International Carbon
Action Partnership

Options and Issues for Restricted Linking of Emissions Trading Systems

Michael Lazarus, Lambert Schneider, Carrie Lee and Harro van Asselt
Stockholm Environment Institute

September 2015, Berlin, Germany

This paper was produced as an input to the ICAP Technical Dialogue on linking emissions trading systems. The authors would like to thank Michel Frerk, Constanze Haug, Aki Kachi and Marissa Santikarn from the ICAP Secretariat as well as several ICAP members for their constructive comments on a draft version of this paper. Bianca Sylvester and Luca Taschini also provided valuable feedback. An earlier version of the paper was presented at the workshop "Comparison and Linkage of Mitigation Efforts in a New Paris Regime", organized by the International Emissions Trading Association, the Harvard Project on Climate Agreements and the World Bank Group's Networked Carbon Markets Initiative. Findings and opinions expressed in this study are those of its authors and do not necessarily reflect the views of ICAP or its members.

Contents

Executive Summary 4

1. Experience with linking ETS, challenges and prospects 6

1.1 Introduction 6

1.2 Why link? 6

1.3 Status of linking 7

1.4 Challenges of full linking 7

1.5 Alternatives, or pathways, to full linking 9

1.6 Objectives and roadmap of this paper 10

2. Restricted linking: options and issues 10

2.1 Quotas and one-way linking 11

2.2 Exchange rates and discount rates 12

3. Assessing the options 15

3.1 Assessment criteria 15

3.2 Using a simplified modeling analysis to assess implications 17

3.3 The reference cases: no linking and full linking 17

3.4 Quotas 19

3.5 One-way linking 22

3.6 Exchange rates 22

3.7 Discount rates 29

3.8 Implications of key ETS design features and conditions 30

3.9 Implications of restricted linking across more than two jurisdictions 32

3.10 Comparison of options 33

4. Findings and conclusions 35

References 37

Executive Summary

Linking emissions trading systems (ETSs) offers many potential benefits. These include economic benefits such as; increased access to abatement options; a potential reduction of carbon leakage risks; enhanced market liquidity and increased resilience to market shocks. There are also political benefits, as linking can signal enhanced cooperation and influence, limit competitiveness concerns, and enable jurisdictions to adopt more ambitious targets. In addition, linking can offer administrative and institutional benefits through shared administrative infrastructure, procedures, and costs. Linking may ultimately be necessary for an ETS to achieve its economic and political objectives, especially for smaller systems.

As the limited number of existing links attests, however, linking faces many challenges. Full linking requires a certain level of harmonization of, and thus compromises on, key ETS design elements. Full linking may also risk reduced ambition if one linked jurisdiction faces significant over-allocation of allowances, as well as real or perceived loss of regulatory autonomy. Jurisdictions may further have unequal institutional and technical capacities, and competing domestic objectives may need to be reconciled during the linking process.

Given the mix of benefits, obstacles, and risks, successful linking requires matching jurisdictions with compatible ETS designs and domestic policy objectives. It can also require finding the right level at which to engage, and tailoring the link to suit the contexts and compatibility of ETSs and their jurisdictions. In many cases, the suitable approach may fall well short of full linking, at least in the near term.

This paper examines alternatives, short of full linking, that jurisdictions could pursue to capture some of the political, economic, and environmental benefits associated with linking. It focuses in particular on the potential advantages and drawbacks of "restricted linking": options that enable the flow of units among jurisdictions, but with specific constraints such as quantity limits ("quotas"), or conditions such as exchange rates, to help address concerns that full harmonization might create. While these options have been discussed in many venues, they have yet to be examined thoroughly in the literature. This paper aims to help fill that gap.

We use four broad criteria (environmental benefit, economic benefit, political feasibility, and other practical and overarching considerations) and a simple model to help assess and compare options for restricted linking – in particular quotas, one-way linking, exchange rates and discount rates – with the situations of not linking or full linking.

From a political perspective, quotas offer some attractive features. Assuming allowances are not over-allocated, quotas have no effect on overall abatement across jurisdictions while enhancing cost-effectiveness, and even low quotas may create significant economic benefits that can be shared among jurisdictions. Perhaps most importantly, quotas allow policy-makers within a jurisdiction to retain a certain level of control on the extent of unit flows and related impacts, assuming that the quotas are set by regulators in the linking jurisdictions. Quotas have already been used in several ETSs, such as the EU ETS quantitative limits on the use of credits from the Kyoto Protocol's flexibility mechanisms. One-way linking was implemented before the Norwegian ETS became fully linked to the EU ETS, and was once proposed for the Australia ETS for accessing

allowances from the EU ETS. Like other restricted linking approaches, one-way linking offers a way to "test the waters", short of committing to full bilateral linking. At the same time, if the allowed direction of unit flow is from lower- to higher-price jurisdictions, then the outcome from a cost-effectiveness perspective would be very similar to that of full linking.

The implications of exchange rates are more complex than those of quotas or one-way linking. First, in contrast to full linking, quotas or one-way linking, exchange rates can affect the total abatement across the two jurisdictions. Second, the implications for cost-effectiveness and total abatement are highly sensitive to the value at which the exchange rate is set. Third, the overall implications depend on: (1) the rationale and policy objectives pursued when setting the exchange rate (e.g. enhancing market liquidity, gaining economic benefits, ensuring the same or a better environmental outcome); (2) the frequency and manner in which the exchange rate might be updated (fixed, fixed with updates, or floating); and (3) who sets the rate. Exchange rates could create considerable environmental and economic benefits or have major (unintended) adverse implications, even compared with no linking, due to information asymmetries between the regulated entities and policy-makers setting the exchange rate, as well as uncertainties about future developments. In practice, setting exchange rates in ways that avoid such unintended consequences can be rather difficult. These challenges would become particularly acute in a context where more than two schemes link.

Discount rates can avoid some of the challenges of exchange rates. Exchange rates are inherently symmetrical: if two jurisdiction A units are required for each ton of emissions in jurisdiction B, then one jurisdiction B unit would cover two tons of emissions in jurisdiction A. In contrast, discount rates can be asymmetrical; one or both jurisdictions would require more than one imported unit to cover a ton of domestic emission. Discount rates could work in a similar manner as price containment mechanisms and be set up in ways that ensure that both cost-effectiveness and total emissions abatement are enhanced, possibly creating a net mitigation benefit as called for in market-based mechanisms by the Cancun Agreements. They implicitly link flexibility and mitigation ambition – the more the regulated entities use the flexibility to import units from another scheme, the more emissions are reduced. Applying discount rates reciprocally (both jurisdictions discount unit imports at the same rate) might address perceptions about valuing reductions differently across jurisdictions. They could therefore be further explored as an alternative or complementary approach to quotas or one-way linking.

While restricted linking options do not, in principle, achieve the potential benefits of full linking, they can lessen some of the potential pitfalls. They offer easier off-ramps to "de-link" and levers to adjust (e.g. quotas or discount rates), should linking concerns prove to be more significant than anticipated. Restricted linking options may thus represent a cautious approach that can be more easily implemented and explored where full linking is either infeasible in the near term or incompatible with the objectives of the jurisdictions involved. To use a "match-making" analogy for the process of finding the right linking partner, restricted linking could be viewed as a step like "moving in together" before – or without the intention of – getting married. With this paper, we hope to offer some insights that can help prospective partners decide whether this step is worth taking.

1. Experience with linking ETS, challenges and prospects

1.1 Introduction

National, subnational, and regional greenhouse gas emissions trading schemes (ETSs) continue to proliferate. Seventeen distinct ETSs are now operational across four continents, with another 15 in preparation or under consideration (ICAP 2015). The emergence of these new schemes, the recent success in connecting the California and Québec ETSs, and recognition of the challenges of building a global carbon market have led to renewed interest in options for linking different trading schemes.

The existing literature on linking ETSs offers useful insights into the economic, legal, and institutional options for, and challenges to, linking (Mehling and Haites 2009; Flachsland et al. 2009a; Jaffe and Stavins 2007; Tuerk et al. 2009; Haites et al. 2001). Much of this literature examines full linking, where units remain unique to each system, but are fully interchangeable, have the same compliance value, and flow freely among linked schemes. Yet, notwithstanding the California-Québec success, efforts to pursue full linking have proven challenging.

This paper examines alternatives, short of full linking, that jurisdictions could pursue to capture some of the political, economic, and environmental benefits associated with linking. It focuses in particular on the potential advantages and drawbacks of what we term "restricted linking": options that enable the flow of units among jurisdictions, but with specific conditions (e.g. exchange rates, discount rates) or limitations (e.g. quotas, one-way linking) to help address near-term concerns that full harmonization might create. While these options have been discussed in many venues, they have yet to be examined thoroughly in the literature. This paper aims to help fill that gap.

1.2 Why link?

Full linking of ETSs can yield multiple benefits. Expanding markets can result in economic benefits such as increased access to abatement options, which lower overall costs; a potential reduction of carbon leakage risks; a larger number of market participants, which may reduce the potential for market manipulation; enhanced market liquidity; and increased resilience to market shocks. There can also be political benefits, as linking can signal enhanced cooperation and influence. Some have also suggested that by lowering overall costs, linking can generate domestic support and encourage jurisdictions to adopt more ambitious targets (Flachsland et al. 2009b; Bodansky et al. 2015).

Linking can also provide administrative and institutional benefits through shared administrative infrastructure, procedures and costs – such as the joint allowance auctions now held by California and Québec (Burtraw et al. 2013). It can reduce management and transaction costs by providing a common set of rules and prices that entities can respond to. Ultimately, linking to other ETSs may be necessary for a jurisdiction to achieve its economic and political objectives.

1.3 Status of linking

The full linking of California and Québec's programs in 2014 marked a major step forward in ETS linking. For the first time, two jurisdictions negotiated a full linking agreement that enabled mutual acceptance of compliance instruments and has since resulted in joint auctions and exchange of units. The foundation for this link was built through the Western Climate Initiative, a multi-jurisdictional collaboration that developed common design recommendations for emissions trading in late 2008. Since then, California and Québec have established their respective ETSs and worked together to align program elements in order to facilitate linking. Both ETSs began operating in 2013, and formally linked fully in 2014.

The provincial ETSs in Japan – in Tokyo and Saitama – are fully linked as well, though trade among systems has yet to occur. The programs were developed at roughly the same time and share similar program elements. The Australian ETS, which was repealed in 2014, included plans for one-way linking to the European Union (EU) ETS, with longer-term ambitions of fuller linking.

Some indirect linking has also occurred, as discussed further in Section 1.5. For example, both the EU ETS and the New Zealand ETS initially allowed the use of common offset units: Certified Emission Reductions (CERs) from the Clean Development Mechanism (CDM) and Emission Reduction Units (ERUs) from Joint Implementation (JI). In China, all seven pilot ETS programs can draw from the same pool of domestic project-based offset credits – China Certified Emission Reductions.

There are other examples of programs that have joined and merged with existing ETSs, which we consider distinct from linking. When programs merge, they no longer have distinct caps, ETS programs, or units. Both the EU ETS and the Regional Greenhouse Gas Initiative (RGGI) have been joined by jurisdictions that intentionally designed their programs to align with the larger ETS's design elements and facilitate merging. For example, the Norway ETS came into force in 2005, the same year as the EU ETS, and was designed to be compatible with the EU ETS. Its first phase provided for a one-way link by allowing EU allowances in the Norwegian system. In 2007, Norway officially joined the EU ETS. Similarly, the Switzerland ETS was designed to match the provisions in the EU ETS, and is now in negotiations to formally link. In the case of RGGI, the ETS was developed through a regional collaborative effort. Each state has its own statutory and/or regulatory authority to participate, but there is one regional cap and unit type for the entire RGGI ETS.

