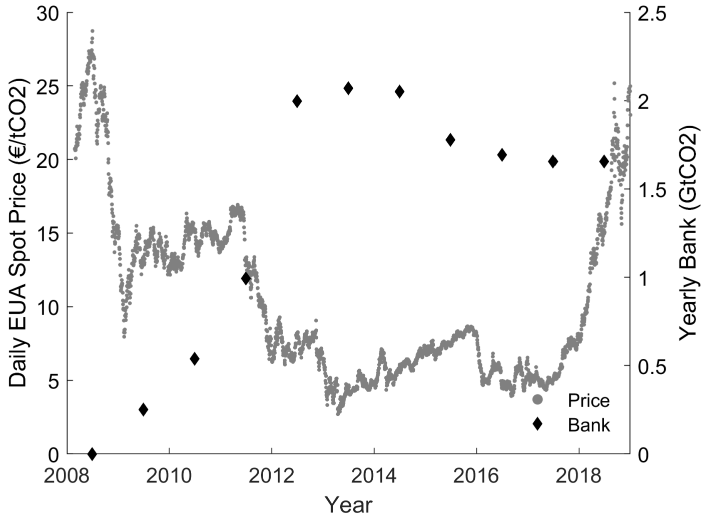

Note: EUA = EU emission Allowance. Data compiled from ICE and the EU Transaction Log.

seen as a way of addressing increasing uncertainty and costly informational requirements (e.g. Easley & Spulber, 1981; Grüne et al., 2015). Additionally, rolling planning is extensively used by firms in their routine production and supply management processes (see Sahin et al., 2013, for a review) and was proposed in the similar context of exhaustible resource extraction to help rationalize the lack of empirical support for the Hotelling rule (Spiro, 2014; van Veldhuizen & Sonnemans, 2018).[4] As another source of bounded rationality, we consider that firms may face cognitive or computational limitations in optimally adjusting their decisions to supply control. We believe that introducing limited responsiveness is an alley worth exploring as it fits well into the context of the EU ETS where there is evidence that, even in the absence of supply control, firms do not fully appreciate the trading and profit opportunities created by the market (e.g. Martin et al., 2015; Venmans, 2016; Karpf et al., 2018; Baudry et al., 2020). Additionally, the firm-level implications of a banking corridor such as the one introduced by the 2018 reform may arguably not be as transparent and straightforward as those of a price signal conveyed by a price collar (e.g. Perino, 2018; Wettestad & Jevnaker, 2019).[5]

As a second contribution, we tailor the model to the EU ETS design and parametrize permit supply and demand using historical emissions and allocation data in alignment with prevailing

[4]See Section 2.1 for a literature overview on rolling planning and its applications, and see Section 2.2 for anecdotal evidence hinting at participants' limited farsightedness in the context of the EU ETS.

[5]In the words of Perino (2018), «the [control] rules should be simple and stable and their impacts predictable so that both market participants and regulators can understand them readily and respond accordingly. Such mechanisms do exist – but the new rules for [EU ETS] Phase IV are not among them».

regulation (e.g. ETS Directive, EU renewable or energy efficiency targets). We calibrate the market-wide discount rate and planning horizon by fitting simulated outputs to past annual banking and price dynamics and compare the merits of two alternative assumptions – infinite vs. rolling horizon – in how well they can rationalize observations. A rolling horizon of a dozen years can (1) reconcile annual banking levels over 2008–2017 with discount rates inferred from futures' yield curves (3% on average) where an infinite horizon can only do so with 'too high' rates, (2) reproduce average annual price changes over 2008–2017 twice better than an infinite horizon (both in size and sign), and (3) pick up the 2018 price rally and regime shift induced by the market reform where an infinite horizon falls short of it (Figure 1).

It is important to recognize one caveat that must be applied to interpreting this result based on the resultant of individual firms' behaviors. Specifically, we offer an interpretation in the spirit of Friedman's (1953) black box approach. That is not to say that all firms actually use rolling horizons, nor a fortiori the same horizon, but this representation has the comparative advantage of reproducing past market outcomes more satisfactorily than the infinite horizon. Additionally, despite the existing heterogeneity in firms' risk preferences and strategic abilities (e.g. Obara & Zincenko, 2017; Hortaçsu et al., 2019),[6] our representative firm approach allows us to develop a parsimonious model within the canonical Rubin–Schennach framework.[7] In this context, our calibration estimates are a novel and valuable contribution to the empirical literature on the EU ETS, which identifies firms' degrees of optimization, foresight levels and horizons as open research questions (e.g. Ellerman et al., 2016; Hintermann et al., 2016; Fuss et al., 2018). Our calibration exercise also enriches the methodology for the ex-post analysis of banking behavior proposed by Ellerman & Montero (2007) in the context of the US Acid Rain Program and it exemplifies how rolling horizons can be key in understanding price and banking dynamics as well as evaluating ETS' performances.[8]

If firms (behave as if they) focus more on the short to mid term than on the long term, this has

[6]Because it does not matter who holds the bank of permits, one could think that for efficient intertemporal behavior to obtain in aggregate it would suffice to have a few competitive rational arbitrageurs doing the intertemporal optimization while all we need from other regulated firms is static cost minimization. However, as we elaborate in Section 2.2, intertemporal optimization opportunities at large are restrained due to financial constraints, missing markets for long-term derivative contracts and (unhedgeable) regulatory uncertainty.

[7]Formally embedding different types of firms in the model is feasible but requires more complex equilibrium concepts (e.g. a cognitive hierarchy model as in Hortaçsu et al. (2019)) or less transparent assumptions for our calibration exercise in Section 4 (i.e. the higher the number of types of firms, the higher the number of risk and time preference parameters that need to be set exogenously).

[8]In assessing the banking efficiency ex post, Ellerman & Montero (2007) compare the fit between observed and simulated banking levels for various given pairs of discount and expected demand growth rates to guess at which pair may have governed the dynamics. We augment this approach by endogenizing changes in firms' expectation about future demand, allowing for rolling horizons and utilizing futures' yield curves.

important ramifications for policy design and outcomes. As a third contribution, we thus use our calibrated model to assess the impacts of the 2018 market reform (European Parliament & Council, 2018) and compare them based on the firms' horizon and responsiveness degree. We decompose the impacts of reform's three main elements, viz. (1) an increase in the annual Linear Reduction Factor of the cap from 1.74% to 2.20% from 2021 on, (2) the introduction of the Market Stability Reserve (MSR) in 2019, a banking corridor that adjusts current auctions downwards (resp. upwards) when the bank level in the previous year was above (resp. below) some threshold, and (3) its reinforcement by the cancellation mechanism (CM) in 2023, which cancels permits stored in the MSR in excess of auctioned volumes in the previous year.

We highlight three key results for policy design. First, quantitative MSR impacts can differ considerably depending on the firms' horizon and responsiveness. For instance, cumulative cancellations vary by a factor of two under an infinite vs. rolling horizon (5 vs. 10 $GtCO_2$).[9] Second, under a rolling horizon, the MSR can compensate for the firms' limited foresight and reduce the costs of meeting a given cumulative emissions target relative to a sole equivalent linear cap. That is, by postponing auctions (i.e. by frontloading abatement efforts), the MSR makes the long-term stringency implied by the linear cap more tangible early on (i.e. it raises the cap's transitional stringency), steering the abatement path more in line with the long-term target.[10] Note that the CM is instrumental in realizing this cost-effectiveness improvement as it prevents fixed, inflexible reinjections of MSR permits into the market in the future. Third, MSR-induced cumulative supply adjustments in response to demand shocks are limited in size and time, which is more marked under a rolling horizon. This suggests a limited stabilizing potential in reflecting the interactions with complementary energy policies and responding to unexpected shocks such as the Covid-19 pandemic or major recessions.

In this respect, our paper relates and adds to the burgeoning literature on the MSR initiated by Richstein et al. (2015), Fell (2016) and Perino & Willner (2016).[11] This literature finds that market outcomes and MSR impacts are sensitive to model parametrization and calibration as well as underlying assumptions about complementary abatement policies. Our finer-grained

[9]This is not surprising as the firms' horizon and discount rate are central to defining banking strategies, thereby affecting the MSR impacts, whose sensitivity to model inputs is notably discussed in Fell (2016).

[10]With a rolling horizon under a declining cap path, the price is 'too low' early on (as firms underestimate the actual stringency of the cap in the long term) and 'too high' later on (catch-up effect). The MSR partially corrects for this by backloading auctions while it is largely irrelevant with an infinite horizon, see Section 5.2.

[11]This includes Perino & Willner (2017), Perino (2018), Bocklet et al. (2019), Carlén et al. (2019), Gerlagh & Heijmans (2019), Gerlagh et al. (2019) and Beck & Kruse-Andersen (2020). For other dynamic supply adjustment mechanisms based on firms' banking similar in the spirit of the MSR, see Kollenberg & Taschini (2016, 2019), Gerlagh & Heijmans (2018) and Lintunen & Kuusela (2018).

calibration thus confers a comparative advantage to our simulation results.[12] Other related papers introduce risk aversion on the part of firms (Kollenberg & Taschini, 2016, 2019; Tietjen et al., 2021) or an output market (Chaton et al., 2018), typically electricity, or they account for firms' investment decisions (Mauer et al., 2019; Bruninx et al., 2020; Tietjen et al., 2021). In these contexts, as with a rolling horizon, the mere postponement of auctions via the MSR (irrespective of the CM) has noticeable impacts on market outcomes, either because it affects the temporal availability of permits and variability thereof, and thus the permit-specific risk premium and equilibrium price path, or because it affects equilibrium investment decisions and thus the speed and path of low-carbon capacity deployment.

The remainder is structured as follows. Section 2 introduces the concept of rolling planning horizons in relation with the literature and provides anecdotal evidence for why it is relevant in the context of the EU ETS. Section 3 embeds rolling horizons and limited responsiveness to supply control into an archetypal permit market modeling framework. Section 4 describes the model parametrization and calibration to the EU ETS. Section 5 analyzes the impacts of the 2018 reform, focusing on the interplay between market design elements and the firms' horizon and responsiveness degree. Section 6 concludes.

2 Planning with rolling horizons

2.1 Concept and literature

The farther away one looks into the future the more that future is uncertain in terms of possible outcomes, their probabilities and how to incorporate them into one's planning today. As a rational way of dealing with increasing uncertainty, informational constraints or cognitive limitations, agents may conceivably use heuristics or rules of thumb in their decision making to contain the associated computational complexity and informational requirements, all the while performing comparably to more complex procedures (see e.g. Simon, 1955; Baumol & Quandt, 1964; Radner, 1996; Gigerenzer & Todd, 1999). As a case in point, rolling horizons are widely used by firms in their decision making. Essentially, this involves making the most immediate decisions (those in the first period) by optimizing over a finite number of periods ahead for which they can make reasonably good forecasts about relevant exogenous factors.

[12] Specifically, if and where others calibrate their models by fitting parameters so as to replicate the market price the year before their simulations, our calibration is based on both annual price and banking observations over 2008–2017. Moreover, a rolling horizon offers a better calibration (both in and out of sample) compared to the infinite (or similar) horizon typically assumed in other models.

Only the first-period optimal decisions are implemented and the procedure is sliding, i.e. it repeats every period thereafter, each time over the rolling horizon. Firms can also revise or update their forecasts each time they initiate a new planning cycle.

A rolling horizon can thus be seen as a particular form of decision making under uncertainty with limited information, either because information is scarce (if not unavailable) or because agents can but imprecisely respond to it. In other words, forecasts can either be too costly or unreliable or both, and the more distant the future, the costlier or less reliable the forecasts. It may even be that forecasts beyond a certain point are simply unavailable. Because of these constraints, agents may rationally choose to postpone decisions regarding the distant future now and only address these later on when things become clearer.

Admittedly, sliding truncated horizons may be seen as a crude way of modeling the behavior of agents facing uncertainty. Micro-foundations for similar types of behavior may involve, inter alia, ambiguity aversion especially through the maxmin decision rule (Gilboa & Schmeidler, 1989), sparsity-based bounded rationality (Gabaix, 2014) or rational inattention in the face of information acquiring and processing costs (Reis, 2006). Although these modeling approaches also restrict informational requirements, a rolling horizon has the comparative advantage of being the most simple expedient for capturing the main driving mechanism that plans are finite without inducing other biases.

Goldman (1968) first delineated and formalized the concept of continual finite planning revision. It was later extended by Easley & Spulber (1981) and Sethi & Sorger (1991) to account for stochasticity and stationarity, by Jehiel (1995) to account for strategic interactions, and by Kaganovich (1985) in the context of capital accumulation, where rolling finite plans are approximately optimal in the infinite horizon sense. More recently, Grüne et al. (2015) have proposed a procedure called 'nonlinear model predictive control' which leverages an iterative solution of optimal control problems with a receding horizon – a widely used procedure in control engineering. Grüne et al. provide a convergence result for discounted optimal control problems and characterize when a rolling horizon approximately yields the infinite horizon optimal paths depending on the length of the horizon and the discount rate.

Rolling horizons have also been extensively developed in the production planning and supply chain management literature, see e.g. Sahin et al. (2013) for a review. Importantly, they have helped rationalize quantitative puzzles in terms of saving behaviors (Caliendo & Aadland, 2007), social security choices (Findley & Caliendo, 2009), and more in line with our paper, the long-run price dynamics of exhaustible resources such as oil (Spiro, 2014; van Veldhuizen & Sonnemans, 2018). In the latter context, when the resource stock is large enough, a rolling

Figure 2: Normalized transacted volumes of EUA futures by maturities (2008–2018)

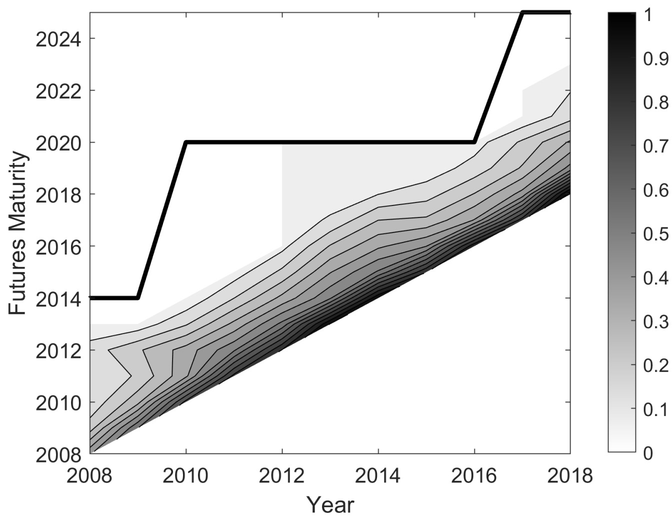

Note: Data compiled from the IntercontinentalExchange platform (ICE). For each year (x-axis), we measure the annual trading volume by maturity (y-axis), normalized by the maximal historical value. The thick black line represents the longest available maturity in each year.

horizon suppresses long-term scarcity constraints so that the dynamic nature of the problem vanishes and resource extraction does not conform to Hotelling's rule (see Section 3.3).

Additionally, existing literature on expectation formation indicates that agents use heuristics to forecast future factors relevant to their decision problems. For instance, as an alternative to rational expectations à la Muth (1961), Brock & Hommes (1997) consider heterogeneity in expectations and adaptive switching between heuristics. Such models of behavioral expectations perform well in representing expectation dynamics based on experimental or survey data (Branch, 2004; Hommes, 2011; Hommes et al., 2019). Relatedly, there is a large experimental body of literature documenting violations of rational behaviors in dynamic decision problems (e.g. Carbone & Hey, 2001; Johnson et al., 2002) or limitations on how far ahead people can plan (e.g. Hey & Knoll, 2007). In all these cases, agents typically attach insignificant if not no (salvage) value to outcomes beyond their horizon. Finally, experiments in asset markets show that traders can be myopic (Smith et al., 1988) and base their expectations on the extrapolation of past trends (Haruvy et al., 2007).

2.2 Anecdotal evidence in the EU ETS

The empirical literature indicates that firms regulated under the EU ETS behave consistently with intertemporal cost minimization, although the exact extent thereof is uncertain (see e.g.

