Linking Cap And Trade Schemes

Roweno J.R.K. Heijmans*
Reyer Gerlagh†

Tilburg University

For the most recent version, click here.

June 14, 2021

Abstract

Recent years have witnessed a rapid increase in the number of cap and trade schemes to mitigate greenhouse gas emissions. With multiple coexisting schemes, their linkage has become a topic of interest. This paper offers a theory of optimal linking under uncertainty. We show that efficient linkage adjusts the joint cap in response to observed trade between the schemes. It lowers allowance price volatility and increases global welfare by efficiently adjusting total emissions in response to private information. Interestingly, while asymmetric information generally harms welfare, asymmetric uncertainty can be exploited to increase welfare. Optimally linked cap and trade schemes expand the range of model parameters for which cap and trade is favored over a carbon tax.

JEL codes: D82, D83, H23, Q52, Q54

Keywords: asymmetric information, policy updating, asymmetric uncertainty, carbon markets, linking

1 Introduction

The number of cap and trade schemes to mitigate greenhouse gas emissions has grown steadily.[1] Reduced to its core, a cap and trade scheme caps CO_2 emissions by allocating

*Email: r.j.r.k.heijmans@uvt.nl

†Email: r.gerlagh@uvt.nl.

[1]The first major emissions trading system (ETS) for greenhouse gases – the European Emissions Trading System (EU ETS) – was established in 2005. To date, there are 20 ETSs in place across five

allowances to emitters who are then allowed to trade their permits; if a firm emits CO_2, it must surrender an equivalent amount of allowances. The policy combines a conservative certainty on emissions offered by direct command-and-control measures with an efficient allocation of abatement efforts realized through a carbon tax.

Cap and trade schemes can link. A linkage between two schemes reciprocally enables the use of permits issued in one scheme to meet compliance obligations pursuant to another. Linking is seen as a promising development in cap and trade regulation (Mehling et al., 2018). Article 6 of the Paris Agreement expressly provides for the possibility of linking. Linking has become increasingly prominent in recent years. California's cap and trade system linked with Quebec's on 1 January 2014 and the schemes organize joint auctions. On 1 January 2020, a link between the European Union's Emissions Trading System (EU ETS) and the Swiss Emisions Trading System came into force. Linkages between the Regional Greenhouse Gas Initiative (RGGI) and the Emissions Trading Systems of Virginia and Pennsylvania are currently on their way, as are implicit linkages between California's ETS and Washintong State's Clean Air Rule.[2]

Linking is efficient because it leads to an equalization of marginal abatement costs across jurisdictions. An additional benefit may be that, through their increased cooperation, local planners are less likely to choose their policies noncooperatively (Mideksa and Weitzman, 2019), leading to a more efficient total emission levels.

This paper proposes a theory of optimal linking under abatement cost uncertainty. Linking two trading systems means that a firm can fulfill compliance obligations pursuant to one jurisdiction by means of emission allowances issued by the other. In practice this means that polluters in either jurisdiction are allowed to trade emission allowances freely and on a one-to-one basis. Free inter-scheme trade of allowances ensures that firms equate marginal abatement costs both within and across schemes. This is necessary for an optimal linkage because as long as marginal abatement costs are not equal across jurisdictions, mutually beneficial exchanges of allowances can be made, contradicting the notion of efficiency. The intuition is essentially the same as that favoring cap and trade over more direct command and control policies in a single jurisdiction. The literature on linking has often emphasized this channel of potential efficiency gains (Carbone et al., 2009; Flachsland et al., 2009b; Doda and Taschini, 2017; Mehling et al., 2018; Doda et al., 2019; Holtsmark

continents and covering 27 jurisdictions which produce almost 40 % of global wealth (GDP). With over a dozen more governments considering or having already scheduled an ETS, emissions trading has emerged as a key instrument to cost effectively decarbonize our economies. (ICAP, 2020)

[2]The latter link is mostly hypothetical at this point, as Washinton State's Clean Air Rule was suspended after a 2018 court ruling. Though contested, the ruling was largely upheld by the Washington State Supreme Court on Jan. 16, 2020.

and Weitzman, 2020).

In addition to marginal abatement cost equalization, an optimal linkage exploits a second channel to boost the performance of the linked schemes: learning. Allowances are traded between the schemes when marginal abatement costs differ across jurisdictions. The number of allowances traded between the schemes thus provides a sufficient statistic for the gap in marginal abatement costs under the initial allocation. Knowledge of this gap allows the planners to learn something about realized abatement costs in both jurisdictions. Since this posterior will generally be different from their prior, the planners thus learn that the initial allocation of emission allowances was inefficient for two reasons. First, the distribution of allowances between the schemes, given the global cap, was suboptimal. This inefficiency is dealt with through trading. But second, the cap itself may turn out ex post inefficient from the point of view of global welfare. An optimal linkage attempts to deal with both of these inefficiencies.

To see how a flow of allowances between jurisdictions signals in inefficient global cap, note that the level of the ideal cap balances the cost of climate change and the cost of abatement. When posterior beliefs about abatement costs deviate from planners' expectations, that balance is lost. In this sense, emission levels are ex post inefficient and the planners know it, not because they possess perfect information, but because they update their beliefs based on net trade between the schemes. We therefore propose to adjust the global cap in response to trade between the jurisdictions. To our knowledge, we are the first to study such endogenous cap adjustments in the context of linked cap and trade schemes.

Interestingly, what drives cap adjustments is not the amount of uncertainty per se, but rather the degree to which uncertainty is asymmetric between jurisdictions. This has a clear intuition. Trade flows signal a wedge in jurisdictional marginal abatement costs. Learning about the *absolute* level of abatement costs occurs when planners update beliefs, mostly for the least predictable jurisdiction. Updating, in turn, is done by anchoring beliefs about the least predictable jurisdiction on beliefs about the most predictable one. In the extreme case where the planners are uncertain about abatement costs in one jurisdiction only, trade flows allow them to pin down abatement costs in the uncertain jurisdiction exactly. Since in this case there is de facto complete information about abatement costs, the planners are able to implement the first-best social optimum. That is, by smartly linking emission trading schemes, aggregate uncertainty can be reduced to match the least uncertain scheme.

Indeed, optimal linking deviates from emissions traded one-for-one (or, in terms of climate change, ton-for-ton), when jurisdictions have asymmetric uncertainty (Holland and Yates, 2015). But the efficiency gain associated with cap flexibility comes with a substantial loss of allocative efficiency. To illustrate, suppose we were to contemplate an exchange rate such that 1 allowance issued by jurisdiction N can be traded against 2 allowances issued in jurisdiction S. Then firms in N and S will trade allowances until the marginal cost of reducing emissions in N equals *half* the marginal cost of reducing emissions in S. A non-unitary exchange rate thus drives firms' incentives away from an efficient distribution of abatement efforts.[3]

Cap adjustments based on observed allowance prices is another alternative policy. Compared to quantities, prices are highly efficient information aggregators. Indeed, when the planners observe both emissions and the market price for allowances, they can perfectly back out the abatement cost functions in both jurisdictions. The possibility of using price information hence allows for implementation of the unconstrained best cap on emissions. Our first and foremost recommendation is therefore to adjust the caps in response to carbon prices. Importantly, in practice proposed price-based interventions are often discrete and based on price thresholds; the implementation then remains imperfect and our quantity-based optimal linkage can perform better as it allows for a continuous processing of market information.[4]

Linking cap and trade systems across jurisdictions is related to integrating cap and trade markets over time (Yates and Cronshaw, 2001). Dynamic linking was studied in Heutel (2020) and Pizer and Prest (2020) for flow pollutants, and in Gerlagh and Heijmans (2020) for stock pollutants.[5]

Although we use the language of multiple cap and trade schemes, our analysis also applies to situations where a previously uncovered industry is newly added to an existing scheme, or to cases where multiple unregulated industries are combined into a newly formed cap and trade scheme. Our results are therefore relevant in discussions on such topics as the inclusion of road transport among the industries covered by the EU ETS, or on the extension of RGGI beyond the electricity sector. Similarly, our analysis motivates the question whether clearly identifiable jurisdictions within existing schemes – member

[3]The same need not apply to local pollutants like NO_X or lead pollution.

