July 2017 ■ RFF WP 17-16

WORKING PAPER

Linking Carbon Markets with Different Initial Conditions

Dallas Burtraw, Clayton Munnings, Karen Palmer, and Matt Woerman

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

Linking Carbon Markets with Different Initial Conditions

Dallas Burtraw, Clayton Munnings, Karen Palmer, and Matt Woerman

Abstract

Linkage of emissions trading systems theoretically minimizes total abatement costs by allowing fungibility of emissions reductions across jurisdictions. We develop a theoretical framework to investigate the implications of linking systems with unique designs. We qualitatively assess the California and the Regional Greenhouse Gas Initiative systems, which we find to be nearly ready to link despite some differences in their initial conditions, including design and stringency. We use a simulation model of regional electricity markets to investigate market outcomes under such a linked system. We consider possible exchange rates for allowances to adjust for differences in program stringency, and we examine how they interact with price floors and ceilings while explicitly representing other program features (e.g., leakage policies, companion policies, and allowance allocation). We find that aggregate emissions and emissions in each jurisdiction change in ways predicted by theory but that efficiency gains can be distributed in nuanced and nonintuitive ways.

Key Words: greenhouse gas, climate change, climate policy, policy coordination

JEL Classification Numbers: Q58, H77

Contents

Linking Carbon Markets with Different Initial Conditions

Dallas Burtraw, Clayton Munnings, Karen Palmer, and Matt Woerman*

1. Introduction

The environmental consequences of greenhouse gas emissions are felt around the globe, regardless of where those emissions originate. Correspondingly, in the 1990s, numerous economists heralded a single international carbon market as the cost-effective solution to climate change. Such a market would, in principle, lead to a single global carbon price through the trade of emissions allowances, which would serve to identify and realize emissions reductions at the lowest possible cost and yield the cost-effective geographic distribution of abatement. Despite the logic of this approach, international policymakers were unable to implement this vision and climate governance has taken a different path. Today twenty international, national, regional, state, provincial and municipal emissions trading systems are in operation, instead of the single international carbon market that was once imagined (World Bank and Ecofys 2016). Moreover, domestic policymakers seem to have followed suit; for example, the US Environmental Protection Agency's Clean Power Plan delegates a host of decisions to states, including whether to start an emissions trading market and, if so, whether to link to other states. This fragmentation leaves important opportunities for improved cost-effectiveness on the table and coordination could enable greater environmental stringency at lower total costs. A central way to improve the cost-effectiveness of this patchwork is to aggregate through bilateral or multilateral linking, a process in which the regulatory authorities in each system mutually allow their regulated firms to use emissions allowances from any of the linked jurisdictions in order to meet compliance obligations (Jaffe, Ranson, and Stavins 2009).[1]

A wealth of qualitative literature describes the potential advantages of linking in economic terms. For example, in principle, bilateral or multilateral linking achieves a unified

* Burtraw is the Darius Gaskins Senior Fellow, Palmer is senior fellow and research director, Munnings is a research associate at Resources for the Future. Woerman is a graduate student at the University of California, Berkeley. Early work on this project was supported by the Energy Foundation, the Merck Family Fund, and Mistra's INDIGO Project. The authors greatly benefited from collaboration with Paige Weber and Lars Zetterberg at the early stages of this project. Anthony Paul provided valuable assistance with the modeling. All errors and opinions remain the responsibility of the authors.

