An endogenous emissions cap

SUMMARY

The European Union's Emissions Trading System (EU ETS) is complemented by a Market Stability Reserve (MSR). After a major revision of the EU ETS in 2018, the MSR effectively makes the supply of allowances responsive to demand. In this paper, we show that a cap-and-trade scheme with an endogenous cap, such as the EU ETS produces a green paradox. Abatement policies announced early but realized in the future are counter-effective because of the MSR, they increase cumulative emissions. We present the mechanisms in a two-period model, and then provide quantitative evidence of our result for an annual model disciplined on the price rise in the EU ETS that followed the introduction of the MSR. Our results point to the need for better coordination between different policies, such as the "European Green Deal." We conclude with suggestions to improve the workings of an endogenous cap, ahead of the MSR review scheduled for 2021.

JEL codes: D59; E61; H23; Q50; Q54; Q58

—Reyer Gerlagh, Roweno J.R.K. Heijmans andKnut Einar Rosendahl

Economic Policy 2021 Printed in Great Britain

An endogenous emissions cap produces a green paradox

Reyer Gerlagh,* Roweno J.R.K. Heijmans † and Knut Einar Rosendahl ‡

*Department of Economics, Tilburg University, the Netherlands and Oslo Center for Research on Environmental Friendly Energy (CREE); †Department of Economics, Tilburg University, the Netherlands and Oslo Center for Research on Environmental Friendly Energy (CREE); ‡School of Economics and Business of the Norwegian University of Life Sciences, Norway, Statistics Norway, and the Oslo Center for Research on Environmental Friendly Energy (CREE), Norway

Corresponding author: Roweno J.R.K. Heijmans. E-mail: r.j.r.k.heijmans@uvt.nl

1. INTRODUCTION

In order to reduce greenhouse gas emissions economists have long advocated carbon pricing, either as a tax or via an emissions trading system (ETS) (c.f. Aldy et al., 2010; Golosov et al., 2014). Where a tax fixes the price of emissions, an ETS sets overall emissions while leaving the price endogenous to forces in the market. The typical ETS in addition allows for banking and, sometimes, borrowing between periods. With banking and borrowing, short-run emissions levels can flexibly adjust to changing market conditions even if the short-run supply of emissions allowances is fixed. Long-run emissions levels are still given, however, as long as the long-run supply of allowances is exogenous.

Economic Policy 2021 pp. 1–34 Printed in Great Britain

Emissions targets are a natural focal point of policy making and, perhaps for this reason, policy makers around the world have generally favored ETSs over emissions taxes. The aim of any climate policy is to halt global warming by reducing greenhouse gas emissions. In this sense, an emissions cap like the European Union's ETS (EU ETS) or the Regional Greenhouse Gas Initiative (RGGI) is the most direct instrument toward the given goal of limiting emissions.[1] Although an emissions tax can, in the end, also achieve a reduction in emissions, the effect is indirect. In addition, a cap offers certainty on emissions whereas a carbon tax leaves the realized amount of emissions reduction to the market, which can be politically undesirable. Finally, in the context of the European Union, an emissions cap can be imposed after simple majority voting whereas an EU-wide tax requires unanimous consent.[2]

Due to uncertainty, the realized ETS price may exceed, or fall short of, prices expected when the system is set-up (Weitzman, 1974). To avoid sustained unexpected deviations of emissions prices, supplementary measures have been proposed or even implemented, such as price collars (Roberts and Spence, 1976; Abrell and Rausch, 2017; Borenstein et al., 2019) or endogenous allocation of allowances to individual firms (e.g. output-based allocation, cf. Fowlie et al. (2016); Böhringer et al. (2017)). The EU ETS, the world's largest operating carbon market, recently implemented involving market-induced cancelation of allowances. Hence, the long-run supply of allowances is no longer fixed – the emissions cap is endogenous by construction (Perino, 2018; Gerlagh and Heijmans, 2019).

In this paper, we show, first analytically in a simple two-period model and then numerically in the context of the EU ETS, that an ETS with a quantity-based endogenous cap produces a green paradox. More precisely, we show that there exists an abatement policy which reduces the demand for allowances but at the same time increases aggregate emissions. When calibrating and simulating a model of the EU ETS, we find that the green paradox may be substantial, especially if demand for emissions allowances is reduced only several years from now but anticipated already today. Our results clearly show that the announcement of future abatement policies can invert the long-run effects from a reduction to an increase in emissions.

The endogenous supply of allowances in the EU ETS is itself endogenous to its history. Over the years from 2008 to 2012, the net supply of allowances increased through a large inflow of certified emissions reductions from clean development mechanism projects. Besides, and due to the economic slowdown that started in 2008, demand for allowances decreased and a large amount of banked (i.e. unused and saved for later)

1 The Regional Greenhouse Gas Initiative is "a cooperative effort among the states of Connecticut, Delaware, Maine, Maryland, Massachusetts, New Hampshire, New Jersey, New York, Rhode Island, and Vermont to cap and reduce CO_2 emissions from the power sector." (retrieved from www.rggi.org).

2 That is a carbon tax would fall under each national government's sovereignty whereas an ETS can be established under supra-national EU law.

allowances accumulated. The large bank exercised a downward pressure on the price of emissions allowances (EUAs), which dropped below 10 €/tCO_2 from 2012 onward. Perceiving these prices as too low, the EU implemented a Market Stability Reserve (MSR) in 2015. The ideas of this MSR is that if aggregate banking in the market exceeds a certain threshold, part of next year's allowances enter the MSR rather than the market (Fell, 2016; Kollenberg and Taschini, 2019). These MSR-held allowances are to be "backloaded" in the future, when demand is higher. Importantly, though, note that this initial MSR only reduced the short-run supply of allowances – the long-run, cumulative cap on emissions remained untouched.

Leaving the cumulative cap untouched, the backloading of allowances did not create more scarcity and thus did not succeed to push up prices. In response, the EU adapted the following MSR-mechanics in 2018: when the size of the MSR exceeds the annual level of auctioned allowances, all allowances above this threshold are permanently canceled. With this adjustment to the MSR, the cumulative cap was effectively reduced, supporting higher prices. Importantly, the amount of canceling has been made endogenous; it depends on the allowances that are banked and subsequently flow into the MSR. Intertemporal supply and demand now find themselves in a delicate balance.

Whereas abatement policies had no effect on cumulative emissions under the old regime, Perino (2018) finds that the new rules (as intended) leave them some leverage. A 1-ton demand reduction in 2018 reduces cumulative emissions (i.e. the long-run emissions cap) by 0.4–0.8 tons according to his calculations. The reasoning is that reduced demand in 2018 increases banking and a bigger inflow into the MSR, which eventually cancels more allowances. That is, the new MSR rules have "punctured the waterbed".[3] The magnitude of the effects depends on the timing of the demand reduction, and the time window over which the MSR takes in allowances.

Gerlagh and Heijmans (2019) extend the analysis by Perino (2018) and consider changes in equilibrium prices as well as second-order effects on banking and allowance cancelation. In the present paper, we add one more element and examine the effects of demand reductions in any period, differentiating between surprise policies and those anticipated before implementation.

With respect to unanticipated policies, our analysis underscores Perino's (2018) finding that emissions-reducing policies are (partially) effective. A surprise policy that reduces the demand for allowances also reduces cumulative emissions if it is announced and implemented at the same time. This conclusion changes drastically when considering anticipated future policies instead (cf. Rosendahl, 2019). If the policymaker announces a complementary emissions-reducing policy today that will be implemented some years from now, firms anticipate the associated reduction in demand to come, and reduce the

3 If the emissions cap is fixed and binding, any additional policies will not affect total emissions, but only shuffle emissions around. This is often referred to as the waterbed effect: Sitting on a waterbed changes the distribution of water inside the bed, but not the total amount of water the bed contains.

amount of allowances they bank: why keep them for a future in which they are not needed? This is important because, by construction of the EU ETS, a decrease in banking leads to less inflow into the MSR. Furthermore, only allowances in the MSR can be canceled. It follows that fewer emissions allowances get written off. Hence, while the complementary policy reduces future emissions indeed, it also starts an unintended chain of events through which, in the years leading up to the reduction, emissions increase. If the policy is anticipated long enough in advance, this causes cumulative emissions to be higher compared with the case in which no policy had been implemented.

The mechanism we described is somewhat reminiscent of the green paradox (Sinn, 2008; Bauer et al., 2018); anticipated future climate policies incentivize fossil fuels producers to speed up extraction, increasing current but not cumulative emissions. In our context, it is not the timing of emissions but cumulative emissions that increase following well-intended climate policies (cf. Gerlagh, 2011). The green paradox we consider is therefore stronger than the classic one, and caused by an artificial market intended to support climate policies.

Importantly, our result does not warrant the conclusion that an endogenous emissions cap like in the EU ETS is a bad idea per se. Rather, it illustrates the cost of leaving out all price information from the MSR design.[4] Quantity-related targets are politically manageable, while price-related regulation is politically sensitive, particularly in the EU. But the resulting pure quantity-based regulation does not interact well with overlapping policies, especially when these are announced in advance of actual implementation. Constructively, we therefore suggest several relatively straightforward fixes to the EU ETS that remedy our green paradox. Most simply, the European Commission could complement any demand-reducing policy with a proportionate decrease in the future supply of allowances. Alternatively, the EU ETS could be enhanced with a price mechanism, such that the supply of allowances is reduced when the allowance price falls.[5] We discuss the relative merits of these solutions at some length in Section 5.

Only a few published studies exist quantifying the impacts of the cancelation rules in the MSR, and none of them consider the green paradox we demonstrate. The first quantitative study is probably Perino and Willner (2017), who simulated the impacts on EUA prices of the (then) proposal to cancel allowances, extending the model in Perino and Willner (2016). Silbye and Sørensen (2019) use a quantitative model similar to ours, concluding that demand-reducing policies in early years reduce cumulative emissions. They find bigger quantitative impacts than the Perino (2018), as the MSR takes in allowances for a much longer period (we find similar results in this paper). Bruninx et al. (2020) use a more detailed model of the EU ETS and investigate the impacts of the

4 Note that the intuition we presented above for the green paradox, caused by the MSR, did not refer to prices.

5 This does not need to be a rigid price floor as in the Regional Greenhouse Gas Initiative. The auctioned volume can continuously decrease with decreasing prices.

MSR on EUA prices and cumulative emissions. Gerlagh et al. (2020) apply the same model as in this paper, examining the impacts of COVID-19 on EUA prices with and without the cancelation rules.