[4]An important example to which our theory applies is any linkage involving the EU ETS. Despite repeated calls for a European price floor (Flachsland et al., 2020), the European Union uses only information on quantities for cap-adjustments. Linkages between any cap and trade scheme and the EU ETS will hence benefit from our analysis, "key features for compatibility for linking" being "complications around intervention in price" (EU Commission, 2015, EU ETS Handbook)."

[5]For EU ETS, Gerlagh and Heijmans (2019) and Gerlagh et al. (2021) illustrate several unexpected side-effects of endogenous intertemporal emission caps. This offers an important warning: endogenous cap-adjustments can be efficient, but details matter.

states in the EU ETS, or countries in RGGI – are currently "linked" in the most efficient way possible. These are highly relevant policy questions that deserve greater attention from policymakers and academics alike.

The paper is organized as follows. Section 2 introduces our model and the building blocks for welfare analysis. Section 3 discusses different types of (integrated) cap and trade policies and develops our theory of optimal linking. At the end, we revisit the perennial question of instrument choice, prices versus quantities (Weitzman, 1974). Section 4 concludes. Proofs and lengthy derivations are in the Appendix.

2 Model

Given are two jurisdictions, North and South, each operating its own cap and trade scheme. The assumption of two jurisdictions is not restrictive. One may simply consider North a representative jurisdiction for two linked jurisdictions, East and West, where trade between North and South is essentially a reduced-form way of writing trade between East, West, and South. The term jurisdiction, too, should not be narrowly interpreted.

Each jurisdiction i is populated by firms that produce a composite good the production of which causes emissions. We may write benefits $B_i(\tilde{e}_i; \theta_i)$ as a function of emissions $\tilde{e}_i$, given by:

$$B_i(\tilde{e}_i \mid \theta_i) = (p^* + \theta_i)(\tilde{e}_i - e_i^*) - \frac{b_i}{2}(\tilde{e}_i - e_i^*)^2. \tag{1}$$

Emissions yield benefits because they allow firms to produce goods and save on the cost of abatement. As an umbrella term, we refer to $B_i(\tilde{e}_{|}\theta_i)$ as abatement costs in jurisdiction i, but other interpretations are possible. For notational convenience, we normalize benefits relative to the ex-ante optimal allocation e_i^* and prices $p_i = p^*$. The parameter θ_i is a fundamental of jurisdiction i's economy and is *private information* of its constituent firms, though it is common knowledge that $\mathbb{E}[\theta_i] = 0$, $\mathbb{E}[\theta_i^2] = \sigma_i^2$, and $\mathbb{E}[\theta_N \theta_S] = \rho\sigma_N\sigma_S$. One way to think about this parameter and the fact that it is unobserved by the policymakers is in terms of uncertainty in the (residual) demand for emission allowances (Borenstein et al., 2019). The variance σ_i^2 is a measure for the uncertainty about jurisdiction i's economy. We say that uncertainty is asymmetric if $\sigma_N \neq \sigma_S$.

Emissions accumulate in the atmosphere and cause climate change as a global externality. The cost of climate change is given by:

$$C(\tilde{e}_N + \tilde{e}_S) = p^*(\tilde{e}_N + \tilde{e}_S - e_N^* - e_S^*) + \frac{c}{2}(\tilde{e}_N + \tilde{e}_S - e_N^* - e_S^*)^2. \tag{2}$$

Costs are also measured relative to the ex-ante optimum e_i^*. We make the simplifying assumption that emissions have local benefits but global costs. One could imagine more complicated settings where emissions also have strictly regional costs like local air pollution (Caplan and Silva, 2005; Antoniou and Kyriakopoulou, 2019). We abstract from these aspects and focus on the optimal linking of cap and trade schemes to regulate a global externality.

Subtracting the costs of climate change from the benefits due to emissions yields welfare:

$$W = B_N(\tilde{e}_N \mid \theta_N) + B_S(\tilde{e}_S \mid \theta_S) - C(\tilde{e}_N + \tilde{e}_S). \tag{3}$$

Policies are set to align polluters' incentives with the global social cost of carbon (Kotchen, 2018). We interpret optimal linking as the integrated policy for e.g. an international climate treaty like the Paris Agreement or its successor. Yet even strictly local planners may more or less explicitly consider global climate damages when deciding on mitigation policies. There are at least two pieces of suggestive evidence to support our claim. First, most if not all existing cap and trade schemes are policymakers' attempts at meeting mitigation obligations implied by the Paris Agreement. Since the Paris Agreement is global in scope and intention, cap and trade schemes set up to satisfy the associated pledges explicitly or implicitly operate with a measure of global climate welfare in mind.[6] Second, actual linkages often start from a certain amount of cooperation and mutual agreement on the cap. Indeed, among the "essential criteria" to ensure "compatibility between the systems" mentioned up in Annex I of the *Agreement between the European Union and the Swiss Confederation on the linking of their greenhouse gas emissions trading systems* are the "ambition and stringency of the cap". Similarly, the independent but linked cap and trade schemes of California and Quebec organize joint auctions of allowances, implying a high degree of cooperation in determining the linked caps.

The timing of the game is as follows:

1. The planner of each jurisdiction chooses its local cap;
2. Firms in each jurisdiction i observe their benefit curve (θ_i) and choose their emissions subject to the local cap and whether or not schemes are linked;

[6] As an example, "[...] the [European] Commission proposed in September 2020 to raise the 2030 greenhouse gas emission reduction target, including emissions and removals, to at least 55% compared to 1990. [...] This will enable the EU to move towards a climate-neutral economy and implement its commitments under the Paris Agreement by updating its Nationally Determined Contribution." Moreover, "[t]o achieve the EU's overall greenhouse gas emissions reduction target for 2030, the sectors covered by the EU ETS must reduce their emissions by 43% compared to 2005 levels." Retrieved from https://ec.europa.eu/clima/policies/strategies/2030_en.

3. The planners observe the number of allowances surrendered in each jurisdiction and, depending on the type of policy in place, may buy back or sell additional allowances;

4. Emissions are realized and the game ends.

Observe that our game has one period and is static in that sense. However, within this one period decisions are taken in order so that we essentially study a dynamic framework that is not repeated. We also assume that that between the moment firms observe their benefit functions and the time the game ends, there are no further innovations in the fundamentals θ_i; that is, we study periods of a duration short enough to focus only on the asymmetric information problem without worrying about future uncertainty.

We write $e_i = \tilde{e}_i - e_i^*$ the deviation in emission levels from the ex-ante optimum; similarly we denote $p_i = \tilde{p}_i - p_i^*$. We let $\Delta e := \tilde{e} - e^{SO}$ be the difference between the value of e and its ex post (after observing θ_i) socially optimal value (see subsection 2.1). For total emissions, we write $\tilde{E} = \tilde{e}_N + \tilde{e}_S$ and for deviations from the ex-ante optimum $E = e_N + e_S$. We use superscripts for policy rules or scenarios.

Firms are profit maximizers. Once a policy k caps emissions at the level e_i^k, individual firms trade allowances until marginal abatement costs for all are equal to:

$$p_i^k = -b_i e_i^k + \theta_i, \tag{4}$$

which is firms' inverse demand for allowances. In a competitive market for allowances, the price at which permits are traded will be p_i^k in equilibrium.