1.4 Challenges of full linking

As the limited number of existing links suggests, linking ETSs can be challenging. Full linking requires a certain level of harmonization of many ETS design elements – the allocation of allowances; monitoring, reporting and verification (MRV) standards; cost containment measures (price floors or ceilings); offsets rules and limits – that may differ based on the political and economic context. The task is easier, though by no means simple, where systems are designed in coordination or evolve from similar design considerations, as was the case for California and Québec.

Further challenges to full linking include assessing and comparing the ambition of the ETS caps set by the jurisdictions; the potential for a sense of reduced ambition if linking to an ETS that is viewed as less stringent; real or perceived loss of regulatory autonomy; unequal capacities among jurisdictions; and potentially competing domestic objectives. As Green et al. (2014) observe, the global climate negotiations have faced very similar obstacles.

To the extent that designs and ambition differ among jurisdictions, linking will alter carbon prices and the location and extent to which long-term investment in low-carbon technologies and any associated co-benefits occur (Flachsland et al. 2009a; Bodansky et al. 2015). While in principle these shifts – from higher- to lower-cost emission reduction opportunities – will improve overall economic efficiency, it may be difficult for policy-makers to justify these shifts to constituencies that may see reduced investment and co-benefits or face higher allowance prices. Furthermore, linking creates financial transfers across jurisdictions, which can be particularly challenging if the direction of flow is from smaller or less affluent jurisdictions toward larger or more affluent ones. Jurisdictions that see a decline (increase) in allowance prices from linking will also see a reduction (increase) in investment in abatement activities, a decrease (increase) in allowance action revenues, and a net flow of resources to (from) other jurisdictions. In addition, linking may create political uncertainty, as another jurisdiction's future withdrawal could undermine the system. Some even argue that linking can create a potential disincentive to maintaining or increasing ambition in an individual jurisdiction (Helm 2003).

Therefore, linking decisions must consider whether such tradeoffs are compatible with the existing balance of domestic policy objectives. Whether linking results in advantages or disadvantages for a specific system, however, ultimately depends on that system's policy priorities. For example, one jurisdiction may prioritize cost minimization and thus allow for unrestricted use of offsets, while a second jurisdiction may prioritize domestic emission reductions and thus restricts the use of offsets (i.e. ensuring "supplementarity"). Linking these two ETSs without further provisions would effectively enable a relatively unlimited supply of offsets across both jurisdictions, thus compromising the second jurisdiction's policy objective of supplementarity (Flachsland et al. 2009b).

Given the mix of benefits, obstacles and risks, successful linking requires matching jurisdictions with compatible ETS designs and domestic policy objectives (Comendant and Taschini 2014), and "navigating tradeoffs between efficiency and political feasibility" (Green et al. 2014, p.1066). It can also require finding the right level at which to engage, and tailoring the link to suit the contexts and compatibility of ETSs and their jurisdictions. For instance, Doda and Taschini (2015) suggest that the difference in (emissions) size as well as the existence of unilateral tax distortions (e.g. a tax on allowance transactions) may affect otherwise mutually advantageous linking arrangements. In many cases, the suitable approach may fall well short of full linking, at least in the near term.

1.5 Alternatives, or pathways, to full linking

The challenges to full linking have increased interest in alternative approaches. Indeed, there is a wide range of options for policy-makers to consider, from no linking to full linking:

1. Full linking, or unrestricted *mutual unit recognition* and fungibility.[1]
2. Restricted linking, which involves *partial, conditional, or restricted unit recognition.* Examples include quantity limits or quotas, one-way linking, exchange rates, discount rates, and fees (a premium on units from other jurisdictions).
3. Indirect linking, which implies *no mutual unit recognition*, but connections or links to the same "third" system (e.g. offsets from the CDM or another common program).
4. Incremental alignment, which suggests *no mutual unit recognition*, but gradual alignment of program elements (carbon prices, coverage, MRV, etc.); this is also referred to as "linking by degrees" (Burtraw et al. 2013).
5. Connection to non-ETS markets and actions, which involves the recognition of greenhouse gas mitigation programs – including non-ETS programs such as carbon taxes – for compliance purposes (Bodansky et al. 2014).
6. No linking, or the continued development and operation of ETSs without explicit or intended efforts to harmonize or align systems, or recognize one another's units or other units in common (indirect linking).

In choosing along this range, policy-makers must consider which option best balances a number of potentially competing policy objectives, such as: reducing global emissions, reducing emissions "at home" (i.e. ensuring supplementarity), creating co-benefits from within-jurisdiction abatement activities (e.g. reducing local air pollution, creating jobs and new businesses), inducing technological change in the long run, minimizing costs to entities and/or consumers, and using proceeds of allowance sales or auctions for other public benefit. We examine these and other tradeoffs further in Section 3.

As illustrated by Figure 1, restricted linking options can be viewed as part of an evolving process – as short-term or interim measures towards a longer-term goal of fuller linking. Alternatively, they can be the approach chosen by a jurisdiction for the foreseeable future, in order to best balance among competing objectives.

[1] Linking with mutual unit recognition and flow (whether full or restricted) benefits from a certain level of similarities in ETS design and operation among linking jurisdictions: MRV, registry systems, banking and borrowing provisions, target type, compliance periods, and cost containment measures (including offset types and limits). The suitability of ETS designs for linking is a key consideration addressed in other literature (e.g. Tuerk et al. 2009), and is not explored further in this paper.

Figure 1: Linking approaches and possible pathways among them

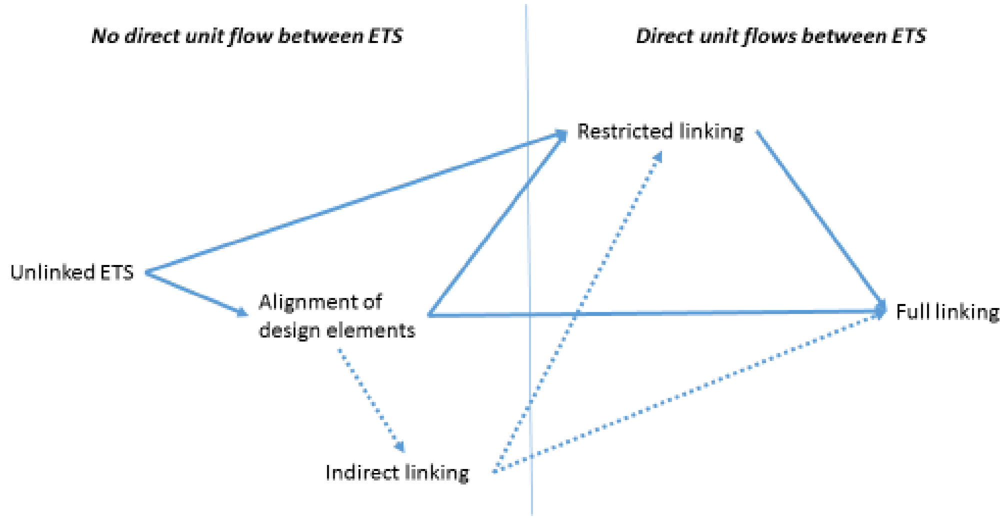

Whether systems are linked fully or only on a limited basis, it is important to align certain design elements – at a minimum, the allowance tracking systems, in particular the registries for tracking serial numbers of units and common units (e.g. metric tons).

1.6 Objectives and roadmap of this paper

The International Carbon Action Partnership (ICAP) has identified linking ETSs as an overarching issue of priority for its 2014–2015 Technical Dialog. As part of this work stream, ICAP members have commissioned this study on alternatives to full linking. As noted here, while studies on linking have been released since the early 2000s, most focus on full linking. Yet there may be cases where the full harmonization of key design features needed for full linking is not feasible, but restricted linking could be considered. This study maps out how alternative models for restricted linking compare with full linking and no direct linking. We assess these options with regard to their environmental benefit (i.e. abatement outcome), economic benefit (i.e. cost-effectiveness), political feasibility, and other practical and overarching considerations.

In the next section, we describe the options for restricted linking, their state of play, insights from the existing literature, and some key questions to consider. In Section 3, we conduct a deeper assessment of these options – focusing primarily on quotas and exchange rates – against a suite of criteria, and use a simple numerical model to illustrate some of their implications. In Section 4, we summarize the findings and offer conclusions.

2. Restricted linking: options and issues

This section explores restricted linking options in more detail, summarizing the literature and identifying general questions that each option raises, some of which we address in the next section. Because in their basic conception, they share some common features – in particular maintaining parity of unit value – we discuss one-way linking in conjunction with quotas.

Similarly, we discuss discount rates together with exchange rates (the fully symmetric and fungible version of discount rates). It is important to emphasize at the outset that these options are not necessarily mutually exclusive, and in fact could be combined (e.g. quotas with exchange rates). Finally, we do not present an exhaustive list here; other restricted linking options are conceivable. For instance, fees for the transfer of an allowance between two ETSs – or permits for the use of another jurisdiction's allowances – could be viewed as other forms of restricted linking.[2]

2.1 Quotas and one-way linking

Quantity limits or quotas can restrict the amount of (specific types of) allowances from other jurisdictions that can be used for compliance. Quotas that limit the availability or use of tradable units are hardly a new concept. Most ETSs apply some type of quantity restriction on the use of offsets. The EU ETS has long adopted the principle of supplementarity, limiting the overall use of international offset credits – CERs from the CDM and ERUs from JI. In many jurisdictions, including in Phase II of the EU ETS, supplementarity limits have been designed with the intent of achieving at least half of abatement through emission reductions at covered sources. California and Québec adopted a similar principle in limiting entities' use of offsets to no more than 8% of their compliance obligations.[3] Chinese ETS pilots, RGGI, and the South Korean and Swiss ETSs all limit offsets use in a similar way. Therefore, the notion of limits or quotas on the use of tradable units from other systems is a familiar concept and would not, in principle, be difficult to implement.

A quota would be typically designed to limit overall net transfers (most likely imposed on imports rather than exports). The limit could be set in order to ensure that a certain fraction of emission reductions are achieved domestically, akin to the supplementarity principle that several jurisdictions have used for offset limits. There are several possibilities for formulating and implementing quotas. They can be expressed as a fraction of total compliance obligations that an entity can surrender, similar to offset limits in a number of ETSs.[4] In fact, policy-makers might elect to combine a limit on allowances from other systems with a limit on offsets with the aim of ensuring that a certain fraction (e.g. half) of emission reductions are achieved by covered sources within the jurisdiction. Quotas could be adjusted over time as an incremental step as the harmonization of instruments increases.

[2] Fees are a novel concept, not discussed further in the remainder of this paper… While the ideas of fees or permits for imported allowance use have been raised informally in meetings, very little has been written about them. Fees or tradable permits for the use of imported allowances could offer regulatory agents a means to capture some of the rent (economic surplus) resulting from access to lower-priced imported units. One challenge of a fee approach would be setting it an appropriate level that achieves the intended outcome with respect to cost containment and other objectives. A permit approach raises a parallel question of how many permits to issue, as well as the question of whether the added administrative and transactions costs of another permit system would be warranted.

[3] Based on a supplementarity principle, the limit was to be set at 4%, but that was increased to account for additional allowance units subsequently placed into the allowance reserve.

[4] Such a fraction can be derived from an absolute limit, divided by the total number of allowable units (inclusive of both domestic and imported units).

As with offset limits, a limit or quota on the use of imported allowances may or may not effectively limit unit flows; it can be difficult to ascertain in advance whether unit flows might otherwise exceed quota levels. If the quota is set larger than the unit flow that would occur under full linking, then the outcome would be the same as under full linking. The quota could nonetheless serve the purpose of guarding against levels of allowance flow between jurisdictions that might be politically unacceptable. As we discuss further below, the effect of a quota depends on the relative costs of reducing emissions among jurisdictions – or in economists' terms, the relative marginal abatement cost (MAC) curves – that are subject to uncertainties as well as information asymmetries between the regulator and entities with regard to cost and feasibility of abatement options.[5]

Quotas raise several policy questions, including how the limit is determined (e.g. through supplementarity principles ex ante, or advance estimates of likely trade flows), how it might be updated, and whether and how it would be integrated with any offset limits.