Our results invite particular concern in light of recent policy developments. In December 2019, the European Commission presented its "European Green Deal", promising 50–55% reductions in greenhouse gas emissions compared with 1990 levels in 2030, and a carbon neutral economy by 2050 (EU Commission, 2019).[6] Although setting an ambitious agenda, the European Green Deal is precisely the kind of demand-reducing policy, announced and anticipated years in advance of actual implementation, to which our findings speak. Unless changes to the EU ETS are implemented in parallel, the announced demand-reductions in future decades may backfire. They may reduce the inflow into the MSR in the near future, reducing cancelation of emissions allowances, eventually increasing cumulative emissions within the EU ETS.

On the other hand, a fairly simply remedy to our green paradox result exists. If the policy maker, upon announcing a demand-reducing policy, simultaneously reduces the supply of allowances, this can undo the green paradox effect. Such a supply-reduction can be implemented either at the EU level through a more rapid reduction in the annual supply of allowances, or at the national level through unilateral cancelation of allowances.[7] Such an adaptation retains the efficiency benefits of an endogenous supply scheme as implemented into the EU ETS yet mitigates the problems due to a green paradox identified in the present paper. We come back to ETS policies complementing demand-reducing measures when discussing our findings in the Section 4.

The structure of the paper is as follows. We first present a stylized two-period model in which we lay out the mechanisms that lead from the endogenous cap to a green paradox. The next section adds the details of the EU ETS. A particular element of the EU ETS cum MSR is that it exhibits multiplicity of equilibria, and that the green paradox specifically arises for abatement policies that reduce future demand. We showcase these elements through a numerical calibration to the model, in which we calculate the size of the green paradox in the EU ETS. In Section 4, we present discuss the policy implications of our results, as well as several fixes. The final section concludes.

2. MODEL

Let there be two periods $t = 1$ and 2, and let e_t denote the emissions in period t. Because firms have to surrender allowances, or permits, for their emissions, we save on notation

6 In December 2020, the EU Heads of State approved an emissions reduction target of at least 55% for 2030, cf. https://ec.europa.eu/commission/presscorner/detail/en/mex_20_2389.

7 The German government is planning to cancel allowances along with the country's phase-out of coal power. At the time of writing, it is not decided how many allowances will be canceled, but the government stated it will take into account cancellation of allowances via the MSR when making the decision (Szabo and Garside, 2020).

and let e_t denote the demand for allowances as well. Permits are traded at a price p_t and we allow for banking and borrowing of allowances between the two periods.[8]

Demand for allowances e_t is decreasing in the price p_t and we describe this relationship via the demand function $f_t(p_t)$. We also leave room for demand reductions that are external to the ETS itself, such as shifts in consumer preferences toward less emission-intensive products or complementary policies affecting the demand for emissions altogether (e.g. the European Green Deal, phasing out of coal power or supporting zero emissions technologies). Let λ_t denotes these external effects, such that $\lambda_t < 0$ describes a *reduction* in demand. We will refer to λ_t as a complementary demand policy in our narrative. With these elements, we obtain the following demand for allowances:

$$e_t = f_t(p_t) + \lambda_t. \tag{1}$$

Aggregate emissions, E, are equal to the sum of emissions in the two periods: $E = e_1 + e_2$. We are particularly interested in $dE/d\lambda_t$, that is, the effect of a demand-reducing policy on aggregate emissions. If $dE/d\lambda_t = 0$, we have the standard waterbed effect, and the policy can be regarded as ineffective. If $0 < dE/d\lambda_t < 1$, the policy is (slightly) effective, while there is a green paradox if aggregate emissions *increase* in response to a demand-reducing policy, $dE/d\lambda_t < 0$.

The demand function given by Equation (1) can be inverted to yield an inverse demand function $\psi_t(.)$:

$$p_t = \psi_t(e_t - \lambda_t). \tag{2}$$

Although it is not necessary to assume price-taking behavior, for the sake of analytical convenience we make the standard assumption that the price rises by the interest rate r:

$$p_2 = (1 + r)p_1. \tag{3}$$

This condition is known as Hotelling's rule and follows from free banking of allowances over time (see footnote 8) combined with unrestricted access by outsider firms to ETS allowances. It describes intertemporal arbitrage between investment opportunities (Hotelling, 1931). If the price would rise at a pace above the interest rate, investors would have an incentive to buy allowances in the first period, and sell them in the second period at a positive net return. But this would lead to a rise in the first period price and a decrease in the second period price, and equilibrium would not be reached until the allowances price rises by the interest rate. A similar reverse mechanism prevents prices from rising below the interest rate.[9]

8 In the EU ETS, borrowing from a future period is not allowed. However, this constraint is currently not binding, and will probably not be binding in the foreseeable future (nor in our simulations).

9 That is, free access for outsiders to the allowances market reduces the feasibility of strategic price-distorting behavior by firms in the market.

Note that a complementary demand-reducing policy in a given period ($\lambda_t < 0$) suppresses the price of emissions in that period. The implication of Hotelling's rule is then that the price of emissions in the other period ($s \neq t$) should fall as well, which means that emissions in period s will rise. Without the MSR, the emissions reductions in period t would be completely undone through increased emissions in period s, that is, the waterbed effect. With the MSR, things are not as straightforward as we shall see below.

Hotelling's rule is intimately related to the ease with which a more ambitious climate policy can be implemented today compared with the future. To see this, consider a 1-unit reduction in cumulative emissions, $dE = de_1 + de_2 = -1$. The change in prices per unit additional emissions reduction in period 1 is given by ψ'_1, the slope of the inverse demand function in the first period (cf. Equation (2)). Similarly, bringing about 1-unit additional emissions reductions in the second period would change prices in that period by ψ'_2.

For a given additional tightening of emissions, the ratio between these price effects can be viewed as a measure for the relative difficulty of reducing emissions in the first period compared with the second. In economic terms, we are interested in the ratio of the elasticity of demand, as a measure of the relative effort of a first-period reduction vis-a-vis a second-period reduction:

$$\eta = \frac{\psi'_1/\psi_1}{\psi'_2/\psi_2} = \frac{(1+r)\psi'_1}{\psi'_2}. \tag{4}$$

An efficient allocation of the climate ambition splits the additional emissions reduction between the two periods, such that the marginal costs rise by the interest rate (the relation to Hotelling's rule suggested earlier). If $\eta < 1$, an efficient policy reduces emissions mostly in the first period. If $\eta > 1$, increased climate ambitions will mostly reduce demand for allowances in the second period.

Although so far, our focus has been on the demand for emissions allowances, these must be matched by supply in the ETS. Let supply be given exogenously by $\bar{s}_t$ in t. If a firm has more allowances in the first period than it uses, it can bank these allowances for use in the second period. Let this bank, aggregated over all firms, be denoted b, so that $b = \bar{s}_1 - e_1$. The level of banking is crucial for the operation of the MSR, and a detailed explanation of this is given in the next section. For our model, we rely on a stylized representation. If the bank exceeds some given threshold $\bar{b}$, so $b > \bar{b}$, then the supply of allowances in the second period is reduced by an amount δb; a fixed fraction of the total bank.[10] Importantly, firms keep their (private) bank of allowances. Banked allowances are not canceled and so Hotelling's rule, Equation (3), is maintained. This, in turn,

10 Note that the EU ETS without MSR-driven canceling of allowances is effectively described by setting $\delta = 0$ in our model. Note also the discontinuity around $\bar{b}$ - supply drops by $\delta\bar{b}$ units as banking crosses $\bar{b}$.

means that firms' banking incentives are not directly affected.[11] Rather, the supply of *new* allowances in the second period is adjusted; supply drops to $\bar{s}_2 - \delta b$.

Supply and demand of allowances must balance in an ETS, which leads to the following set of conditions:

$$e_1 + b = \bar{s}_1, \tag{5}$$

$$e_2 = \bar{s}_2 + b \text{ iff } b \leq \bar{b}, \tag{6}$$

$$e_2 = \bar{s}_2 - \delta b + b \text{ iff } b > \bar{b}. \tag{7}$$

For now, we assume that the bank exceeds the cancelation threshold: $b > \bar{b}$. We return to this assumption below, though we note that it is clearly met in the real EU ETS. As total emissions equal total supply as well, we then obtain from Equations (5) and (7):

$$E = e_1 + e_2 = \bar{s}_1 + \bar{s}_2 - \delta b. \tag{8}$$

Equation (8) highlights an important implication of the MSR: cumulative emissions *decrease* proportionally with banking. As banking decreases in first period demand, this means that low demand in the first period leads to decreased supply in the second, and therefore to a decrease in cumulative emissions. This is the essence of the EU ETS' stability mechanism.

Observation 1. *The change in cumulative emissions equals the change in banking in the first periods, multiplied by the cancelation parameter δ (as long as $b > \bar{b}$):*

$$dE = -\delta db = \delta de_1 \tag{9}$$

We can now derive our key results as an implication of the simplified model developed in this section. Imagine that the government enacts a complementary demand-reducing policy in the first period: $\lambda_1 < 0$. This policy, by its nature, reduces the demand for emissions in period 1 ($e_1 \downarrow$). The decline in demand mechanically leads to more banking ($b \uparrow$), which in turn leads to less supply of allowances in the second period through the MSR ($\bar{s}_2 - \delta b \downarrow$). Consequently, emissions overall will fall ($E \downarrow$). This result is as intended: a policy reducing the demand for emissions leads to lower cumulative emissions.

The more counter-intuitive case arises when the government enacts a complementary demand-reducing policy in the second period, $\lambda_2 < 0$, which is announced (or at least anticipated) in the first. Anticipating a lower demand for allowances in the second

11 If a policy were to touch the allowances held by private firms, it would disturb abatement incentives away from Hotelling's Rule and that is inefficient. Quantity-based policies to remedy a perceived over-supply of allowances can therefore be efficient only if they reduce the net supply ot allowances directly, without touching allowances held by private firms.

period ($e_2 \downarrow$), firms will bank fewer allowances in the first period for use in the second ($b \downarrow$). This decreased banking implies a lower reduction of supply in the second period through the MSR ($\bar{s}_2 - \delta b \uparrow$). Aggregate emissions rise accordingly ($E \uparrow$). This, at first, is a counter-intuitive result: a policy reducing demand for emissions *in the second period* leads to higher emissions overall.

Proposition 1 formalizes these discussions.

Proposition 1 (Green Paradox). *Assume that $b > \bar{b}$. Then we have:*
The MSR retains but dampens the effect of demand-reducing policies in the first period:

$$0 < \frac{dE}{d\lambda_1} = \frac{\delta\eta}{\eta + 1 - \delta} < 1. \tag{10}$$

The MSR reverses the effect of demand-reducing policies in the second period:

$$\frac{dE}{d\lambda_2} = -\frac{\delta}{\eta + 1 - \delta} < 0 \tag{11}$$

The proofs of this and the next proposition are found in Appendix A.

Recall that our stylized model of the EU ETS *with* MSR can also describe a situation *without* MSR by setting the cancelation-parameter $\delta = 0$. In this case, demand-reducing policies in either period have no effect on emissions: the waterbed effect.