2.1 Global Social Optimum

In a perfect world not plagued by an asymmetric information problem – the Social Optimum – a welfare-maximizing planner allocates emissions in each jurisdiction $i = N, S$ so that they satisfy:

$$\frac{\partial W}{\partial \tilde{e}_i} = \underbrace{\frac{\partial B_i(\tilde{e}_i \mid \theta_i)}{\partial \tilde{e}_i}}_{MB_i} + \underbrace{\frac{\partial B_i(\tilde{e}_j \mid \theta_j)}{\partial \tilde{e}_i}}_{0} - \underbrace{\frac{\partial C(\tilde{e}_i + \tilde{e}_j)}{\partial \tilde{e}_i}}_{MC} = 0, \tag{5}$$

where $j \neq i$ denotes the jurisdiction that is not i. Equation (5) immediately implies that marginal emission benefits should be equal across jurisdictions in an efficient allocation. Since both jurisdictions by assumption operate a cap and trade scheme, marginal benefits of emissions are also equal to the carbon price, so $p_N^{SO} = p_S^{SO} = p^{SO} = MB^{SO}$. Next, since the level of emissions if efficient if and only if marginal climate costs equal marginal benefits in either jurisdiction, we obtain the following conditions for the socially optimal emission levels:

$$c \cdot (e_N^{SO} + e_S^{SO}) = p^{SO} = -b_i e_i^{SO} + \theta_i. \tag{6}$$

Equation (6) characterizes the Social Optimum and represents three equalities in three unknowns: e_N^{SO}, e_S^{SO}, and p^{SO}. Solving for these, we obtain:

$$p^{SO} = \frac{c(b_S\theta_N + b_N\theta_S)}{cb_N + cb_S + b_N b_S}, \tag{7}$$

$$e_i^{SO} = \frac{b_{-i}\theta_i + c(\theta_i - \theta_{-i})}{cb_N + cb_S + b_N b_S}, \tag{8}$$

$$E^{SO} = \frac{b_S\theta_N + b_N\theta_S}{cb_N + cb_S + b_N b_S}, \tag{9}$$

where $i = N, S$ and $-i$ refers to "the other jurisdiction that is not i". As is intuitive, socially optimal emission levels are higher when abatement is more expensive; E^{SO} increases in θ_i. This basic observation will be useful later on, when we illustrate that trades of allowances between (linked) schemes signal private information about abatement costs to the planners. An optimal linkage exploits this information by adjusting the global cap in response.

For future reference, we note that the variance of prices is given by:

$$\mathbb{E}\left[\left[p^{SO}\right]^2\right] = \left(\frac{c}{c \cdot b_N + c \cdot b_S + b_N b_S}\right)^2 \left[b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S \rho \sigma_N \sigma_S\right]. \tag{10}$$

Thus, abatement cost uncertainty translates into price volatility. We will return to this later.

2.2 Welfare Losses

We rank policies according to their expected welfare levels. Suppose a policy k induces emission levels $\tilde{e}_i^k$ in jurisdiction i. From firms' equilibrium behavior (4), we see that deviations in emissions from the social optimum scale with prices:

$$\Delta^k p_i = -b_i \Delta^k e_i. \tag{11}$$

Expected welfare losses relative to the social optimum are then given by:

$$L^k = \mathbb{E}\left[\Delta^k B_N + \Delta^k B_S - \Delta^k C\right]$$

$$= \frac{c}{2}\mathbb{E}\left[\left(\Delta^k E\right)^2\right] + \sum_i \frac{b_i}{2}\mathbb{E}\left[\left(\Delta^k e_i\right)^2\right]. \tag{12}$$

Throughout the analysis, we interpret a policy as the implementation of a constrained expected welfare maximization problem. We note that the level of welfare in the Social Optimum, W^{SO}, is hypothetical and unaffected by the policy implemented. Thus, expected welfare maximization subject to a set of policy constraints can be treated as the dual problem of expected welfare loss minimization relative to the Social Optimum, subject to the same constraints. Given the equivalence between these two approaches, we may use them interchangeably for convenience.

Equation (12) shows that there are essentially two sources of welfare losses. One derives from an inefficiently high or low level of global emissions; this relates to the balance between the benefit of emissions and the cost of climate change. The second derives from an inefficient distribution of emissions given a global cap; this relates to the balance between the benefit of emissions in each jurisdiction. Conditional on any level of the global cap, whether efficient or not, welfare W is maximized if and only if

$$\frac{\partial W}{\partial \tilde{e}_i} = \frac{\partial W}{\partial \tilde{e}_j}, \tag{13}$$

for $i, j = N, S$ and $i \neq j$. Equation (13) implies that given any level of aggregate emissions, global welfare can be increased whenever marginal emission benefits are not equal across jurisdictions. If a policy leads to marginal benefit equalization, it guarantees that the resulting allocation of emission levels is at least second-best: it yields the highest level of global welfare given a (possible inefficient) global cap on emissions.

Linking policies as we study here exhibit equal prices across jurisdictions in equilibrium, as we shall explain below. By (11), emissions (both local and global) under such policies can therefore be expressed in terms of the common price gap Δp. Plugging this into (12), expected welfare losses from a policy featuring equal prices across jurisdictions can be written as:

$$L^k = \frac{1}{2}\frac{(cb_N + cb_S + b_N b_S)(b_N + b_S)}{b_N^2 b_S^2}\mathbb{E}\left[\left(\Delta^k p\right)^2\right]. \tag{14}$$

3 Policies

3.1 Regional cap and trade

The simplest policy operates two separate cap and trade schemes. In this case, the planner of jurisdiction i sets a cap e_i to maximize (3). The first-order condition this policy should implement equates expected marginal benefits in each scheme to marginal climate damages. Thus, emissions e_i are set such that

$$\mathbb{E}[MB_i \mid e_i] = MC, \tag{15}$$

resulting in emission levels capped at the ex-ante optimum $e_N = e_S = 0$. Plugging these, through (8)-(9), into (12), we obtain expected welfare losses when both jurisdictions operate a regional cap and trade scheme:

$$L^R = \frac{1}{2}\frac{(c+b_S)\sigma_N^2 + (c+b_N)\sigma_S^2 - 2c\rho\sigma_N\sigma_S}{cb_N + cb_S + b_N b_S}. \tag{16}$$

We note that the marginal damages (the RHS of (15)) are perfectly known when caps are set, whereas marginal abatement costs are stochastic variables due to the unobserved fundamentals θ_i. Thus, regional cap and trade implements socially optimal emission levels if and only if abatement costs turn out exactly as expected ($\theta_N = \theta_S = 0$). If abatement costs deviate from expectations, regional cap and trade is inefficient for two reasons. First, the allocation of abatement efforts between the jurisdictions may be ex post inefficient (abatement costs may differ between them). Second, the level of emissions is ex post inefficient. The first of these is remedied by linking schemes across jurisdictions.

3.2 Linking

When North and South link their cap and trade schemes, each planner sets the expected optimal cap for its jurisdiction but allowances issued in one scheme may be used to fulfill abatement obligations pursuant to another. Thus polluters are free to trade their allowances as long as global emissions are not affected:

$$e_N + e_S = E = 0. \tag{17}$$

Linking has attracted a lot of attention in recent years (Doda and Taschini, 2017; Mehling et al., 2018; Doda et al., 2019; Holtsmark and Weitzman, 2020). Just as trade in emission allowances between firms within a jurisdiction is an efficient way to achieve a given amount of local abatement, so the trade of allowances between jurisdictions is an efficient way to achieve a pre-determined amount of global abatement. The intuition is essentially the same: while linking does not affect global emissions, and therefore climatic damages, it does lead to a closer alignment of the benefits and costs of abatement across jurisdictions. In fact, as long as marginal abatement costs in the jurisdictions are *not* the same, firms can and will profitably exchange permits. The integrated market for emission allowances therefore reaches equilibrium only once the price of an allowance is the same in North and South:

$$p_N = p_S. \tag{18}$$

Moreover, since a firm in either jurisdiction is willing to sell (buy) an allowance against the going market price as long as the price is above (below) its private marginal abatement cost, marginal benefits are also the same across firms and jurisdictions when schemes are linked:

$$-b_N e_N^L + \theta_N = -b_S e_S^L + \theta_S, \tag{19}$$

where e_N and e_S are chosen by the firms conditional on θ_N and θ_S, subject to (17). Plugging this into our welfare loss-function (12), we obtain:

$$L^L = \frac{1}{2}\frac{1}{b_N + b_S}\frac{b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S \rho\sigma_N\sigma_S}{cb_N + cb_S + b_N b_S}. \tag{20}$$