One could also envision a situation where one jurisdiction allows unlimited allowance imports, while another linked jurisdiction sets a quota. One-way linking, where units in system A would be recognized in system B, but not vice-versa, could be considered to be the extreme version of this, with a quota of zero in one direction. For example, the now-defunct Australian ETS would have allowed limited imports of EU ETS allowances, but without any provision for the EU ETS to import Australian units. One-way linking can also serve as a transitional instrument, with the intent of fuller linking when design elements (and resulting allowance prices) are better aligned in future compliance periods. As discussed above, this happened in the first phase of the Norway ETS, before the country officially joined the EU ETS.

2.2 Exchange rates and discount rates

With exchange rates or trading ratios, units from one system can be used for compliance in another, but their value is adjusted by a conversion factor. Perhaps most examined through the World Bank's Globally Networked Carbon Markets initiative,[6] the concept has received limited attention in the carbon market literature. Researchers have tended to focus on trading ratios, the more common term for exchange rates in the academic literature, in the context of non-uniformly mixed pollutants such as SOx or NOx, where the ultimate damage is strongly affected by the location of emission (Mendelsohn 1986). There, the case for trading ratios is clearer – in an ETS for local or regional air pollutants, a trading ratio can make a ton of emission reduction in a more populated air shed more valuable than one in a remote area, increasing both health benefits and economic efficiency.

For carbon dioxide (CO_2) and most other greenhouse gases (GHGs), however, damage is largely unaffected by the location of emission. Nonetheless, some economists have suggested other

[5] As discussed further below, Holland and Yates (2014) suggest that these asymmetries are in fact a rationale for using exchange rates.

[6] See http://www.worldbank.org/en/topic/climatechange/brief/globally-networked-carbon-markets.

economic efficiency arguments for imposing trading ratios for GHGs (Holland and Yates 2014).[7] The nature and rationale for trading ratios (or exchange rates) suggested by the International Emissions Trading Association (IETA), the World Bank initiative, and others have more to do with expanding markets, improving liquidity, and enabling links that might not otherwise be feasible.

Indeed, some existing and proposed ETSs have employed a variant of trading ratios, typically in the form of "discounts" on certain unit types. The proposed American Clean Energy and Security Act of 2009 (commonly referred to as the Waxman-Markey bill), which was never adopted, stated that for every ton of an entity's compliance obligation covered by an offset it would need to turn in 1.25 tons of offsets, effectively a 20% discount. Similarly, France introduced a general 10% discount on domestic JI projects (only 90% of the emission reductions being issued as ERUs).

In the next section, we distinguish **discount rates**, such as these, from fuller unit **exchange rates** or trading ratios. Exchange rates operate in a symmetrical fashion, much the way that openly traded currencies do. If an exchange rate is set such that two units from system A could be used in place of one unit in system B, then in system A, one system B unit would be worth two system A units. This symmetry creates fungibility and enables units to flow readily back and forth among systems. In contrast, discount rates can be asymmetrical. Jurisdiction A can decide on a ratio to apply for units from system B; unlike with an exchange rate, jurisdiction B need not automatically accept units from jurisdiction A using the inverse of that ratio. It could, for example, apply the same discount rate such that each jurisdiction places a greater value on units from its own system. We return to this particular notion below; under certain conditions, particularly when future abatement costs are highly uncertain, it may offer some advantages over exchange rates as a restricted linking approach.

Burtraw et al. (2013) examine the implications of linking the RGGI and California ETSs with an exchange rate where three RGGI units are equivalent to one California unit. This example is illustrated in Figure 2, which contrasts the exchange rate with a situation of full linking.

Each block in Figure 2 represents the number of tons of CO_2 that the holder of an allowance from a given jurisdiction can emit in each jurisdiction (red blocks are jurisdiction A allowances; blue blocks are jurisdiction B allowances). With an exchange rate of 3:1, as shown, if an entity in jurisdiction B acquires a unit from jurisdiction A, it can emit only 1/3 of a ton of CO_2, and overall emissions decrease across the two systems by 2/3 of a unit transferred. Conversely, if an entity in jurisdiction A acquires a unit from jurisdiction B; it can emit 3 tons of emissions (versus 1 ton in the originating jurisdiction); overall emissions increase by 2 tons.

As illustrated in Section 3, the direction in which most units flow will determine whether emissions increase or decrease. Whether *total* emissions (abatement) would increase or decrease

[7] Holland and Yates (2014) suggest that abatement cost uncertainties and the related asymmetry of abatement cost information between regulators and polluters provide a rationale for trading ratios. Due to the information asymmetry, policy-makers cannot set the overall cap in an efficient way ex ante, and ex post adjustments could adjust to the cap and increase overall efficiency, balancing damages and mitigation costs. However, as current caps fall significantly short of a cost-efficient pathway to address climate change (from the perspective of total social costs), these considerations may be more academic.

compared with no linking (or linking without an exchange rate) depends on how an exchange rate is set (relative to the underlying cap and marginal abatement costs) and the resulting net flow of allowances between the systems.

Figure 2: Schematic of exchange rate example compared to full linking

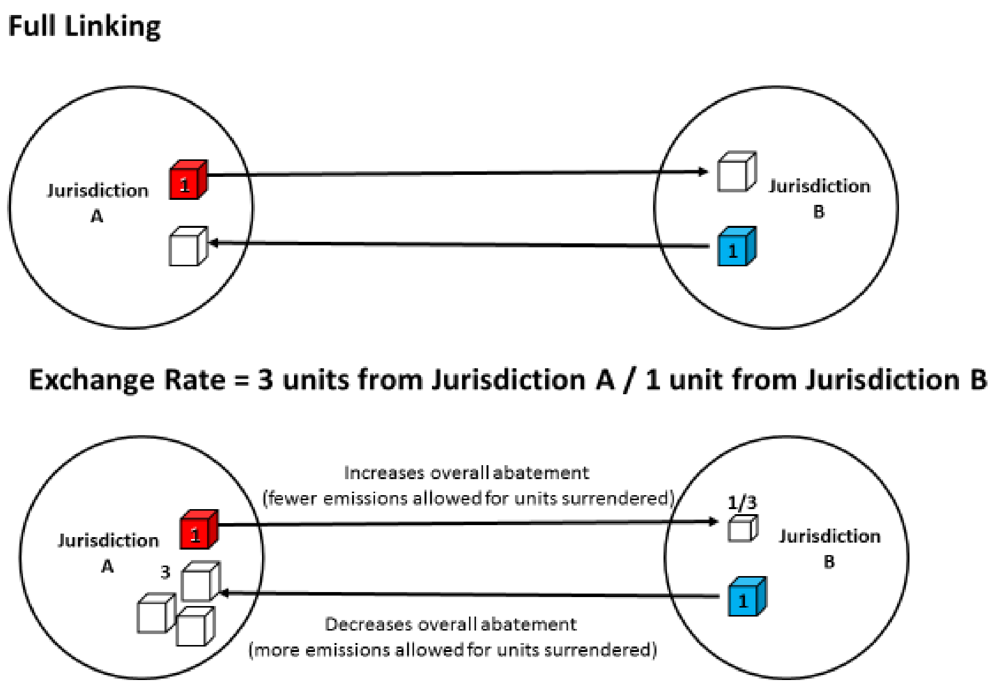

Exchange rates raise several additional policy questions, such as who sets the rates and how; whether the rates are fixed or floating; and how rates can be updated while retaining the integrity of allowance markets. The World Bank Globally Networked Carbon Markets program has begun exploring some of these questions, as well as the concept of a central aggregator or reserve that could help in rating unit types (e.g. based on relative price, mitigation value, environmental integrity, etc.) and creating a "blended unit" (i.e. that, much like a mutual fund, reflects holdings of multiple unit types), how to translate ratings into rates, and how to factor in implications of other policies. We do not delve deeper into those issues in this paper; they are of considerable importance to the feasibility and outcome of exchange rates and deserve further study.

Having outlined in this section some of the key options for restricted linking, we provide in the next section a preliminary assessment against a set of four criteria.

3. Assessing the options

In this section we assess and compare options for restricted linking – in particular one-way linking, quotas, exchange rates, and discount rates – with the situations of no linking or full linking. We first introduce a suite of four broad criteria that policy-makers can use to help evaluate the merits and drawbacks of each option. Under each criterion we describe specific implications that can be used to inform them – such as cost-effectiveness to inform the criterion of economic benefits, or allowance flows to inform the criterion of political feasibility. We use a simple representation of two emission trading schemes in two jurisdictions to illustrate some of the environmental and economic implications of various restricted trading options. We apply this model, first for reference, to the cases of no linking and full linking, and then to quotas, one-way linking, exchange rates and discount rates, and relate the findings as well as other considerations to common criteria. We then briefly discuss how certain ETS features, namely differing price containment mechanisms, offset rules and over-allocation – as well as the possibility of multiple linked jurisdictions – can make restricted linking more challenging to assess and implement.

3.1 Assessment criteria

There are many issues for policy-makers to consider when linking ETSs. For the purposes of comparing the options for linking, we group them into four broad criteria:

1. ***Environmental benefit (abatement outcome).*** There are several ways to look at the environmental, or emission reduction, benefits of linking. One is to look at the direct implications: how an option affects overall abatement across linked jurisdictions, given existing emissions caps (expected abatement outcome). It is also important to consider how linking might affect the level of abatement in each individual jurisdiction. Some stakeholders may be reluctant to see a significant reduction in local abatement activity, especially if that might also reduce any environmental, economic, or social co-benefits, which in turn can affect political feasibility (below). In addition, by affecting allowance prices, and changing the context in which future targets are set by linked jurisdictions (i.e. more interdependently), linking could either encourage or discourage increased ambition (deeper reduction targets) in future periods, an implication that is more difficult to assess without significant speculation.

2. ***Economic benefit.*** One of the core rationales for linking is the classic "gains from trade" that accrues when markets are opened, increasing access to lower-cost production, in this case of emission reductions. To assess economic benefits, we consider effects on total abatement costs (across and within individual jurisdictions), on liquidity, on unit fungibility and flows, on within-jurisdiction investment in abatement options, and on cost-effectiveness.[8] We also assess the risk for carbon leakage – i.e. the risk of production

[8] For a given abatement outcome, we define cost-effectiveness as the total abatement costs in both jurisdictions with no linking divided by the total abatement costs in both jurisdictions with restricted (or full) linking. In cases where a

shifts between jurisdictions as a result of carbon markets affecting competitiveness – which depends on several factors, including differences in carbon prices, how allowances are allocated to entities, whether allowances prices are reflected in the cost of products, and the extent to which barriers may limit the trade of products among jurisdictions.[9] We do not assess potential macroeconomic impacts, such as effects on employment, economic growth or changes in environmental co-benefits, since they strongly depend on the specific context of the jurisdictions and the nature of abatement options pursued or not pursued as the result of linkage.[10]

3. *Political feasibility.* Political feasibility is closely related to the first two criteria or, perhaps more importantly, to how economic and environmental outcomes are perceived by influential actors. For instance, a linking decision may be more attractive if it can be shown to maintain or increase global emissions benefits, and thus helps demonstrate global leadership (Bodansky et al. 2014), as could be the case with some forms of exchange rates or discount rates, as discussed below. The impact of linking on a jurisdiction's emissions goals, investment in low-carbon technologies, abatement costs for covered entities, or consumer costs, and how such impacts are assessed and communicated, are key drivers of political feasibility. Linking, whether restricted or full, will generally involve tradeoffs among policy objectives and thus, as with most difficult public policy decisions, can be controversial among stakeholders. To complicate matters, the outcomes of linking options can be quite uncertain in some cases, both to understand ex ante and to detect "real-time". The fact that one jurisdiction today has a lower allowance price than another does not guarantee that this situation will not change due to external (e.g. global fuel prices) or internal (e.g. economic growth, other energy or climate policy measures) factors, thus changing the expected flow of units and financial transfers among jurisdictions. Linking options can be examined with respect to how they affect or potentially erode the sovereignty and regulatory autonomy of a jurisdiction's decision-makers, or the extent to which linking demonstrates "momentum" that increases support from domestic stakeholders.