With positive canceling of allowances, $0 < \delta < 1$, Proposition 1 tells us that demand-reducing policies in the first period indeed lower cumulative emissions: the waterbed is punctured with respect to early supplemental climate policies (cf. Perino, 2018). A green paradox arises when the government enacts demand-reducing policies in the second period that are anticipated in the first, for then cumulative emissions increase.

In the special case of complete cancelation, $\delta = 1$, the waterbed is not just punctured for early demand reductions, it is leaking altogether. For, as Proposition 1 makes clear, in this case early demand-reducing policies are fully translated in aggregate emissions reductions ($\frac{dE}{d\lambda_1} = 1$). On the downside, complete cancelation also leads to a sizable green paradox. Indeed, if the cost of achieving increased climate ambitions today is relatively low compared with achieving the same ambitions in the future ($\eta < 1$), our green paradox exceeds 100%; reducing demand in the second period by 100 ton of CO_2 leads to a more than 100 ton increase in aggregate emissions.

We note that our green paradox is not due to an accidental and unfortunate combination of factors in the EU ETS. It is a fundamental feature of any endogenous emissions cap that works through quantities (i.e. some $\delta \neq 0$), rather than through price information. This suffices to understand the economics behind the green paradox that arises in our simulations. For completeness we graphically illustrate Proposition 1 in the Appendix A, see the left panel of Figure 6.

There is one thing left to be discussed. Our simple model also illustrates an unwanted side effect of *discrete* supply-adjustments in response to trigger events. In the EU ETS, supply in the second period is not necessarily continuously reduced in response to banking. Rather, the marginal effect of banking on supply reductions jumps *discretely* when banking crosses the cancelation threshold $\bar{b}$. To be more precise, for all banking levels $b < \bar{b}$, there is no cancelation of allowances in the second period, whereas for all banking levels $b > \bar{b}$, supply in the second period is reduced by an amount δb. Hence, when banking crosses the threshold $\bar{b}$, the cancelation of allowances in period 2 jumps up from zero to some amount at least $\delta\bar{b}$. This discrete adjustment of supply may lead to unexpected problems of uniqueness or existence. See the Appendix, Figure 6, right panel, for a graphical representation. The next proposition formalizes.

Proposition 2 (Multiplicity). *If an equilibrium exists with banking sufficiently close to the threshold,* $|b - \bar{b}| < \varepsilon$ *and* ε *small, then at least two distinct equilibria exist. These equilibria are supported by distinct price-paths* $(p_1^*, p_2^*) < (p_1^{**}, p_2^{**})$*, and different levels of cumulative emissions* $E^* > E^{**} + \delta\bar{b}$.

The problem with equilibrium multiplicity is the inherent unpredictability of the market it implies. A policy-maker expects firms to behave according to the equilibrium of the (implicit) game they are playing. But if there is more than one equilibrium, which outcome should the policy-maker expect? Worse still, what should firms expect other firms to do? This leads to an intricate system of expectations with no clear outcome. Firms are essentially forced to act by guess and by golly, which may lead to coordination failure and inefficiency (Van Huyck et al., 1990).

In addition to equilibrium multiplicity and coordination failures, discrete supply-adjustments are undesirable as they invite participating firms to engage in strategic gambling. Although the consequences of such behavior are hard to assess without clear data, there is an experience with it in other domains; currency attacks reveal the potential for private gains from exploiting policy interventions triggered by market indicators (Morris and Shin, 1998). In the context the EU ETS, suppose the expected bank size at the end of the year is slightly below the threshold that triggers a flow into the MSR. A large firm could then buy a substantial number of allowances, driving up the price by a small amount. This leads to a reduction in demand for other firms and could thus push the bank above its MSR-threshold, inducing a large write off of allowances through the MSR. After the market switches to a new equilibrium, the large firm can then sell its allowances at a higher price and book a substantial gain. In this sense, discrete trigger events build a kind of strategic complementarities (Bulow et al., 1985) into the game.

The EU ETS is not the only system where trigger events lead to discrete adjustments in supply. RGGI admits a similar property: the supply of allowances is reduced by a *discrete* amount when prices fall below a specific level. RGGI, too, may therefore be susceptible to equilibrium multiplicity.

Propositions 1 and 2 are not intended to constitute a criticism of endogenous emissions caps altogether. For a pollutant with the characteristics of climate change, where damages dependent on *cumulative* emissions, a reduction of future supply in response to lower current demand yields substantial welfare improvements (Gerlagh and Heijmans, 2018). Rather than suggesting that the EU ETS abandons its MSR, we therefore argue the MSR mechanics should be adopted to preempt the possible problems we identify. We return to this point later in the paper.

3. QUANTITATIVE ASSESSMENT

3.1. EU ETS model

In this section we develop and simulate a stylized, dynamic model of the EU ETS that captures the mechanics of the MSR in detail. We first briefly revisit the EU ETS and the rules of the MSR.

EU ETS is the largest market for carbon to date and as one of the first such instruments, it has experienced many difficulties since its conception. Firms under the EU ETS at risk of relocating have led the EU to adopt (too) generous compensation mechanisms (Martin et al., 2014). The price of allowances has been consistently low and highly volatile, carrying along some counter-intuitive implications for firms' profit (Bushnell et al., 2013). The low price of carbon in the EU ETS can be traced back to interactions with supplemental climate policies as well as the general economic recession during part of its existence. The cap on emissions has been considered set too loosely, as evidenced by a strong accumulating "bank" of unused allowances, privately stored by firms for future use, despite the low prices.[12]

In response, the EU introduced a MSR and set the new rules in 2018. From 2019 the MSR takes in allowances that are otherwise auctioned, the amount of which equals 24% (12% as of 2024) of banked allowances, every year the (cumulative) bank exceeds 833 $MtCO_2$.[13] These allowances, *not* taken from the private bank of allowances but from the volumes otherwise auctioned, will return to the market later; in years when the bank has shrunk to below 400 $MtCO_2$, an additional 100 $MtCO_2$ is auctioned from the MSR. When too many allowances end up in the MSR, however, all MSR-held allowances in excess of the volume auctioned in the previous year are canceled permanently (starting in 2023). In this sense, the MSR with canceling effectively makes the cap on emissions in the EU ETS endogenous. The MSR reforms have been documented in Perino (2018) and Gerlagh and Heijmans (2019). The equations used for our simulations are provided in the Appendix, Section B.

12 Hintermann et al. (2016) suggests that the emissions cap in phase II of the EU ETS was not binding, as thus the nonzero price toward the end of the phase" reflected expectations of a cap on overall emissions that is binding in the long term, given the opportunity to bank allowances".

13 The EU has introduced the term" Total number of allowances in circulation (TNAC)" (EU (2019)), which for our purpose is equivalent with private banking of allowances.

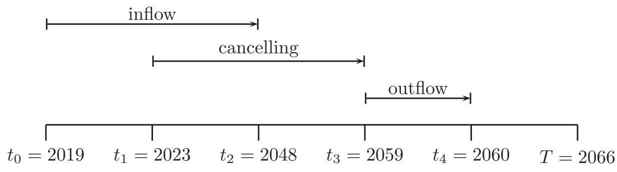

Figure 1. Time line for MSR.

Figure 1 presents the timeline for the MSR in our calibrated model.[14] From 2019 (t_0) to 2048 (t_2), the MSR takes in allowances, reducing the amount auctioned (as mentioned above, the intake rate is reduced from 24 to 12% in 2024). The intake stops in 2048 as the bank drops below 833 $MtCO_2$. From 2023 (t_1) to 2059 (t_3), allowances in the MSR are canceled when they exceed the volume auctioned in the previous year. In 2059, the bank has dropped below 400 $MtCO_2$, and the MSR returns the remaining allowances into the market, for 1 year (t_4). The ETS lasts till 2066 (T) in our calibrated model. Below, throughout this section, we will use general notation t_i when we emphasize the mechanisms at work. When presenting quantitative numbers, we will refer to years.

We assume allowances have constant assets return $1 + r$ leading to the Hotelling's rule for prices (i.e. generalization of Equation (3)):

$$p_{t+1} = (1 + r)p_t. \qquad (12)$$

The ETS is in equilibrium when there are no left-over unused allowances. As the MSR is emptied before the end of the ETS, cumulative emissions are given by cumulative supply minus canceled allowances. Given the stages displayed in Figure 1, all additions to the MSR before t_2 become canceled one-to-one. In other words, if some policy or other economic changes move demand from early to late periods, so that banking in early periods increases, such a change in the demand path reduces cumulative emissions. We replicate Observation 1 in the context of the EU ETS:

Observation 2. *The change in cumulative emissions equals the change in banking in periods before t_2, multiplied by the shaving parameter (24% before 2024, 12% after).*

The observation tells us why the timing of demand shocks is important for the final effect on emissions. The mechanism is the same as in Section 2. Early demand reductions, while not 100% effective, still lead to an increase in banking and a strictly positive fall in emissions in the aggregate. Late shocks, on the other hand, when anticipated today, lead to a decrease in banking, and thus to an increase in emissions. The increase in demand in early periods is effectively taken from the allowances otherwise canceled from the MSR. Thus, net emissions increase in relation to the case where no future reduction in

14 Whereas the mentioned years are specific to our model, other quantitative assessments of the MSR also tend to find a similar timeline, i.e., an inflow phase partly overlapping with cancellation, followed by an outflow phase (e.g., Bruninx et al. (2020), Silbye and Sørensen (2019) and Perino and Willner (2017)).

emissions demand had occured. A green paradox arises. We will come back to the importance of anticipation in the next section (see Figure 4).

One important condition for our green paradox is that the demand-reducing policy implemented in the future is anticipated today, so a forward-looking agent takes the future drop in demand (or prices) into consideration when making decisions on banking. For "surprise demand reduction" in the second period, the result does not hold. This insight highlights the importance of policy announcement. Although the timing of a policy matters, the timing of its announcements matters as well.

3.2. Model calibration

We now calibrate a stylized, dynamic model of the EU ETS. We then use this model to quantify the effects discussed in Section 2.

The model is given by the equations in Appendix B, and is conceptually similar to the analytical model in Section 2 (but with more periods and more detailed modeling of the MSR). Here we focus on the specification and calibration of the demand function, which we specify as follows:

$$e_t(p_t;\lambda_t) = (a - bp_t)(1 + ct) + \lambda_t, \tag{13}$$

where a, b and c are the three parameters to be calibrated. a/b is the (constant) choke price (i.e. the price at which demand equals zero), $1/b$ is the initial slope of the inverse demand function and c determines how the demand function changes over time (for a given price). Parameter specifications are shown in Table 1 in the Appendix. Here we give a brief explanation of how the model is parameterized.