Comparing welfare losses, we can now formally state our first substantive result.[7]

Proposition 1. *Linking cap and trade schemes increases global welfare.*

In contrast to regional cap and trade, linking guarantees that marginal abatement costs are equal in both jurisdictions. For this reason, linking is always weakly better for welfare than regional cap and trade. The planners' problem consists of two steps, each with its own an intuitive meaning. First, the planners of North and South cap local emissions at levels that, in expectations, maximize global welfare. When the local caps are set, allowances can be traded on a one-to-one basis between schemes, as long as emissions overall remain fixed at the sum of the two jurisdictional caps. By linking their schemes, the planners of North and South effectively delay the determination of local emission caps until after θ_N and θ_S are known, guaranteeing an ex post efficiency gain through marginal

[7]See appendix for derivations and proof.

abatement cost equalization

$$MB = MB_N = MB_S \tag{21}$$

In expectations thes equal marginal damages. Cumulative emissions $e_N + e_S$ is chosen such that:

$$MB = MB_N = MB_S \tag{22}$$

$$\mathbb{E}[MB \mid e_N + e_S] = MC. \tag{23}$$

Linking benefits global welfare (Proposition 1) but the effects on individual jurisdictions are ambiguous. To see this, note that prices in equilibrium equate marginal benefits, so that the volatility of prices is equal to the volatility of marginal benefits. Hence if abatement costs in South are much less predictable than in North, $\sigma_S > \sigma_N$, North may import part of the price volatility to which South is subject. If this effect is strong enough, North may be harmed by its linkage with South (Holtsmark and Weitzman, 2020; Habla and Winkler, 2018).

Linking suffers from another, more implicit type of inefficiency. Suppose that after trading the planners observe emissions levels e_N^L and e_S^L in jurisdictions N and S, respectively. Since firms will trade allowances between the jurisdictions until marginal abatement costs are equal everywhere, we know that (19) must be satisfied under the observed emission levels e_N^L and e_S^L. Rewriting (19), the planners therefore learn $\mu = \theta_N - \theta_S$. Conditional on μ, planners update their beliefs on the true marginal abatement cost curve in both of the jurisdictions. Since this posterior will generally deviate from their priors, the expected optimal cap on emissions should ideally respond to the observed trade of allowances. It does not under standard linking.

3.3 Optimal Linking

We will now construct our optimal linking policy. We proceed in two steps. First, we derive the expected optimal emission level conditional on the trade flows between the jurisdictions. Second, we formulate a mechanism that is known to all firms and allows the planners to implement the expected optimal cap for any observed trade of allowances.

As we discussed, when post-trading emissions levels are e_N and e_S in jurisdictions N and S, respectively, the planners can back out μ, the vertical distance between regional marginal abatement cost functions:

$$\mu := b_N e_N - b_S e_S = \theta_N - \theta_S. \tag{24}$$

A key observation is that μ contains more information than the relative position of jurisdictions' abatement cost functions alone – it also signals something about the *absolute* location of the curves. To see this most simply, suppose that the planners are uncertain only about θ_N whereas θ_S is perfectly known (i.e. suppose that $\sigma_S = 0$). In this hypothetical case, observing μ is clearly equivalent to observing θ_N directly. In the more general case where both θ_N and θ_S are unknown, such sharp posteriors are not possible. Still the planners know all the combinations of θ_N and θ_S consistent with μ. Depending on the regional uncertainties σ_N and σ_S, some of these combinations will be more likely than others. The planners can therefore calculate the expected marginal abatement cost in both jurisdictions, conditional on μ and/or post-trade emissions:

$$\begin{aligned}
\mathbb{E}[MB \mid \mu] &= \mathbb{E}[\theta_N \mid \mu] - b_N e_N = \mu \frac{\mathbb{E}[\mu\theta_N]}{\mathbb{E}[\mu^2]} - b_N e_N \\
&= \mu \frac{\sigma_N^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S} - b_N e_N \\
&= \frac{\sigma_N^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S} b_S e_S - \frac{\sigma_S^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S} b_N e_N.
\end{aligned} \tag{25}$$

When $e_i \neq 0$, $i = N, S$, expression (25) is troubling: the initial cap $e_N + e_S = E = 0$ is set at the level that is efficient conditional on $\theta_N = \theta_S = 0$, or $MB_N = MB_S = 0$. Upon learning $\mu \neq 0$, however, the planners no longer think that $MB_i = 0$ since their posterior belief $\mathbb{E}[MB \mid \mu]$ is 0 only if $e_N = e_S = 0$, see (25). Indeed, having learned μ the planners hold the following posterior beliefs on (θ_N, θ_S):

$$\begin{aligned}
\hat{\theta}_N^\mu &= \mathbb{E}[\theta_N \mid \mu] = \frac{\sigma_N^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}\mu = \frac{\sigma_N^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}(b_N e_N - b_S e_S) \\
\hat{\theta}_S^\mu &= \mathbb{E}[\theta_S \mid \mu] = -\frac{\sigma_S^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}\mu = \frac{\sigma_S^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}(b_S e_S - b_N e_N).
\end{aligned} \tag{26}$$

It is immediate that $\hat{\theta}_i^\mu = 0$ only if $\mu = 0$. To see why this signals an inefficient initial cap, recall that marginal climate damages due to emissions are equal to $c \cdot (e_N + e_S)$, which is 0 when $e_N + e_S = 0$. At the same time, conditional on μ expected marginal benefits from emissions are given by (25), and these 0 only if $\mu = 0$, meaning that for all $\mu \neq 0$ the initial cap is inefficient. Indeed, conditional on beliefs $(\hat{\theta}_N^\mu, \hat{\theta}_S^\mu)$ the expected optimal emission levels e_N^O and e_S^O are determined by:

$$c \cdot (e_N^O + e_S^O) = \hat{\theta}_N^\mu - b_N e_N^O = \hat{\theta}_S^\mu - b_S e_S^O, \tag{27}$$

which says that marginal damages from pollution must be equal to expected marginal benefits in both jurisdictions conditional on the planners' posterior beliefs $(\hat{\theta}_N^\mu, \hat{\theta}_S^\mu)$.

We seek a policy that implements (27) for all μ; that is, we want to solve the asymmetric information problem using aggregate market signals (Kwerel, 1977; Dasgupta et al., 1980). Moreover, like standard Linking the policy must guarantee that marginal benefits from emissions are equal across jurisdictions whether or not the planners' beliefs $(\hat{\theta}_N^\mu, \hat{\theta}_S^\mu)$ turn out correct:

$$\theta_N - b_N e_N^O = \theta_S - b_S e_S^O \quad \text{for all } \theta_N, \theta_S. \tag{28}$$

The policy that implements both (27) and (28) is called Optimal Linking.

We make two observations. First, if an Optimal Linking policy implements (27), then the information on marginal benefits contained in jurisdictions' choices of emissions levels is again summarized by $\mu = b_N e_N^O - b_S e_S^O = \theta_N - \theta_S$. Compared to standard Linking, though, the emission levels are optimal not with respect to the prior belief that $\theta_N = \theta_S = 0$ but with respect to the posterior belief $(\theta_N, \theta_S) = (\hat{\theta}_N^\mu, \hat{\theta}_S^\mu)$. Thus, an Optimal Linking policy adjusts the (global) cap in response to the private information revealed in emission levels.

Second, an Optimal Linking policy allows firms to exchange emissions allowances "ton-for-ton". This requirement is necessary for allocative efficiency: any other trading basis creates incentives for firms to exchange allowances beyond the point where the marginal benefit of emissions is equal across all firms and jurisdictions, contradicting (28). Thus, a "trading ratio" on allowances cannot be part of an Optimal Linking policy.