---

restricted linking option changes the overall abatement outcome, we use this changed abatement level in comparing the costs with no linking; to this end, we assume that the abatement under no linking would change in each jurisdiction proportionally to their emission reduction targets.

[9] We assess the risk for carbon leakage only with regard to the first aspect, the exposure to allowances prices, but do not consider other aspects, as these may strongly depend on the specific circumstances of sectors and jurisdictions.

[10] Some abatement options, such as investment in improved building efficiency, may have significant macroeconomic benefits to the extent that they directly stimulate job creation and reduce purchases of imported energy, and re-spending of any cost savings stimulates the local economy. Other abatement options, in contrast, can involve substitution of a locally produced commodity (e.g. coal) with an imported one (e.g. natural gas), with potentially negative macroeconomic implications.

4. ***Other practical and overarching considerations.*** A number of additional practical considerations may play a role in determining the suitability of restricted linking options. These include administrative cost and complexity: some options may be harder to implement or to communicate effectively, and thus may be harder to "sell" to target audiences. A key overarching consideration is how a linking decision could enhance or degrade the vulnerability or resilience to unexpected developments, such as an economic recession or the growth of emissions-intensive industries, in one or both of the jurisdictions. Such developments can have economic, environmental, and political implications, as the recent history of the EU ETS suggests.

3.2 Using a simplified modeling analysis to assess implications

In our model, we define each jurisdiction by two assumptions: (1) the total abatement, reflected by the difference between its emission cap and business-as-usual (BAU) emissions, and (2) a simplified, linear MAC curve that represents how its cost of reducing emissions increases with the level of abatement. We use this simplified model to assess environmental and economic implications of the different options.

As an illustrative example, we set up the model such that, in the absence of any linking, one jurisdiction has an allowance price (per ton of CO_2) three times higher than the other. We refer alternately here to the "lower-price" jurisdiction (or jurisdiction A) and the "higher-price" jurisdiction (or jurisdiction B). We assume that each jurisdiction's cap leads to similar levels of abatement. Thus, the reason for the difference in price is that jurisdiction A has more tons CO_2 of abatement opportunities at lower cost.

While we construct this simplified example, the direction of environmental and economic implications can be generalized to other two-jurisdiction relationships,[11] where marginal abatement costs increase in a roughly linear fashion, targets are set below BAU levels, and the cost curve is largely unaffected by price containment mechanisms (such as offsets, reserves, floors, caps and/or triggers). Indeed, many ETSs have price containment mechanisms and in some cases targets above BAU (over-allocation), as well as more complex abatement cost functions. These features create well-recognized challenges to linking. We highlight some potential implications of these features in Section 3.8 below.

3.3 The reference cases: no linking and full linking

The more familiar cases of unlinked and fully linked ETSs create a helpful backdrop for examining restricted linking options. The implications and outcomes of these two "reference cases" are relatively straightforward and well-understood. In the absence of linking, differing emission caps, marginal abatement costs, and other design elements lead to different allowance prices among systems. If systems are fully linked (with design elements aligned), their allowance prices would

[11] For example, see the parameterization of California and RGGI ETSs in Burtraw et al. (2013).

converge, and abatement investment and activity would shift from the higher-price jurisdiction (B) toward the lower-price jurisdiction (A).

Figure 3: Implications of no linking and full linking (simplified model)

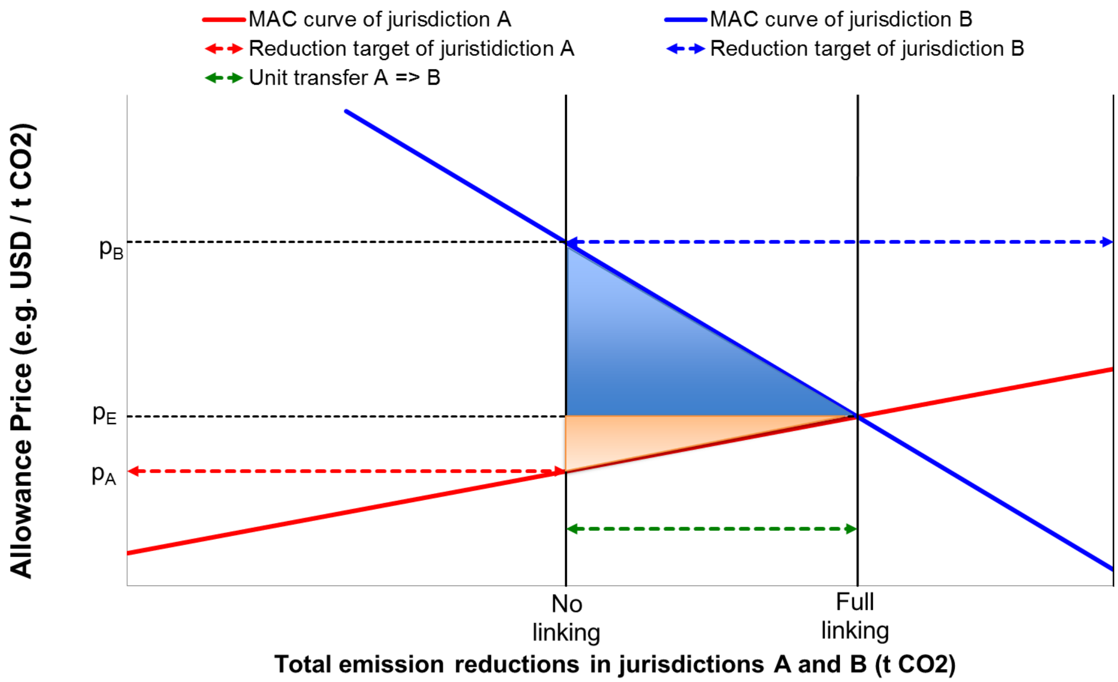

Figure 3 depicts this dynamic. It shows the contribution of each jurisdiction to total emission reductions or abatement (the width of the x-axis). Assuming no linking, abatement activity in the lower-price jurisdiction (A) moves from left to right along an abatement cost curve (red line) until its targeted emission reductions are achieved (width of the red dotted line). Similarly, abatement activity in the higher-price jurisdiction (B) moves along its abatement cost curve (blue line) but starting from the right side of the chart until its targeted emission reductions are achieved (width of the blue dotted line). The resulting allowances prices in the two systems (p_A and p_B) differ, as shown.

If these systems are fully linked, the allowance prices in the two systems converge at a common price (p_E). More emissions are abated in the lower-price jurisdiction (A) and fewer in the higher-price jurisdiction (B). Units are transferred from the lower-price jurisdiction (A) to the higher-price jurisdiction (B) (green dotted line), with a net financial transfer from the higher-price to the lower-price jurisdiction (the amount of units transferred, shown by the green dotted arrow, times the linked-system allowance price).

Full linking leads to several economic benefits. Total costs to achieve the overall abatement outcome are reduced, and these benefits from increased cost-effectiveness are shared among the two jurisdictions. While there is greater abatement in the lower-price jurisdiction (A), the financial transfer (by entities) from the higher-price jurisdiction (B) in acquiring allowances exceeds the cost of the added abatement, providing an economic surplus, as shown by the red shaded area in Figure 3. For the higher-price jurisdiction (B), overall costs decrease because the reduction in abatement costs exceeds the transfer payment made to the lower-price jurisdiction (A), creating the economic surplus shown by the blue shaded area.

Furthermore, linking enhances market liquidity by increasing the number of allowances and actors across the linked market, creating an added economic benefit. Full linking may also reduce carbon leakage risks. To the extent that businesses in the two jurisdictions compete for the same markets, the risks of carbon leakage between the two jurisdictions may decrease, because the regulated entities are exposed to the same, rather than differing, carbon prices. Effects on potential carbon leakage outside the jurisdictions, likely to be a more significant concern, may increase or decrease, depending on whether linking increases or decreases allowance prices relative to no linking.

The primary beneficiaries of the economic surplus from linking are the entities covered by the ETSs. To the extent that allowances are auctioned, linking will, other things equal, tend to reduce overall auction revenues across the jurisdictions. Auction proceeds will increase in the lower-price jurisdiction (A), and decrease in the higher-price jurisdiction (B). Where auction proceeds are already earmarked for specific uses (tax rebates or investments), the decline, if significant, could create a political barrier to linking for jurisdiction B.

3.4 Quotas

Quotas will only have implications if they effectively limit the transfer of allowances. If a quota is set larger than, or equal to, the allowance transfer that would occur with full linking, the outcome should, in principle, be the same as under full linking. The same holds if the lower-price jurisdiction (A) establishes a quota for allowance imports. A quota is thus only effective if it limits imports to the higher-price jurisdiction (B), and if it is set lower than the allowance flow in the market equilibrium under full linking.

Figure 4 illustrates the implications of a quota that enables 50% of the unit transfer that would occur with full linking (orange lines). Quotas imply that different prices persist in the two jurisdictions, though with a smaller difference ($p_{A,Quota}$, $p_{B,Quota}$). As under full linking, abatement increases in the lower-price jurisdiction (A) and decreases in the higher-price jurisdiction (B), but to a lesser extent than under full linking. The amount of allowances transferred (green line) corresponds to the quota limit.

Figure 4: Implications of a quota set at 50% of the net allowance transfers under full linking (simplified model)

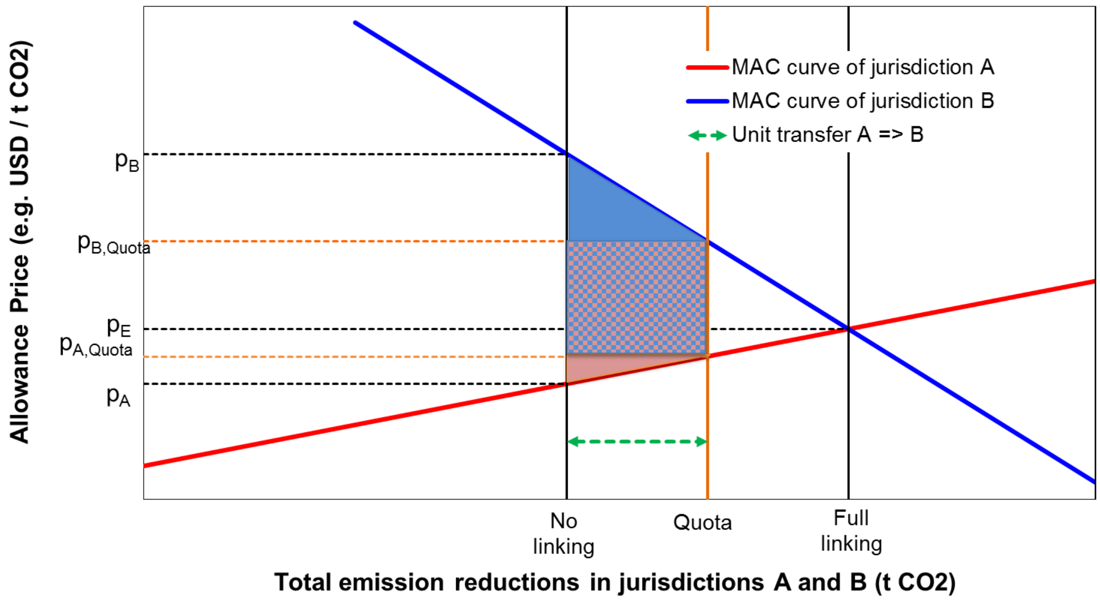

Generally, the economic implications of a quota will tend to lie between those of no linking and full linking. As Figure 4 shows, economic gains (shaded area) per net unit transferred among jurisdictions decrease as one moves from no linking (where the price disparities are highest) toward full linking (where price disparities disappear). For example, with linear MAC curves, such as in Figure 4, a 50% quota yields 75% of the cost-effectiveness gains from full linking.