Figure 3: Regulatory timeline and trading phases in the EU ETS

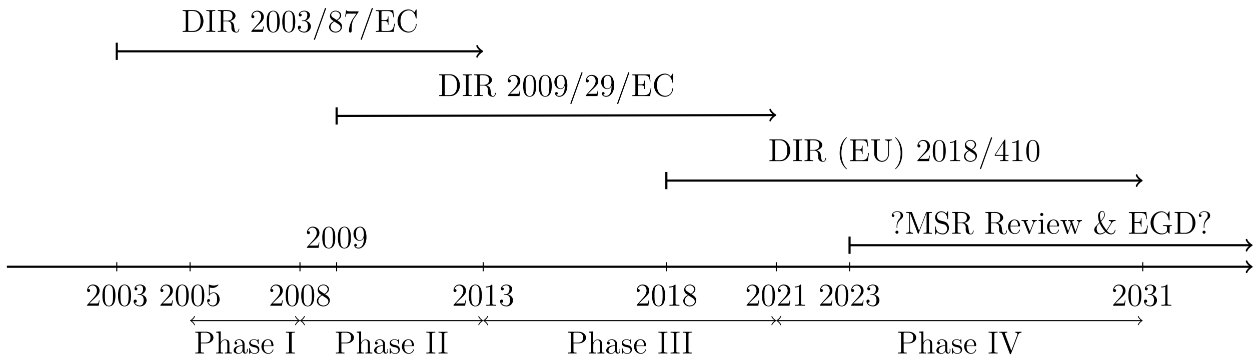

Note: The two ends of an arrow ↦ respectively indicate the year when the Directive is passed and the horizon associated with the quantitative objectives the Directive should help attain (viz. the 2012 Kyoto targets for DIR 2003/87/EC, the 2020 objectives for DIR 2009/29/EC and the 2030 objectives for DIR (EU) 2018/410. An overlap between arrows means that the more recent Directive supersedes the older. We include the market reform to come (expected in 2023 or 2024) following the 2021 review of the Market Stability Reserve and the implementation of other ETS-based measures as part of the EU Green Deal (EGD). For visual clarity, we do not include intermediary regulatory changes such as Regulations and Decisions, though they may affect the ETS design and functioning (e.g. Regulation 176/2014 that specifies the short-term backloading measure).

Ellerman et al., 2016; Hintermann et al., 2016; Fuss et al., 2018). Indeed, firms' degrees of optimizing behavior, levels of foresight and planning horizons largely remain open empirical questions. Below we list and discuss various reasons leading us to think that in practice firms are, whether intentionally or incidentally, likely to utilize rolling finite horizons.

Firstly, business and production plans are, as a rule, formulated over some finite horizon and updated on a regular basis. This is relevant for emission allowances as a factor of production part of overarching supply chain management strategies that typically rely on rolling horizon procedures (see e.g. Sahin et al., 2013; Zhang & Xu, 2013). Additionally, large industrial and power companies routinely hedge future sales to lock in profits. In the EU ETS for instance, the latter's carbon hedge is usually formulated as a percentage of future emissions generated by forward power sales, viz. 80%, 40% and 20% of expected emissions in the following year, two and three year's time, respectively (e.g. Eurelectric, 2009). Hedging against fluctuations (and the expected long-term increase) in carbon prices can be achieved through the banking of permits or the purchase of derivative contracts or both.

However, the intertemporal dimension of the EU ETS is characterized by a 'missing market' issue and by firms' constrained resources and limited capability to optimize over long horizons. The missing long-term markets is a well-identified issue in power markets, potentially causing under-investment in generation capacities (e.g. Newbery, 2018). It has a direct parallel with the carbon market, in which the time span of futures contracts can be construed as an upper bound on firms' horizon. As Figure 2 shows, maturities extend from the end of the current year up to ten years in the future, and liquidity quickly decreases with time-to-maturity. For instance, trade in EU emission allowances (EUA) futures contracts beyond a 3-year maturity is marginal, which aligns with the hedging horizon of future power sales.

Firms' horizon and ability to cost minimize over time can also be limited by financial resources and behavioral aspects. For instance, financial constraints on leverage and willingness to tie up capital in banked permits are discussed in Bredin & Parsons (2016) and can affect firms' compliance and market behaviors (Dardati & Riutort, 2016). Additionally, one may expect shorter horizons and less sophisticated intertemporal abatement and trading strategies from smaller firms which lack in-house expertise and trading desks, see Baudry et al. (2020) for a review. For instance, firms can have little leeway in deviating from standard risk management procedures (Schopp et al., 2015) while others perceive the ETS as a command-and-control policy rather than a compliance market (Venmans, 2016). This restrains intertemporal considerations and efficiency gains (e.g. Martin et al., 2015; Schleich et al., 2020).

More fundamentally, regulatory uncertainty prevails in the EU ETS, which has a bearing on market outcomes (e.g. Koch et al., 2016; Salant, 2016).[13] As a result, market participants may excessively if not exclusively focus on the short term, because of longer-term credibility issues or simply owing to the fact that regulation is never set in stone and subject to government interventions (Kydland & Prescott, 1977). Figure 3 depicts the EU ETS regulatory timeline, illustrating that the ETS Directive (the central legal instituting piece) is an evolving process undergoing important amendments on a regular basis, every 5 to 10 years at most. It is also noteworthy that the prevailing Directive establishes market design over an horizon over of a dozen years or so, suggesting a de facto upper bound on firms' horizon of a regulatory nature. Credibility issues may also arise when participants doubt the regulator's ability to intervene to 'fix the market',[14] or when regulatory language is vague.[15] This is crucial in a compliance

[13]Analyzing regulatory and policy announcements as price drivers in the EU ETS, Koch et al. (2016) detect a less pronounced sensitivity to the associated longer-term implications. They note that the relatively limited economic significance of announced long-term targets might be due to two interrelated aspects: the difficulty for regulators to commit to such long-term targets and the short foresight horizon of market participants.

[14]A clear example is a 2014 sentiment survey of market participants which found that 30% of respondents thought the ETS would no longer exist after 2020 (Thomson Reuters, 2014). See also The Economist (2013).

[15]A clear example of vagueness in language in our case is that the add-on cancellation mechanism might in fact not cancel permits or even be enforced. The first source of vagueness relates to the practical interpretation of the notion of permit validity vs. cancellation as the official terminology is 'no longer be *valid*'. The second source of vagueness relates to the implementation of the measure itself, as it might be '*decided otherwise* in the first review' in 2021, see (23) in Directive 2018/410 (European Parliament & Council, 2018). Relatedly, there are discrepancies within the Directive itself ('*should* no longer be valid' in (23) vs. *shall* in Art. 2(5a))

market where the value of emissions rights is purely extrinsic and imposed by the regulator.

3 Model

We consider a competitive emissions trading system with full banking and limited borrowing of emissions permits over time. Time is discrete and indexed by $t \geq 1$. The system starts at date 1 and compliance is required at each date t. The regulator sets a cap on system-wide emissions for each date t, which consists of freely allocated and auctioned permits f_t and a_t. Some fixed quantity of offset credits O may also be surrendered for compliance over a given time period and o_t denotes the volume of offsets used at date t with $\sum_t o_t \leq O$. At each date t, we assume that regulated firms fully acquit their compliance obligations, i.e. they remit as many permits and offsets as needed to cover their current emissions.[16]

The decentralized market equilibrium can be characterized indirectly as the solution to joint compliance cost minimization among all firms (Montgomery, 1972; Rubin, 1996).[17] Because our focus is on the temporal dimension of the system, we abstract from its spatial trading dimension and take the perspective of the regulated perimeter as a whole, hereafter the firm. This representative firm approach has been used extensively in the literature (e.g. Schennach, 2000; Ellerman & Montero, 2007; Fell, 2016; Lintunen & Kuusela, 2018; Kollenberg & Taschini, 2019). We let e_t and u_t denote the firm's levels of realized and unregulated (baseline) emissions at date t, respectively. End-of-pipe abatement $u_t - e_t \geq 0$ is costly and we let C_t denote its minimum total abatement cost function at date t with $C'_t, C''_t > 0$.[18]

Baseline emissions u_t hinge on business cycle fluctuations (Koch et al., 2014; Bel & Joseph, 2015; Chèze et al., 2020) and the variable performances of companion policies, i.e. complementary climate and energy policies which can affect counterfactual emission levels independently of the permit price (Burtraw & Keyes, 2018; Borenstein et al., 2019). On the supply side, there exist small discrepancies between announced and realized cap levels, future supply $\{f_t\}_t$ and $\{a_t\}_t$ can be affected by regulatory changes, and $\{o_t\}_t$ depends on external offset market conditions (de Perthuis & Trotignon, 2014; Ellerman et al., 2016). In the following, we use

and with a more recent Communication which resorts to *will* (European Commission, 2018a).

[16]Penalties incurred for various permit and self-reporting violations are implicitly assumed to be adequately designed and enforced. See Stranlund et al. (2005) for a discussion.

[17]With competitive trading, this is warranted only when firms' baseline emissions are private information, which can be aggregated via the market by a sufficient statistic (the equilibrium price), while informational efficiency breaks down when firms' abatement costs are private information (Cantillon & Slechten, 2018).

[18]Formally, the representative firm's abatement cost function C_t is the envelope of all firm-level abatement cost functions $C_{i,t}$, i.e. its derivative C'_t obtains by horizontal summation of the $C'_{i,t}$'s (Montgomery, 1972).

the tilde notation to signify those quantities that are uncertain in nature for the firm.

We next describe the properties of the competitive intertemporal market equilibrium before introducing in turn rolling horizon planning, supply-side control and bounded responsiveness to the control, along with the associated resolution procedures. We also discuss the relative implications and intertemporal behaviors under infinite and rolling finite horizons.

3.1 Competitive intertemporal equilibrium

The firm's annual demand for permits has two components: one for contemporaneous compliance, the other reflecting intertemporal arbitrage.[19] Specifically, at any date t, given the prevailing permit price p_t and realized baseline u_t, the firm's emission level, or demand for permits for compliance $e_t^\star(p_t, u_t)$, satisfies the usual first-order necessary condition

$$C_t'(u_t - e_t^\star(p_t, u_t)) - p_t = 0. \tag{1}$$

In addition, the firm can over-comply relative to total contemporaneous available supply and carry over (i.e. bank) unremitted permits into future dates. Likewise, it can under-comply and front-load (i.e. borrow) yet-unallocated permits from its future self and use them to achieve full compliance today. While banking is unlimited, we consider that borrowing at date t is authorized up to a limit $l_t \geq 0$. As the firm cost minimizes over time, limited intertemporal trading opportunities imply a no-arbitrage condition closely following the rationale of competitive commodity storage with negligible storage costs and no stock depreciation over time (e.g. Wright & Williams, 1982; Deaton & Laroque, 1992).[20] Permit banking, whose level at date t is denoted b_t, thus constitutes the second determinant of permit demand and satisfies the following two conditions with complementary slackness

$$b_t + l_t \geq 0 \quad \perp \quad p_t - \beta \mathbb{E}_t\{p_{t+1}\} \geq 0, \tag{2}$$

where $\mathbb{E}_t\{\cdot\}$ denotes expectation conditional on all information available to the firm at date t and $\beta = (1+r)^{-1}$ is the firm's discount factor with r the discount rate, possibly inclusive of a

[19] We do not explicitly account for forwards and futures as the aggregate demand for such bilateral contracts is nil in equilibrium, so p_t is de facto the date-t spot price (Laffont & Tirole, 1996; Seifert et al., 2008). Note, however, that (2) can legitimately hold provided an active market for futures exists (Pindyck, 1993). Besides, prices on the primary and secondary markets have no reason to differ given the uncertainty structure assumed here, specifically u_t, f_t, a_t and o_t are revealed at the beginning of date t (see e.g. Kling & Rubin, 1997).

[20] Schennach (2000) first pointed out the tight connection between commodity storage and permit banking. One key difference is that stockouts, which here coincide with positive borrowing ($b_t < 0$), are not feasible.

permit-specific risk premium.[21] As long as $\beta\mathbb{E}_t\{p_{t+1}\} > p_t$, banking is profitable and increases date-t permit demand, which raises p_t and lowers $\mathbb{E}_t\{p_{t+1}\}$ until all arbitrage opportunities are exhausted and the firm breaks even, i.e. when the cost-of-carry price coincides with the spot price grown at the discount rate ($\beta\mathbb{E}_t\{p_{t+1}\} = p_t$). Conversely, as long as $\beta\mathbb{E}_t\{p_{t+1}\} < p_t$, borrowing is profitable but only authorized up to l_t. As soon as this constraint is binding, the connection between current and expected future prices ceases to hold and the price rises at a rate less than the discount rate. In sum, the price can rise at a rate at most as high as the discount rate in a rational expectations equilibrium.

Note that banked (resp. borrowed) permits add to (resp. subtract from) future permit supply, i.e. total available supply at date t amounts to $f_t + a_t + o_t + b_{t-1}$. Market clearing at date t, which implies that total supply equalizes total demand, thus reads

$$f_t + a_t + o_t + b_{t-1} = e_t + b_t. \tag{3}$$

Combining the compliance, no-arbitrage, and market-clearing conditions in (1), (2) and (3) then leads to two regimes in the equilibrium price and emission dynamics

$$p_t = \max\{\beta\mathbb{E}_t\{p_{t+1}\};\, C_t'(u_t - (l_t + f_t + a_t + o_t + b_{t-1}))\}, \tag{4a}$$

$$e_t = \min\{e_t^\star(\beta\mathbb{E}_t\{p_{t+1}\}, u_t);\, l_t + f_t + a_t + o_t + b_{t-1}\}, \tag{4b}$$

$$\text{with } b_t = f_t + a_t + o_t + b_{t-1} - e_t \geq -l_t. \tag{4c}$$

The first regime features intertemporal flexibility in the firm's emission stream over time in line with Hotelling's rule in expectation ($b_t > -l_t$).[22] In the second regime the borrowing constraint is binding ($b_t = -l_t$) so that permit demand becomes solely determined by annual compliance requirements in (1) and emissions are pegged to the contemporaneous amount of permits on hand. Constrained borrowing implies an asymmetric demand shock dampening potential for the market,[23] and induces a non-linearity so that no closed-form solution to (4) can be derived.[24] Our resolution procedures are described in the following sections.

[21]In the spirit of the Capital Asset Pricing Model, Schennach (2000), Perino & Willner (2016) and Kollenberg & Taschini (2016, 2019) introduce an exogenous risk premium on top of the risk-free rate to capture the impacts of future emissions uncertainty on the firm's decisions. See Tietjen et al. (2021) for a derivation of an endogenous premium formally accounting for firms' hedging demand and investment decisions.

[22]This regime is always finite in time but may not be unique under uncertainty (see Schennach, 2000).

[23]In principle, the firm can entirely smooth out the price impact of a downward demand shock (temporary glut) by stockpiling more permits while it can only be so for a symmetric upward shock (temporary shortage) to the extent that the bank is not too small. See also Kollenberg & Taschini (2019) for a discussion.

[24]Deaton & Laroque (1992) use a fixed-point approach in a similar commodity storage problem, for which

3.2 Introducing rolling finite horizons

At the beginning of date t, u_t, f_t, a_t and o_t are given and known to the firm, which also keeps track of the state variable, i.e. the bank b_{t-1} (initialized with $b_0 = 0$).[25] The firm selects its date-t emission e_t and implied bank b_t by minimizing its expected present value of compliance costs. That is, with an infinite horizon, the firm solves the following program

$$\min_{\{e_\tau\}_{\tau \geq t}} \mathbb{E}_t \Big\{ \sum_{\tau \geq t} \beta^{\tau - t} C_\tau(\tilde{u}_\tau - e_\tau) \Big\} \tag{5a}$$

$$\text{subject to } \; 0 \leq e_\tau \leq \tilde{u}_\tau, \tag{5b}$$

$$\text{and } \; b_\tau = b_{\tau-1} + \tilde{f}_\tau + \tilde{a}_\tau + \tilde{o}_\tau - e_\tau \geq -l_\tau, \tag{5c}$$

where (5b) contains feasibility constraints for the emission path and (5c) describes the law of motion for the state variable (i.e. annual market clearing), where the constraint on the bank ensures the cumulative emissions cap is fulfilled (i.e. overall market clearing).