Combining these two observations, one can show that an Optimal Linking policy can be implemented through:

$$e_N^O + e_S^O = E^O = (1 - \delta) e_N^O, \tag{29}$$

subject to the constraint that allowances can be traded one for one. Equation (29) says that global emissions $e_N^O + e_S^O$ are capped at the level E^O, which itself is endogenous to regional emission levels through $E^O = (1 - \delta) e_N^O$. The parameter δ is endogenous to the structure of our model and given by:

$$\delta = \frac{b_N[\sigma_S^2 - \rho\sigma_N\sigma_S] + c[\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S]}{b_S[\sigma_N^2 - \rho\sigma_N\sigma_S] + c[\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S]}. \tag{30}$$

We refer to δ as the *cap-adjustment rate* since it prescribes how the global cap on emissions should be adjusted in response to the demand for emissions in both jurisdictions. Interestingly, the cap-adjustment rate δ may be negative. When this happens, higher emissions in one jurisdiction translate into higher emissions in the other jurisdiction as well. The reason is intuitive: if abatement costs in one jurisdiction are very unpredictable yet strongly correlated to those in the other, high abatement costs in the latter are likely matched by equally high costs in the former.

Practically, Optimal Linking can be thought of as a policy that proceeds in three simple "steps":

1. The planners issue a total number $E = 0$ of allowances;
2. Firms exchange emission on a one-to-one basis both within and across jurisdictions;
3. Conditional on the net number of allowances traded, the planners buy or auction extra allowances until the total number of allowances available for use is E^O.

Proposition 2. *Optimal Linking is the best possible cap and trade policy using only information on quantities. Expected welfare losses under an Optimal Linking policy are given by:*

$$L^O = \frac{1}{2}\frac{b_N + b_S}{cb_N + cb_S + b_N b_S}\frac{(1-\rho)(1+\rho)\sigma_N^2\sigma_S^2}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}. \tag{31}$$

Proof. The first part of the proposition is true by construction, i.e. see (27) and (28). The expression for L^O requires a lengthy derivation and is relegated to the Appendix. □

Note that although an Optimal Linking policy operates via the pre-defined "rule" (29), it does not require us to assume that planners commit to this rule (*c.f.* Boleslavsky and Kelly, 2014). Once the planners observe μ, their incentive is to adjust the global cap to the level prescribed by our Optimal Linking policy since this is expected to be optimal conditional on their posterior beliefs. Indeed, if the planners had somehow agreed upon another emissions-based updating principle, they would want to deviate to our Optimal Linking rule.[8]

We emphasize that Optimal Linking policy does not involve a "trading ratio" on allowances (Holland and Yates, 2015). Trading ratios will perform strictly worse than Optimal Linking. To understand why, suppose the planners would impose a trading ratio $\alpha \neq 1$ on allowances such that an allowance worth one ton of emissions in North is

[8]Evidently we do assume that the policymakers commit to updating their policy using information on emissions only; within this class of policy instruments, no further commitment-assumptions are needed.

worth α tons on emissions in South. Under such a regime, firms would continue to trade allowances across jurisdictions up until the point where $\theta_N - b_N e_N = \alpha(\theta_S - b_S e_S)$, i.e. until the ratio marginal benefits from emissions between North and South is α. This is in direct contradiction with (28) and the notion of efficiency, which says that marginal benefits should be exactly equal across jurisdictions. Thus, even if a trading ratio could, in principle, implement the same expected optimal global cap E^O as an Optimal Linking policy, the resulting distribution of abatement efforts will be strictly less efficient.

An optimal linkage performs remarkably well. When uncertainty about abatement costs is strongly asymmetric across jurisdictions ($\sigma_N/\sigma_S \to 0$ or ∞), or when abatement costs are highly correlated ($\rho \to \pm 1$), optimal linking allows for welfare levels very close to the full information Social Optimum. This reflects the planners' great scope for learning in these cases.

Corollary 1. *Where the planners have perfect information about one of the two linked jurisdictions ($\sigma_i = 0$ for $i \in \{N, S\}$), or when abatement costs are perfectly correlated ($\rho = \pm 1$), Optimal Linking implements the first best levels of emissions.*

Even though asymmetric *information* leads to welfare losses, asymmetric *uncertainty* compensates for part (and, in extreme cases, all) of these losses.[9]

Another way to see the great advantage of an optimal linkage is to compare its welfare performance with that of a classic linking policy:

$$\frac{L^O}{L^L} = \frac{(b_N^2 + b_S^2 + 2b_N b_S)(1+\rho)\sigma_N\sigma_S}{b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S \rho\sigma_N\sigma_S} \cdot \frac{(1-\rho)\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}. \tag{32}$$

As we expect from Corollary 1, an optimal linkage performs much better than classic linking ($L^O/L^L \to 0$) when uncertainty is highly asymmetric or when abatement costs are strongly correlated ($\rho \to \pm 1$). If we believe that abatement costs are driven by macroeconomic conditions and available technologies they likely are strongly correlated across jurisdictions indeed – this makes a compelling case for optimal (rather than a standard) linking of the cap and trade schemes. Interestingly, though the absolute levels of welfare losses under both classic an optimal linking depend on climate damages through c, the relative performance of an optimal linkage is independent of the slope of climate damages.

[9]For a simple illustration of this point, suppose $b_N = b_S = b$ and $\rho = 0$. (i) When uncertainty is perfectly symmetric, i.e. $\sigma_N = \sigma_S = \sigma$, expected welfare losses under an Optimal Linking regime are strictly positive: $L^O = (1/2)(1/(b+2c))\sigma^2 > 0$. (ii) When uncertainty is highly asymmetric, i.e. $\sigma_N/\sigma_S \to 0$ or $\sigma_N/\sigma_S \to \infty$, expected welfare losses under an Optimal Linking regime vanish: $L^O \to 0$.

We recall that, though global welfare increases after establishing an optimal linkage, individual jurisdictions may still be worse off. A jurisdiction in which abatement costs are relatively predictable may, through linking, expose itself to the volatile abatement costs in the other jurisdiction and therefore import a variable allowance price (Holtsmark and Weitzman, 2020). To the extent that such concerns are important for real world linkages, our optimal linking policy offers some relief.

Proposition 3. *Optimally linked cap and trade schemes admit lower price volatility than classically linked cap and trade schemes:*

$$\mathbb{E}\left[\left(p^{OL}\right)^2\right] \leq \mathbb{E}\left[\left(p^{L}\right)^2\right]. \tag{33}$$

Though it is still possible that an individual jurisdiction experiences higher allowance price volatility after optimally linking cap and trade schemes, this effect (if it occurs) will be less severe than under standard linking. Note that with inter*temporal* trading of permits, an endogenous cap also reduces price volatility, see Gerlagh et al. (2020).

Optimal Linking is the best the planners can do when using only information on quantities to update their beliefs. From a theoretical viewpoint, the limitation to quantity-information is arbitrary. If the planners are willing to use both (post-trade) emissions and allowance prices, they can perfectly back out the marginal abatement cost function in each jurisdiction i (whether linked or not). To see this, we rewrite firms' inverse demand function (4) to obtain:

$$\theta_i = p_i + b_i e_i. \tag{34}$$

An ideal policy uses information on both prices and quantities to pin down θ_i in each jurisdiction $i = N, S$ and, given these, adjusts the caps so that emission levels end up in the Social Optimum. Our first and foremost recommendation is therefore to implement a policy along those lines.

It is important to note that the ideal instrument is continuous in prices and emissions. A simple price collar – often proposed in the context of cap and trade policies – will perform far worse. To our knowledge, there is no literature on optimal price collars. However, we conjecture that an Optimal Linking policy outperforms (optimal) price-collar-based cap adjustments when uncertainty is highly asymmetric or when abatement costs are strongly correlated (which likely they are since abatement costs are largely driven by technological developments and macroeconomic conditions). In these cases, an optimal linkage implements welfare levels very close to the Social Optimum, see Corollary 1.