To estimate the demand function, the three parameters are disciplined using historic evidence. We require that the following three conditions are met: (i) the level of demand should be consistent with the observed price and demand combination in 2018; (ii) the simulated Base Case scenario, which includes the MSR, should have an initial price of 21.0 €/tCO_2 in 2019; and (iii) a simulated scenario that does not include the MSR should have an initial price of 7.5 €/tCO_2 in 2019. In other words, the model should be able to reproduce both the current ETS prices but also those before the new MSR rules were introduced. We take the real interest rate to be 5%.[15]

The calibration leads to a choke price of 221.5 €/tCO_2. Further, the annual shift in the demand function is –2.1% (of initial demand). Taken together, this means that

15 It is difficult to know what the appropriate interest rate should be. On the one hand, there exists a market for future allowances from which one can derive the discount rate, revealing low returns (for instance, at the time of writing, the future price in December 2026 is 7.2% higher than the future price in December 2021, indicating a nominal interest rate of 1.4% per year, cf. https://www.barchart.com/futures/quotes/CK*0/futures-prices). On the other hand, the future of the EU ETS is uncertain, suggesting a higher rate. Our choice of 5% is "middle of the road' compared with the literature.

demand (i.e. emissions) becomes zero by 2066. Annual gross supply (s_t) becomes zero after 2057, assuming a continuation of the linear reduction rate after 2020. For this reason, we calibrate the final year in which the EU ETS is operative to be 2066 in our calibrations.

3.3. Quantitative results: Baseline scenario

The model described above can easily be simulated to derive the EU ETS market equilibrium for the period 2019–66.[16] The outcome is shown in Figures 2 and 3.[17] Note that this should not be taken as a forecast of the EU ETS market. The purpose of this analysis is to examine the effects of demand-reducing policies at different points of time, given a possible but fairly realistic scenario for the future EU ETS market.

Figure 2 shows that supply exceeds demand until 2050 – which then reverses. Annual demand is equivalent with annual emissions, while supply refers to gross supply (s_t), that is, before taking into account interaction with the MSR. Initially, net supply is significantly below gross supply (see Figure 2), and also well below demand, owing to a large inflow into the MSR.

Figure 3 shows the stocks of allowance reserves, both privately held ("banking") and in the MSR. It also displays how allowances enter into, or are taken out of, the MSR, as well as the canceled allowances. There is a notable change in 2023 − 24, due to two important factors in those years: cancelation of allowances begins in 2023 (t_1) and the withdrawal rate drops from 24 to 12% in 2024. The latter explains the decline in allowances entering the MSR in 2024 (labeled "MSR-in" in Figure 3, labeled m_t in Equation (19)), corresponding to the increased net supply (Figure 2). In this scenario, the MSR stops taking in allowances after 2048 (t_2), increasing net supply the next year (Figure 2). Cancelation of allowances is clearly the biggest in 2023, but continues for more than three decades in this scenario. In total, 6.9 Gt of allowances are canceled until cancelation ends in 2059 (t_3), of which 3.6 Gt are canceled by 2030.[18]

3.4. Quantitative results: Effects of demand-reducing policy

We now turn to the main purpose of the numerical analysis, which is to examine the effects on cumulative emissions of a demand-reducing policy. We consider policies that reduce demand in a given year t ("reduction year") by 1 million EUAs (corresponding

16 The model is simulated using the MCP solver in GAMS (Brooks et al., 1996). The GAMS program is provided in Appendix D.

17 By assumption, the ETS price starts at 21 Euro per ton in 2019, and reaches 208 Euro in 2066 (due to Equation (12)).

18 As a comparison, RefinitivCarbon (2018) expects 3.3 Gt to be canceled by 2030, and a total surplus of allowances of 1.6 Gt in 2030 (banking in the market plus MSR) implying further cancellation post-2030, especially since that study predicts a rising surplus in the market between 2025 and 2030.

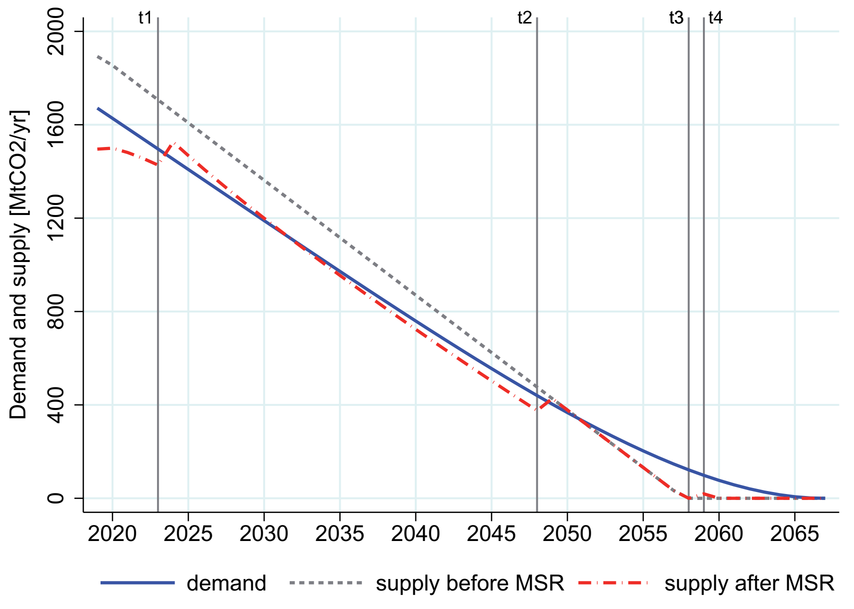

Figure 2. Market balance in Baseline scenario. Annual figures for the period 2019–66.

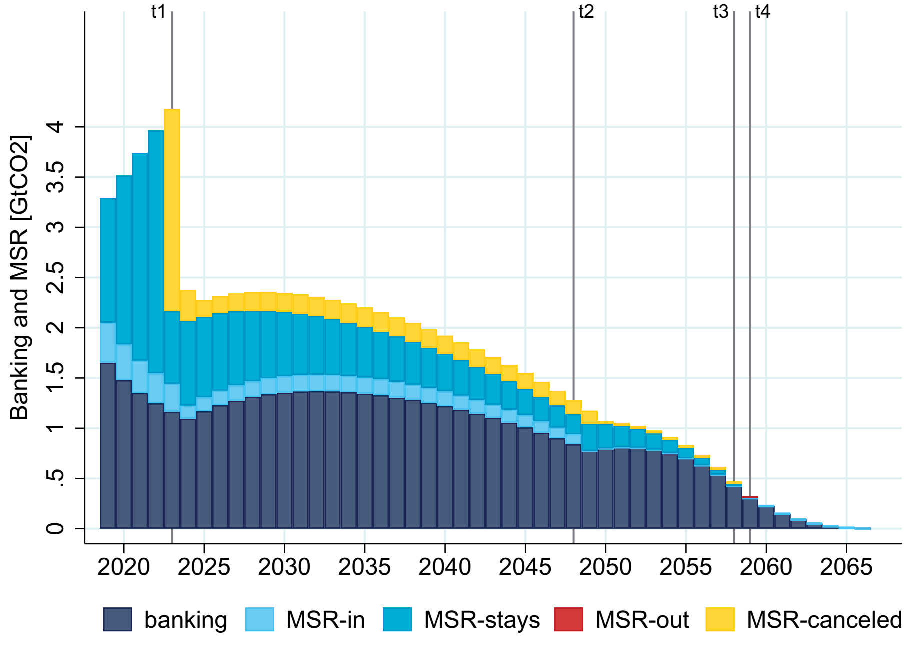

Figure 3. Stocks of allowances. The MSR is divided into the following four contents (cf. Equation 19): Input of allowances into MSR this period (m_t, "MSR-IN"); other allowances that remain in the MSR next period ("MSR stays"); allowances that leave MSR next period (n_t,"MSR-OUT"); and allowances that are canceled ("MSR Canceled"). Annual figures for the period 2019–66 in Baseline scenario. For the meaning of year labels $t1$, $t2$, $t3$, $t4$, see Figure 1.

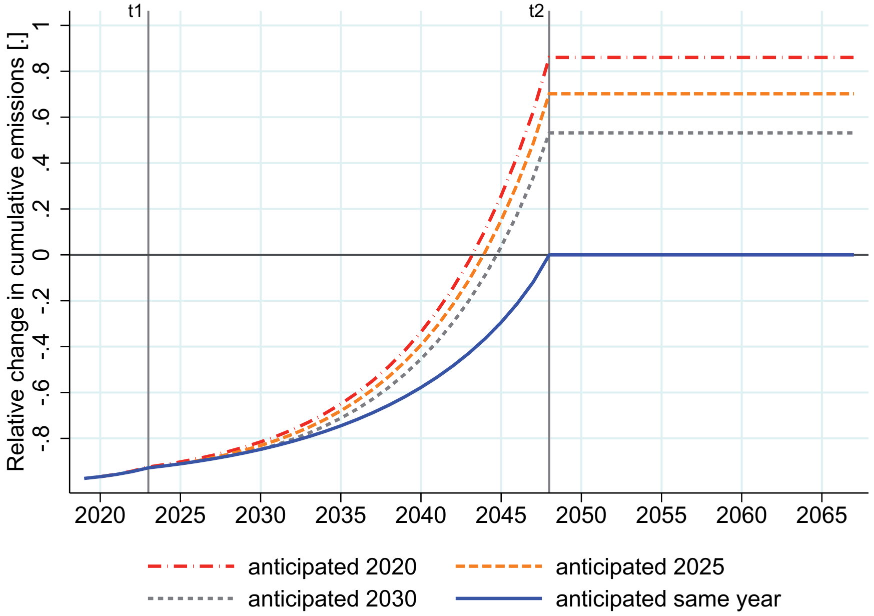

Figure 4. Effects on cumulative emissions of a demand-reducing policy that reduces demand by 1 million EUAs (1 $MtCO_2$) in the"reduction year" t, with the policy anticipated in year $s \leq t$. Years $t1, t2$ refer to start of canceling in the MSR and the end of the intake, respectively.

to 1 $MtCO_2$). Moreover, the announcement of the policy can take place in any year s ("announcement year") up to the year when the demand reduction takes place ($s \leq t$).

Figure 4 shows the effect on cumulative emissions of such a demand-reducing policy. On the horizontal axis, we have the reduction year t. The curve "Announcement 2020" shows the effects on cumulative emissions of announcing the policy in 2020 ($s = 2020$), and we have similar curves for $s = 2025$ and $s = 2030$. The fourth curve shows the effects of announcing the policy the same year ($s = t$).