To deal with uncertainty the firm may choose to optimize over a finite horizon $h \geq 0$ within which it is confident about its forecasts of future exogenous variables (see Section 2).[26] Hence it selects its date-t emission e_t and implied bank b_t by solving the following program

$$\min_{\{e_\tau\}_{\tau=t}^{t+h}} \sum_{\tau=t}^{t+h} \beta^{\tau-t} C_\tau(\hat{u}_\tau^t - e_\tau) \tag{6a}$$

$$\text{subject to } \; 0 \leq e_\tau \leq \hat{u}_\tau^t, \; b_\tau = b_{\tau-1} + \hat{f}_\tau^t + \hat{a}_\tau^t + \hat{o}_\tau^t - e_\tau \geq -l_\tau, \tag{6b}$$

$$\text{and } \; \sum_{\tau=t}^{t+h} \Big[\hat{u}_\tau^t - e_\tau \Big] \geq \sum_{\tau=t}^{t+h} \Big[\hat{u}_\tau^t - (\hat{f}_\tau^t + \hat{a}_\tau^t + \hat{o}_\tau^t) \Big] - b_{t-1}, \tag{6c}$$

where $\hat{x}_\tau^t$ denotes the date-t forecast for $x = \{u, f, a, o\}$ at date $\tau \geq t$. Although redundant, (6c) is added for clarity: it specifies the firm's date-t assessment of the system stringency over the horizon h, i.e. the sum of forecasted individual raw abatement efforts $\hat{u}_\tau^t - (\hat{f}_\tau^t + \hat{a}_\tau^t + \hat{o}_\tau^t)$ corrected for the initial bank b_{t-1}. Then, (6a) dictates that this forecasted overall abatement effort be spread over the horizon in accordance with the equimarginal value principle.

they prove that with time-independent linear consumption demand (which translates to time-independent linear marginal abatement cost functions here) there is a unique stationary rational expectations equilibrium.

[25] Offset usage is assumed exogenous to the firm's problem. In Section 4.3 we explain how we tackle offset usage for the ex-post model calibration and why this assumption is innocuous for our analysis in Section 5. See Koch et al. (2017) for a treatment of joint permit and offset usage decisions.

[26] We treat the horizon length h as given in the model and calibrate it in Section 4.4 so as to best replicate annual 2008-2017 market outcomes. As discussed in Section 2, h could also endogenously emerge as a result of informational barriers and computational cost containment (see e.g. Reis, 2006).

Under a rolling horizon, the firm solves for the equilibrium path from date t up to $t+h$ given its current forecasts $\{\hat{x}_\tau^t\}_\tau$, but only implements the first date of the plan, which pins down the state variable for the next date. At date $t+1$ the firm revises its forecasts based on new information about x and initiates a new planning cycle from date $t+1$ to $t+h+1$ taking the state variable b_t as given.[27] This date-on-date solving and updating of finite plans and the sequential execution of the first date of these plans then unfolds over time. Technically, a rolling horizon shrinks the firm's problem dimensionality from uncertainty to certainty given appropriate forecasts. Note also that the current plan is not contingent on future plans, as this would otherwise call for high informational requirements and computational complexity, perhaps even more so than for an infinite horizon.

Last but not least, note that terminal conditions matter with a finite horizon. In our setting, they can be construed as the belief about the continuation value of banked permits as valued by the market at the end of the horizon. Because of informational constraints, this value can reasonably only be based on an educated guess at best, which is unlikely to correspond to the equilibrium permit price at this date that would obtain under an infinite horizon (see Section 2.1). Thefore, program (6) assumes a zero salvage value for simplicity and without loss of generality for our results.[28]

3.3 Infinite versus rolling finite horizons

Intuitively, as h grows, the equilibrium paths obtained with a rolling horizon should converge to those with the infinite horizon. In our setting, we need two additional assumptions to arrive at this convergence result and ensure a fair comparison between rolling and infinite horizon outcomes. First, we derive the expected equilibrium paths under the infinite horizon invoking a first-order approximation along certainty-equivalent paths for the exogenous variables x.[29] This shrinks the dimensionality of the infinite horizon problem, making it similar to that of a rolling finite horizon. Second, we impose that any date-t certainty-equivalent paths coincide

[27] We consider that the firm can revise its forecasts and initiate a new planning cycle at each date, although in practice the re-planning periodicity can be a choice variable for the firm. See Sections 4.2 and 4.3 for a description of future demand and supply forecast heuristics and how they are updated.

[28] A non-zero salvage value affects our quantitative results in a predictable way: a positive value implies higher price and bank paths, and vice versa. As a result, when assessing the 2018 reform impacts in Section 5, a positive salvage value would induce larger permit withdrawals by the market stability reserve.

[29] This approximation was suggested but not operationalized by Schennach (2000). The certainty equivalence property entails that up to a first-order approximation optimal decisions at date t coincide with those under full information provided that random variables equal their date-t expected values. It naturally comes about with linear marginal abatement cost functions.

with the corresponding forecasts over the relevant time span, i.e. $\hat{x}_{\tau}^{t} = \mathbb{E}_t\{\tilde{x}_{\tau}\}$.

Specifically, by chaining (4a) over time and invoking the tower rule, the date-t expected price path under an infinite horizon satisfies $\mathbb{E}_t\{p_{t'} - \lambda_{t'}\} = \beta^{t-t'} p_t$ for any $t' > t$, where $\lambda_{t'} \geq 0$ is the Lagrange multiplier associated with the limited borrowing constraint at t', i.e. $b_{t'} \geq -l_{t'}$. When there is a zero probability of a binding constraint in the future, i.e. $\mathbb{E}_t\{\lambda_{t'}\} = 0$ for all $t' \in [\![t_1; t_2]\!]$, the date-t expected price grows at rate r over $[\![t_1; t_2]\!]$. When this probability is positive, i.e. $\mathbb{E}_t\{\lambda_{t'}\} > 0$ for some $t' \in [\![t_1; t_2]\!]$, the expected price rises at a rate less than r (which is not uniquely pinned down) with the downward offset rising over $[\![t_1; t_2]\!]$. When it is unity, the expected price is uniquely determined by $\mathbb{E}_t\{p_{t'}\} = \mathbb{E}_t\{C_{t'}(\tilde{u}_{t'} - (\tilde{f}_{t'} + l_{t'} + \tilde{a}_{t'} + \tilde{o}_{t'}))\}$. An *expected* equilibrium path thus exhibits the above three-regime dynamics while the *actual* equilibrium path only features the two regimes described in (4).

Schennach (2000) proposed to invoke the certainty equivalence property and solve (5) assuming that the expected equilibrium paths follow the same two-regime dynamics as the actual paths in (4). We follow the suggested approach. That is, we by construction only consider the size of the bank in expectation. Therefore, a second-order bias between our approximate and the exact expected paths arises as we do not capture the possibility that $\mathbb{E}_t\{\lambda_{t'}\}$ may become positive for some $t' > t$ although the bank is still expected to satisfy the limited borrowing condition by that time. Schennach (2000) also showed that the approximate expected price path is slightly biased downward in the first regime and early on in the second one, biased upward for the rest of it, and unbiased in the third regime.

To develop insights into the relative implications of infinite versus rolling finite horizons, we describe and compare the associated market outcomes in a simple setup shutting down other components that might come into play. Specifically, we assume perfect foresight on the part of the firm, no borrowing ($l_t = 0$), and time invariant abatement cost functions to single out the impacts of different time horizons and discount rates. When emission caps are declining over time and increasingly binding relative to baseline emissions (Figure 4c), it is rational for the cost-minimizing firm to curtail emissions below the cap early on and accumulate a bank of permits which is drawn down later on when the cap becomes more stringent (Figure 4b). As Figure 4a shows, the associated price path rises at the discount rate as long as the bank is positive, and then it equals the marginal cost of meeting the cap (no borrowing).

The optimal price, bank and emission paths depend on the combination of the firm's horizon and discount rate. Specifically, for h (resp. r) given, the larger r (resp. smaller h) the shorter the banking period and the lower the banked volumes at all points in time. Accordingly, the permit price is lower early on but rises at a higher rate over time during the banking period

Figure 4: Qualitative implications of infinite vs. rolling horizons

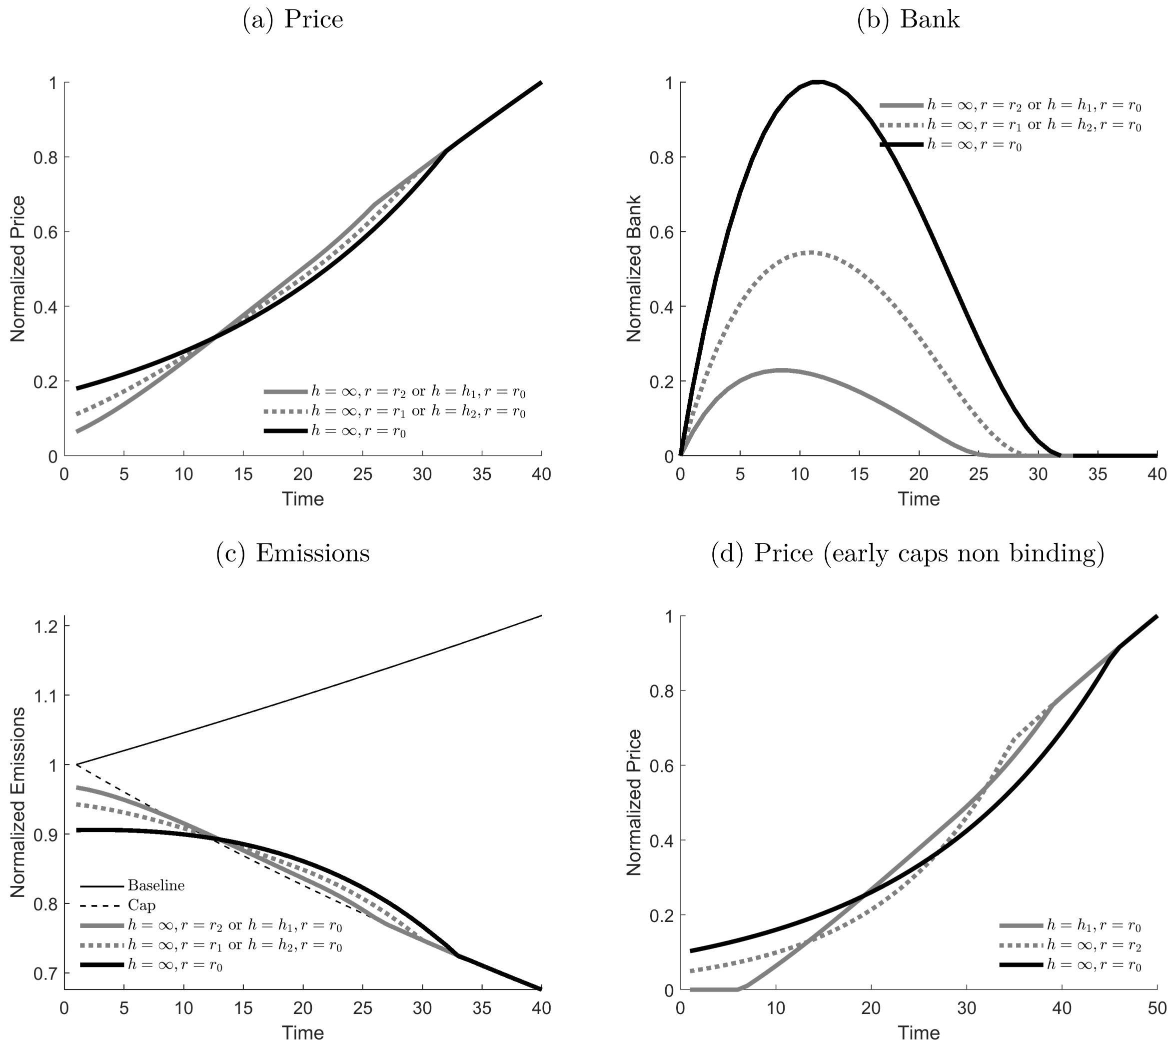

Note: All cases assume $r_2 > r_1 > r_0 > 0$, $0 < h_1 < h_2 < \infty$, $l_t = 0$ and $C''_t = c$ for all t for some positive scalar c. Figs. 4a–4c only feature binding caps, while early caps are not binding in Fig. 4d.

as it is the vehicle that equalizes supply and demand over the horizon. This exemplifies that the intertemporal behavior of a firm using a rolling finite horizon is observationally equivalent to that of a firm using an infinite horizon with a higher discount rate, all else equal.

Crucially, however, this qualitative equivalence result only holds when (1) caps are binding and (2) there is no supply-side control. Regarding (1), Figure 4d shows how the equivalence breaks down when early caps are not binding: when the horizon is sufficiently short, the firm may perceive no overall constraint on emissions implying that the price is constant and nil for some time, while with an infinite horizon and a larger discount rate the price exhibits the same behavior as in Figure 4a.[30] Regarding (2), we will explore the interaction between the firm's horizon and discount rate and a specific supply control in Section 5.

3.4 Introducing supply control via a banking corridor

The regulator can embed some resilience into its system via supply-side control (e.g. Roberts & Spence, 1976; Pizer, 2002). These controls adjust supply schedules based on the value of a given market indicator, typically the permit price, relative to some predefined thresholds.[31] This usually requires the creation of a reserve of permits whose stock at date t is denoted by $s_t \geq 0$. We here consider a banking corridor automatically adjusting current auctions a_t based on banking history $\{b_\tau\}_{\tau<t}$ according to

$$a_t \longleftarrow a_t - \min\Big\{a_t; R \cdot \sum_{\tau<t} \mathbb{1}\{b_\tau > \bar{b}\} x_\tau b_\tau\Big\} + \min\{I; s_{t-1}\} \cdot \sum_{\tau<t} \mathbb{1}\{b_\tau < \underline{b}\} x_\tau, \tag{7}$$

where $\mathbb{1}\{\cdot\}$ is the indicator function, $\underline{b} > 0$ and $\bar{b} > \underline{b}$ lower and upper bank thresholds, $I > 0$ an injection quantity, $R \in [0;1]$ an absorption rate and historical weights $\{x_\tau\}_{\tau<t}$ are such that $x_\tau \in [0;1]$ for all $\tau < t$ and $\sum_{\tau<t} x_\tau = 1$. In parallel, the stock of permits stored in the reserve follows the complementary dynamics

$$s_t = s_{t-1} + \min\Big\{a_t; R \cdot \sum_{\tau<t} \mathbb{1}\{b_\tau > \bar{b}\} x_\tau b_\tau\Big\} - \min\{I; s_{t-1}\} \cdot \sum_{\tau<t} \mathbb{1}\{b_\tau < \underline{b}\} x_\tau. \tag{8}$$

In words, when $b_{\tau<t}$ is above $\bar{b}$, a predefined share $x_\tau R$ thereof is withheld from auctions at date t and placed in the reserve. Symmetrically, when $b_{\tau<t}$ is below $\underline{b}$ and the current stock of the reserve allows, a fixed quantity of reserve permits $x_\tau I$ is added to auctions at date t. Otherwise, the banking corridor is inactive. Because the change in auctions is determined by the banking history it is fixed once and for all at the beginning of each date.

The banking corridor defined in (7–8) may be expected to preserve the cumulative emissions cap as it essentially rearranges the auction schedule over time. In principle, after a certain period of time the bank will pass below $\underline{b}$ (see Figure 4b) implying that the reserve should gradually empty. That is, all reserve permits should eventually return to the market so that

[30] See Spiro (2014) for an extensive discussion in the context of exhaustible resource extraction.

[31] A standard approach to controlling supply consists in introducing steps in otherwise vertical (infinitely price inelastic) supply schedules and typically takes the form of a price corridor, i.e. a combination of a price floor and ceiling (see inter alia Grüll & Taschini, 2011; Fell et al., 2012; Abrell & Rausch, 2017).

cumulative emissions are unchanged (Perino & Willner, 2016).[32] This ceases to be the case if the corridor is equipped with an add-on cancellation mechanism that invalidates permits stored in the reserve in excess of a certain threshold. In this case, the reserve stock is further adjusted such that

$$s_t \longleftarrow s_t - \max\{0; s_t - k_t\}, \tag{9}$$

where $k_t \geq 0$ is the maximum number of permits that can be stored in the reserve at date t. The cumulative emissions cap is endogenized and becomes a market outcome. Whether and by how much it is reduced depends on the thresholds $\{k_t\}_t$ and the equilibrium dynamics of the reserve stock, which itself hinges on the initial bank and reserve stock conditions and is ultimately governed by the firm's horizon, discount rate and responsiveness to the control.

3.5 Introducing bounded responsiveness to the control

The key decision quantity that the firm has to appraise at each date is its required cumulative abatement effort over its planning horizon in (6c). It hinges on the firm's future permit supply and demand forecasts and, crucially, further interacts with supply control. We consider two polar degrees of responsiveness on the part of the firm in foreseeing and understanding this interaction, viz. zero and full responsiveness. In the former case, the firm does not account for control impacts on future supply when evaluating (6c). At each date, the unresponsive firm simply discovers and factors in control impacts on present auction volumes but remains completely oblivious as to what future control impacts will be.[33] In the latter case, the fully responsive firm comprehends the interplay between its intertemporal decisions and associated control impacts over time, which yields the rational competitive equilibrium (Muth, 1961).