3.4 Prices vs. Quantities

We saw how jurisdictions can optimally link their cap and trade schemes. As an alternative, each jurisdiction could instead levy a carbon tax. We now revisit the classic question of Weitzman (1974) on instrument choice.

We assume that emissions are taxed at an expected optimal rate – that is, taxes are set to minimize (12) subject to the constraint that, in equilibrium, firms will emit until the marginal benefit of emissions equal the tax. Recall that p^* is defined as the expected welfare-maximizing carbon price in each jurisdiction, so the expected optimal tax sets $p_N = p_S = 0$. We may therefore invoke (10) and (14) to derive expected welfare losses when both jurisdictions tax emissions:

$$L^{tax} = \frac{1}{2}\left(\frac{c}{b_N b_S}\right)^2 \frac{b_N + b_S}{cb_N + cb_S + b_N b_S}(b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S \rho\sigma_N\sigma_S). \tag{35}$$

All else equal, the expected welfare loss when jurisdictions tax emissions is increasing in c, the marginal climate damage. We can now compare (35) and (31), which yields the following proposition.

Proposition 4. *Optimally Linked cap and trade schemes are favored over Taxes if and only if:*

$$\frac{(1+\rho)\sigma_N\sigma_S}{b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S\rho\sigma_N\sigma_S}\frac{(1-\rho)\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S} < \left(\frac{c}{b_N b_S}\right)^2 \tag{36}$$

Inequality (36) is likely satisfied when both regions have strongly asymmetric uncertainty ($\sigma_N/\sigma_S \to 0$ or ∞) or when abatement costs are highly correlated ($\rho \to \pm 1$). This should not come as a surprise: Corollary 1 showed that precisely in these cases two optimally linked cap and trade schemes implement welfare levels very close to the complete information social optimum. Proposition 4 once again underlining the real-world relevance of optimal linking regime.

4 Discussion and Conclusions

4.1 Contributions and Limitations

This paper studies the problem of linking cap and trade schemes under uncertainty. We design a novel Optimal Linking policy that can greatly improve the global welfare performance of cap and trade schemes. Although the benefits and costs of linking cap and trade schemes are frequently discussed in the literature (Flachsland et al., 2009a,b; Doda and Taschini, 2017; Mehling et al., 2018; Doda et al., 2019; Holtsmark and Weitzman, 2020; Holtsmark and Midttømme, 2021), ours is the first paper to explore what an optimal linkage might look like.

The core principle of an Optimal Linking policy is to extract private information from firms' emission decisions and, in response, adjust the (global) cap on emissions according to a simple pre-defined rule. Our policy thus uses aggregate market signals to solve the asymmetric information problem (Kwerel, 1977; Dasgupta et al., 1980). Importantly, Optimal Linking does not use "trading ratios" on allowances (Holland and Yates, 2015) – though effectively causing an adjustment of the global cap indeed, a trading ratio incentivizes firms to trade allowances beyond the mere equalization of marginal benefits and this is inefficient in the case of a global pollutant. Moreover, even though our policy operates through a pre-defined and commonly known rule, it does not rely on commitment (*c.f.* Boleslavsky and Kelly, 2014).

An important limitation of our analysis is the assumption that policymakers use only information on emissions to update the (global) cap. Using both price- and quantity-information, the planners could do far better and implement the first-best level of welfare. Importantly, such a policy would have to be continuous in prices and emissions. A simple price collar (Fell, 2016; Holt and Shobe, 2016; Flachsland et al., 2020) will continue to exhibit equilibrium inefficiency. It is not clear whether a price collar performs better than an Optimal Linking policy; we hypothesize that it will not when emission benefits are strongly correlated between the jurisdictions since, in that case, Optimal Linking implements welfare levels very close to the first-best.

Cap and trade schemes are typically dynamic, i.e. they regulate emissions in multiple periods and allow for the use of allowances issued in one (earlier) period to cover emissions emitted in another (later). It is not obvious how two dynamic cap and trade scheme should ideally be linked. For example, how would the aggregate cap respond when firms in one jurisdiction buy allowances from the other and then keep them for future use? In this case, the mere number of allowances traded is not informative about the regional benefit functions (though note that the number of allowances surrendered per jurisdiction still is; that is, the equivalence of allowance trading and emissions in a given period breaks down in a dynamic setting). Similarly, it is not clear how a dynamic cap and trade policy should respond to the composition of allowances surrendered or stored. A complete theory of linking should account for the fact cap and trade schemes can be dynamic.

The Optimal Linking policy is derived under strong assumptions regarding functional forms. While we do not investigate the robustness of our results to more general specifications, we nevertheless believe that our key insights generalize. Whenever firms are allowed to trade allowances, the resulting allocation of emissions must be such that marginal benefits are equal for all firms (if not, firms can continue to make mutually beneficial trades, which we may reasonably assume they will). The fact that an observed distribution of allowances implies marginal benefit equalization allows a policymaker to construct posterior beliefs on the true marginal benefit function in each jurisdiction. When this posterior deviates from the prior belief that was used to set the global cap, the initial cap may turn out inefficient, which motivates policy updating. Similarly, the idea that asymmetric uncertainty can improve the performance of a well-designed cap and trade scheme with an endogenous cap seems fairly robust: belief updating is always easier when there is a strong (relative) “anchor” to update beliefs upon.

We pursue a top-down approach toward linking. Our Optimal Linking policy is derived from global welfare maximization. Unfortunately, while expected global welfare is strictly higher under an Optimal Linking regime compared to local cap and trade, regional welfare may be lower since a relatively stable jurisdiction may, through linking, expose itself to imported uncertainty from another. Since global welfare is strictly higher it is of course possible to construct side payments such that each jurisdiction expects to benefit from an Optimal Linkage; if, however, side payments are (politically) infeasible, individual jurisdictions may opt out of an Optimal Linkage. Our analysis is therefore best seen as an attempt to formulate the kind of linking policy an international agreement like the Paris Agreement or its successors might stimulate.

Finally, we treat uncertainty in a very particular way. Our analysis considers the case in which the intercept of the marginal abatement function is private information of the firms; however, the planners nevertheless know all other parameters of the emissions benefit function. Similarly, we assume that the environmental damage function to be perfectly known. Future work might relax these restrictive assumptions.

4.2 Policy Implications

Cap and trade schemes have become a major policy instrument in the fight against climate change. In Europe alone, roughly 45% of greenhouse gas emissions are regulated by EU ETS, the world’s largest market for carbon. As more and more cap and trade schemes are erected, linking has become a prominent policy issue – in fact, multilateral linking is explicitly suggested by Article 6 of the Paris agreement and linkages between local schemes already exist. The are linkages between EU ETS and the Swiss ETS, between RGGI and Quebec, between Quebec and California. The up-and-coming carbon markets of China and the post-Brexit UK will create new possibilities for linking. Given the large amount of money and CO_2 involved in cap and trade, and given the increasing prevalence of linkages between jurisdictional schemes, constructive ideas on optimal linking are called for. This paper offers some initial thoughts. We have three key messages for efficient policymaking.

First, allowances should be traded "ton-for-ton" between linked schemes to guarantee an efficient distribution of emissions given the aggregate cap. The reason for this requirement is that a trading ratio on allowances (Holland and Yates, 2015) would incentivize firms to distribute emissions between jurisdictions beyond the point at which marginal benefits of emissions are the same. Since climate change is indifferent to the source of emissions, such an allocation of emissions would be inefficient.

Second, the aggregate cap on emissions should be adjusted in response to allowance trading – or the demand for emissions – between the jurisdictions. The idea here is that the choice to buy or sell allowances, given an initial cap, reveals information about the marginal benefits due to emissions of the firms involved in the transaction. This information can be used to update the policymakers' beliefs about an efficient global cap, which may result in cap adjustments.