Another interesting implication is how the relative economic gains for the two jurisdictions could differ from those under full linking. With quotas, the economic benefit shown by the hatched rectangle in Figure 4 could in principle accrue to either jurisdiction A or B, depending on how the quota is implemented. If the quota is expressed as a percentage of the compliance obligations by entities in the higher-price jurisdiction (B), it would limit the *demand* for allowance transfers from the lower-price jurisdiction (A). As the entities in the lower-price jurisdiction would compete to sell allowances to the higher-price jurisdiction (B), the price for such transfers would likely settle at the allowance price in the lower-price jurisdiction ($p_{A,Quota}$), and the economic benefit shown by the hatched area in Figure 4 would confer to the entities in the higher-price jurisdiction (B).

However, a quota could also be set on the *supply* side, as an export quota. To implement an export quota, the lower-price jurisdiction (A) could limit the number of allowances that can be used by other jurisdictions (e.g. B), for example, through auction licenses or permits that entities from other jurisdictions would need to hold and surrender when using its allowances outside the jurisdiction. Entities in the higher-price jurisdiction (B) would then compete for the rights to use a limited amount of allowances from the lower-price jurisdiction (A), and the price for using transferred allowances would likely settle at the allowance price in the higher-price jurisdiction ($p_{B,Quota}$). In other words, the lower-price jurisdiction (A) would capture the scarcity rent (hatched area in Figure 4). In general, quotas that limit *supply* will create greater economic benefits for the exporting, lower-price jurisdiction (A), while quotas that limit *demand (or use)* will create greater economic benefits for the importing, higher-price jurisdiction (B). Jurisdictions could also devise approaches designed to share the scarcity rent,[12] which might be quite substantial if the price differences between jurisdictions remain significant after implementing the quota.

Relative to no linking, quotas will increase market liquidity as they will increase the number of potential buyers and holders of units. Without full linking, however, jurisdiction A and jurisdiction B allowances would not be fully fungible and would therefore continue to be traded as different commodities at commodity exchanges. The implications of quotas for carbon leakage risks vary in a proportionate manner with the changes in prices, between no linking and full linking.

From a political perspective, quotas offer some attractive features. Quotas should have no effect on overall abatement across jurisdictions. In cases of over-allocation in one of the schemes, the overall ambition could indeed be lowered, but less so than with full linking. Assuming little or no over-allocation, to the extent they enable allowance flow among jurisdictions (as compared to no linking), quotas can enhance cost-effectiveness, and even low quotas may create significant economic benefits that can be shared among jurisdictions, as compared with no linking. Perhaps most importantly, quotas allow policy-makers within a jurisdiction to retain a certain level of control on the extent of unit flows and related impacts, assuming that the quotas are set by regulators in the linking jurisdictions. Limits akin to quotas have already been used in several ETSs, such as the EU ETS quantitative restrictions on the use of CDM and JI credits.

In terms of practical feasibility, the use of quotas still raises several key questions:

At what level should the quota be set (and on what basis)? As noted above, quotas are only effective if set lower than the unit flow in the market equilibrium under full linking. However, that flow cannot be known with full certainty in advance. Analysts could employ a modeling approach such as that shown schematically in Figure 4, using estimated marginal abatement potentials and costs, relative to a projection of BAU emissions. But history has shown that abatement potentials, abatement costs, and emissions trajectories can often diverge significantly from expectations. The political context of any given jurisdiction will influence how restrictive the quota policy-makers might aim for. For example, where linking is expected to lead to significantly higher allowance prices, covered entities (that need to acquire allowances through trading or auctions) may favor a lower quota to limit any price increase.

How often should the quota be updated (and under which conditions will it be updated)? Updating the quota may make sense in case of unexpected developments (e.g. a substantial increase in abatement opportunities in jurisdiction B may warrant a lowering of the quota or even obviate the need for one). However, updating may also undermine investment certainty, unless changes in quota levels are predictable. Revising quotas may also have implications for the other trading scheme, and may require the explicit agreement of regulators in the other jurisdiction; such a requirement may be laid down in a linking agreement, but it could also be made subject to a

[12] For example, the two jurisdictions could each issue and auction an agreed share of the permits that allow entities in the higher-priced jurisdiction to use the lower-priced jurisdiction's units.

simplified amendment procedure (cf. Mehling and Haites 2009). One option could also be an automatic adjustment of the quota, triggered by observed allowance prices. If allowance prices exceed a threshold over a defined time period, an import quota may be enlarged, and conversely, it may be adjusted downwards if prices fall below a floor. Such an approach might be similar to the price thresholds for importing (offset) units originally included in RGGI.

Who should set the quota? Quotas could most easily be introduced in an agreement between both jurisdictions to set quotas, taking into account that circumstances (e.g. abatement opportunities and costs) may change over time.

3.5 One-way linking

With one-way linking, units would flow in only one direction. One jurisdiction would recognize the use of another's units for compliance, without the other jurisdiction reciprocating. If under a one-way link, a jurisdiction recognized units from a jurisdiction with clearly lower prices, most outcomes would still resemble those of full linking: allowance prices in the two ETSs would converge at the same, lower level, and abatement would shift toward the lower-price jurisdiction in the same manner as full linking. One important difference with full linking is that liquidity would not increase significantly, as the two unit types would not be fully fungible and therefore would continue to be traded as different commodities.

Although one-way linking may suggest that systems may not need to be as fully harmonized as would be required for full linking, certain design elements may still require altering. For instance, Australia had planned a one-way link with the EU ETS that would have enabled Australian ETS entities to acquire and surrender a limited number of EU allowances. The intent was to consider mutual recognition and fuller linking at a later time, thus postponing the need to address a number of questions that might otherwise complicate full linking negotiations. However, Australia did agree to not implement a price floor and to lower the amount of offsets that could be used.[13] The proposed link also raised concerns about the current over-allocation in the EU ETS; while the one-way link would have lowered compliance costs in Australia and helped to reduce unit surpluses in the EU ETS, it might have appeared to weaken the ambition of the Australian system, while creating financial flows from Australia to the EU.

3.6 Exchange rates

The implications of exchange rates are more complex than those of quotas. First, in contrast to full linking or quotas, exchange rates can affect the total abatement across the two jurisdictions. Second, the implications for cost-effectiveness and total abatement are highly sensitive to the range of values in which the exchange rate is set. Third, the overall implications are linked to how exchange rates are set, depending on: (1) the rationale and policy objectives pursued when setting the exchange rate (e.g. enhancing market liquidity, gaining economic benefits, ensuring

[13] See http://europa.eu/rapid/press-release_IP-12-916_en.htm.

the same or better environmental outcome relative to no linking); (2) the frequency and manner in which the exchange rate might be updated (fixed, fixed with updates, floating); and (3) who sets the rate. As we will show below, exchange rates could either create considerable environmental and economic benefits or have major (unintended) adverse consequences, even compared with no linking, due to information asymmetries between the regulated entities and policy-makers setting the exchange rate, as well as uncertainties about future developments. We therefore discuss the implications for different exchange rate levels and explore different options for setting exchange rates.

For the purpose of our analysis, we define the exchange rate as the number of units from the jurisdiction with the *expected* lower price absent linking (A) that are needed for compliance in the jurisdiction with the *expected* higher price absent linking (B). As illustrated in Figure 2 above, an exchange rate of 3 would imply that three units from jurisdiction A are required to emit one more ton of CO_2 in jurisdiction B, or that with one unit from jurisdiction B, three more tons could be emitted in jurisdiction A.

Three ranges of exchange rate values and their implications

Using the model and parameters described in Section 3.2, Figure 5 illustrates how key indicators for environmental and economic benefits – total abatement outcome and cost-effectiveness – fluctuate with the exchange rate. For ease of comparison, the figure also illustrates the outcome for full linking and quotas. The blue line reflects different exchange rate levels (R) that increase clock-wise on the curve.

Figure 5: Implications of varying exchange rates on total abatement and cost effectiveness

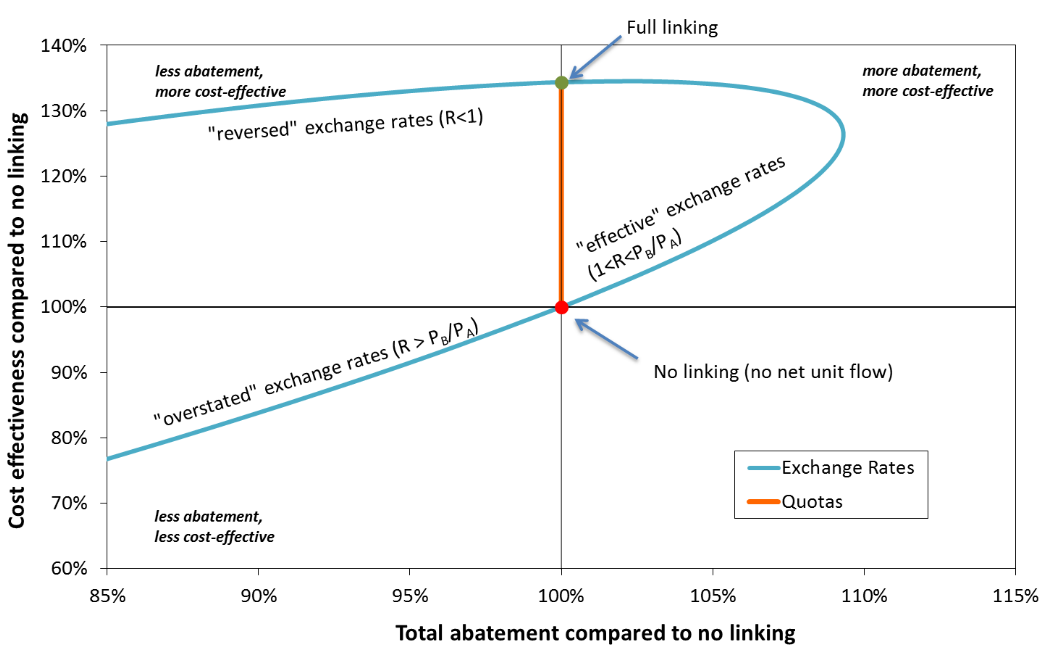

The implications of exchange rates differ for three ranges of possible values:

"Effective" exchange rates: We define *effective* exchange rates, found in the upper right quadrant of Figure 5, as rates at which both the total abatement outcome and the cost-effectiveness are higher than with no linking. Exchange rates are *effective* if set above 1 (which corresponds to full linking, the green dot in Figure 5), but lower than the ratio of allowance prices under no linking ($1 < R < p_B / p_A$).[14] An exchange rate that reflects this ratio (which corresponds to 3 in our example) has the same implications as no linking (the red dot in Figure 5), since regulated entities in the two jurisdictions would not have any economic benefits from net transfers of allowances. In the range of *effective* exchange rates (between 1 and 3 in the example here), the net flow of allowances is from the lower-price jurisdiction (A) to the higher-price jurisdiction (B), and as a result fewer emissions (one-third as many) are allowed in the higher-price jurisdiction (B) for every allowance from the lower-price jurisdiction (A) (see upper arrow in exchange rate case shown in Figure 2). As exchange rates move from the no linking value (3 here) to parity (the full linking value), the flow of allowances from A to B increases (see Figure 6 below), while the net emissions benefit per unit transferred decreases, leading to the behavior shown in Figure 5, where the overall abatement outcome reaches a maximum at a middle range rate (in our example the maximum occurs at a rate of 1.74).[15]

"Reversed" exchange rates: We define *reversed* exchange rates, found in the upper left quadrant in Figure 5, as below 1 ($R < 1$). We refer to them as *reversed* exchange rates because they reverse in which jurisdiction allowances are valued higher: while units still flow from the lower-price jurisdiction (A) to the higher-price jurisdiction (B), their value is effectively "inflated": less than one allowance from the lower-price jurisdiction (A) is required to emit one more ton in the higher-price jurisdiction (B). As illustrated in Figure 6, abatement decreases in each jurisdiction relative to full linking, and as a result total abatement declines relative to both full linking and no linking. Cost-effectiveness decreases relative to full linking as well, though it remains above the no linking level until exchange rates reach low levels (not shown in Figure 6).