Our indirect approach to solving for the recursive competitive equilibrium as the outcome of a planning problem for a representative firm is valid under laissez faire (Samuelson, 1971). However, Salant (1983) showed in the related context of commodity storage that this could mischaracterize the rational expectations equilibrium in the presence of supply-side controls.[34] That is, forward-looking rational agents may take advantage of the control rules which the regulator adheres to, which may engender speculative attacks on the scheme and result in

[32]In our simulations in Section 5, this conventional result that cumulative emissions are preserved is not recovered – it requires several assumptions that are not satisfied herein, see e.g. footnote 59.

[33]Such an extreme form of bounded rationality could reflect the difficulty for an individual firm to fathom and respond to the future supply-side implications of a banking corridor, which are based on aggregate bank levels, at least in comparison with the more straightforward price signals conveyed by a price corridor.

[34]See Hasegawa & Salant (2015) for a transposition of Salant's original point to emissions trading.

policy failure instead of price stabilization. Under a banking corridor for instance, firms would collectively like the policy handle (the bank) to fall below the intake threshold for the control to induce a minimal contraction in overall supply, if any.[35] Individually, however, they have a negligible impact on the bank and cannot coordinate their banking decisions to 'game the system'.[36] In a competitive equilibrium, therefore, supply control cannot alter intertemporal efficiency (i.e. the equalization of discounted expected marginal abatement costs across firms and periods) although it does have an impact on market clearing at each date.

To derive the recursive competitive equilibrium in the full responsiveness case, we solve for the fixed point of the mapping between the firm's beliefs about the control impact profile and optimal beliefs.[37] That is, the equilibrium obtains when the firm's belief coincides with the actual law of motion for the control actions generated by the firm's intertemporally effective choices induced by this belief. At each step in this procedure, the firm has a forecast for both the control action and supply profiles, evaluates its required cumulative abatement effort, and cost minimizes over time. Notice that the associated first-order conditions are congruent with those obtained in a competitive equilibrium.

Specifically, starting from the zero responsiveness case at any date t with horizon h, the firm first derives the date-t expected equilibrium path by recursively solving the program

$$\min_{\{e_\tau\}_{\tau=t'}^{t'+h}} \sum_{\tau=t'}^{t'+h} \beta^{\tau-t'} C_\tau(\hat{u}_\tau^t - e_\tau) \tag{10a}$$

$$\text{subject to } \; 0 \leq e_\tau \leq \hat{u}_\tau^t \text{ and } b_\tau = b_{\tau-1} + \hat{f}_\tau^t + \hat{a}_\tau^t + \hat{o}_\tau^t - e_\tau \geq -l_\tau. \tag{10b}$$

Program (10) is solved for each $t' \in [\![t; t+h]\!]$ iteratively: only the first-date optimal outputs for each t' are implemented, assuming that (1) the firm's date-t forecasts are kept unchanged throughout and materialize as forecasted; and (2) the initial bank condition at t' (i.e. $b_{t'-1}$) is set by the previous optimization round. This iterative process captures that the responsive firm is aware of its rolling horizon and accounts for it when computing the date-t expected equilibrium path. Throughout this process, the control affects contemporaneous auctions as per (7), although future control impacts remain unforeseen to the firm.

[35] Specifically, banking levels below intertemporally efficient ones could yield a smaller supply tightening and hence lower compliance costs, but would not conform to a competitive equilibrium. See Stocking (2012) for an example of manipulation under a hard price ceiling where purchasing permits at the ceiling price when the market price is lower is not cost efficient but may reduce the overall compliance cost by relaxing the cap.

[36] See Carlén et al. (2019) for an example where firms are assumed to be able to coordinate. Such strategic behavior cannot occur in our competitive framework (here under full responsiveness).

[37] Such a fixed-point procedure is hardly new. Lucas & Prescott (1971) were the first to implement one to determine a rational expectations equilibrium in a Bellmanized indirect planning optimization program.

This process yields a path of permit flows into and out of the reserve for all t' as per (8) and (9). The firm then adjusts its beliefs about the future auctions stream based on the obtained flow path, and repeats the above process. If this is not neutral vis-à-vis the previously optimal emission and bank paths, the firm will revise them. This changes the control impact profile, which in turn again affects the firm's intertemporal decisions, and so forth. At each step in this procedure the firm holds beliefs about future control impacts, behaves rationally with respect to these beliefs, and updates them in between each step. Gradually, a fixed point is attained, yielding the date-t competitive equilibrium in the full responsiveness case.

4 Calibration to the EU ETS

In this section, we first set the model parameters to align with the EU ETS design. We next describe how we use historical data and regulatory texts to construct the permit demand and supply schedules, as well as the corresponding forecasts the firm forms and revises over time. We finally calibrate the firm's discount rate, planning horizon and marginal abatement cost parameter to fit the observed annual banking and price dynamics over 2008–2017.

4.1 Market design

In the EU ETS, compliance is required on a calendar-year basis, by April 30$^{\text{th}}$ of the following year. Therefore, each date t in the model coincides with a calendar year. Banking within and across trading phases is authorized and there is no vintage restriction on the temporal validity of issued permits for future compliance since 2008 (European Parliament & Council, 2003). Year-on-year borrowing is tacitly authorized since freely-allocated year-t vintage permits are issued two to four months prior to the year-$(t-1)$ compliance deadline. Hence we set $l_t = f_{t+1}$ as firms can effectively tap into their allotment for two consecutive years.

We parametrize the banking corridor to conform with the features of the Market Stability Reserve as adopted in 2018 (European Parliament & Council, 2018). The MSR begins operations in 2019 and is initially seeded with backloaded permits over 2014-2016, i.e. $s_{2018} = 900$ million (European Commission, 2014). Non-allocated Phase-III permits up to 2017 are also placed in the reserve in 2021, the number of which we estimate at 581 million.[38] The MSR

[38]According to the European Commission (2015) the expected number of unallocated Phase-III permits is between 550 and 700 million. Discrepancies between announced and realized annual supply arise mostly due to non-distributed permits from the New Entrants Reserve and Article 10(c), or from plant closures and production capacity changes. We are thus in the lower range of the estimate as we assume 2018–20 allocations

thresholds are set at $\bar{b} = 833$ million and $\underline{b} = 400$ million. Moreover, we set $I = 100$ million and $R = 0.24$ until 2023 with $R = 0.12$ afterward (European Commission, 2018a).[39] Finally, we set the historical weights such that $x_{t-2} = 2/3$ and $x_{t-1} = 1/3$ to account for the mismatch between the MSR and compliance calendars.[40]

A cancellation mechanism as defined in (9) is active from 2023 on and operates such that any permits stored in the reserve in a given year in excess of the number of auctioned permits in the previous year are invalidated (European Parliament & Council, 2018). Hence we set $k_t = a_{t-1}$, where a_{t-1} is endogenously determined through the MSR as per (7). Finally, since the reform was finalized in late 2017 and enacted in early 2018, we consider that the impacts of the MSR and the cancellation mechanism on annual auction volumes can be anticipated and factored in by the responsive firm from 2018 on.

4.2 Permit demand

Our aim is to construct a counterfactual scenario for baseline CO_2 emissions for the EU ETS perimeter, i.e. emissions as they would be in the absence of the scheme but accounting for industrial production growth and all complementary energy and climate policies in place. To that end we use a simple decomposition of baseline CO_2 emissions into three Kaya indexes

$$CO_2\,emissions = \underbrace{Production}_{economic\ activity} \times \underbrace{\frac{Energy}{Production}}_{energy\ intensity} \times \underbrace{\frac{CO_2\,emissions}{Energy}}_{carbon\ intensity}. \tag{11}$$

Figure 5a depicts the reconstructed and projected trajectories of these three indexes between 1990 and 2050. We assume that the permit price has negligible impacts on both production and energy intensity.[41] Specifically, we compute the production index ex post from Eurostat sector-level production data and consider a 1% p.a. production growth from 2018 on.[42] Next, we compute the energy intensity index ex post from total final energy consumption time series for a proxy of regulated sectors (viz. electricity, heat, industry, and energy industry own use

to coincide with the announced levels but note that the exact number of non-issued permits or date at which they enter the reserve has negligible impacts on our quantitative results.

[39] Since $R \ll 1$, a dozen iterations typically suffice for the fixed-point procedure in Section 3.5 to converge.

[40] The official bank value, or 'total number of allowances in circulation', for year $t-1$ is published in May of year t, and is used for MSR operations over a twelve-month period from September 1st of year t onwards.

[41] Both assumptions are reasonable. Regarding production, there is no ex-post evidence of carbon leakage in the EU ETS (Joltreau & Sommerfeld, 2019; Naegele & Zaklan, 2019). Regarding energy intensity, Figure 5a shows that it declines less over 2005–2015 with the ETS in place than prior to the ETS (1990–2005).

[42] Link to Eurostat data (`http://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=sts_inpr_a`).

Figure 5: Kaya indexes, baseline emissions and total cap on emissions

(a) Kaya Indexes (b) Supply & Demand

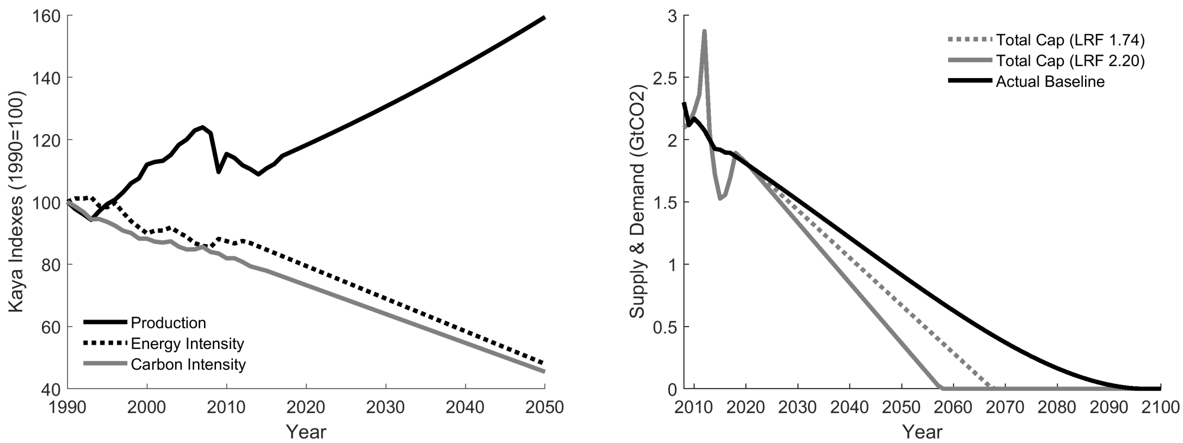

Note: The amounts by which the cap declines yearly correspond to the LRF multiplied by the 2010 emissions of the covered perimeter in Phase III: 38.3 and 48.4 million under a LRF of 1.74% and 2.20%, respectively.

and losses) obtained from the 1990–2015 balance sheets of the International Energy Agency.[43] From 2016 on we utilize a linear interpolation ensuring that EU energy efficiency targets for 2020 and 2030 are met, and we assume this linear trend to be valid afterward.

Next, we compute the carbon intensity index ex post by reconstructing regulated CO_2 emissions over 1990–2004 based on IEA primary energy consumption data and standard EU-level emission factors.[44] From 2005 on, we need to calculate emissions as they would have been absent the ETS. We assume that the permit price may have driven some fuel switching (thus impacting the carbon content of energy) but not the development of renewable energy.[45] To isolate and account for renewable deployment while neutralizing fuel switching in the baseline emissions, we fit a linear relationship between renewable deployment and the carbon content of energy prior to Phase II. We extrapolate this first-pass relationship for later years using observed renewable deployment for 2008–2015 and then assuming that EU renewable targets set for 2020 and 2030 (and their continuation afterward) are attained linearly. Computing the ratio of CO_2 emissions to energy consumed finally yields the carbon intensity index.

Thus, baseline emissions at any point in time are by construction independent of the history of permit prices.[46] Graphically, the black line in Figure 5b depicts the resulting baseline path,

[43] Link to IEA data (`https://www.iea.org/statistics/?country=EU28`).

[44] Link to IEA data. Emission factors are 4.2, 3.1 and 2.4 tCO_2/toe for coal, oil and gas, respectively.

[45] This is supported by evidence that the EUA-price equivalent of renewable subsidies has been significantly higher than EUA prices (Marcantonini & Ellerman, 2015; Marcantonini & Valero, 2017; Abrell et al., 2019).

[46] This independence assumption becomes less tenable when prices reach higher levels than those observed

Table 1: Forecasted trends of baseline emissions

Forecast period	Climate Energy Package	$\bar{u}_{2050}/e_{2008}$	$\bar{u}_t = 0$ in
2008–2013	CEP#1	57.5%	2115
2013–2017	CEP#2	50.7%	2105
2018–2100	Reinforced CEP#2	39.7%	2096

Note: The 2050 baseline trend $\bar{u}_{2050}$ is expressed as a percentage of 2008 emissions ($e_{2008} = 2.12$ GtCO$_2$).

which we obtain by plugging the evolution of the three Kaya indexes in (11). It is downward sloping, which is in line with the steady decline in ETS perimeter emissions observed prior to the launch of the ETS. Regarding demand forecasts, we assume that the firm uses a simple heuristic congruent with the deterministic part of an AR(1) process, allowing for growth and a varying trend.[47] That is, the year-t forecast for demand in year $t+1$ is defined by

$$\hat{u}_{t+1}^{t} = \varphi(1+\gamma_t)u_t + (1-\varphi)\bar{u}_{t+1}^{t}, \tag{12}$$

where u_t is the realized baseline in year t, $\varphi \in [0;1]$ captures some persistence in baselines and γ_t is the expected annual growth rate in year t for future years. We allow the trend $\bar{u}$ to vary over time. Typically, it can be thought of as declining over time due to the achievements of complementary policies outside the system's perimeter. Importantly, notice that realized baselines u_t differ from their past forecasts $\hat{u}_t^{\tau<t}$. This induces the firm to revise its demand forecasts on a yearly basis and to adjust its intertemporal decisions accordingly.

We assume that the trend in year t for some future year $t' > t$ (i.e. $\bar{u}_{t'}^{t}$) is set to align with the attainment of the Climate Energy Package objectives prevailing in year t. The computed trends are reported in Table 1. Additionally, we set γ_t in line with GDP growth forecasts by the European Commission over 2008–2017 and consider a 1% p.a. growth rate afterward.[48] For the persistence parameter, we finally follow Fell (2016) and use $\varphi = 0.9$.[49]

up to now. Note, however, that endogenous baselines are not considered in similar modeling approaches. Fell (2016) and Perino & Willner (2016, 2017) assume given baseline paths, respectively increasing and constant over time. Similarly, baselines in Beck & Kruse-Andersen (2020) are decreasing over time due to increasing renewable deployment, which slows down over time but nonetheless remains independent of the permit price.

[47] To facilitate calibration, the firm is assumed to use one practically relevant heuristic and to stick to it. In general, models of behavioral expectations consider a set of forecast rules with agents rationally switching from one to another based on their relative performances in the recent past (e.g. Hommes et al., 2019).

[48] Link to EC forecasts published in spring of year t for year $t+1$.

[49] In related contexts, Heutel (2012) and Lintunen & Kuusela (2018) use $\varphi = 0.95$ and $\varphi = 0.8$ respectively. Roughly speaking, the lower φ the more the firm expects future baselines to coincide with the trend. As the trend happens to be relatively close to actual baselines, a lower φ thus implies 'better foresight'.

4.3 Permit supply

In year t the firm observes annual permit supply $f_t + a_t$ and it forecasts future permit supply to coincide with the cap trajectory as given in currently prevailing regulatory texts (e.g. EU Directives, Decisions or Communications).[50] As soon as regulation is revised or upon release of actual supply data (e.g. EC Carbon Market Reports), the firm corrects its forecast. Hence yearly forecast updates and year-on-year discrepancies between forecasted and actual supply naturally come about. For instance, the firm factors in a total cap path from Phase IV (2021) based on the currently effective linear reduction factor, i.e. LRF=1.74% before vs. 2.20% after the 2018 reform, with 57% of the total cap auctioned off.