Third, the demand for emissions in a less predictable jurisdiction should have a relatively stronger effect on global emissions in an efficient cap and trade scheme. This condition follows from the observation that if firms choose to re-allocate allowances in a way unforeseen by the policymakers, so their prior beliefs about the true benefit functions were off, then the likeliest explanation is that benefits were "most off" in the unpredictable jurisdiction. Cap adjustments therefore respond more strongly to demand for emissions in the unpredictable jurisdiction.

Our narrative focuses on linking of cap and trade schemes at the level of a jurisdiction. Another, perhaps more natural interpretation of our model is in terms of covering emissions in different sectors, industries, or even countries with a single cap and trade scheme. Consider the aviation sector. Passenger flights outside the European Economic Area are not covered by the EU ETS. To nevertheless reduce emissions in the aviation industry, "the International Civil Aviation Organization (ICAO) agreed on a Resolution for a global market-based measure to address CO2 emissions from international aviation as of 2021. The agreed Resolution sets out the objective and key design elements of the global scheme, as well as a roadmap for the completion of the work on implementing modalities. The Carbon Offsetting and Reduction Scheme for International Aviation, or CORSIA, aims to stabilize CO2 emissions at 2020 levels by requiring airlines to offset the growth of their emissions after 2020."[10] Though it is still unclear how exactly CORSIA and EU ETS will interact, our analysis suggests that a direct incorporation of CORSIA into the EU ETS may be suboptimal. Similarly, our results suggest that trade of allowances between sectors or clearly identified jurisdictions within an existing cap and trade scheme – countries in the EU ETS, states in RGGI, industries in the South Korea ETS – can be made more efficient by incorporating a policy along the lines of our optimal linking regime.

4.3 Concluding Remarks

We propose a simple theory of optimal linking. Our results have the potential to greatly increase welfare compared to current practices. The core of an optimal linkage boils down to a basic observation: trading in allowances between schemes signals valuable information about abatement costs in the jurisdictions. An efficient policy aims to incorporate this information and adjusts the linked global cap in response to allowance trading. We pin down a precise analytic formulation for such endogenous policy updating.

There are various ways to adjust the global cap in response to trade flows, but under an optimal linkage, firms are allowed to exchange allowances one-for-one both within and across schemes. Thus, our optimal linking policy cannot rely on "trading ratios" for emissions allowances (Holland and Yates, 2015). While trading ratios on allowances do indeed endogenize the (global) cap in response to trading (Holland and Yates, 2015), they also disturb individuals firms' incentives away from an exact equalization of marginal abatement costs. A straightforward way to achieve this is to either inject new permits or buy back already issued ones as called for by the observed trade flows. This is not too complicated; Hintermayer (2020) analyzes a buyback policy in a dynamic model of the EU ETS.

An important concept is asymmetric uncertainty. When two schemes trade allowances, information on *relative* abatement costs is revealed. But when the schemes are asymmetrically uncertain, this information on relative abatement costs can be used to make (sharp) predictions about *absolute* abatement costs as well. The same is not possible when trade between symmetrically uncertain schemes is observed. Our results suggest that the study of asymmetric uncertainty deserves a more prominent place in environmental economics.

Our theory is preliminary and does not address several aspects of real life emission trading. Most large-scale cap and trade schemes are dynamic and allow covered industries to bank (and sometimes borrow) allowances across periods. How two dynamic cap and

[10] Retrieved from https://ec.europa.eu/clima/policies/transport/aviation_en

trade schemes should optimally be linked will likely depend on details of the dynamic policy. Similarly, it remains unclear how two cap and trade schemes, each with their own price collar on allowances, should best be linked. We leave an exploration of these and other aspects for future research.

References

Antoniou, F. and Kyriakopoulou, E. (2019). On the strategic effect of international permits trading on local pollution. *Environmental and Resource Economics*, 74(3):1299–1329.

Boleslavsky, R. and Kelly, D. L. (2014). Dynamic regulation design without payments: The importance of timing. *Journal of Public Economics*, 120:169–180.

Borenstein, S., Bushnell, J., Wolak, F. A., and Zaragoza-Watkins, M. (2019). Expecting the unexpected: Emissions uncertainty and environmental market design. *American Economic Review*, 109(11):3953–77.

Caplan, A. J. and Silva, E. C. (2005). An efficient mechanism to control correlated externalities: redistributive transfers and the coexistence of regional and global pollution permit markets. *Journal of Environmental Economics and Management*, 49(1):68–82.

Carbone, J. C., Helm, C., and Rutherford, T. F. (2009). The case for international emission trade in the absence of cooperative climate policy. *Journal of Environmental Economics and Management*, 58(3):266–280.

Dasgupta, P., Hammond, P., and Maskin, E. (1980). On imperfect information and optimal pollution control. *The Review of Economic Studies*, 47(5):857–860.

Doda, B., Quemin, S., and Taschini, L. (2019). Linking permit markets multilaterally. *Journal of Environmental Economics and Management*, 98:102259.

Doda, B. and Taschini, L. (2017). Carbon dating: When is it beneficial to link etss? *Journal of the Association of Environmental and Resource Economists*, 4(3):701–730.

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

Flachsland, C., Marschinski, R., and Edenhofer, O. (2009a). Global trading versus linking: Architectures for international emissions trading. *Energy Policy*, 37(5):1637–1647.

Flachsland, C., Marschinski, R., and Edenhofer, O. (2009b). To link or not to link: benefits and disadvantages of linking cap-and-trade systems. *Climate Policy*, 9(4):358–372.

Flachsland, C., Pahle, M., Burtraw, D., Edenhofer, O., Elkerbout, M., Fischer, C., Tietjen, O., and 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):133–142.

Gerlagh, R. and Heijmans, R. (2020). Regulating stock externalities.

Gerlagh, R. and Heijmans, R. J. (2019). Climate-conscious consumers and the buy, bank, burn program. *Nature Climate Change*, 9(6):431–433.

Gerlagh, R., Heijmans, R. J., and Rosendahl, K. E. (2020). COVID-19 tests the market stability reserve. *Environmental and Resource Economics*, 76(4):855–865.

Gerlagh, R., Heijmans, R. J., and Rosendahl, K. E. (2021). An endogenous emission cap produces a green paradox. *Economic Policy*.

Habla, W. and Winkler, R. (2018). Strategic delegation and international permit markets: Why linking may fail. *Journal of environmental economics and management*, 92:244–250.

Heutel, G. (2020). Bankability and information in pollution policy. *Journal of the Association of Environmental and Resource Economists*, 7(4):779–799.

Hintermayer, M. (2020). A carbon price floor in the reformed EU ETS: Design matters! *Energy Policy*, 147:111905.

Holland, S. P. and Yates, A. J. (2015). Optimal trading ratios for pollution permit markets. *Journal of Public Economics*, 125:16–27.

Holt, C. A. and Shobe, W. M. (2016). Reprint of: Price and quantity collars for stabilizing emission allowance prices: Laboratory experiments on the eu ets market stability reserve. *Journal of Environmental Economics and Management*, 80:69–86.

Holtsmark, B. and Weitzman, M. L. (2020). On the effects of linking cap-and-trade systems for CO2 emissions. *Environmental and Resource Economics*, 75(3):615–630.

Holtsmark, K. and Midttømme, K. (2021). The dynamics of linking permit markets. *Journal of Public Economics*, 198:104406.

ICAP (2020). Icap brief 3: Emissions trading at a glance. *International Carbon Action Partnership*.

Kotchen, M. J. (2018). Which social cost of carbon? A theoretical perspective. *Journal of the Association of Environmental and Resource Economists*, 5(3):673–694.

Kwerel, E. (1977). To tell the truth: Imperfect information and optimal pollution control. *The Review of Economic Studies*, 44(3):595–601.

Mehling, M. A., Metcalf, G. E., and Stavins, R. N. (2018). Linking climate policies to advance global mitigation. *Science*, 359(6379):997–998.