We first notice that a demand-reducing policy announced and realized in 2020 will reduce cumulative emissions quite substantially (relatively speaking). A decrease in emissions in 2020 by 1 $MtCO_2$ will reduce cumulative emissions by 0.97 Mt. The intuition is, as explained by Perino (2018), that less emissions in 2020 lead to more banking over many years, which further increases the inflow into the MSR, and subsequently more cancelation of allowances.

Next, we see from the fourth (solid) curve that we get a similar but less pronounced effect as long as the demand-reducing policy is announced in the same year, that is, until 2048 (t_2). Afterwards, the MSR does not take in more allowances (in our scenario, cf. Figure 3), which means that from 2048 onwards the supply of allowances is fixed. The reason why the effect on cumulative emissions is the biggest in the early years is that there are more years with additional inflow into the MSR when banking is increased early on.

If the demand-reducing policy is announced years before it is realized, the effects are quite different though. For instance, if the policy is announced in 2020, but realized in 2048 or later ($t \geq t_2$), the net effect of the policy is to increase cumulative emissions by 0.86 Mt (according to our simulations). That is, the policy has quite the opposite effect of what is intended as it increases rather than decreases total emissions. Hence, a green paradox. The intuition is that when agents in the ETS market foresee a less tight market in the future, it becomes less profitable than before to bank allowances from the preceding periods. With less banking, fewer allowances enter the MSR, and thus fewer allowances become canceled. Moreover, when fewer allowances are taken out of the market, this further reduces the market tightness – hence, there is a multiplier effect which is bigger the longer the MSR is taking in allowances.

If the announcement is made in 2025 (or 2030), the effects on cumulative emissions are still perverse, but to a lesser degree as banking before 2025 (or 2030) is not affected. This illustrates the importance of policy announcement. It is not only the timing of the policy that matters, but also the timing of announcement.

We also see from the figure that if the demand-reducing policy takes place in year $\hat{t}$, where $\hat{t}$ is only a few years before the MSR stops taking in allowances (t_2), it can still have a perverse effect on cumulative emissions (if it is announced several years in advance). In this case, there will be less banking before, and more banking after, year $\hat{t}$. Hence, fewer allowances enter the MSR before year $\hat{t}$, whereas more allowances enter after year $\hat{t}$. If year $\hat{t}$ is quite close to t_2, the first effect dominates, and hence the net effect on cumulative emissions is positive.

3.5. Quantitative results: Multiple equilibria

In Proposition 2 we noted that distinct equilibria can exist, given the trigger points and discrete jumps in supply. Here we want to investigate this issue in the context of the numerical model of the EU ETS. As we will see, the calibrated demand function indeed supports three distinct equilibria. One equilibrium has been used in the subsections above (i.e. the calibrated baseline scenario), the others have a slightly higher price path.

When looking into this, it is useful to consider the level of banking at the end of the last period, considering different starting prices producing Hotelling-consistent price paths. The outcome of the exercise is shown in Figure 5 for the first-period price interval 20–22 euro per tCO_2. In equilibrium final banking must equal zero.

At first thought, we would expect net banking to be a monotonically increasing function of the price, as a higher price increases abatement and hence reduces demand for allowances. We see from the figure that net banking is only piecemeal increasing in the price, however, and then drops down at certain price levels. Moreover, we notice that there are three distinct first period prices where net banking at the end of the last period is zero, one at 21.0, one at 21.3 and one at 21.4 euro per tCO_2 (marked with small

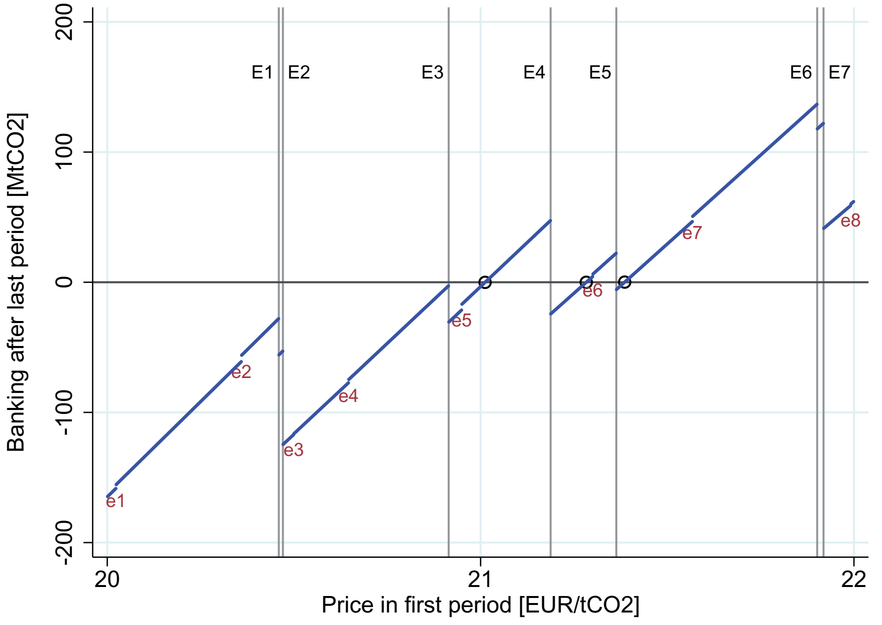

Figure 5. This figure shows banking after the last period as a function of the initial price, and illustrates the multiplicity of equilibria generated by the MSR. By definition, an equilibrium is characterized by the intersection of the banking curves with the horizontal line at 0. In this particular case, three equilibria exist: one at a first-period price of 21.0, the other at a price of 21.3 and 21.4. The seven vertical lines at discrete jumps in the banking function indicate major MSR-events. The eight minor discontinuities labeled by e1–e8 indicate minor MSR-events. More details in Appendix C.

circles). In other words, all these three prices (price paths) are feasible equilibria given the calibrated demand function.

When comparing the distinct equilibrium price paths, we observe that these are rather close to one another, suggesting equilibrium multiplicity is not an important problem in terms of magnitude. The minor difference in (initial) prices is somewhat deceptive, though. A slightly higher initial price (e.g. 21.0 versus 21.3) leads to slightly more banking, which can cause a substantial jump in cumulative cancelation of allowances and therefore emissions. Indeed, cumulative cancelation is close to 200 Mt higher in the equilibrium supported by an initial price of 21.4 euro per tCO_2 compared with the equilibrium with a starting price of 21.0 euro per tCO_2. The net decrease in emissions of nearly 200 Mt is roughly equal to Dutch CO_2 emissions in 2019 or four times Norway's. We provide a more detailed discussion in Appendix C, where we also show the impacts on cancelation.

4. POLICY IMPLICATIONS

What are the implications of our results for the design of cap-and-trade schemes? On the one hand, it is intuitively desirable to allow for endogenous cap-adjustments in cap-and-trade schemes – the EU experience with a large oversupply of allowances serves as a good illustration. On the other hand, the present paper establishes that a cap-and-trade scheme with an endogenous cap suffers from a green paradox. What avenues are there to reconcile these two observations?

First, and most intuitively, the system could match any demand-reducing policy with a (sufficient) decrease in the future supply of allowances. This reduction in supply directly avoids the green paradox. Such a solution is not without complications. The MSR was intended to avoid the need for discretionary policy-making through manual adjustment of supply. For years, the European Commission had been worried about the steadily increasing bank of unused emissions allowances and understood something had to be done about it. The MSR was introduced as a solution to the perceived over-supply of allowances without the need for ad hoc supply-adjustments and the political difficulties involved. Hence, simply complementing reduced future demand by a reduction in future supply, while an academically proper solution indeed, may well be politically difficult. Moreover, the ETS remains sensitive, in the counter-intuitive direction, to expectations about future demand driven by, for example, drifting consumer preferences.

Second, as a more drastic change, the European Commissions might add a price targeting mechanism to the MSR. Revealed preferences of EU policy makers suggest that they are not fond of too low allowance prices, as that makes it obvious that the ETS does not significantly contribute to EU climate policy. On the other hand, they seem to be afraid that tightening the allowance supply "too much" will lead to a carbon price that is unacceptably high for voters and firms in the EU. If those are the basic political economy forces shaping the design of the ETS, how could we make the best of the system?

In a simple price-focused setting, canceling can be triggered when prices fall below a floor price, like in RGGI or as has been proposed for the EU ETS. As discussed in our theoretical model, such discrete events introduce multiplicity and thus unpredictability when the equilibrium comes close to a trigger event. As a fix, one could devise more sophisticated (continuous) rules that implement an upwards sloping "marginal damage curve" for climate change under uncertainty (Gerlagh and Heijmans, 2018). Under such a policy, canceling would decrease, and cumulative emissions would rise, continuously with prices. A well-designed hybrid price-quantity policy along those lines prevents the green paradox. A drop in demand, independently of when it occurs and whether or not it is anticipated, lowers the price of allowances and increases canceling. This policy therefore reduces cumulative supply unambiguously. It establishes a negative feedback loop between demand and supply and thereby maintains effectiveness of complementary climate policies. As an additional benefit, it reduces price volatility substantially.

Although the above suggestions concern canceling *within* the MSR, a future revision of the ETS must also consider the exchange *between* the market and the MSR, that is, the intake and outflow. We raise two points in this regard. First, we see no clear benefits from discrete jumps, while we do see important disadvantages. We therefore propose a change toward continuous rules rather than moving discrete lumps of allowances in and out of the MSR in response to trigger events. Second, we believe that the flows of allowances between the market and the MSR serve a different purpose than the cancelation rules. Their set-ups should therefore be guided by a different principle. Cancelation is meant to insure an efficient balance of supply and demand. Subject to our proposed reforms of the MSR, we think the EU ETS would indeed achieve this balance. The flows of allowances, in addition, have an effect on market liquidity. We believe these should also be considered by the policy maker. On the one hand, a large bank of privately held allowances turns price volatility into asset risks.[19] On the other hand, a small bank of privately held allowances causes a collapse of intertemporal trade, which *causes* price volatility. The latter type of induced price volatility is illustrated by experiences in the South Korean ETS. The rules for the flows of allowances between the market and the MSR should thus aim at sufficient but not excessive market liquidity. To try and reach this balance, flows could be made responsive to the ratio between the amount of reserve allowances held by firms versus those surrendered.

In the end, our results may leave one wonder whether the endogenous cap in the EU ETS is a good thing at all. This is an issue regarding the relevant counterfactual of the policy and can be split into two subquestions. One is whether cancelation through the MSR reduces emissions in general, the other whether the EU ETS interacts properly with other emissions-reducing policies. Our answer to the first of these is affirmative: canceling allowances in the MSR reduces cumulative emissions compared with the functioning of the EU ETS before 2018. The green paradox identified in this paper expressly pertains to the interaction of the EU ETS with other policies intended to bring down emissions. It is those policies that, when interacted with the MSR, may become counterproductive. Future revisions of the EU ETS should therefore pay careful attention to the interaction of the rules for cancelation with other policies. Price-based cancelation rules would make for a useful addition to the ETS toolkit.