Usage of Kyoto offsets (viz. CERs and ERUs) is authorized in Phases II and III, up to the cumulative limit $O \approx 1.6$ $GtCO_2$.[51] Within that period, we assume that the firm does not decide how many offsets to surrender, but note this simplification is innocuous for our ex-ante analysis (i.e. from Phase IV on) in Section 5. Specifically, o_t is given at the beginning of year t and equal to observed offset usage. Moreover, the firm forecasts in year t that the remaining allowed quantity of offsets that can be surrendered from year $t+1$ on (i.e. $O - \sum_{\tau=2008}^{\tau=t} o_\tau$) is equally split across the remaining years of the period. Again, discrepancies between realized and forecasted offset usage naturally arise, which causes forecasts to change over time.

Graphically, the grey line in Figure 5b depicts total annual supply $\{f_t + a_t + o_t\}_t$. The peak in 2011–2012 is due to a massive use of offsets, totalling about 1 $GtCO_2$ in Phase II (Trotignon, 2012; de Perthuis & Trotignon, 2014). The following dip is due to the 900 $MtCO_2$ backloading in 2014–2016 and to non-issued Phase-III permits, totalling about 600 $MtCO_2$ in 2013–2017 (European Commission, 2015). From Phase IV on, supply coincides with the announced cap which is set to decline at a yearly linear reduction factor of 2.20%, implying that supply is nil from 2058 on. Figure 5b also shows the pre-reform cap path with an LRF of 1.74%.

4.4 In-sample calibration (2008–2017)

To calibrate the firm's discount rate and planning horizon, we choose to restrain our sample to annual price and banking data over 2008–2017 for three reasons. First, we leave aside the trial Phase I (2005–2007) as banking and borrowing across Phases I and II were not authorized,

[50] We only consider stationary sources and exclude the aviation sector (intra-EEA flights) as it was brought under the ETS in 2012, represents a small fraction of the 'regular cap' (∼2%) and regulatory uncertainty lingers pending ICAO's adoption of the relevant CORSIA instruments (European Commission, 2018b).

[51] EU legislation specifies qualitative and maximum limits on offset usage (European Parliament & Council, 2003; European Commission, 2013). Our estimate for O is based on aggregated individual entitlements.

Table 2: Best-fit calibration results (2008–2017 banking and price data)

Horizon type	Horizon & discount rate	Marginal abatement cost
Infinite	$h = \infty^{\star}$ $r = 7.06\%$ (std.dev = 52.9 $MtCO_2$)	$c = 5.53 \cdot 10^{-8}$ €/$(tCO_2)^2$ (std.dev = 4.04 €/tCO_2)
Rolling	$h = 13$ $r = 3\%^{\star}$ (std.dev = 64.9 $MtCO_2$)	$c = 5.72 \cdot 10^{-8}$ €/$(tCO_2)^2$ (std.dev = 2.12 €/tCO_2)

Note: Parameters are calibrated by minimizing the distance between annual simulated and observed banking (center column) and price (right column) levels over 2008–2017. A $\star$ indicates parameters fixed exogenously.

de facto truncating the firm's horizon, e.g. with regard to usage of Phase-I permits.[52] Second, limiting the calibration sample to 2017 implies we need not take a stance on the firm's degree of responsiveness for the calibration (as the reform passed in 2018, we assume that its effects can only be anticipated from this point on). Third, this allows us to exploit the regulatory change offered by the reform to compare how in-sample calibrated infinite and rolling horizons fare out of sample, i.e. in capturing the induced regime shift and price rally.

We consider that permit demand is linear in the permit price, which is a standard assumption (e.g. Ellerman & Montero, 2007; Kollenberg & Taschini, 2019). We thus assume that $C_t''' \approx 0$, which can be viewed as a local Taylor approximation of more general functional forms. We also assume that the slope of the linear marginal abatement cost functions is time invariant, i.e. $C_t'' = c$ for all t. We do so for three reasons. First, it is a conservative assumption given that we have limited empirical and theoretical guidance as to how marginal abatement costs evolve over time (e.g. Bréchet & Jouvet, 2008). Second, as a fixed scaling parameter, c only affects the level, but not the shape, of simulated price paths: it is neutral vis-à-vis both the firm's intertemporal decision making and our analysis in Section 5. Third, it ensures that the two-step calibration approach described below is legitimate as a constant c does not influence the firm's banking strategies, which only depend on its discount rate and horizon. That said, note that the linear intercept of the marginal abatement cost curve is gradually lowered over time as the actual baseline path is downward sloping (see Figure 5b).

We calibrate the model parameters following a two-step procedure in the spirit of a standard least squares maximum likelihood estimation with one free parameter. In the first step, we select r given h or h given r so that the simulated banking path deviates the least from the observed banking path over 2008–2017. In the second step, given r and h, we select c so that the simulated price path deviates the least from the yearly-averaged spot price path over

[52] As a result of overallocation and no banking, the Phase-I permit price dropped to zero in 2007.

Figure 6: Best-fit calibration results (2008–2017 banking and price data)

(a) Banking (b) Price

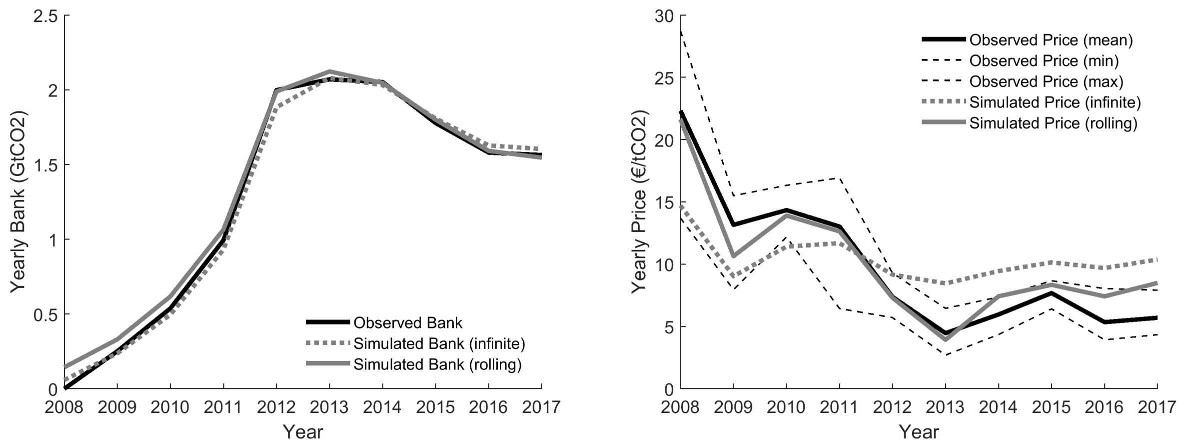

2008–2017. In each step, the free parameter is fitted by minimizing the distance between simulated and observed paths. Table 2 reports the best-fit parameters and Figure 6 depicts the observed and best-fit simulated paths over 2008–2017 for visual comparison.

With the infinite horizon (we set $h = \infty$), we find that a discount rate $r \approx 7\%$ best replicates past banking with a fit of 53 $MtCO_2$/year. This aligns with general rates of return on risky assets (see e.g. Jordà et al., 2019) but is in the higher range of the rates that can be inferred from futures' yield curves since Phase II (see Table 3).[53] Moreover, because permits can be banked for hedging purposes, one might argue that required returns should be less than those for standard risky assets. With a rolling horizon, we thus set $r = 3\%$ which is a central value for inferred discount rates,[54] and let h be the free parameter. We find that $h \approx 13$ years best replicates past banking with a similar fit of 65 $MtCO_2$/year. Next, given these select couples (r, h) for the infinite and rolling horizons, the best price fit values for c are similar and in the order of $5.5 \cdot 10^{-8}€/(tCO_2)^2$, which is line with dedicated studies, viz. $4.3 \cdot 10^{-8}€/(tCO_2)^2$ in Böhringer et al. (2009) and $5.7 \cdot 10^{-8}€/(tCO_2)^2$ in Landis (2015).[55]

[53] Ellerman & Montero (2007) analyze the permit-specific CAPM beta to select appropriate values for the discount rate. We choose a different approach and impute implicit rates from futures' yield curves. Indeed, futures markets can provide information about the discount rates applied by market participants in valuing present vs. future permits. Bredin & Parsons (2016) and Trück & Weron (2016) find similar values.

[54] Fell (2016) and Kollenberg & Taschini (2016) pick the same value while Beck & Kruse-Andersen (2020) and Perino & Willner (2016, 2017) use 5% and 10%, respectively. In terms of sensitivity, our calibrated h is increasing with the select value of r: 12, 14, 15 and 16 years for an r of 1 (or 2), 4, 5 and 6%, respectively. For r=7.06%, any horizon $h \geq 27$ years yields the same calibration output as the infinite horizon.

[55] From Böhringer et al. (2009), we obtain 1 $GtCO_2$ of abatement in 2020 for 43 €/tCO_2, which pins down the marginal abatement cost slope. See Table 4 for 2020 in Landis (2015).

Table 3: Discount rates implied from daily futures' yield curves (2008–2017)

Daily yield curve	Mean	Median	Std.Dev	Min	Max
Fut. Dec Y+1 / Spot	2.4%	2.5%	1.5%	0.2%	7.0%
Fut. Dec Y+1 / Fut. Dec Y	2.9%	2.6%	1.8%	0.3%	8.7%
Fut. Dec Y+2 / Fut. Dec Y+1	3.6%	3.7%	2.0%	0.2%	8.7%
Fut. Dec Y+3 / Fut. Dec Y+2	4.1%	2.5%	2.0%	0.6%	9.2%

Note: Summary statistics for four year-on-year futures' yield curves with daily price data. With t_1 the day's date (spot) or maturity (futures) of asset a with price $p_a^{t_1}$ and $t_2 > t_1$ the maturity of futures b with price $p_b^{t_2}$ the implied discount rate is given by $\ln(p_b^{t_2}/p_a^{t_1})/(t_2 - t_1)$ since storage costs are nil.

Our calibration exercise shows that a rolling finite horizon is able to reconcile the past banking dynamics with implicit discount rates where the infinite horizon cannot. Moreover, the rolling horizon is better able to pick up the past yearly-averaged price changes, both in sign and size (Figure 6b), and the price fit is twice as good as with the infinite horizon (Table 2).

In the absence of conclusive evidence on how forward-looking firms plan, we use an approach in the spirit of Friedman's (1953) black box model and compare the merits of two alternative assumptions – infinite vs. rolling horizons – in how well they replicate past market outcomes. Our results lend more support to the latter but do not imply that market actors actually use a rolling 13-year horizon with a 3% discount rate. Indeed, our representative firm model erases prevailing heterogeneity in individual behavior and risk preferences. As such, our calibration results should be taken as a first-pass ex-post assessment at the market level, which will also prove crucial for policy design and ex-ante evaluation as the next section shows.

5 Simulations: The 2018 EU ETS reform

In this section, we use our calibrated model to provide a quantitative assessment of the 2018 EU ETS reform with a twofold objective. First, we separate out the impacts of the three main reform design elements on market outcomes, viz. the increase in the Linear Reduction Factor from 1.74% to 2.20% from 2021 on, the introduction of the Market Stability Reserve from 2019 on, and its reinforcement with the Cancellation Mechanism from 2023 on. Second, we characterize and quantify how market outcomes depend on the firm's planning horizon and degree of responsiveness to supply control. We begin with an analysis of the general impacts of the reform on equilibrium price and banking paths and on cumulative emissions. We then appraise the MSR-induced effectiveness in achieving the associated cumulative emissions, the interaction between the various reform elements, and the MSR potential to adjust cumulative emissions in the face of exogenous permit demand shocks.

Specifically, we compare market outcomes under the infinite and rolling horizons as calibrated in Section 4.4 and under two polar assumptions about the firm's responsiveness to supply-side policies as described in Section 3.5, namely zero and full responsiveness. Thus, the scenarios we consider are: the status quo, i.e. the continuation of Phase III rules without reform (NO REF), a sole LRF increase (NO MSR) and the same scenario with the MSR (MSR). The MSR scenario is divided into four sub-scenarios, depending on whether the cancellation mechanism is on (C) or off (N) and whether the firm is fully responsive to the supply control (F) or not at all (Z). For instance, the scenario MSR F+C features a fully responsive firm and the MSR augmented by the cancellation mechanism (CM). We present our results until 2100 since all permits are used and emissions are zero in all scenarios by that time.

5.1 General impacts of the 2018 EU ETS reform

The left (resp. right) hand side of Figure 7 depicts the equilibrium price, bank and cumulative withdrawal (MSR stock + cancellation volume) paths with the infinite (resp. rolling finite) horizon.[56] Comparing NO REF with NO MSR, we see that a sole LRF increase induces higher price levels throughout and a shorter banking period, albeit with higher banked volumes early on. With the infinite horizon, all post-reform price paths reach a peak when the bank becomes empty. This is also when emissions become nil since the cap had already shrunk to zero (in 2058 with a 2.20% LRF). After the peak, there are no permits left in circulation (both the cap and bank are nil) so the firm can no longer emit and has to abate its baseline emissions. As baseline emissions gradually decline to zero, so do yearly abatement efforts and associated costs at the margin, hence the declining price paths. We observe similar though less clear-cut patterns with the rolling horizon.

Introducing the MSR on top of the larger LRF further hikes the price and reduces the bank. Crucially, the firm's horizon matters so we begin with the infinite horizon. A noticeable result is that price and bank paths with the MSR are much alike irrespective of both the CM and the firm's responsiveness, at least until reserve permits may be released into circulation in the 2050's. Logically, with zero responsiveness, price and bank paths are identical with and without the CM before reinjections actually occur, and the CM then sustains higher prices. With full responsiveness, price and bank paths are always higher with the CM. Indeed, as the firm foresees larger future reinjections without the CM, it forecasts lower overall abatement

[56] Prices are in current EUR, using annual inflation rates over 2008–2018 and 1.5% p.a. afterward.

Figure 7: MSR impacts as a function of the horizon, responsiveness and the CM

(a) Price ($h = \infty$, $r = 7.06\%$)

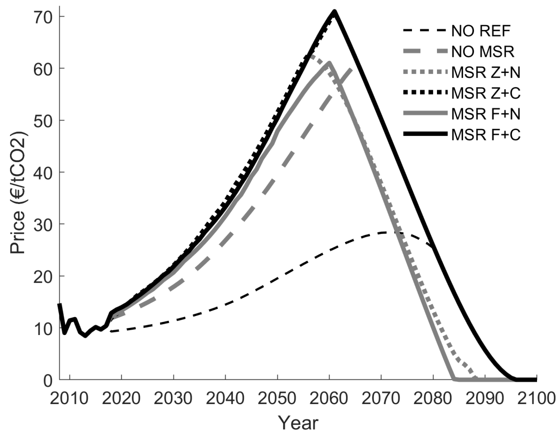

(b) Price ($h = 13$, $r = 3\%$)

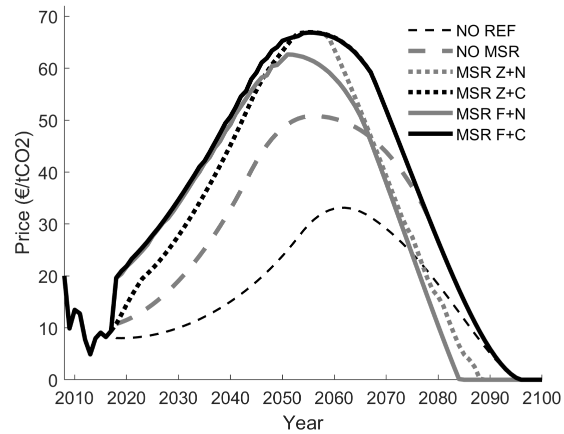

(c) Banking ($h = \infty$, $r = 7.06\%$)

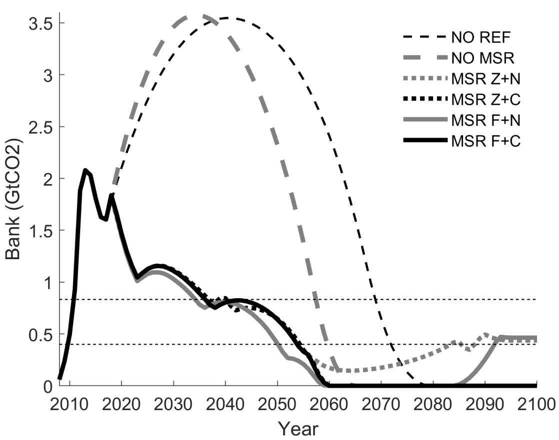

(d) Banking ($h = 13$, $r = 3\%$)