"Overstated" exchange rates: We define *overstated* exchange rates, found in the lower left quadrant in Figure 5, as rates set above the ratio of allowance prices under no linking ($R > p_B / p_A$). We refer to them as *overstated* because they exaggerate the expected price difference between both jurisdictions with no linking. Remarkably, they reduce both total abatement and cost-effectiveness relative to no linking. In other words, *overstated* exchange rates lead to worse outcomes than no linking. Under *overstated* exchange rates, the direction of allowance flows and transfer payment reverses, with entities in the lower-price jurisdiction (A) buying allowances from the higher-price jurisdiction (B). As a consequence, less abatement occurs in the jurisdiction with the lower-cost abatement opportunities (A), while more abatement occurs in the jurisdiction with higher-cost abatement opportunities (B), decreasing the cost-effectiveness. While not calling

[14] We use p_A and p_B to refer to the allowance prices that would prevail were the jurisdictions not linked. As discussed below, were systems to be linked, then these relative prices would no longer be observable.

[15] The cost-effectiveness is highest in the case of full linking and decreases with increasing exchange rates. The relative losses in cost-effectiveness are smaller with an increase in the exchange rate close to 1 than with a further increase close to the situation of no linking (Figure 6 below).

them *overstated* exchange rates per se, Burtraw et al. (2013) discuss such rates in the context of linking California-Québec and RGGI.[16]

Table 1 compares these three types of exchange rates.

Table 1: Possible ranges of exchange rates

	Exchange Rates		
	Effective	**Overstated**	**Reversed**
Exchange rate (units of lower price jurisdiction (A) per higher-price jurisdiction (B) unit)	Set between 1 and relative unlinked prices (pB/pA)	Set above ratio of relative unlinked prices (pB/pA)	Set below 1
Direction of net allowance flow	From lower-price (A) to higher-price jurisdiction (B)	From higher-price (B) to lower-price jurisdiction (A)	From lower-price (A) to higher-price jurisdiction (B)
Conditions that could lead to (unintentionally) setting the exchange rate at overstated or reversed ranges		The likely unit price in the absence of linking was estimated too high in B and/or too low in A. This could occur, for example, if: • Abatement is much more (less) costly in A (B) than expected; • BAU emissions rise much more (less) rapidly in A (B) than expected.	The likely unit price in the absence of linking was estimated too low in B and/or too high in A. This could occur, for example, if: • Abatement is much more (less) costly in B (A) than expected • BAU emissions rise much less (more) rapidly in B (A) than expected

In all three cases, assuming no constraints on allowance flow or use, the exchange rate will determine the ratio of allowance prices: the price in jurisdiction B is always R times higher than the one in jurisdiction A. This relationship has important implications for the setting of exchange rates: the regulators or other entities entrusted with setting exchange rates have considerable policy and political impact, as the exchange rate establishes relative allowance prices – which also affect levels of domestic abatement – among the linked jurisdictions.

In comparison with quotas, exchanges rates allow the two jurisdictions' allowances to be fully fungible. Therefore, any market developments, such as changes in demand for allowances due to changes in economic growth, climatic conditions, technological developments, or target adjustments, will affect both systems equally, as under full linking. For the same reasons, exchange rates offer the same liquidity as full linking; a single financial product could be developed for the two commodities.

[16] Burtraw et al. (2013) evaluated an exchange rate of 3 between California-Québec and RGGI. The ratio of current allowances prices, adjusted by their different metrics (short ton and metric ton), is lower. However, the findings are similar to the simplified example in Section 3, though our example does not account for the presence of floor prices.

Not surprisingly, the potential risks for carbon leakage between the two systems resemble those of full linking with exchange rates close to 1, and increase with exchange rates moving away from 1 (below or above).

The financial transfer from the higher-price jurisdiction (B) to the lower-price jurisdiction (A) is highest under full linking and decreases with lower or higher exchange rates, and reverts to financial transfers from the lower-price jurisdiction to the higher-price jurisdiction for *overstated* exchange rates (see Figure 6). The financial implications for any revenues from auctioning allowances are proportional to the change in allowance prices (see Figure 6).

Figure 6: Impacts of varying exchange rates on allowance prices, abatement, and transfer payments (using simplified model assumptions)

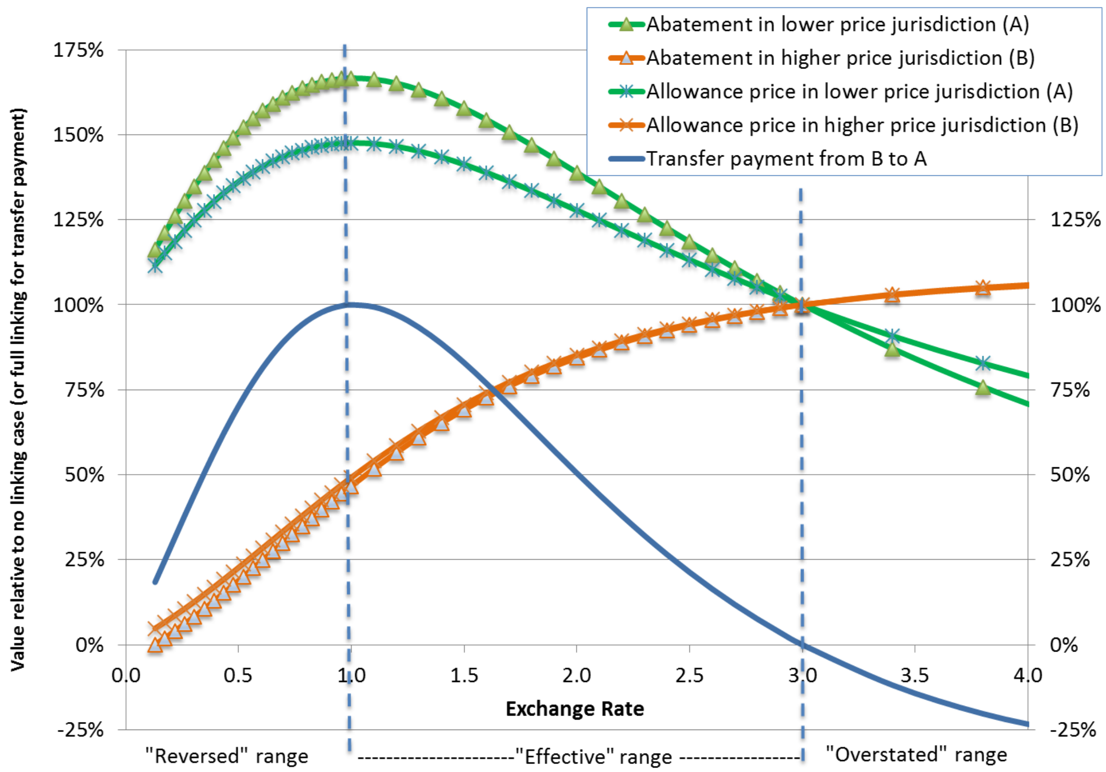

How can exchange rates be set?

Exchange rates could be introduced to pursue different policy objectives, and could be set at different values depending on which policy objectives are prioritized. If the primary policy objective is enhancing market liquidity while avoiding any other impacts, such as a relocation of abatement, financial transfers, or changes in allowance prices and auctioning revenues, one could argue that an exchange rate set at the ratio of the allowance prices under no linking ($R = p_B / p_A$) would best fit that purpose. Exchange rates close to that ratio provide for enhanced liquidity, while limiting any other implications. If a key policy objective is enhancing cost-effectiveness, values closer to 1 may provide more benefits. If policy-makers intend to ensure that total abatement is increased rather than decreased, moderate exchange rates, set between 1 and the ratio of allowance prices under no linking, would best ensure that this policy objective is met (aside, of course, from tightening the cap). Hence, approaching the "best" exchange rates may require balancing different policy objectives. That said, for a broad range of policy objectives, policy-makers will likely want to set the exchange rate within the spectrum of *effective* exchange rates ($1 < R < p_B / p_A$), rather than *overstated* or *reversed* exchange rates, which lead to several adverse impacts.

In practice, setting the exchange rate at the "ideal" level along this *effective* range and ensuring that it remains within the *effective* range is challenging. Doing so requires that those who set exchange rates have good information, as well as foresight, for each jurisdiction, on BAU emissions and on the abatement potential and costs. Yet, BAU emissions and abatement potential and costs may face major uncertainties related to future economic growth, technological development, international fuel prices, or changes in climatic conditions. Furthermore, there may be information asymmetries between regulators (including those who establish exchange rates) and the regulated entities with respect to abatement opportunities and costs. Because of these uncertainties and information asymmetries, a regulator might set an exchange rate *ex ante* that might, based on best available information at the time, appear to be in the *effective* range, but turn out to be a *reversed* or *overstated* exchange rate.

For example, an unanticipated decline in international energy prices could mean that a (coal-dependent) jurisdiction that was previously expected to be the higher allowance price jurisdiction could now become the lower price one (if for example, the coal-gas price spread were to narrow). Such a development could convert what was designed as an *effective* exchange rate into a *reversed* exchange rate. Alternatively, if economic growth in the higher-price jurisdiction were weaker than expected and lowered demand (and prices) for allowances significantly, an exchange rate could become *overstated*, over-valuing jurisdiction B units, and resulting in both decreased abatement and cost-effectiveness compared with a no linking situation. Inadequate information about the abundance of low-cost abatement opportunities in the higher price jurisdiction (due e.g. to information asymmetry) could also lead to a similar outcome.

Adjusting exchange rates over time

Regulators could aim to adjust exchange rates to mitigate these risks. However, in practice, this too could be challenging. If there is no quota or constraint on the flow and use of units, once systems are linked through an exchange rate, only a single price signal will remain (mediated by the exchange rate). It then becomes impossible to observe, and difficult to impute, what the relative allowances prices in different jurisdictions would have been were the systems not linked, the ratio of which (p_B / p_A) is essential for understanding whether an exchange rate is *effective*, *overstated*, or *reversed*. Regulators can track information on the magnitude and direction of allowance flows, but such data would provide only a rough and time-lagged indicator of relative prices under no linking. This is true regardless of whether there is effective real-time tracking of transactions or holdings.