[1] We primarily focus on bilateral links, although a variety of other linking types exist, including incremental alignment of carbon policies, which Burtraw et al. (2013) refer to as "linking by degrees", unilateral linking, and various forms of restricted links (Mehling and Hates 2009; Lazarus et al. 2015).

carbon price across the newly linked system that is expected to lower overall abatement costs. The potential efficiency gains are greater the greater are differences in pre-linked allowance prices (Flachsland, Marschinski, and Edenhofer 2009). Linking also can dampen allowance price volatility caused by regional variations in the demand or supply of allowances because typically the factors that influence emissions such as weather or economic activity are imperfectly correlated across jurisdictions (Burtraw et al. 2013). In some circumstances, linking can ameliorate concerns over competitiveness impacts by explicitly addressing the possibility for leakage of economic activity between jurisdictions that may result from differences in program stringency (Jaffe, Ranson, and Stavins 2009). Moreover, there are other potentially significant benefits to linking that are not economic in nature. From an environmental perspective, the reduction in abatement costs achieved by linking could make it easier to enhance ambition (Bodansky et al. 2015). From a political perspective, linking starts to dispel the free-rider narrative that can prevent individual jurisdictions from pricing carbon in the absence of an international carbon price (Flachsland, Marschinski, and Edenhofer 2009).

There is also a significant qualitative literature that outlines the potential costs of linking. First and foremost, established links between trading systems have required significant negotiations between jurisdictions in order to harmonize the design of the systems; the time and resources spent on this process of harmonization can be thought of as a fixed cost of linking. In addition, the efficiency gains achieved by linking may come with associated costs. For example, linking requires ceding some control over domestic allowance prices, which might be regarded as a political cost (Ranson and Stavins 2016), or a virtue when it insulates policymakers from narrow interest groups within their jurisdiction (Burtraw et al. 2013). While linking may reduce overall abatement costs, it may have negative economic impacts on particular actors in each jurisdiction (Newell, Pizer, and Raimi 2013).[2] Moreover, linking can exacerbate allowance price volatility in certain cases (Doda and Taschini 2016). From an environmental perspective, linking could increase emissions leakage if allowance prices increase in the system that is more susceptible to leakage (Jaffe, Ranson, and Stavins 2009) and may alter incentives for cap setting, encouraging systems to set lower caps to achieve lower prices and therefore export more allowances, thereby resulting in higher emissions than would occur without linking (Bohm 1992, Helm 2003). Linking also might provide an incentive to introduce companion policies, such as

[2] In jurisdictions where allowance prices increase due to linking, compliance entities or consumers who purchase goods from these entities will experience greater costs. Conversely, in jurisdictions where allowance prices decrease due to linking, any agent holding excess permits will experience a reduction in the value of these assets and governments will receive less revenue from allowance auctions.

technology support policies, that reduce local demand for allowances, in order to increases allowance exports and associated government revenues.

Weighing the advantages and disadvantages of a specific link requires an accounting of the unique designs of each of the involved trading systems and how they would interact under a particular linking architecture. Quantitative approaches are useful in this regard. One vein of the quantitative literature on linking utilizes models to provide estimates of the efficiency gains achieved by linking (a selection of which are reviewed by Springer 2003, or the emissions outcomes of different coalitions of linked trading systems (e.g., Paltsev 2001). A second vein of the quantitative literature on linking takes an analytical approach to investigate the impact of different linking architectures (e.g., a link between mass and rate-based trading programs or a restricted one-way link that discounts incoming allowances (Fischer 2003, Lazarus et al. 2015)), or the impacts of unique program design features (e.g., market size) on the economic implications of linking (Doda and Taschini 2016).

Jurisdictions considering a potential link have some control over the domestic costs and benefits of the link through the use of an allowance exchange rate, which denominates the value of an emissions allowance (i.e., the quantity of emissions per allowance) differently in each system. That is, an exchange rate mandates that an allowance from one system is worth more or less in terms of compliance (allowable tons per allowance) than is an allowance from another system. While economists typically discuss exchange rates in the context of pollutants that impose local damages that vary by the source of emissions (Hung and Shaw 2005), the interest in applying exchange rates in the context of greenhouse gas emissions has increased recently (Fischer 2003, Metcalf and Weisbach 2012, Holland and Yates 2015). Greenhouse gas allowance exchange rates have also been included in recent policy discussion, including efforts by the World Bank's Networked Carbon Market Initiative[3] (Macinante 2016) and China's stated intentions to discount allowances from regional emissions markets when its national trading system launches (Carbon Pulse 2015).