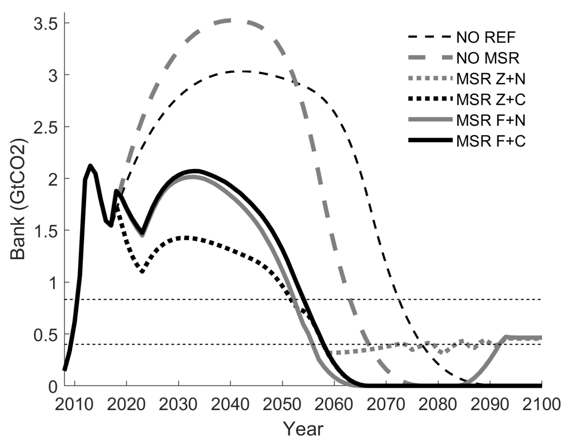

(e) Cumul. withdrawals ($h = \infty$, $r = 7.06\%$)

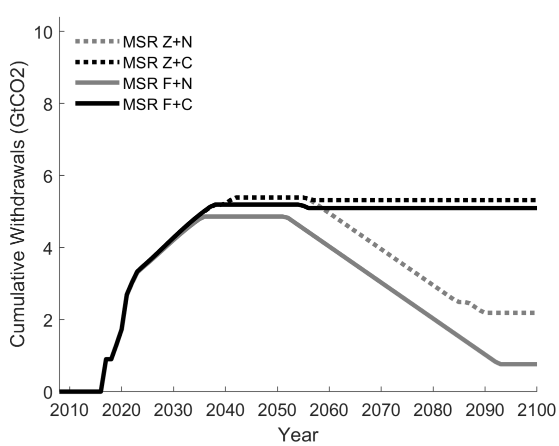

(f) Cumul. withdrawals ($h = 13$, $r = 3\%$)

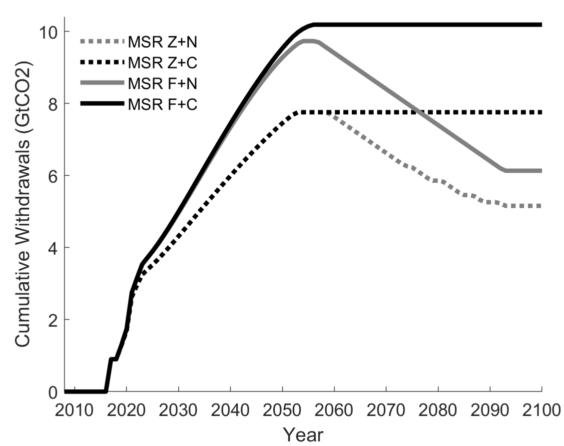

Note: Except NO REF for which the LRF is 1.74%, all other scenarios have an LRF of 2.20% and feature full (F) or zero (Z) responsiveness to the control with (C) or without (N) the CM alternatively.

efforts, which leads to lower bank and price levels throughout. This in turn entails that the MSR withdraws less permits (in volume and duration) and starts reinjecting sooner.

The MSR eats away at the bank but it takes about two decades for banking to fall below the intake threshold $\bar{b}$.[57] As Figure 7e shows, the MSR withdraws permits each year until then. The bank then remains above the outtake threshold $\underline{b}$ for 15 years (see the plateau in Figure 7e) and reinjections do not occur before the 2050's.[58] Because most of the permits set aside by the MSR have been cancelled by then, only about 100 million permits are reinjected with the CM in place and cumulative emissions are reduced by about 5 $GtCO_2$. Without the CM, cumulative emissions are also reduced, albeit by a lower amount. Non-conservation of cumulative emissions without the CM occurs because the MSR does not have time to empty before the market terminates.[59] Finally note that, intuitively, cumulative MSR intakes (and ensuing cancellations when applicable) are smaller when the firm is responsive.

With the rolling finite horizon, price and bank paths are higher, i.e. the MSR induces a more pronounced price increase and reduces the bank less sharply than with the infinite horizon. Importantly, with full responsiveness, the rolling horizon captures the observed 2018 price rally (see Figure 1), where the infinite horizon falls short of it (the simulated price increase is three times bigger with the rolling horizon). Therefore, in addition to better in-sample price and discount rate fits, the rolling horizon also yields a better price fit out of sample.

Price and banking levels are higher when the firm is responsive and, crucially, irrespective of the CM. The underlying mechanism is the following. With the rolling horizon, MSR-driven supply cuts have a larger relative impact on the firm's perceived overall abatement effort and more of it is abated early on due to the lower discount rate.[60] These two effects concur to yield

[57] As the MSR takes in permits and cuts back on yearly auctions, it forces the firm to tap into its private bank of permits to compensate for reduced contemporaneous supply.

[58] The slight bank upticks occurring after the intake and outtake thresholds are passed arise as the MSR suddenly stops withdrawing or starts reinjecting permits, respectively. Without the CM, annual reinjections of 100 million permits induce a second banking period when the firm is responsive (Perino & Willner (2016) find a similar effect) while when the firm is not responsive the bank never drops to zero in the first place and fluctuates around the outtake threshold.

[59] Note that we do not impose a terminal condition on the MSR stock. Compared with Perino & Willner (2016), conservation of cumulative emissions under the sole MSR does not occur in our simulations because the firm (1) stops needing permits, i.e. baselines become nil, before the reserve could empty; (2) has imperfect foresight, for otherwise the responsive firm would not waste permits this way. When the firm is responsive: (1) at the end of the century, reserve permits return to the market in excess of the firm's actual compliance needs so that banking increases and passes above the outtake threshold (see Figure 7c), and the MSR stops emptying; (2) the MSR reinjects more permits than effectively withdrawn as the terminal MSR stock is lower than the amount of permits the MSR is initially seeded with (900 + 581 million).

[60] Initially, reinjections are far off into the future and beyond the firm's horizon so the CM is irrelevant for the responsive firm early on. As with the infinite horizon, price and bank paths are identical with and without the CM when the firm is not responsive before reinjections kick in.

higher price and bank levels than with the infinite horizon. In addition, as the responsive firm foresees shorter supply over its horizon, it drives up abatement and banking, which in turn inflates future MSR intakes. This raises the firm's overall abatement forecast, which leads to higher banking and future MSR intakes, and so forth.[61] As a result, MSR intakes are larger (due to higher banking), last longer (intakes stop 15 years later), and no reinjections occur with the rolling horizon. This translates into a larger contraction in cumulative emissions, in the order of 6 to 10 $GtCO_2$ when the firm is responsive, without and with the CM.

5.2 Focus on cumulative emissions and cost effectiveness

The introduction of the MSR and CM endogenizes the cumulative cap on emissions, which is now a function of the firm's horizon, discount rate and responsiveness. Table 4 contains the 2008–2100 cumulative emissions that obtain with the MSR and a 2.20% LRF in each possible scenario. It also reports the equivalent LRFs (LRF_{eq}) used from 2021 on that would, without the MSR, deliver the same cumulative emissions. For instance, with the rolling horizon, full responsiveness and the CM, the LRF_{eq} can be as high as 2.95%. This corresponds to a linear emissions cap path declining by 64.9 $MtCO_2$ per annum which becomes nil in 2048 (compared to 48.4 $MtCO_2$ and 2058 with a 2.20% LRF).

On the face of it, relative to the MSR, the LRF_{eq} specifies a clear supply trajectory without inducing a priori distortions in the firm's optimal intertemporal allocation of abatement. We investigate this issue by computing the additional total costs of achieving a given cumulative cap with the MSR in place relative to the corresponding LRF_{eq}, reported in the second-to-last column of Table 4. With the infinite horizon, the MSR-induced effectiveness loss is negligible with the CM and in the order of 10% without. Similar results obtain with the rolling horizon, with the notable exception that the MSR and CM can improve cost effectiveness relative to the LRF_{eq} with a responsive firm, achieving a 2% reduction in total costs.

How can we explain this result? For visual comparison, Figure 8 depicts the simulated price paths with the MSR vs. the LRF_{eq} in the four cases where the firm is responsive. Without the CM, we see that by postponing auctions, the MSR induces too much abatement early on and too little later on (because of fixed, inflexible yearly reinjections) with both the infinite and rolling horizons, hence the higher costs. This 'distortion' or 'imbalance' is partly remedied by the CM, showing that it is a key tool to flank the MSR with. In particular, with the infinite

[61] This adjustment procedure eventually converges to a fixed point, see Section 3.5. As time goes by and reinjections enter the horizon, the firm forecasts less of an abatement effort without the CM, implying lower bank and price levels than with the CM.

Table 4: MSR impacts on cumulative emissions and compliance costs

Horizon	Respons.	CM	Cumul. emissions	LRF_{eq}	Effic. loss	Interaction
Infinite	Zero	Off	57.86 $GtCO_2$	2.28%	9.0%	16.4%
		On	55.16 $GtCO_2$	2.48%	0.2%	4.6%
	Full	Off	59.26 $GtCO_2$	2.18%	11.5%	11.1%
		On	55.39 $GtCO_2$	2.46%	0.2%	3.1%
Rolling	Zero	Off	54.87 $GtCO_2$	2.50%	9.0%	11.7%
		On	52.72 $GtCO_2$	2.70%	0.6%	0.6%
	Full	Off	53.88 $GtCO_2$	2.59%	7.9%	1.7%
		On	50.29 $GtCO_2$	2.95%	-2.2%	-5.2%

Note: LRF_{eq} is the linear reduction factor from 2021 on yielding the same cumulative emissions over 2008–2100 without MSR as those with the MSR and a 2.20% LRF. Effic. loss measures the additional cumulative compliance costs implied by the MSR with a 2.20% LRF relative to the sole LRF_{eq}. Interaction measures the reductions in cumulative emissions due to the MSR with a 1.74% LRF that are additional to those obtained with the MSR and a 2.20% LRF, expressed as a percentage of the decrease in cumulative emissions due to the LRF increase from 1.74% to 2.20% (9.0 $GtCO_2$).

horizon, price paths are quasi-identical with the MSR and the LRF_{eq} since the MSR-driven change in the supply path is essentially irrelevant for the firm (Figure 8c).[62]

Crucially, the MSR-driven auction backloading is no longer neutral with the rolling horizon, and the MSR has potential to compensate for the firm's limited foresight. Indeed, with no MSR, the firm partially factors in the long-term cap stringency implied by the LRF_{eq}. This results in 'too low' abatement levels early on and 'too high' abatement levels later on, when the firm realizes it underestimated the 'true' cap stringency. In comparison, by frontloading abatement efforts, the MSR makes the long-term stringency embedded in the linear cap path more tangible early on, steering the abatement path more in line with the long-term target.

5.3 Focus on the interaction between reform elements

We analyze the nature of the interaction between the LRF increase from 1.74% to 2.20% and the implementation of the MSR. That is, we assess whether these elements are complements, substitutes or independent in terms of resulting cumulative emissions. Independently of the MSR, the LRF increase lowers cumulative emissions by 9 $GtCO_2$ (from 68 down to 59 $GtCO_2$ over 2008–2100). We use this as a reference to evaluate the LRF-MSR interaction considering a hypothetical case where the LRF stayed at 1.74%. That is, we quantify how much of these 9

[62]Negligible differences occur due to changes in the stringency of the limited borrowing constraint. Indeed, a higher LRF implies reduced free allocations, thus less borrowing opportunities. This affects intertemporal arbitrage constraints and hence market equilibrium paths at the margin.

Figure 8: Comparative price paths with the MSR and an equivalent LRF without MSR

(a) MSR F+N ($h = \infty$, $r = 7.06\%$)

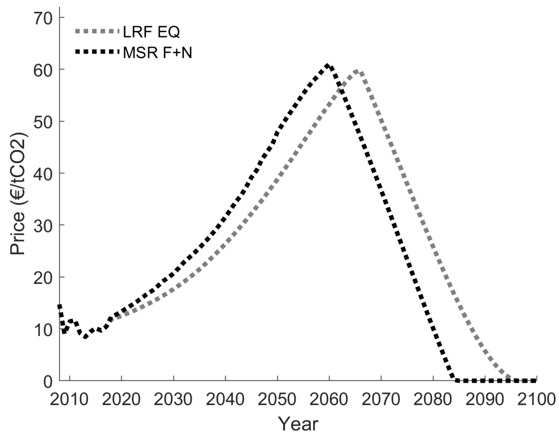

(b) MSR F+N ($h = 13$, $r = 3\%$)

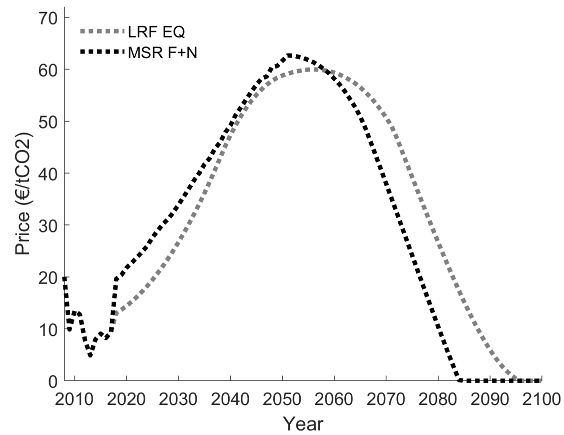

(c) MSR F+C ($h = \infty$, $r = 7.06\%$)

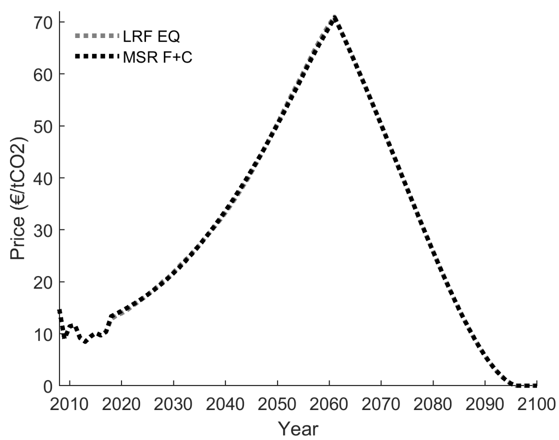

(d) MSR F+C ($h = 13$, $r = 3\%$)

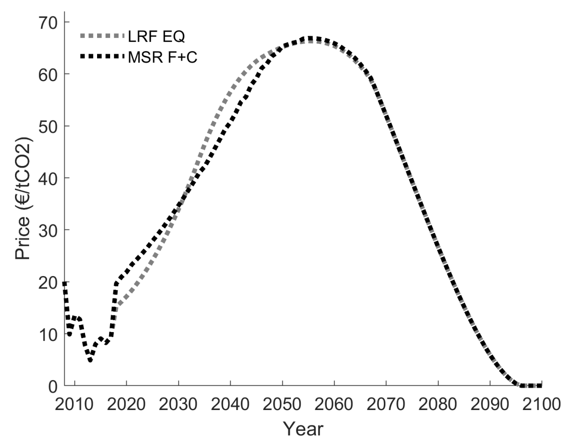

Note: MSR scenarios use a 2.20% LRF. LRF EQ scenarios use the corresponding LRF_{eq} in Table 4.

$GtCO_2$ would be permanently withdrawn from circulation by the MSR with a 1.74% LRF (i.e. on top of the total withdrawals achieved by the MSR with a 2.20% LRF). The last column of Table 4 reports our results. The second-order magnitude of these additional cumulative MSR withdrawals suggests that, to a large extent, the LRF increase and the MSR are independent elements as far as cumulative emissions are concerned, especially with the CM in place.[63]

In terms of transitional stringency, however, and in light of the discussion in Section 5.2, we wish to underscore that the LRF increase, the MSR and the CM complement each other well in the case of a rolling finite horizon. Specifically, while the LRF increase reinforces the long-

[63] See Quemin (2020) for a detailed analysis of the synergies between the MSR parameters and the LRF.

term cap stringency, the MSR makes it more visible early on by increasing the transitional cap stringency (and eventually the CM further raises the cumulative cap stringency by emptying the MSR). In other words, the LRF commands the long-term policy aspects while the MSR is designed to address those in the short to mid term. To see this, observe that despite the reduction in cumulative supply associated with the LRF increase (9 GtCO2) is similar to that induced by the MSR and CM (Figure 7f), the price rise attributable to the sole LRF increase represents only 25% of that witnessed in conjunction with the MSR and CM (Figure 7b). By contrast, with the infinite horizon, the sole LRF increase represents 70% of the simulated 2018 price rise obtained in conjunction with the MSR and CM (Figure 7a).