For example, in principle a reversal in the expected direction of net flow of allowances could be used as an indication of an unintentionally *overstated* exchange rate, and trigger a policy intervention to adjust the exchange rate. However, the ultimate use of allowances, and hence the overall *net* flow of allowances between jurisdictions in a compliance period would only be known for certain at the time of allowance surrender, which might be as seldom as once a year, depending on ETS compliance requirements, and adjusting the exchange rate at this point in time might be too late to avert an exchange rate error. In the interim, purchases and holdings of units could be tracked, but this information also has practical limitations: allowances may temporarily be held in accounts in other countries than where they will be ultimately surrendered. For example, financial institutions often hold and trade allowances on behalf of regulated entities. The financial institutions may have accounts in other countries than those of the regulated entities surrendering the units for compliance. Moreover, allowances could be swapped and banked into future periods. Modeling allowance prices in the absence of linking could be another approach, but models may fall short in reflecting the behavior of market participants in real-time and identifying the real abatement opportunities and costs due to information asymmetries. Hence, adjusting exchange rates over time faces several practical constraints; it might help to prevent the worst outcomes – e.g. highly *overstated* exchange rates – but requires careful consideration and further analysis.

When adjusting exchange rates over time, key questions include the frequency of any adjustments and whether the way in which the exchange rate is adjusted is pre-determined (e.g. through an algorithm, as proposed for the market stability reserve of the EU ETS), or is decided based on the information available. The frequency may need to balance responding to short-term effects and long-term trends. As pointed out above, in a fully fungible market, the net flow of allowances between jurisdictions may change over hours, days, or weeks and may not be related at all to scarcity in either jurisdictions, but rather to where market players, such as financial institutions, are based. In this regard, only the use of allowances for compliance could provide insights on long-term trends with regard to the direction of *net* unit flows. However, in most ETSs, allowances are surrendered only annually, and hence any adjustments may come significantly later than an unforeseen market shock that changes the effects of the exchange rate.

Who should set exchange rates?

It is conceivable that exchange rates could be negotiated and agreed by the jurisdictions involved. Alternatively, an independent international body could set exchange rates. However, given that exchange rates can have significant, unintentional adverse impacts on economies and climate mitigation, jurisdictions may be reluctant to delegate the authority of setting exchange rates to a third party. Others have suggested that financial markets themselves could set exchange rates, and they could indeed "float", as they do for currencies and other financial products.[17] However, it remains unclear how this would work in practice. Unlike other products and services and the currencies used for their exchange, emissions allowances have no real "value" outside the markets created by the regulators themselves.

---

[17] See, e.g., the World Bank's Globally Networked Carbon Markets Initiative, http://www.worldbank.org/en/topic/climatechange/brief/globally-networked-carbon-markets

Overall feasibility of exchange rates

Exchange rates may be viewed as a valuation of a jurisdiction's climate mitigation actions; this perception makes this option particularly politically sensitive. Ensuring that an exchange rate actually delivers the intended policy objectives is challenging, as noted above. This predicament is not unique to exchange rates; one can point to other policies in place which have not always achieved their intended objectives, such as the CDM and its failure to deliver sustainable development benefits, especially in its early years. However, the novel and complex nature of exchange rates may make them somewhat more difficult for policy-makers to embrace. It also requires negotiating the terms of exchange rates, including the basis for any adjustments over time, with other jurisdictions – or entrusting this process to a third party – which could be as challenging or more than negotiating the terms for harmonization of rules and full linking.

3.7 Discount rates

Discount rates can be regarded as a variation on exchange rates, with a view to addressing some of the challenges inherent to exchange rates. They could help avoid setting rates unintentionally at levels that lead to adverse outcomes, and thus help ensure that the envisaged policy objectives are actually achieved. In other words, discount rates could be designed more easily in ways that ensure that rates are *effective*, and neither *reversed* nor *overstated*.

While exchange rates inherently require a symmetrical relationship in the value of jurisdictions' allowances (e.g. 3:1 from A to B and 1:3 from B to A), discount rates do not. Jurisdictions could apply one discount rate in one direction of allowance flow (e.g. 3:1 from A to B) and parity (1:1) or a rate of different magnitude in the other direction (e.g. 2:1 from B to A).

To ensure that they are set at an *effective* range – i.e. they lead to both higher cost-effectiveness and higher overall abatement outcome compared with no linking – discount rates can be implemented in either of two ways:

1. *One-way discount rates.* One jurisdiction could allow imports of allowances from another jurisdiction, converted at a discount rate above 1, i.e. more than one allowance from the other jurisdiction is required to emit one more ton in the importing jurisdiction. No unit exports would be allowed to the other jurisdiction.
2. *Two-way discount rates*. Both jurisdictions could allow importing units, each using a discount rate that increases the abatement outcome. In both directions of allowance flows, more than one imported allowance unit would be required to emit one more ton in the importing jurisdiction.

With both options, policy-makers would ensure that they always "land" at an *effective* exchange rate, in the right upper quadrant of Figure 5 above, even under information asymmetry and uncertainty with regard to BAU emissions, abatement potential, and abatement costs. They would guarantee that abatement could only increase as a result of allowance trade among jurisdictions, and cost-effectiveness could only increase as well: without an economic benefit to accompany the added abatement, entities would have no incentive to acquire and surrender imported allowances.

With discount rates, allowances are only transferred between jurisdictions if the difference in allowance prices without linking exceeds the discount rate. For example, if both jurisdictions applied a 20% discount to allowances imported from the other jurisdiction (1.25 imported allowances required per ton emitted), trade would only start if the price difference between the jurisdictions (with no linking) is larger than this differential. If the price difference is smaller, the outcome would be the same as for no linking. In other words, the price difference between the two jurisdictions is capped to a maximum of 20%. In this regard, such discount rates can work in part as price containment mechanisms.

When units are transferred, the implications are practically the same as those for *effective* exchange rates described above. However, unlike exchange rates, discount rates do not provide for full fungibility of allowances and hence, like quotas or one-way linking, they fail to provide for significantly enhanced market liquidity.

3.8 Implications of key ETS design features and conditions

In this section, we briefly discuss how some key ETS design features (price containment and offsets) and market conditions (over-allocation or multiple links) might affect linking options. In general these considerations are complex, and the implications are likely to be highly context-specific, dependent upon, for instance, exactly how price floors are set, how offset restrictions differ, or how long-lived any over-allocation may be.

Price containment

Several jurisdictions have introduced or are considering mechanisms that affect allowance prices, including price floors, price ceilings, or mechanisms to lower or expand allowance availability, such as the allowance reserve in place in California or the "market stability reserve" planned for the EU ETS. Such mechanisms can have considerable impacts when linking ETSs. The impacts depend not only on the design of the price mechanisms but also on the size of the ETSs, their caps, and their marginal abatement cost curves. A detailed assessment is beyond the scope of this study, but we briefly assess the implications of the most common of these features, price floors, for different options for (restricted) linking.

Price floors are generally established as minimum prices for allowance auctions (as done for example in RGGI, California, and Québec). Generally, price floors only affect market prices if, in the market equilibrium, prices would normally fall below the floor price. However, linking two ETSs changes the allowances prices and could thus make a price floor "kick in" where it would not otherwise.

If two ETSs have a price floor, full linking will generally, but not always, impose the higher price floor on both schemes. Similarly, if only one of the two ETSs has a price floor, it will become applicable in both schemes. However, the outcome of linking jurisdictions with differing price floors, price caps, and allowance reserves is difficult to assess without reference to specific design features, sizes, and market conditions. If linking does change the level of effective floor prices, it could affect the amount of allowances that might remain unauctioned, which in turn could affect the amount of allowances available for compliance, and thus total emissions (abatement) across the two systems. These and other questions related to price containment become even more challenging and important to assess if exchange rates are involved.

In principle, linking with exchange rates will transmit the price floor level times the exchange rate to the other system, and the higher of the two price floors becomes effective. For example, if jurisdiction A had a price floor of $5 and jurisdiction B one of $12, an exchange rate of 3 would increase the effective price floor in jurisdiction B from $12 to $15, while the price floor would remain at $5 in jurisdiction A. With an exchange rate of 2, the effective price floor in jurisdiction A would increase from $5 to $6, while it would remain at $12 in jurisdiction B. As this relatively simple example suggests (no reserves are included, for example), the implications of price floors may not always be intuitive; a lower, absolute price floor in the lower price jurisdiction could implicitly impose a higher price floor in the higher price jurisdiction (B).

In conclusion, restricted, just as full, linking could affect how price containment features function; and in that sense, policy-makers will need to consider the extent to which they may be willing to compromise their price control objectives or modify these features.

Offsets

Many ETSs allow entities to use offsets to meet their compliance obligations. Most ETSs impose both quantitative and qualitative restrictions, limiting the number and the type of offsets eligible for compliance. As with price floors, the implications of linking depend on the specific context of the two ETSs and the potential for, and prices of, offsets.

Quantitative limits on offsets are similar to linking two ETSs with a quota, as illustrated in Figure 4, but where the lower cost marginal abatement cost curve (jurisdiction A) represents the offset supply curve. If a quantitative limit effectively restricts unit flow, a price spread between offsets and allowances emerges, as shown in Figure 4 (and as observed in the EU ETS). Otherwise, offsets and allowances would trade at the same price, as under full linking represented in Figure 3.

Linking can affect the amount and type of offsets used, depending on the context. For example, if jurisdiction A allows unlimited use of offsets while jurisdiction B restricts offset use, the limit in jurisdiction B would become ineffective under full linking because an unlimited number of offsets could be acquired by entities in jurisdiction A, freeing up jurisdiction A allowances that could then be transferred to jurisdiction B entities. Assuming that offsets can be acquired in this manner at costs lower than allowance prices in jurisdiction B, linking would increase the overall use of offsets. The same holds for qualitative limits: any quality restrictions implemented by only one of the jurisdictions would become less effective. It is also conceivable that linking reduces the overall use of offsets. As a result of linking, entities in the higher-price jurisdiction (B) might, for example, rather import allowances from the lower price-jurisdiction (A) than using some higher-cost offsets which they would have otherwise acquired. Alternatively, linking may have no effect on overall offset use, if, for example, both ETSs had quantitative limits that would be fully met with and without linking.

Similar dynamics would be present whether linking were full or restricted. The extent to which quotas or one-way linking would restrict the net flow of allowances might also limit the extent to which offset flows might increase (or decrease). The effects of exchange rates and discount rates on offset use are much harder to generalize, and would depend on how they were set, and whether offsets might also be subjects to discount rates. Policy-makers will need to evaluate the implications of offset provisions when linking, since linking could allow a way around restrictions in offset type or amounts that may be present in one jurisdiction and not another. If set up appropriately, restricted linking options could potentially reduce the implications of different provisions for offset use, depending on the context of the ETS and the options for restricted linking.

Over-allocation (oversupply) of allowances

Several ETSs have been subject to over-allocation or oversupply of allowances, where administrators issue more allowances than ultimately needed to cover BAU emissions. This situation can create excess allowances. Unless these excess allowances are directly accounted for in setting future emissions caps, they present a risk to environmental integrity if traded into other ETSs or otherwise used. If one system is over-allocated and the other not, then linking them could reduce overall emissions abatement.

Restricted linking could reduce, but not fully mitigate, this risk. For example, effective exchange rates or discount rates would reduce the overall amount of allowances available, whereas quotas might reduce the amount of excess allowances imported from the over-allocated system.

3.9 Implications of restricted linking across more than two jurisdictions

Thus far, we have considered potential outcomes of restricted linking options with only two jurisdictions involved. Establishing links with additional jurisdictions is rather straightforward under full linking. However, with restricted linking, some options may be more challenging than others. Implementing quotas or discount rates seems relatively straightforward to implement. Jurisdictions could apply a quota or discount rate to any imports from, or exports to, other jurisdictions, possibly in combination with quotas or discount rates for offsets. The level of any quotas or discount rates may not necessarily need to be mutually agreed among all jurisdictions but could be set and adjusted by each jurisdiction. Exchange rates appear more challenging to implement, as they would require mutual agreement among all involved jurisdictions on their level. Similarly, any adjustments over time would need to be centrally managed. As highlighted above, exchange rates face considerable risks of unintended consequences if the exchange rate is set wrongly. This risk increases significantly if more than two jurisdictions get involved. If abatement costs are higher or lower than expected in only one jurisdiction, the overall outcome in terms of abatement and cost-effectiveness may be affected for the entire system.