5. CONCLUSIONS

This paper establishes that a cap-and-trade system with an endogenous emissions cap like in the EU ETS suffers from a strong green paradox: cumulative emissions may increase in response to overlapping policies that reduce demand. Our analysis highlights the importance of anticipation, an expected shift in consumer preferences, or currently announced policies aiming to reduce emissions in the future, run the risk of being severely impaired if not more than overturned, whereas surprise policies may still be (somewhat) effective. Our green paradox is even stronger than the one previously pointed to in the economics literature (Sinn, 2008).

19 This is particularly relevant as it adds a layer of firms' interests to future changes in ETS regulation that is not so easy to gauge.

That preannounced policies may be less effective than the "surprise" ones, is not a new insight, nor is it limited to the case of environmental policy. In fiscal policy, for instance, preannouncement of policies has been found to substantially decrease their *net* effect (Auerbach and Gorodnichenko, 2012; Mertens and Ravn, 2012). Monetary policy is another such example (Sheehan, 1985). Our finding of a strong green paradox only underlines further the importance of carefully considering new policies, including how and when to communicate them, especially so if this communication takes place in advance of actual implementation.

A particular case in point to which our result applies is the European Green Deal. Presented by the European Commission in December 2019, this policy pledges to a 50–55% reduction in greenhouse gas emissions by 2030, increasing to 100% by 2050. The mechanism highlighted in our paper speaks directly to this proposal. Market participants, anticipating a policy-induced plunge in demand for allowances by 2030, attach less value to allowances beyond then, let alone 2050. Consequently, more allowances will be used today, leading to a reduced bank. This automatically reduces inflow into the MSR, and thus leads to less cancelation. In the coming decade, fewer allowances may be permanently canceled, increasing aggregate ETS emissions as compared with the situation where no Green Deal had been enacted. One can come up with several solutions to this dismal result, as discussed in the previous section.

One crucial assumption behind our analysis is that the market has perfect foresight about the future ETS market. This is a strong assumption, but we believe that the mechanism underlying our result is highly relevant also with imperfect foresight. Still, an important question is to what degree market participants let expectations about the future affect their current decisions (Fabra and Reguant, 2014; Kollenberg and Taschini, 2019). Incorporating different forms of expectations into our model framework would be one interesting avenue for future research.

ACKNOWLEDGEMENTS

We are grateful for help from two anonymous referees and the editors for their constructive suggestions on the structure of the paper and discussion section. We acknowledge funding from The Research Council of Norway through CREE (Grant 209698 to R.G. and K.E.R.) and the NorENS project (Grant 280987 to K.E.R.).

CONFLICT OF INTEREST

There are no conflict of interest.

A. PROOFS AND FIGURE FOR THE TWO-PERIOD MODEL

PROOF OF PROPOSITION 1:

Proof. Totally differentiating the price Equation (2) gives

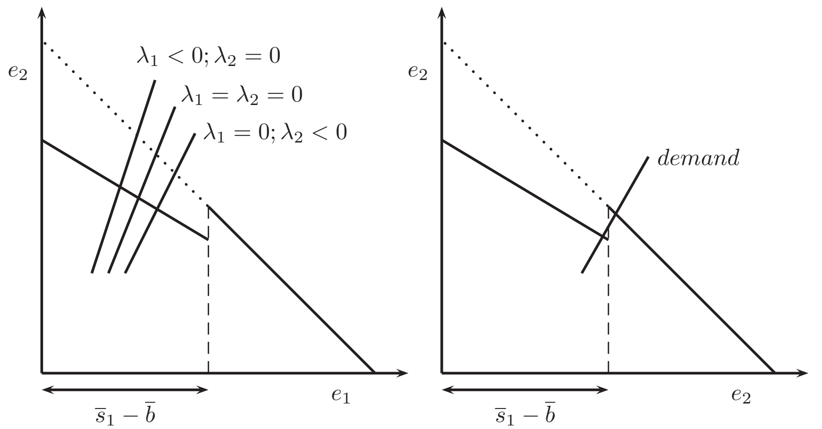

Figure 6. Equilibrium. Left panel presents demand shock dampening and green paradox as in Proposition 1. Right panel shows multiplicity of equilibrium as in Proposition 2. Upwards sloping lines represent demand satisfying Hotelling's rule (3). Downwards sloping lines represent supply as in Equations (5)–(7).

$$dp_1 = \psi_1'(de_1 - d\lambda_1), \tag{14}$$

$$dp_2 = \psi_2'(de_2 - d\lambda_2). \tag{15}$$

We cancel p_1 and p_2 and merge both equations into one, by Hotelling's rule (3).

$$(1 + r)\psi_1'(de_1 - d\lambda_1) = \psi_2'(de_2 - d\lambda_2). \tag{16}$$

Then we substitute e_1 for e_2 through aggregation of the allowances balances Equation (5), (7) over both periods, and taking differences,

$$de_2 = -(1 - \delta')de_1, \tag{17}$$

resulting in

$$((1 + r)\psi_1' + (1 - \delta')\psi_2')de_1 = (1 + r)\psi_1' d\lambda_1 - \psi_2' d\lambda_2. \tag{18}$$

Together with Equation (4), (9), this gives the proposition's equations.Q.E.D.

PROOF OF PROPOSITION 2:

Proof. Without loss of generality, assume that an equilibrium exists just below the canceling jump, $b = \bar{b} - \varepsilon$, supported by p_1^*, and with cumulative emissions E_1^*. If we slightly raise prices, first-period demand goes down and banking goes up. When $b = \bar{b}$ is reached, aggregate supply drops discretely, by $\delta\bar{b}$, resulting in strict excess demand. We have to further raise prices to find a new equilibrium. With those higher prices, banking in the first period further increases, thus cumulative supply further decreases. This proves that if we find a new equilibrium, it must satisfy $E^* < E^{**} - \delta\bar{b}$. Assuming supply does not become negative with rising prices, a second equilibrium with the stated properties must exist. Q.E.D.

FIGURES SUPPORTING THE TWO-PERIOD MODEL

For readers interested in a rigorous yet intuitive understanding of the above two proofs of propositions, below we add a graphical representation with emissions in the two periods on the axes. Upwards sloping lines represent demand satisfying Hotelling's rule (3). That is, the lines represent demand (e_1, e_2) for a set of possible prices (p_1, p_2) that satisfy $p_2 = (1 + r)p_1$. If prices go up, demand in both periods goes down. If prices go down, demand in both periods goes up. Hence, the curve is upwards sloping. The downwards sloping lines represent supply as in Equations (5)–(7). If banking falls short of the threshold, $b < \overline{b}$, so that $e_1 > \bar{s}_1 - \overline{b}$, then cumulative supply is fixed by $\bar{s}_1 + \bar{s}_2$. If banking exceeds the threshold $(e_1 > \bar{s}_1 - \overline{b})$, cumulative supply drops by a discrete amount and decreases with reductions in first-period demand, that is, the slope of the supply curve is decreasing at $<45°$.

In the left panel, the central line of the three demand lines is the benchmark, with no policies, $\lambda_1 = \lambda_2 = 0$. Demand shifts left if demand is reduced in the first period $(\lambda_1 < 0)$, and this reduces cumulative emissions since the supply curve is decreasing at $<45°$. A demand-decreasing policy in the second period, $\lambda_2 < 0$, that is a shift of the demand curve down or to the right, must increase emissions, for the *same* reason: the supply curve is decreasing at $<45°$.

B. EU ETS MODEL DETAILS

B.1 MODEL STRUCTURE

For our quantitative model of the EU ETS, we consider time periods $t \in \{1, \ldots, T\}$ and refer to the entire time window if not stated otherwise. We use capitals for stocks at the end of a period (so that a stock at the start of the first period has index 0), and lower-case variables for flows. The stock of allowances in the MSR is defined through the following mechanical rule:

$$M_t = \min(\beta s_{t-1}, M_{t-1}) + m_t - n_t, \tag{19}$$

where

$$(m_t, n_t) = \begin{cases} (0, \min(M_{t-1}, \Gamma)) \quad \text{if} B_{t-1} < \underline{B} \, (0,0) \text{if} \underline{B} \leq B_{t-1} < \overline{B} (\alpha B_{t-1}, 0) \text{if} \underline{B} \leq B_{t-1}, \\ \end{cases} \tag{20}$$

with s_t the maximum (exogenous) number of allowances issued in period t, β the share of these auctioned, B_t banking from period t to $t+1$, and m_t and n_t flows into and out of the MSR, respectively. The model can be parameterized to the EU ETS by setting $\beta = 0.57$,[20] $\Gamma = 100$, $\underline{B} = 400$, $\overline{B} = 833$, $\alpha = 0.24$ (0.12 from 2024). If M_t exceeds βs_t, the difference is shaved off, and these allowances are canceled permanently.

Equilibrium is characterized through demand e_t, supply s_t, and flows into and out of the MSR. Excess supply is added to the bank of allowances available for future use B_t.

$$B_t - B_{t-1} = s_t - e_t(p_t; \lambda_t) - m_t + n_t \tag{21}$$

As before, the one-dimensional parameter λ_t is a demand shifter, through which we study comparative dynamics. It captures the structure of the economy, also describing changes brought about by climate-oriented or other policies. We keep the same notation as in the general model and normalize the parameter λ_t such that $\partial e_t / \partial \lambda_t = 1$.[21] By means of notation, we will abbreviate $\partial e_t / \partial p_t$ as d'_t, so $d'_t < 0$. We define cumulative emissions as $E = \sum_t e_t$.

As allowances are complementary mostly to fossil fuel use, demand is bound from above and well-defined for zero prices. We also set a finite choke price, where no emissions are profitable anymore (e.g. fossil fuels are replaced by renewables).[22]

Given the above structure, the full EU-ETS model is characterized by the parameters presented below.

B.2 MODEL PARAMETRIZATION

Table 1 displays the specification of parameter values in the model. Several of the parameters are either specified by the policy, or based on historic observations (i.e. emissions and banking). The last four parameters in Table 1 are uncertain but important. The main text explains the calibration procedure. Here some more details are provided.

First, for the (real) interest rate, 5% is chosen. There are arguments for both higher and lower rates. Looking at futures prices of EUAs suggest a lower interest rate, even in nominal terms.[23] At the time of writing, the annual futures prices increase by 3–4% in the period 2020–25. On the other hand, the future of the EU ETS is uncertain, and recurring regulatory changes enhance the future price uncertainty. This suggests a high market interest rate (or a gradually higher interest rate to reflect that regulatory uncertainty increases over time, especially between phases).[24]

---

20 This is the approximate share of allowances that are auctioned, according to Perino (2018).

21 We could, for example, specify $D(.) + \lambda_t$ as residual demand, but we like to think of policies in a more generic framework.