Mideksa, T. K. and Weitzman, M. L. (2019). Prices versus quantities across jurisdictions. *Journal of the Association of Environmental and Resource Economists*, 6(5):883–891.

Pizer, W. A. and Prest, B. C. (2020). Prices versus quantities with policy updating. *Journal of the Association of Environmental and Resource Economists*, 7(3):483–518.

Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies*, 41(4):477–491.

Yates, A. J. and Cronshaw, M. B. (2001). Pollution permit markets with intertemporal trading and asymmetric information. *Journal of Environmental Economics and Management*, 42(1):104–118.

A Derivations and Proofs

DERIVATION OF (20):

Combining the definition with the firms' FOCs, (4), we find the change in permit use by jurisdiction:

$$\Delta^L e_N = \frac{\theta_N - \theta_S}{b_N + b_S} \tag{37}$$

$$\Delta^L e_S = \frac{\theta_S - \theta_N}{b_N + b_S}. \tag{38}$$

PROOF OF PROPOSITION 1:

Proof. We only need to compare welfare losses under a linked cap and trade regime, equation (20), to those under jurisdictional cap and trade, equation (16). Linking outperforms regional cap and trade iff:

$$\frac{b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_Nb_S\rho\sigma_N\sigma_S}{b_N + b_S} < (c + b_S)\sigma_N^2 + (c + b_N)\sigma_S^2 - 2c\rho\sigma_N\sigma_S$$

$$\Longleftrightarrow$$

$$\frac{2\rho\sigma_N\sigma_S}{\sigma_N^2+\sigma_S^2} < \frac{cb_N+cb_S+b_Nb_S}{cb_N+cb_S},$$

which is always true as the LHS is below and the RHS is above one. □

DERIVATION OF (30):

Regional and global deviations from Socially Optimal permit use are given by:

$$\Delta^O e_N = \frac{b_S}{b_N+\delta b_S}\frac{[\delta b_S-c(1-\delta)]\theta_N+[b_N+c(1-\delta)]\theta_S}{cb_N+cb_S+b_Nb_S} \tag{39}$$

$$\Delta^O e_S = \frac{b_N}{b_N+\delta b_S}\frac{[\delta b_S-c(1-\delta)]\theta_N+[b_N+c(1-\delta)]\theta_S}{cb_N+cb_S+b_Nb_S} \tag{40}$$

$$\Delta^O Q = \frac{b_N+b_S}{b_N+\delta b_S}\frac{[\delta b_S-c(1-\delta)]\theta_N+[b_N+c(1-\delta)]\theta_S}{cb_N+cb_S+b_Nb_S}. \tag{41}$$

Define

$$\xi := \frac{b_N+c(1-\delta)}{b_N+\delta b_S} \implies 1-\xi := \frac{\delta b_S-c(1-\delta)}{b_N+\delta b_S}. \tag{42}$$

Welfare losses can now be written as:

$$\begin{aligned} L^O &= \frac{1}{2}\frac{c(b_N+b_S)^2+b_N^2b_S+b_Nb_S^2}{(cb_N+cb_S+b_Nb_S)^2}\mathbb{E}\left[(1-\xi)\theta_N+\xi\theta_S\right]^2 \\ &= \frac{b_N+b_S}{2}\frac{(1-\xi)^2\sigma_N^2+\xi^2\sigma_S^2+2\xi(1-\xi)\rho\sigma_N\sigma_S}{cb_N+cb_S+b_Nb_S}. \end{aligned} \tag{43}$$

If for notational convenience, we define:

$$\psi := \frac{1}{2}\frac{b_N+b_S}{cb_N+cb_S+b_Nb_S}, \tag{44}$$

it is straightforward to derive:

$$\frac{\partial}{\partial\xi}\frac{L^O}{\psi} = 2\xi\sigma_S^2-2(1-\xi)\sigma_N^2+2(1-\xi)\rho\sigma_N\sigma_S-2\xi\rho\sigma_N\sigma_S. \tag{45}$$

The welfare-maximizing ξ^* therefore satisfies:

$$\xi^* = \frac{\sigma_N^2-\rho\sigma_N\sigma_S}{\sigma_N^2+\sigma_S^2-2\rho\sigma_N\sigma_S}. \tag{46}$$

From the definition of ξ, the optimal cap-adjustment rate δ^* follows:

$$\delta^* = \frac{(b_N + c)[\sigma_S^2 - \rho\sigma_N\sigma_S] + c[\sigma_N^2 - \rho\sigma_N\sigma_S]}{(b_S + c)[\sigma_N^2 - \rho\sigma_N\sigma_S] + c[\sigma_S^2 - \rho\sigma_N\sigma_S]}, \tag{47}$$

as given.

PROOF OF PROPOSITION 2:

Proof. Plugging (46) in (43), we find:

$$\begin{aligned}
\frac{L^O}{\psi} &= \left[\frac{\sigma_S^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}\right]^2 \sigma_N^2 + \left[\frac{\sigma_N^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}\right]^2 \sigma_S^2 \\
&\quad + \left[\frac{\sigma_S^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}\right]\left[\frac{\sigma_N^2 - \rho\sigma_N\sigma_S}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S}\right]\rho\sigma_N\sigma_S \\
&= \frac{(1-\rho^2)\sigma_N^2\sigma_S^2}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S} \\
&\Longrightarrow \\
L^O &= \frac{1}{2}\frac{b_N + b_S}{cb_N + cb_S + b_N b_S}\frac{(1-\rho^2)\sigma_N^2\sigma_S^2}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S},
\end{aligned}$$

as stated. This is strictly lower than the welfare loss under traditional Trading if and only if:

$$L^L - L^O \geq 0$$

$$\Longrightarrow$$

$$\frac{1}{b_N + b_S}\frac{b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S\rho\sigma_N\sigma_S}{cb_N + cb_S + b_N b_S} - \frac{b_N + b_S}{cb_N + cb_S + b_N b_S}\frac{(1-\rho^2)\sigma_N^2\sigma_S^2}{\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S} \geq 0$$

$$\Longrightarrow$$

$$(\sigma_N^2 + \sigma_S^2 - 2\rho\sigma_N\sigma_S)(b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S\rho\sigma_N\sigma_S) - (1-\rho^2)(b_N^2 + b_S^2 + 2b_N b_S)\sigma_N^2\sigma_S^2 \geq 0$$

$$\Longrightarrow$$

$$[(b_S\sigma_N^2 - b_N\sigma_S^2) + (b_N - b_S)\rho\sigma_N\sigma_S]^2 \geq 0,$$

which is always true. $\square$

PROOF OF PROPOSITION 3:

Proof. We derived quantity derivations under both policies. Prices are equal in both jurisdictions, so without loss of generality we can solve for price deviations in jurisdiction

1:

$$\Delta^L p_N = \frac{b_S\theta_N + b_N\theta_S}{b_N + b_S}$$
$$\Delta^{OL} p_N = \frac{\delta b_S\theta_N + b_N\theta_S}{b_N + \delta b_S}.$$

Thus:

$$\mathbb{E}\left[\left(\Delta^L p\right)^2\right] = \frac{b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2b_N b_S\rho\sigma_N\sigma_S}{b_N^2 + b_S^2 + 2b_N b_S}$$
$$\mathbb{E}\left[\left(\Delta^{OL} p\right)^2\right] = \frac{\delta^2 b_S^2\sigma_N^2 + b_N^2\sigma_S^2 + 2\delta b_N b_S\rho\sigma_N\sigma_S}{b_N^2 + \delta^2 b_S^2 + 2\delta b_N b_S}.$$

Writing these out, we obtain:

$$\mathbb{E}\left[\left(\Delta^{OL} p\right)^2\right] < \mathbb{E}\left[\left(\Delta^L p\right)^2\right] \iff (\delta - 1)\left[b_S\left(\sigma_N^2 - \rho\sigma_N\sigma_S\right) - b_N\left(\sigma_S^2 - \rho\sigma_N\sigma_S\right)\right] < 0.$$

This condition is always satisfied. $\square$