5.4 Focus on the 'puncture of the waterbed' effect

Given a fixed cumulative emissions cap, if 1 tCO_2 under the cap is not emitted today because of some exogenous event or policy, i.e. independently of the permit price, it will necessarily be emitted at some point in the future, and vice versa. This is usually referred to as a 'waterbed effect' over time.[64] As a consequence of the reform, however, the cumulative emissions cap is no longer a key fixed policy element but has become a market outcome, hence uncertain ex ante. As such, it can be affected by exogenous factors that curtail permit demand, typically complementary policies. Importantly, some share of an exogenously abated tCO_2 can be made permanent, i.e. not emitted in the future, and this share decreases with the date when that tCO_2 is initially abated. In other words, the reform partially and temporarily 'punctures the waterbed' (Perino, 2018). This aspect of the reform has received broad attention in the policy debate because it allows some supply-side adjustments reflecting the interaction between the permit market and various exogenous factors.[65]

We contribute to this debate by quantifying and comparing the partial and temporary nature of the puncture with our calibrated infinite and rolling finite horizons. To this end, we analyze the impacts on cumulative emissions of one-off marginal shifts in baseline emissions, which can for instance be attributed to marginal changes in economic activity or renewable deployment. Specifically, we consider unanticipated one-off 10 $ktCO_2$ drops in baseline emissions occurring

[64]A waterbed effect also occurs over space in the form of carbon leakage or fuel switching, but is beyond the scope of this paper. That is, if 1 tCO_2 is abated independently of the permit price in some location within the system, the induced permit price decrease can trigger 1 tCO_2 emission elsewhere, say in a more CO_2-intensive location or via fuel switching as CO_2-intensive technologies become relatively cheaper to operate.

[65]The prolonged price downturn which led to the reform has in large part been attributed to the economic recession and the achievements of overlapping renewable and energy efficiency policies. Some scholars see cap adjustments as a way of partially preserving the environmental integrity of these complementary policies.

Table 5: Long-term effects of one-off marginal shifts in baseline emissions

			Year of shift				
Horizon	Respons.	CM	2020	2025	2030	2035	2040
Infinite	Zero	Off/On	50.1%	39.4%	27.2%	5.3%	0%
	Full	Off	30.7%	24.3%	11.6%	0%	0%
		On	51.9%	40.3%	24.8%	12.3%	0%
Rolling	Zero	Off/On	9.8%	10.3%	11.3%	12.8%	15.3%
	Full	Off	15.7%	17.3%	18.8%	19.6%	22.5%
		On	15.8%	17.4%	18.9%	21.0%	22.5%

Note: Permanent impacts on cumulative emissions as a percentage of the magnitude of the one-off shift.

in different years. The arbitrary size and sign of the shift do not affect our results provided it is small enough in absolute terms so that the MSR intake cut-off year is preserved. This allows us to measure and compare marginal impacts in a meaningful way by ruling out threshold effects (i.e. discontinuities).[66] Table 5 contains the shares of the shift that are made permanent depending on the presence of the CM and the firm's horizon and responsiveness. The higher this share, the more significant the puncture of the waterbed.

Our results with the CM, an infinite horizon and full responsiveness are consistent with other studies (see Carlén et al., 2019; Perino et al., 2019; Beck & Kruse-Andersen, 2020). That is, the puncture is partial (notice we find smaller punctures in size) and temporary, i.e. smaller the closer the shift from the intake cut-off year (a zero puncture obtains for shifts occurring after this year). By contrast, because the MSR never empties completely without the CM, we find that the puncture is not zero in this case (but remains smaller than with the CM). As a novelty, we also provide results in other cases. First, with zero responsiveness, the puncture is independent of the CM as it has no bearing on the bank paths before the 2050's. Second, a smaller puncture (twice as small for early shifts) obtains with the rolling horizon but it is remarkably stable and even slightly increasing with the year of the shift.

How can we explain these results? The one-off decrease in baseline emissions in year t does not translate into a one-to-one increase in banking in year t, which itself does not translate into a one-to-one increase in cumulative MSR withdrawals (terminal MSR stock or cancellations).[67]

[66] We note that the intake cut-off year is a key indicator of the MSR-induced resilience to external demand shocks. See e.g. Quemin (2020) for an analysis of threshold effects and associated oscillatory behavior.

[67] With (resp. without) the CM, these extra withdrawn permits are cancelled (resp. still sitting in the MSR when the program ends) and never return to the market.

Table 6: Decomposition of the effects of one-off marginal shifts in baseline emissions

Horizon	Share	Year of shift 2020	2025	2030	2035	2040
Infinite	Y_t	60%	57%	54%	53%	52%
	W_{cumul}	90%	76%	59%	23%	0%
Rolling	Y_t	24%	28%	31%	34%	40%
	W_{cumul}	92%	87%	85%	80%	70%

Note: Dampening factors Y_t and W_{cumul} are defined in (13). Case with the CM and full responsiveness.

We can represent these two successive dampening effects by

$$\overbrace{-X_t}^{\text{baseline shift}} \longrightarrow \overbrace{Y_t \times X_t}^{\text{bank increment}} \longrightarrow \overbrace{W_{cumul} \times Y_t \times X_t}^{\text{cumulative withdrawals}}, \tag{13}$$

with X_t the absolute value of the size of the negative baseline shift, $Y_t \in [0, 1]$ the share of the shift that is passed into the bank as a result of intertemporal smoothing, and $W_{cumul} \in [0, 1]$ the share of the bank increment converted into cumulative withdrawals via the MSR. Table 6 reports Y_t and W_{cumul} with the CM in place and full responsiveness. The product $Y_t \times W_{cumul}$ yields the overall impact on cumulative emissions displayed in Table 5.

For a given baseline shift year t, Y_t is always larger with the infinite horizon. This is because with the rolling finite horizon the firm has less room to spread the shift over time and because its bank is larger to start with. Both elements concur to yield a lower incentive to bank the freed-up permits, and observe that the wedge in Y_t between the infinite and rolling horizons decreases with the shift year. Conversely, for a given shift year, W_{cumul} is always larger with the rolling horizon. This is because the intake cut-off always occurs later in this case so that the MSR has more time to absorb the initial bank increment.[68] Naturally, W_{cumul} decreases with the shift year as there is less time before the intake cut-off occurs.

[68] In the first order, $W_{cumul} \approx 1 - (1 - 0.24)^x (1 - 0.12)^y$ where x and y are the number of years between the shift year and the intake cut-off year with $R = 24$ and 12%, respectively. Given its exponential property, W_{cumul} does not (1) decrease by much over 2020–2040 with the rolling horizon and (2) differ by much between the infinite and rolling horizons for early shifts.

6 Concluding remarks

We have made three contributions to the literature on emissions trading as a climate policy tool. As a first contribution, we have built a model of competitive emissions trading under uncertainty with supply control departing from the existing literature along two dimensions: firms can utilize rolling finite horizons as a way of addressing uncertainty and exhibit bounded responsiveness to the control. We have motivated these two departures based on anecdotal evidence in the context of the EU ETS (e.g. financial constraints, missing long-term markets, regulatory uncertainty and other behavioral aspects which have a bearing on intertemporal optimization) and by drawing from the existing literature in other related fields.

As a second contribution, we have tailored the model to the EU ETS design and calibrated the market-wide discount rate and horizon to replicate past market developments. Compared to a standard infinite horizon model, a rolling horizon can reconcile the banking dynamics over 2008–2017 with discount rates derived from futures' yield curves, replicate annual price changes over 2008–2017 twice better, and capture the regime shift and price run-up in 2018. Our calibration estimates are a valuable addition to the empirical literature on the EU ETS. They substantiate previous claims that shortsightedness could be crucial in understanding price and banking dynamics (e.g. Ellerman et al., 2015; Fuss et al., 2018) and increase the relevance of our simulation results that quantify MSR impacts.

As a third contribution, we have used our calibrated model to assess the impacts of the 2018 EU ETS reform depending on what is assumed about the firms' horizon and responsiveness. We have decomposed the impacts of its three main elements (viz. the increase in the Linear Reduction Factor from 1.74% to 2.20%, the introduction of the Market Stability Reserve and its reinforcement with the Cancellation Mechanism) and quantified the interactions between them and the firms' horizon and responsiveness. This has important implications for policy design and evaluation. For instance, the LRF (which commands the long-term stringency of the policy) and the MSR (which is designed to raise stringency in the short to mid term) are policy complements when firms are limitedly farsighted. Importantly, the CM is instrumental to ensure cost effectiveness as it balances out MSR-induced distortions in the abatement path. That said, we find that by design the MSR is ill-equipped to provide stability to the market (i.e. adequate supply adjustments) in the face of demand shocks in the future.[69]

More broadly, our framework as well as our calibration and simulation results in the context of

[69]Bruninx & Ovaere (2020) and Gerlagh et al. (2020) further study how the MSR response depends on the shocks' structure. The MSR can also induce oscillations of its own due to threshold effects (Quemin, 2020).

the EU ETS contribute to improving one's understanding of the intertemporal performances of ETSs in general – a topic which is high on both policy and research agendas (Hasegawa & Salant, 2015; Ellerman et al., 2016) – and the interactions between intertemporal trading and supply-side controls. Our framework can also serve as a good basis for an assessment of changes in the MSR parameters for the 2021 review (Quemin, 2020) or more profound design changes such as the introduction of a price floor mechanism in place or on top of the MSR, e.g. via an auction reserve price (Newbery et al., 2019; Flachsland et al., 2020).

Finally, our paper constitutes a first step towards analyzing hybrid emissions trading systems when firms employ rolling finite horizons or exhibit bounded responsiveness to supply-control instruments. Several alleys for future research seem fruitful. A first line of important research would be optimal policy selection when the regulator is aware that firms are shortsighted or limitedly responsive. This would imply revisiting the normative analysis of price vs. quantity vs. hybrid policies and establishing welfare preferences over hybrid policy designs, e.g. price vs. banking collars.[70] Second, follow-up work could endogenize the firms' horizon as a result of a trade-off between compliance cost minimization and forecasting or computational costs, which may depend on supply-side control complexity. A third line of research could envisage a formal treatment of the heterogeneity of firms' time and risk preferences that goes beyond the representative firm approach developed in this paper and the related literature.

Acknowledgments

We would like to thank Jan Ahrens, Anna Cretì, Christian Flachsland, Peter Kruse-Andersen, Michael Pahle, John Parsons, Grischa Perino, Christian de Perthuis, Katheline Schubert, Luca Taschini and Oliver Tietjen for valuable conversations at different stages of this project. We also thank an anonymous referee from FAERE, and benefited from helpful comments at various conferences (FSR Climate 2018, AFSE 2018, IAEE 2019, EAERE 2019) and seminars (College of Europe, LSE, MCC, PIK, Universität Hamburg). All remaining errors are ours.

Simon Quemin acknowledges funding from LABEX Louis Bachelier Finance and Sustainable Growth (ANR 11-LABX-0019 Project) under the 'Investments for the Future' programme, in accordance with Article 8 of the Assigment Agreements for Communication Assistance, and from the Grantham Foundation for the Protection of the Environment, the UK Economic

[70]More generally, since there is no economic rationale for regulating the firms' intertemporal use of permits (i.e. is permit use frontloaded or backloaded compared to MSR thresholds?), it would be worth studying how the MSR behaves relative to optimal supply adjustment mechanisms, e.g. à la Laffont & Tirole (1996).

and Social Research Council (ESRC) through the Centre for Climate Change Economics and Policy. Both authors further acknowledge funding from the Climate Economics Chair, Paris Dauphine University (PSL Research University).

References

Abrell, J., Kosch, M. & Rausch, S. (2019). Carbon Abatement with Renewables: Evaluating Wind and Solar Subsidies in Germany and Spain. *Journal of Public Economics*, **169**, 172–202.

Abrell, J. & Rausch, S. (2017). Combining Price and Quantity Controls under Partitioned Environmental Regulation. *Journal of Public Economics*, **145**, 226–42.

Baudry, M., Faure, A. & Quemin, S. (2020). *Emissions Trading with Transaction Costs.* Working Paper 341, Grantham Research Institute, London School of Economics.

Baumol, W. J. & Quandt, R. E. (1964). Rules of Thumb and Optimally Imperfect Decisions. *American Economic Review*, **54**(2), 23–46.

Beck, U. & Kruse-Andersen, P. K. (2020). Endogenizing the Cap in a Cap-and-Trade System: Assessing the Agreement on EU ETS Phase 4. *Environmental & Resource Economics*, **77**, 781–811.

Bel, G. & Joseph, S. (2015). Emission Abatement: Untangling the Impacts of the EU ETS and the Economic Crisis. *Energy Economics*, **49**, 531–9.

Bocklet, J., Hintermayer, M., Schmidt, L. & Wildgrube, T. (2019). The Reformed EU ETS – Intertemporal Emission Trading with Restricted Banking. *Energy Economics*, **84**, 104486.

Böhringer, C., Löschel, A., Moslener, U. & Rutherford, T. F. (2009). EU Climate Policy up to 2020: An Economic Impact Assessment. *Energy Economics*, **31**, S295–305.

Borenstein, S., Bushnell, J., Wolak, F. A. & Zaragoza-Watkins, M. (2019). Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design. *American Economic Review*, **109**(11), 3953–77.

Branch, W. A. (2004). The Theory of Rationally Heterogeneous Expectations: Evidence from Survey Data on Inflation Expectations. *Economic Journal*, **114**, 592–621.

Bréchet, T. & Jouvet, P.-A. (2008). Environmental Innovation and the Cost of Pollution Abatement Revisited. *Ecological Economics*, **65**(2), 262–5.

Bredin, D. & Parsons, J. (2016). Why is Spot Carbon so Cheap and Future Carbon so Dear? The Term Structure of Carbon Prices. *Energy Journal*, **37**(3), 83–107.

Brock, W. A. & Hommes, C. H. (1997). A Rational Route to Randomness. *Econometrica*, **65**(5), 1059–95.

Bruninx, K. & Ovaere, M. (2020). Estimating the Impact of COVID-19 on Emissions and Emission Allowance Prices Under EU ETS. *IAEE Energy Forum*, 40–2, Covid-19 Issue.

Bruninx, K., Ovaere, M. & Delarue, E. (2020). The Long-Term Impact of the Market Stability Reserve on the EU Emission Trading System. *Energy Economics*, **89**, 104746.

Burtraw, D. & Keyes, A. (2018). Recognizing Gravity as a Strong Force in Atmosphere Emissions Markets. *Agricultural & Resource Economics Review*, **47**(2), 201–19.

Caliendo, F. & Aadland, D. (2007). Short-Term Planning and the Life-Cycle Consumption Puzzle. *Journal of Economic Dynamics & Control*, **31**(4), 1392–415.

Cantillon, E. & Slechten, A. (2018). Information Aggregation in Emissions Markets with Abatement. *Annals of Economics & Statistics*, **132**, 53–79.

Carbone, E. & Hey, J. D. (2001). A Test of the Principle of Optimality. *Theory & Decision*, **50**(3), 263–81.

Carlén, B., Dahlqvist, A., Mandell, S. & Marklund, P. (2019). EU ETS Emissions under the Cancellation Mechanism – Effects of National Measures. *Energy Policy*, **129**, 816–25.

Chaton, C., Cretì, A. & Sanin, M.-E. (2018). Assessing the Implementation of the Market Stability Reserve. *Energy Policy*, **118**, 642–54.

Chèze, B., Chevallier, J., Berghmans, N. & Alberola, E. (2020). On the CO_2 Emissions Determinants During the EU ETS Phases I and II: A Plant-level Analysis Merging the EUTL and Platts Power Data. *Energy Journal*, **41**(4), 153–84.

Dardati, E. & Riutort, J. (2016). Cap-and-Trade and Financial Constraints: Is Investment Independent of Permit Holdings? *Environmental & Resource Economics*, **65**(4), 841–64.

Deaton, A. & Laroque, G. (1992). On the Behaviour of Commodity Prices. *Review of Economic Studies*, **59**(1), 1–23.

Easley, D. & Spulber, D. F. (1981). Stochastic Equilibrium and Optimality with Rolling Plans. *International Economic Review*, **22**(1), 79–103.

Ellerman, A. D., Marcantonini, C. & Zaklan, A. (2016). The European Union Emissions Trading System: Ten Years and Counting. *Review of Environmental Economics & Policy*, **10**(1), 89–107.

Ellerman, A. D. & Montero, J.-P. (2007). The Efficiency and Robustness of Allowance Banking in the U.S. Acid Rain Program. *Energy Journal*, **28**(4), 47–71.

Ellerman, A. D., Valero, V. & Zaklan, A. (2015). *An Analysis of Allowance Banking in the EU ETS*. Working Paper RSCAS 2015/29, European University Institute.

Eurelectric (2009). EU ETS Phase 3 Auctioning – Timing and Futures versus Spot. Brussels.