22 The prices at which emissions become unprofitable may not be as excessively high as previously believed, see for example, Wilson and Staffell (2018) and Gillingham and Tsvetanov (2019).

23 https://www.barchart.com/futures/quotes/CK*0/all-futures.

24 An alternative approach could be to assume (partly) myopic behaviour by the market participants.

Table 1. Specification of parameter values

Parameter	Description	Value	
$\overline{B}$	Threshold for inflow into MSR	833 Mt	
$\underline{B}$	Threshold for outflow from MSR	400 Mt	
$\bar{\alpha}$	Withdrawal rate (pace of inflow into MSR)	0.24	(2019–23)
		0.12	(After 2024)
Γ	Outflow from MSR	100 Mt	
β	Threshold factor for canceling allowances	0.57	
s_{2019}	Supply of allowances in 2019	1,893 Mt	
	Linear reduction factor of supply per year	–0.0174	(Until 2020)
		–0.0220	(After 2020)
B_{2018}	Banking end of 2018	1,654 Mt	
	Size of MSR end of 2018	1,640 Mt	
	First year of cancelation	2023	
a	Maximal demand in first year	1,846 Mt	
b	Demand function slope in first year	8.336 Mt/€	
c	Relative decrease in demand per year	–0.0206	
r	Discount rate	0.05	

Mainly based on data from European Commission (https://ec.europa.edu/climat/policies/ets/reform_en and https://ec.europa.eu/clima/policies/ets/cap_en), Perino (2018) and RefinitivCarbon (2018).

Next, as mentioned in the main text we require three features to be fulfilled when calibrating the demand function. First, the level of demand (emissions) should be consistent with the observed price and demand combination in 2018. The average EUA price in 2018 was 16.0 euro per ton. Emissions in 2018 were 1,749 Mt.[25]

Second, the simulated Base Case scenario, which includes the MSR rules, should have an initial price in 2019 at 21.0 euro per ton. This is equal to the average price in the last quarter of 2018 (when adjusting for the interest rate). The EUA price was rising steadily in the three first quarters of 2018, whereas the price trend afterwards has been quite flat (the price has been volatile though).

Third, a simulated scenario that does not include the MSR rules should have an initial price in 2019 at 7.5 euro per ton. The average price from the start of phase 3 in 2013 to the first half of 2017, that is, just before the price started to take off, was 5.8 euro. Adjusting for the (real) interest rate of 5% and inflation rate of 1.5%, this corresponds to 7.5 euro in 2019.

As mentioned in the main text, the calibration leads to a choke price of 221.5 euro per ton and an annual reduction factor for demand of 2.1%, which is of the same size as the reduction factor the EU applies for supply.

More generally, it is difficult to know how price responsive demand is, and it is hard to foresee how the demand function will change over time. On the one hand, economic growth tends to push the demand upwards. On the other hand, technological progress and supplementary policies related to renewables, energy efficiency and

25 https://ec.europa.eu/clima/news/emissions-trading-emissions-have-decreased-39-2018_en.

coal phase-out, tend to push the demand downwards. The calibration might suggest that market participants in aggregate believe the latter to be dominating the former.

As explained in the main text, the end year of the EU ETS follows from the calibration, and turns out to be 2066.

Regarding the initial size of banking and MSR, 1,654 million allowances were banked in the market from 2018.[26] In total, 900 million allowances were "backloaded" in 2014–16, which means that auctioning of these allowances was postponed (implicitly banked by the regulator). Eventually, it has been decided that they should enter into the MSR, together with expectedly 740 million allowances (RefinitivCarbon, 2018).

C. MULTIPLICITY OF EQUILIBRIA

What are the details behind the multiplicity of equilibria in Section 3.5? It is useful to first consider the drop in net banking at the price of 21.2 (Event 4 (E4) in Figures 5 and 7). When the initial price is around 21.2, the level of banking falls below the threshold of 833 Mt in 2048. Hence, no more allowances enter into the MSR the following year. If the initial price is below 21.2, banking never exceeds the threshold

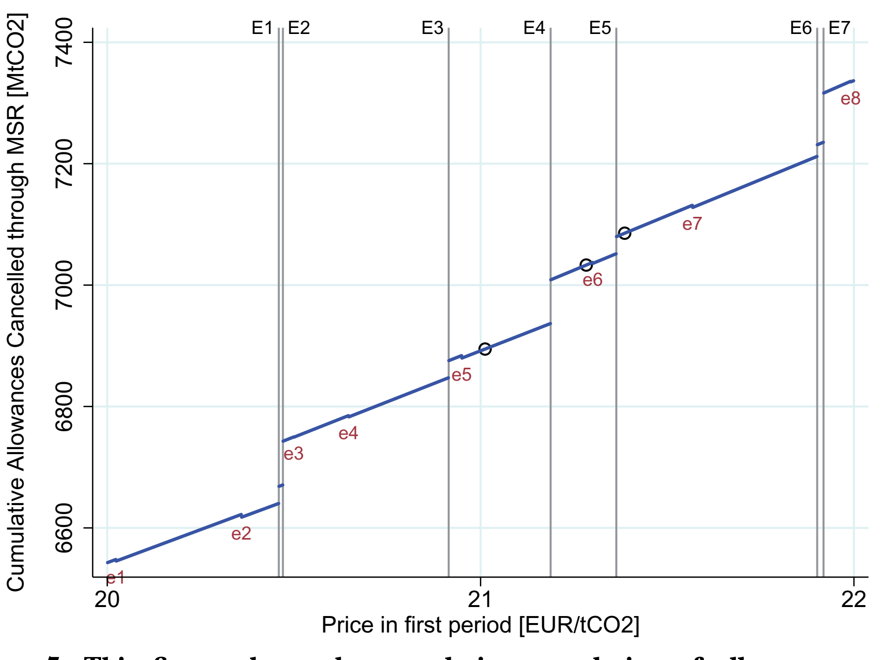

Figure 7. This figure shows the cumulative cancelation of allowances as a function of the initial price, and relates to Figure 5. Whenever one of the MSR thresholds is passed, cumulative cancelation shifts up or down.

26 https://ec.europa.eu/clima/sites/clima/files/ets/reform/docs/c_2019_3288_en.pdf.

again. If the initial price is 21.2, however, banking rises slightly above the threshold once more in 2049. Hence, 100 Mt ($0.12 * 833$) more allowances enter into the MSR (instead of being auctioned) compared with the case where the initial price is just below 21.2, and thus fewer allowances are available in the market. Net banking at the end of the last period is therefore lower even though the price path is (marginally) higher.

As the size of the MSR increases by 100 Mt in 2050 when the initial price is 21.2, but not if it is just below 21.2, it follows that 100 Mt more allowances are (not) shaved off if the initial price is equal to (just below) 21.2. In Figure 7, we see indeed that cumulative cancelation jumps considerably around the initial price of 21.2, but not as much as 100 Mt. The reason is that the other MSR threshold also plays a role here, that is, when allowances should return to the market (400 Mt). If the initial price is just below 21.2, banking in 2057 is above the threshold, while if the price is 21.2, banking is below the threshold. Only in the latter case are allowances released from the MSR the following year, in which case there is no more cancelation. In the former case, the size of the MSR is somewhat above the cancelation threshold, and 28 Mt more allowances are canceled before the threshold is passed the year after. This mitigates to some degree what happens in 2049, and so the net difference in cancelation is 72 Mt ($100 - 28$).

A similar story explains the drop in net banking when the initial price is around 20.5 (Event 2) or 21.9 (Event 7), that is, there is 1 more year of inflow into the MSR when the initial price is marginally above the stated price compared with when it is marginally below (the 400 Mt threshold also plays a role in these cases). For the other and smaller drops in net banking in Figure 5 (E1, E3, E5 and E6), only the outflow threshold plays a role.[27]

D. GAMS PROGRAM

Sets

* EU ETS is simulated for the years 2019-2067. $t = 0$ is 2018, so $t = 49$ is 2067. Both demand and supply are zero from year 2067 according to the calibration.

t Time period/0*49/

t0(t) Period $t = 0$ (before simulation starts)

ts(t) Simulation periods

27 The figures also show eight minor events: small jumps up and down in banking and cancellation, respectively (e1–e8). In these cases, banking first drops below the 833 Mt threshold and then rises above the threshold again 1 or 2 years later, lasting only 1 year implying 1 more year of inflow into the MSR. If the price is marginally *above*, for example, the e1 price, banking is *marginally* above the threshold 1 year later, while if the price is marginally *below* the e1 price, banking is *significantly* above the threshold 2 years later. Hence, there is more inflow and subsequent cancellation in the latter case. Note that in the simulations in Sections 3.4–3.5 there is no such “pause” in the inflow into the MSR.

```
ts2(t) Simulation periods except t=1
tn(t) Last period
;

alias(t,tt);
alias(t,ttt);

t0(t) = yes$(ord(t) eq 1) ;
ts(t) = yes$(ord(t) gt 1) ;
ts2(t) = yes$(ord(t) gt 1 and ord(t) lt card(t)) ;
tn(t) = yes$(ord(t) eq card(t)) ;

Scalars
r Discount rate
beta Threshold for canceling allowances (as a share of s)
p0 Average price in 2018 (t=0)
d0 Demand in 2018 (t=0)
apar Parameter a in demand function
bpar Parameter b in demand function
cpar Parameter c in demand function
;

r = 0.05 ;
* Assumed share of auctioning
beta = 0.57 ;
* Average price in 2018 used to calibrate demand function
p0 = 16 ;
* Demand (incl aviation) in 2018
d0 = 1749 ;
bpar = 1/0.1175 ;
apar = (d0 + p0*bpar) ;
cpar = -0.020566 ;

Parameters
s(t) Fixed allocation of quotas
alpha(t) Withdrawal rate–share of annual auction volume entering into MSR
deltaD(t) Reduced demand for quotas in year t
;

* Supply (incl aviation) from 2019 based on https://ec.europa.eu/clima/policies/ets/cap_en
s(t)$(ord(t) le 3) = 1931 - (ord(t)-1)*38.264;
s(t)$(ord(t) gt 3) = s("2") - (ord(t)-3)*49.216;
alpha(t)$(ord(t) le 6) = 0.24 ;
```

```
alpha(t)$(ord(t) gt 6) = 0.12 ;
deltaD(t) = 0 ;

Positive Variables
p(t) Price
d(t) Demand for quotas
CumD Cumulative demand for quotas
m_in(t) Number of quotas entering into MSR
m_out(t) Number of quotas taken out of MSR and into the ETS market
M(t) Size of MSR
C(t) Cancellation of quotas
CumC Cumulative cancellation of quotas
;

Variables
B(t) Banking of quotas
;

Equations
EQ1(t) Quotas entering into MSR
EQ2(t) Quotas taken out of MSR
EQ3(t) Cancellation of quotas
EQ4(t) MSR stock change
EQ5(t) Market balance
EQ6(t) Price movement
EQ7(t) Demand for quotas

* The following equations sum up cumulative cancellation and demand:
EQ3SUM Cumulative cancellation of quotas
EQ7SUM Cumulative demand for quotas

* The following equation is used in the model without MSR:
EQ3NO(t) Without cancellation of quotas from MSR
;

* Due to the discontinuity of the m_in function, the formulation is somewhat different from the equation in the paper,
* and a marginal number is added to the denominator to avoid division by zero
EQ1(t)$(ts(t)).. m_in(t) =E= MAX(0, alpha(t)*B(t-1)*(B(t-1) – 833))*(B(t-1) – 833) / ((B(t-1) – 833)*(B(t-1) – 833) + 0.01) ;

* Due to the discontinuity of the m_out function, the formulation is somewhat different from eq.2 in the paper,
* and a marginal number is added to the denominator to avoid division by zero
```