3.10 Comparison of options

Table 2 summarizes the results of our analysis, comparing full linking and restricted linking options against the assessment criteria introduced in section 3.1 above. In the next section, we summarize these findings and explore what they might mean for the role of restricted linking.

Table 2: Summary comparison of options against criteria (relative to no linking)

Option	Environmental benefit (total abatement outcome)	Economic benefit (cost-effectiveness and liquidity)	Political feasibility	Practical and overarching considerations
Full linking	• No change (assuming no over-allocation or indirect offset effects) • May influence ambition of future targets	• Maximizes cost-effectiveness • Increases liquidity significantly	• Signals enhanced cooperation and influence • Can face many challenges (price changes, location of abatement and co-benefits, financial flows to other jurisdictions)	• Enables administrative efficiency gains • Requires considerable effort to harmonize (easier if harmonized at outset)
Quotas	• No change (assuming no over-allocation or indirect offset effects) • May influence ambition of future targets	• Enhances cost-effectiveness • Increase liquidity but much less so than full linking	• Can be set in a manner to minimize political risks • Jurisdictions can maintain stronger policy oversight and adjust quota levels	• Less harmonization needed; may be easier to implement than full linking for established ETSs with more disparate features • Hard to determine appropriate level ex ante due to uncertainty and information asymmetry
One-way linking	• No change (assuming no over-allocation or indirect offset effects) • May influence ambition of future targets	• Maximizes cost-effectiveness (if direction of unit flow is the same as under full linking) • Increases liquidity but much less so than full linking and only in one system	• Can be pragmatic first step on path to full linking	• Requires efforts to harmonize (easier if harmonized at outset)
Exchange rates	• Effective rates can increase total abatement; reversed or overstated rates can reduce it • May influence ambition of future targets	• Enhanced cost-effectiveness with effective or moderately reversed rates, decreased cost-effectiveness with overstated rates • Increase liquidity significantly (similar to full linking)	• Mutually agreeing exchange rate values and modalities for future adjustments could be politically challenging	• Hard to determine appropriate level due to uncertainty and information asymmetry • Outcome could be contrary to the policy objectives pursued due to uncertainties and information asymmetries in setting the exchange rate
Discount rates	• Increases total abatement • May influence ambition of future targets	• Enhances cost-effectiveness • Increases liquidity but much less so than full linking	• Jurisdictions maintain stronger policy oversight as discount rates may not necessarily need to be mutually agreed but could be possibly determined and adjusted by each jurisdiction	• Hard to determine appropriate level due to uncertainty and information asymmetry • Less harmonization needed; may be easier to implement than full linking for established ETSs with more disparate features

4. Findings and conclusions

While restricted linking options do not, in principle, achieve the potential benefits of full linking, they can lessen some of the potential pitfalls. They offer easier off-ramps to "de-link"[18] and levers to adjust (e.g. quota levels or discount rates) should linking concerns prove to be more significant than anticipated. Restricted linking options may thus represent a cautious approach that can be more easily implemented and explored where full linking is either infeasible in the near term or incompatible with the objectives of the jurisdictions involved. Borrowing the "match-making" analogy that Comendant and Taschini (2014) invoke for the process of finding the right "linking partner", restricted linking is akin to moving in together, either before (or with no intention of) getting married.

While the benefits (and risks) of linking are often presumed to be economic (cost-effectiveness and liquidity) and political, they can also be environmental in the form of increased or decreased emissions abatement. Over-allocation of allowances in one ETS can reduce the total abatement outcome across jurisdictions when linking, while a price floor in one ETS could increase the aggregated abatement outcome from both ETSs. Linking can change the amount and type of offsets used and thus abatement outcome, depending on the extent of over- or under-crediting by different offset types (Erickson et al. 2014). Among restricted linking options, discount rates can increase the abatement outcome, whereas exchange rates could either increase or decrease it.

In general, restricted linking can reduce, but not wholly avoid, the need for harmonization of design elements required for full linking. Because of the difficulties of "de-linking", full linking requires more careful consideration of differences in design features in order to avoid unexpected or unintended consequences. Restricted linking options that limit the amount of net allowance flows (e.g. quotas and discount rates) provide levers that can be adjusted should negative consequences arise, thus lessening the level of harmonization that might be needed. Yet all of the options considered here create some level of unit flow among jurisdictions and thus would require alignment or coordination of registries and MRV requirements to ensure integrity.

Generally, the implications and feasibility of linking – whether full or restricted – will depend heavily on the design of the ETSs (e.g. allocation rules, price containment measures), the ambition of the caps (e.g. over-allocation or accumulated surpluses), the size of the ETSs (e.g. if one scheme is much smaller than the other), the marginal abatement cost curves, and the use of offsets. It is therefore important for policy-makers to consider the specific context carefully.

The most appropriate option for restricted linking may also depend on the specific policy objectives pursued. If the ultimate goal of policy-makers is indeed full linking, and restricted linkage is viewed as a step toward that end, then quotas or one-way linking may be attractive

[18] The notion of "delinking" has received scant attention in the literature (as well as in linking agreements), but it can be generally seen as terminating a link. Where linking is formalized in a linking agreement, this may require a termination procedure (Mehling and Haites 2009). The way delinking is organized may affect abatement costs as well as subsequent price divergence (Pizer and Yates 2015). Further research could shed light on the economic and environmental implications of pursuing restricted linking options in the process of delinking.

options. In contrast to exchange or discount rates, they do not create different perceptions of the relative value and implied ambition of each jurisdiction's units and ETSs that could ultimately hamper efforts to reach agreement on full linking. They allow for unit parity from the start of unit exchanges, and thus may create an atmosphere of equivalence that can build confidence among stakeholders in both jurisdictions. Yet as illustrated above, quotas, if relatively restrictive, can enable one jurisdiction to capture much of the economic surplus that linking would create. That said, quotas have the advantage of providing a lever (quota levels) that can be adjusted if significant concerns arise. .

Exchange rates could either generate environmental and economic benefits or lead to adverse impacts, depending on how they are set. Their main advantage over other restricted linking options is that they provide for full liquidity, as under full linking. However, information asymmetries and uncertainties in setting the rate could seriously undermine the intended policy objectives. Exchange rates can strongly affect the location, level and cost of abatement, as well as transfer payments, auctioning revenues or co-benefits – to a similar extent as full linking. Especially with more than two jurisdictions linking, the risks of unintended consequences proliferate and setting exchange rates may become politically more difficult. In this regard, one could question whether exchange rates would not mainly add another layer of complexity when negotiating a linking agreement, while not necessarily addressing the concerns with full linking.

By contrast, discount rates can avoid some of the challenges of exchange rates. Discount rates could work in a similar manner as price containment mechanisms and be set up in ways that ensure that both cost-effectiveness and total emissions abatement are enhanced. They implicitly link flexibility and mitigation ambition – the more the regulated entities use the flexibility to import units from another scheme, the more emissions are reduced. Applying discount rates reciprocally (both jurisdictions discount unit imports at the same rate) might address perceptions about valuing reductions differently across jurisdictions. In this regard, they could be further explored as an alternative or complementary approach to quotas or one-way linking.

To conclude, this paper has offered new insights into the options for and issues raised by restricted linking in the absence of, and acknowledging challenges to, fully harmonized ETSs. Our analysis shows that the restricted linking options presented here may be able to reap some of the benefits of full linking, but also offer some important caveats. The outcomes are highly sensitive to the ETS design and context of the prospective linking partners. We recommend that restricted linking options be carefully assessed based on concrete, practical context and policy objectives pursued by the prospective partners.

References

Bodansky, D., Hoedl, S., Metcalf, G. E. and Stavins, R. N. (2014). *Facilitating Linkage of Heterogeneous Regional, National, and Sub-National Climate Policies through a Future International Agreement*. Discussion Paper, Harvard Project on Climate Agreements. Belfer Center for Science and International Affairs, Harvard Kennedy School, Cambridge, MA. belfercenter.ksg.harvard.edu/publication/24568.

Bodansky, D. M., Hoedl, S. A., Metcalf, G. E. and Stavins, R. N. (2015). Facilitating linkage of climate policies through the Paris outcome. *Climate Policy*, (ahead-of-print). 1–17.

Burtraw, D., Palmer, K. L., Munnings, C., Weber, P. and Woerman, M. (2013). *Linking by Degrees: Incremental Alignment of Cap-and-Trade Markets*. RFF Discussion Paper 13-04. Resources for the Future, Washington, DC.
http://www.rff.org/Publications/Pages/PublicationDetails.aspx?PublicationID=22167.

Comendant, C. and Taschini, L. (2014). *Submission to the Inquiry by the House of Commons Select Committee on Energy and Climate Change on 'Linking Emissions Trading Systems'*. Centre for Climate Change Economics and Policy and Grantham Research Institute on Climate Change and the Environment. http://www.lse.ac.uk/GranthamInstitute/wp-content/uploads/2014/10/Comendant-and-Taschini-policy-paper-April-2014.pdf.

Doda, B. and Taschini, L. (2015). Carbon Dating: When is it Beneficial to Link ETSs? *Available at SSRN 2610076*.

Erickson, P., Lazarus, M. and Spalding-Fecher, R. (2014). Net climate change mitigation of the Clean Development Mechanism. *Energy Policy*, 72. 146–54. DOI:10.1016/j.enpol.2014.04.038.

Flachsland, C., Marschinski, R. and Edenhofer, O. (2009a). Global trading versus linking: Architectures for international emissions trading. *Energy Policy*, 37(5). 1637–47. DOI:10.1016/j.enpol.2008.12.008.

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. 358–72. DOI:10.3763/cpol.2009.0626.

Green, J. F., Sterner, T. and Wagner, G. (2014). A balance of bottom-up and top-down in linking climate policies. *Nature Climate Change*, 4(12). 1064–67. DOI:10.1038/nclimate2429.

Haites, E., Mullins, F. and Consultants, M. (2001). *Linking Domestic and Industry Greenhouse Gas Emission Trading Systems*. Margaree Consultants.

Holland, S. and Yates, A. J. (2014). *Optimal Trading Ratios for Pollution Permit Markets*. 19780. National Bureau of Economic Research. http://www.nber.org/papers/w19780.

ICAP (2015). *Emissions Trading Worldwide: ICAP Status Report 2015*. International Carbon Action Partnership, Berlin. https://icapcarbonaction.com/status-report-2015.

Jaffe, J. and Stavins, R. (2007). *Linking Tradable Permit Systems for Greenhouse Gas Emissions: Opportunities, Implications, and Challenges*. International Emissions Trading Association. http://belfercenter.ksg.harvard.edu/files/IETA_Linking_Report.pdf.

Mehling, M. and Haites, E. (2009). Mechanisms for linking emissions trading schemes. *Climate Policy*, 9(2). 169–84. DOI:10.3763/cpol.2008.0524.

Mendelsohn, R. (1986). Regulating heterogeneous emissions. *Journal of Environmental Economics and Management*, 13(4). 301–12. DOI:10.1016/0095-0696(86)90001-X.

Pizer, W. A. and Yates, A. J. (2015). Terminating links between emission trading programs. *Journal of Environmental Economics and Management*, 71. 142–59.

Tuerk, A., Mehling, M., Flachsland, C. and Sterk, W. (2009). Linking carbon markets: concepts, case studies and pathways. *Climate Policy*, 9. 341–57. DOI:10.3763/cpol.2009.0621.