European Commission (2013). *Commission Regulation (EU) 1123/2013*. Brussels, 9.11.2013.

European Commission (2014). *Commission Regulation (EU) 176/2014*. Brussels, 16.2.2014.

European Commission (2015). *Impact Assessment SWD(2015) 135*. Brussels, 15.7.2015.

European Commission (2018a). *Communication C(2018) 2801*. Brussels, 15.5.2018.

European Commission (2018b). *Report COM(2018) 842 final*. Brussels, 17.12.2018.

European Parliament & Council (2003). *Directive 2003/87/EC*. OJ L 275, 13.10.2003.

European Parliament & Council (2018). *Directive (EU) 2018/410*. OJ L 76, 19.03.2018.

Fell, H. (2016). Comparing Policies to Confront Permit Over-Allocation. *Journal of Environmental Economics & Management*, **80**, 53–68.

Fell, H., Burtraw, D., Morgenstern, R. D. & Palmer, K. L. (2012). Soft and Hard Price Collars in a Cap-and-Trade System: A Comparative Analysis. *Journal of Environmental Economics & Management*, **64**(2), 183–98.

Findley, T. S. & Caliendo, F. N. (2009). Short Horizons, Time Inconsistency, and Optimal Social Security. *International Tax & Public Finance*, **16**(4), 487–513.

Flachsland, C., Pahle, M., Burtraw, D., Edenhofer, O., Elkerbout, M., Fischer, C., Tietjen, O. & Zetterberg, L. (2020). How to Avoid History Repeating Itself: The Case for an EU

Emissions Trading System (EU ETS) Price Floor Revisited. *Climate Policy*, **20**(1), 133–42.

Friedman, M. (1953). *Essays in Positive Economics*. University of Chicago Press.

Fuss, S., Flachsland, C., Koch, N., Kornek, U., Knopf, B. & Edenhofer, O. (2018). A Framework for Assessing the Performance of Cap-and-Trade Systems: Insights from the European Union Emissions Trading System. *Review of Environmental Economics & Policy*, **12**(2), 220–41.

Gabaix, X. (2014). A Sparisty-Based Model of Bounded Rationality. *Quarterly Journal of Economics*, **4**, 1661–710.

Gerlagh, R., Heijmans, R. J. R. K., & Rosendahl, K. E. (2019). *Endogenous Emission Caps Always Induce a Green Paradox*. Working Paper 7862, CESifo.

Gerlagh, R. & Heijmans, R. J. R. K. (2018). *Regulating Stock Pollutants*. Working Paper 7383/2018, CESifo.

Gerlagh, R. & Heijmans, R. J. R. K. (2019). Climate-Conscious Consumers and the Buy, Bank, Burn Program. *Nature Climate Change*, **9**, 431–3.

Gerlagh, R., Heijmans, R. J. R. K. & Rosendahl, K. E. (2020). COVID-19 Tests the Market Stability Reserve. *Environmental & Resource Economics*, **76**, 855–65.

Gigerenzer, G. & Todd, P. M. (1999). *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.

Gilboa, I. & Schmeidler, D. (1989). Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics*, **18**(2), 141–53.

Goldman, S. M. (1968). Optimal Growth and Continual Planning Revision. *Review of Economic Studies*, **35**(2), 145–54.

Grüll, G. & Taschini, L. (2011). Cap-and-Trade Properties under Different Hybrid Scheme Designs. *Journal of Environmental Economics & Management*, **61**(1), 107–18.

Grüne, L., Semmler, W. & Stieler, M. (2015). Using Nonlinear Model Predictive Control for Dynamic Decision Problems in Economics. *Journal of Economic Dynamics & Control*, **60**, 112–33.

Haruvy, E., Lahav, Y. & Noussair, C. N. (2007). Traders' Expectations in Asset Markets: Experimental Evidence. *American Economic Review*, **97**(5), 1901–20.

Hasegawa, M. & Salant, S. W. (2014). Cap-and-Trade Programs under Delayed Compliance: Consequences of Interim Injections of Permits. *Journal of Public Economics*, **119**, 24–34.

Hasegawa, M. & Salant, S. W. (2015). The Dynamics of Pollution Permits. *Annual Review of Resource Economics*, **7**, 61–79.

Heutel, G. (2012). How Should Environmental Policy Respond to Business Cycles? Optimal Policy under Persistent Productivity Shocks. *Review of Economic Dynamics*, **15**(2), 244–64.

Hey, J. D. & Knoll, J. A. (2007). How Far Ahead Do People Plan? *Economics Letters*, **96**(1), 8–13.

Hintermann, B., Peterson, S. & Rickels, W. (2016). Price and Market Behavior in Phase II of the EUETS: A Review of the Literature. *Review of Environmental Economics & Policy*,

10(1), 108–28.

Hommes, C. (2011). The Heterogeneous Expectations Hypothesis: Some Evidence from the Lab. *Journal of Economic Dynamics & Control*, **35**(1), 1–24.

Hommes, C., Massaro, D. & Weber, M. (2019). Monetary Policy under Behavioral Expectations: Theory and Experiment. *European Economic Review*, **118**, 193–212.

Hortaçsu, A., Luco, F., Puller, S. L. & Zhu, D. (2019). Does Strategic Ability Affect Efficiency? Evidence from Electricity Markets. *American Economic Review*, **109**(12), 4302–42.

Hotelling, H. (1931). The Economics of Exhaustible Resources. *Journal of Political Economy*, **39**(2), 137–75.

ICAP (2020). *Status Report 2020*. Berlin: International Carbon Action Partnership.

Jehiel, P. (1995). Limited Horizon Forecast in Repeated Alternate Games. *Journal of Economic Theory*, **67**(2), 497–519.

Johnson, E. J., Camerer, C., Sen, S. & Rymon, T. (2002). Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining. *Journal of Economic Theory*, **104**(1), 16–47.

Joltreau, E. & Sommerfeld, K. (2019). Why does Emissions Trading under the EU Emissions Trading System (ETS) not Affect Firms' Competitiveness? Empirical Findings from the Literature. *Climate Policy*, **19**(4), 453–71.

Jordà, O., Knoll, K., Kuvshinov, D., Schularick, M. & Taylor, A. M. (2019). The Rate of Return on Everything, 1870–2015. *Quarterly Journal of Economics*, **134**(3), 1225–98.

Kaganovich, M. (1985). Efficiency of Sliding Plans in a Linear Model with Time-Dependent Technology. *Review of Economic Studies*, **52**(4), 691–702.

Karpf, A., Mandel, A. & Battiston, S. (2018). Price and Network Dynamics in the European Carbon Market. *Journal of Economic Behavior & Organization*, **153**, 103–22.

Kling, C. & Rubin, J. (1997). Bankable Permits for the Control of Environmental Pollution. *Journal of Public Economics*, **64**(1), 101–15.

Koch, N., Fuss, S., Grosjean, G. & Edenhofer, O. (2014). Causes of the EUETS Price Drop: Recession, CDM, Renewable Policies or a Bit of Everything?−New Evidence. *Energy Policy*, **73**, 676–85.

Koch, N., Grosjean, G., Fuss, S. & Edenhofer, O. (2016). Politics Matters: Regulatory Events as Catalysts for Price Formation under Cap-and-Trade. *Journal of Environmental Economics & Management*, **78**, 121–39.

Koch, N., Reuter, W. H., Fuss, S. & Grosjean, G. (2017). Permits vs. Offsets under Investment Uncertainty. *Resource & Energy Economics*, **49**, 33–47.

Kollenberg, S. & Taschini, L. (2016). Emissions Trading Systems with Cap Adjustments. *Journal of Environmental Economics & Management*, **80**, 20–36.

Kollenberg, S. & Taschini, L. (2019). Dynamic Supply Adjustment and Banking under Uncertainty in an Emissions Trading Scheme: The Market Stability Reserve. *European Economic Review*, **118**, 213–26.

Kydland, F. E. & Prescott, E. C. (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, **85**(3), 473–91.

Laffont, J.-J. & Tirole, J. (1996). Pollution Permits and Compliance Strategies. *Journal of Public Economics*, **62**(1), 85–125.

Landis, F. (2015). *Final Report on Marginal Abatement Cost Curves for the Evaluation of the Market Stability Reserve*. Dokumentation 15-01, ZEW Mannheim.

Lintunen, J. & Kuusela, O.-P. (2018). Business Cycles and Emission Trading with Banking. *European Economic Review*, **101**, 397–417.

Lucas, R. E. & Prescott, E. C. (1971). Investment Under Uncertainty. *Econometrica*, **39**(5), 659–81.

Marcantonini, C. & Ellerman, A. D. (2015). The Implicit Carbon Price of Renewable Energy Incentives in Germany. *Energy Journal*, **36**(4), 205–39.

Marcantonini, C. & Valero, V. (2017). Renewable Energy and CO_2 Abatement in Italy. *Energy Policy*, **106**, 600–13.

Martin, R., Muûls, M. & Wagner, U. J. (2015). Trading Behavior in the EU Emissions Trading Scheme. In M. Gronwald & B. Hintermann (eds.) *Emissions Trading as a Policy Instrument: Evaluation and Prospects*, chapter 9, (pp. 213–38). MIT Press.

Mauer, E.-M., Okullo, S. J. & Pahle, M. (2019). *Evaluating the Performance of the EU ETS MSR*. Working Paper, Postdam Insitute for Climate (PIK).

Montgomery, W. D. (1972). Markets in Licenses and Efficient Pollution Control Programs. *Journal of Economic Theory*, **5**(3), 395–418.

Muth, J. F. (1961). Rational Expectations and the Theory of Price Movements. *Econometrica*, **29**(3), 315–35.

Naegele, H. & Zaklan, A. (2019). Does the EU ETS Cause Carbon Leakage in European Manufacturing? *Journal of Environmental Economics & Management*, **93**, 125–47.

Newbery, D. M. (2018). What Future(s) for Liberalized Electricity Markets: Efficient, Equitable or Innovative? *Energy Journal*, **39**(1), 1–28.

Newbery, D. M., Reiner, D. M. & Ritz, R. A. (2019). The Political Economy of a Carbon Price Floor for Power Generation. *Energy Journal*, **40**(1), 1–24.

Obara, I. & Zincenko, F. (2017). Collusion and Heterogeneity of Firms. *RAND Journal of Economics*, **48**(1), 230–49.

Perino, G. (2018). New EU ETS Phase IV Rules Temporarily Puncture Waterbed. *Nature Climate Change*, **8**, 262–4.

Perino, G., Ritz, R. A. & van Benthem, A. (2019). *Overlapping Climate Policies*. Discussion Paper 13569, Centre for Economic Policy Research, revised in November 2020.

Perino, G. & Willner, M. (2016). Procrastinating Reform: The Impact of the Market Stability Reserve on the EU ETS. *Journal of Environmental Economics & Management*, **80**, 37–52.

Perino, G. & Willner, M. (2017). EU ETS Phase IV: Allowance Prices, Design Choices and the Market Stability Reserve. *Climate Policy*, **17**(7), 936–46.

de Perthuis, C. & Trotignon, R. (2014). Governance of CO_2 Markets: Lessons from the EU ETS. *Energy Policy*, **75**, 100–6.

Pindyck, R. S. (1993). The Present Value Model of Rational Commodity Pricing. *Economic Journal*, **103**(418), 511–30.

Pizer, W. A. (2002). Combining Price and Quantity Controls to Mitigate Global Climate Change. *Journal of Public Economics*, **85**(3), 409–34.

Quemin, S. (2020). *Using Supply-Side Policies to Raise Ambition: The Case of the EU ETS and the 2021 Review*. GRI Working Paper 335, London School of Economics.

Radner, R. (1996). Bounded Rationality, Indeterminacy, and the Theory of the Firm. *Economic Journal*, **106**, 1360–73.

Reis, R. (2006). Inattentive Producers. *Review of Economic Studies*, **73**(3), 793–821.

Richstein, J. C., Chappin, E. J. & de Vries, L. J. (2015). The Market (In-)Stability Reserve for EU Carbon Emission Trading: Why It Might Fail and How to Improve It. *Utilities Policy*, **35**, 1–18.

Roberts, M. J. & Spence, A. M. (1976). Effluent Charges and Licenses under Uncertainty. *Journal of Public Economics*, **5**(3), 193–208.

Rubin, J. D. (1996). A Model of Intertemporal Emission Trading, Banking and Borrowing. *Journal of Environmental Economics & Management*, **31**(3), 269–86.

Sahin, F., Narayanan, A. & Robinson, E. P. (2013). Rolling Horizon Planning in Supply Chains: Review, Implications and Directions for Future Research. *International Journal of Production Research*, **51**(18), 5413–36.

Salant, S. W. (1983). The Vulnerability of Price Stabilization Schemes to Speculative Attack. *Journal of Political Economy*, **91**(1), 1–38.

Salant, S. W. (2016). What Ails the European Union's Emissions Trading System? *Journal of Environmental Economics & Management*, **80**, 6–19.

Samuelson, P. A. (1971). Stochastic Speculative Price. *Proceedings of the National Academy of Sciences*, **68**(2), 335–7.

Schennach, S. M. (2000). The Economics of Pollution Permit Banking in the Context of Title IV of the 1990 Clean Air Act Amendments. *Journal of Environmental Economics & Management*, **40**(3), 189–210.

Schleich, J., Lehmann, S., Cludius, J., Abrell, J., Betz, R. & Pinkse, J. (2020). *Active or Passive? Companies' Use of the EU ETS*. Working Paper Sustainability and Innovation S07/2020, Fraunhofer ISI, Karlsruhe.

Schopp, A., Acworth, W., Hauptmann, D. & Neuhoff, K. (2015). *Modelling a Market Stability Reserve in Carbon Markets*. Discussion Paper 1483, DIW Berlin.

Seifert, J., Uhrig-Homburg, M. & Wagner, M. (2008). Dynamic Behavior of CO_2 Spot Prices. *Journal of Environmental Economics & Management*, **56**(2), 180–94.

Sethi, S. & Sorger, G. (1991). A Theory of Rolling Horizon Decision Making. *Annals of Operations Research*, **29**(1), 387–415.

Simon, H. A. (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, **69**(1), 99–118.

Smith, V. L., Suchanek, G. L. & Williams, A. W. (1988). Bubbles, Crashes, and Endogenous

Expectations in Experimental Spot Asset Markets. *Econometrica*, **56**(5), 1119–51.

Spiro, D. (2014). Resource Prices and Planning Horizons. *Journal of Economic Dynamics & Control*, **48**, 159–75.

Stocking, A. (2012). Unintended Consequences of Price Controls: An Application to Allowance Markets. *Journal of Environmental Economics & Management*, **63**(1), 120–36.

Stranlund, J. K., Costello, C. & Chávez, C. A. (2005). Enforcing Emissions Trading when Emissions Permits are Bankable. *Journal of Regulatory Economics*, **28**(2), 181–204.

The Economist (2013). ETS, RIP? 20 April 2013.

Thomson Reuters (2014). Carbon Market Survey – Into Smoother Waters? Point Carbon.

Tietjen, O., Lessmann, K. & Pahle, M. (2021). Hedging and Temporal Permit Issuances in Cap-and-Trade Programs: The Market Stability Reserve under Risk Aversion. *Resource & Energy Economics*, **63**, 101214.

Trotignon, R. (2012). Combining Cap-and-Trade with Offsets: Lessons from the EU ETS. *Climate Policy*, **12**(3), 273–87.

Trück, S. & Weron, R. (2016). Convenience Yields and Risk Premiums in the EU-ETS – Evidence from the Kyoto Commitment Period. *Journal of Futures Markets*, **36**(6), 587–611.

van Veldhuizen, R. & Sonnemans, J. (2018). Nonrenewable Resources, Strategic Behavior and the Hotelling Rule: An Experiment. *Journal of Industrial Economics*, **66**, 481–516.

Venmans, F. (2016). The Effect of Allocation Above Emissions and Price Uncertainty on Abatement Investments under the EU ETS. *Journal of Cleaner Production*, **126**, 595–606.

Wettestad, J. & Jevnaker, T. (2019). Smokescreen Politics? Ratcheting Up EU Emissions Trading in 2017. *Review of Policy Research*, **36**(5), 635–59.

World Bank (2020). *State and Trends of Carbon Pricing 2020*. Washington, DC: World Bank.

Wright, B. D. & Williams, J. C. (1982). The Economic Role of Commodity Storage. *Economic Journal*, **92**(367), 596–614.

Zhang, B. & Xu, L. (2013). Multi-Item Production Planning with Carbon Cap and Trade Mechanism. *International Journal of Production Economics*, **144**(1), 118–27.