EQ2(t)$(ts(t)).. m_out(t) =E= MIN(M(t-1),(MAX(0, 100*(400 – B(t-1)))*(400 – B(t-1)) / ((400 – B(t-1))*(400 – B(t-1)) + 0.01)));

EQ3(t)$(ord(t) gt 4).. C(t) =E= MAX(0 , M(t) – beta*s(t)) ;

EQ4(t)$(ts(t)).. M(t) =E= M(t-1) + m_in(t) – m_out(t) – C(t-1) ;

EQ5(t)$(ts(t)).. s(t) – m_in(t) + m_out(t) =E= d(t) + B(t) – B(t-1) ;

EQ6(t)$(ts2(t)).. p(t + 1) =E= p(t)*(1+r) ;

EQ7(t)$(ts(t)).. d(t) =E= (apar – bpar*p(t))*(1 + cpar*(ord(t)-1)) - deltaD(t) ;

EQ3NO(t).. C(t) =E = 0 ;

EQ3SUM. CumC =E= sum(t,C(t)) ;

EQ7SUM. CumD =E= sum(t,d(t)) ;

* Main model with MSR: Model MSR_YES/EQ1.m_in, EQ2.m_out, EQ3.C, EQ4.M, EQ5.p, EQ6.B, EQ7.d, EQ3SUM.CumC, EQ7SUM.CumD /;

* Model without MSR: Model MSR_NO/EQ1.m_in, EQ2.m_out, EQ3NO.C, EQ4.M, EQ5.p, EQ6.B, EQ7.d, EQ3SUM.CumC, EQ7SUM.CumD /;

* Model to help GAMS find solution, with fixed price and endogenous banking last period Model MSR_EX /EQ1.m_in, EQ2.m_out, EQ3.C, EQ4.M, EQ5.B, EQ7.d, EQ3SUM.CumC, EQ7SUM.CumD / ;

* The initial value of MSR and B:
M.fx(”0”) = 900 + 740 ;
B.fx(”0”) = 1654 ;

* Fixing variables in period 0 (2018):
m_in.fx(”0”) = 0 ;
m_out.fx(”0”) = 0 ;
d.fx(”0”) = 0 ;
C.fx(t)$(ord(t) le 4) = 0 ;

* Last period requirements:
M.fx(t)$tn(t) = 0 ;
B.fx(t)$(ord(t) eq card(t)) = 0 ;

* Ensure that prices must be strictly positive:
p.lo(t) = 0.1;

option iterlim = 100000000;
option reslim = 2000.0;

```
option limrow = 10;

*Help GAMS find the wanted equilibrium (due to multiple equilibria)
P.fx(t) = 20.5*(1 + 0.5*r)*(1+r)**(ord(t)-2);

Solve MSR_EX using mcp;

* Then relax prices and require last period banking to be zero
P.lo(t) = 0.1;
P.up(t) = inf ;
B.fx(t)$(ord(t) eq card(t)) = 0 ;
* Solve the model including MSR: Solve MSR_YES using mcp;

******************************
* Without MSR (and backloading)
alpha(t) = 0 ;
B.fx("0") = B.l("0") + M.l("0");
M.fx("0") = 0 ;

* Solve the model excluding MSR:
Solve MSR_NO using mcp;
```

REFERENCES

Abrell, J. and S. Rausch (2017). 'Combining price and quantity controls under partitioned environmental regulation', *Journal of Public Economics*, 145, 226–42.

Aldy, J. E., A. J. Krupnick, R. G. Newell, I. W. Parry and W. A. Pizer (2010). 'Designing climate mitigation policy', *Journal of Economic Literature*, 48, 903–34.

Auerbach, A. J. and Y. Gorodnichenko (2012). 'Measuring the output responses to fiscal policy', *American Economic Journal: Economic Policy*, 4, 1–27.

Bauer, N., C. McGlade, J. Hilaire and P. Ekins (2018). 'Divestment prevails over the green paradox when anticipating strong future climate policies', *Nature Climate Change*, 8, 130.

Böhringer, C., K. E. Rosendahl and H. B. Storrøsten (2017). 'Robust policies to mitigate carbon leakage', *Journal of Public Economics*, 149, 35–46.

Borenstein, S., J. Bushnell, F. A. Wolak and M. Zaragoza-Watkins (2019). 'Expecting the unexpected: emissions uncertainty and environmental market design', *American Economic Review*, 109, 3953–77.

Brooke, A., D. Kendrick and A. Meeraus (1996). GAMS: A User's Guide. GAMS Development Corporation: Washington DC.

Bruninx, K., M. Ovaere and E. Delarue (2020). 'The long-term impact of the Market Stability Reserve on the EU Emission Trading System', *Energy Economics*, 89, 104746.

Bulow, J. I., J. D. Geanakoplos and P. D. Klemperer (1985). 'Multimarket oligopoly: strategic substitutes and complements', *Journal of Political Economy*, 93, 488–511.

Bushnell, J. B., H. Chong and E. T. Mansur (2013). 'Profiting from regulation: evidence from the European carbon market', *American Economic Journal: Economic Policy*, 5, 78–106.

EU (2019). Communication from the commission. publication of the total number of allowances in circulation in 2018 for the purposes of the market stability reserve under the EU emissions trading system established by directive 2003/87/ec. Technical report. https://ec.europa.eu/clima/sites/clima/files/ets/reform/docs/c_2019_3288_en.pdf.

EU Commission (2019). Communication on the European green deal. https://ec.europa.eu/info/publications/communication-european-green-deal_en.

Fabra, N. and M. Reguant (2014). 'Pass-through of emissions costs in electricity markets', *American Economic Review*, 104, 2872–99.

Fell, H. (2016). 'Comparing policies to confront permit over-allocation', *Journal of Environmental Economics and Management*, 80, 53–68.

Fowlie, M., M. Reguant and S. P. Ryan (2016). 'Market-based emissions regulation and industry dynamics', *Journal of Political Economy*, 124, 249–302.

Gerlagh, R. (2011). 'Too much oil', *CESifo Economic Studies*, 57, 79–102.

Gerlagh, R. and R. Heijmans (2018). Regulating stock externalities. CESifo Working Paper, 7383.

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

Gerlagh, R. and R. J. R. K. Heijmans (2019). 'Climate-conscious consumers and the buy, bank, burn program', *Nature Climate Change*, 9, 431.

Gillingham, K. and T. Tsvetanov (2019). 'Hurdles and steps: estimating demand for solar photovoltaics', *Quantitative Economics*, 10, 275–310.

Golosov, M., J. Hassler, P. Krusell and A. Tsyvinski (2014). 'Optimal taxes on fossil fuel in general equilibrium', *Econometrica*, 82, 41–88.

Hintermann, B., S. Peterson and W. Rickels.(2016). 'Price and market behavior in Phase II of the EU ETS: a review of the literature', *Review of Environmental Economics and Policy*, 10, 108–28.

Hotelling, H. (1931). 'The economics of exhaustible resources', *Journal of Political Economy*, 39, 137–75.

Kollenberg, S. and L. Taschini (2019). 'Dynamic supply adjustment and banking under uncertainty in an emission trading scheme: the market stability reserve', *European Economic Review*, 118, 213–26.

Martin, R., M. Muûls, L. B. De Preux and U. J. Wagner (2014). 'Industry compensation under relocation risk: a firm-level analysis of the EU emissions trading scheme', *American Economic Review*, 104, 2482–508.

Mertens, K. and M. O. Ravn (2012). 'Empirical evidence on the aggregate effects of anticipated and unanticipated US tax policy shocks', *American Economic Journal: Economic Policy*, 4, 145–81.

Morris, S. and H. S. Shin (1998). 'Unique equilibrium in a model of self-fulfilling currency attacks', *American Economic Review*, 88, 587–97.

Perino, G. (2018). 'New EU ETS phase 4 rules temporarily puncture waterbed', *Nature Climate Change*, 8, 262.

Perino, G. and M. Willner (2016). 'Procrastinating reform: the impact of the market stability reserve on the EU ETS', *Journal of Environmental Economics and Management*, 80, 37–52.

—— (2017). 'EU-ETS Phase IV: allowance prices, design choices, and the Market Stability Reserve', *Climate Policy*, 17, 936–46.

RefinitivCarbon (2018). Eua price forecast: a new era for European carbon or calmer waters ahead? Technical report, Refinitiv.

Roberts, M. J. and M. Spence (1976). 'Effluent charges and licenses under uncertainty', *Journal of Public Economics*, 5, 193–208.

Rosendahl, K. E. (2019). 'EU ETS and the waterbed effect', *Nature Climate Change*, 9, 734–5.

Sheehan, R. G. (1985). 'Money, anticipated changes, and policy effectiveness', *The American Economic Review*, 75, 524–9.

Silbye, F. and P. B. Sørensen (2019). National climate policies and the European emissions trading system, in *Climate Policies in the Nordics*. *Nordic Economic Policy Review*, Nordic Council of Ministers, Copenhagen, Denmark, pp. 63–111.

Sinn, H.-W. (2008). 'Public policies against global warming: a supply side approach', *International Tax and Public Finance*, 15, 360–94.

Szabo, M. and B. Garside (2020). 'Germany unveils long-awaited coal phaseout bill after cabinet approval', *Carbon Pulse*, 29 January 2020.

Van Huyck, J. B., R. C. Battalio and R. O. Beil (1990). 'Tacit coordination games, strategic uncertainty, and coordination failure', *The American Economic Review*, 80, 234–48.

Weitzman, M. L. (1974). 'Prices vs. quantities', *The Review of Economic Studies*, 41, 477–91.

Wilson, I. G. and I. Staffell (2018). 'Rapid fuel switching from coal to natural gas through effective carbon pricing', *Nature Energy*, 3, 365.