Both the qualitative and quantitative veins of the literature are useful in characterizing the theoretical benefits and costs of linking but tend to assume that trading systems are nearly identical in design. In reality, however, the array of existing systems exhibits various designs and stringencies. We complement the existing literature by evaluating the linking of systems that

[3] The World Bank's Networked Carbon Market Initiative is focused on facilitating cross-border allowance trades based on a shared understanding of the relative value of different actions, instead of "harmonzing" climate actions so that units can be traded on a one-to-one basis.

have various and different designs (i.e., explicitly different price floors and ceilings, allocation methods, leakage policies and cap ambition and implicitly different offset and companion policies) and considering how different design parameters interact with alternative architectures for linking (e.g., different exchange rates for allowances). In particular, we make two primary contributions with this work. First, we develop an analytical model that formalizes the economic implications and emission market outcomes of linking, both with and without an exchange rate. This model yields novel propositions on the results of linking emission markets, as well as the formalization of results that had previously been described only qualitatively. Second, we test several of these propositions and illustrate other important market outcomes of linking by simulating a link between the California and Regional Greenhouse Gas Initiative (RGGI) trading programs. We use a simulation model of regional electricity markets within the US in order to characterize the particular design features of California's and RGGI's programs, accounting for how they interact with their respective regional electricity markets. We simulate the trading programs under autarky (when they are independent) and under various exchange rates. The electricity market model allows us to consider a wide range of economic implications and emissions outcomes that can arise from linking without losing the detailed designs of the two emissions markets as well as the nuanced and important interactions that might occur between them when linked.

2. An Analytical Model of Linking

The model considers a regional power sector with electricity supplied by a representative electric utility.[4] We first show how this representative utility responds when faced with a CO_2 policy that imposes a price on CO_2 emissions. We next describe the equilibrium outcomes of an emissions trading market in autarky. We then show how the outcomes change when two emissions markets link through the trade of allowances.

This model describes the linking of two regional trading systems, each of which covers the electric power sector in its region. This is only an illustrative example of linking, however, and the insights of this theoretical model apply more broadly to the linking of any two emissions trading systems, such as the linking of different sectors or the linking of two federal economy-

[4] An alternative model of electricity generation has independent power producers (IPPs) coordinated by an independent system operator (ISO) in a wholesale electricity market, rather than a vertically-integrated utility serving both of these roles. The representative utility model yields the same results as this alternative model when the IPPs reveal the true merit order to the ISO, so the ISO can schedule the lowest-cost dispatch of generators. We model the representative utility for simplicity.

wide trading systems, that gives emitters a marginal incentive to equate the marginal abatement cost with the marginal cost of emissions. In discussion we explain that the model also applies to the linking of a broader set of carbon pricing policies, such as a carbon tax, which can be interpreted as an emissions trading system with a price floor that is coincident with a price ceiling.

2.1. Electricity Generation

The power sector in a region is characterized by a representative utility that owns a portfolio of power plants and dispatches these plants to meet electricity demand at lowest cost. The cost to the representative utility of generating electricity is given by $C(Q, E)$ and is a function of the quantity of electricity generated, Q, and CO_2 emitted, E. Generating electricity from low-emitting sources is more costly than from high-emitting sources, so $\frac{\partial C(Q,E)}{\partial E} < 0$ over the relevant range of Q and E considered in this model. The utility is subject to a CO_2 trading program that imposes an opportunity cost of p on each unit emitted. The utility is also required to meet the demand for electricity, which is assumed to be fixed at $\bar{Q}$.[5] The utility selects the level of emissions that minimizes its total cost while meeting its required level of generation:

$$\min_{E} C(\bar{Q}, E) + pE$$

This optimization problem yields the first-order condition:

$$-\frac{\partial C(\bar{Q}, E^*)}{\partial E} = p$$

This is the familiar result that the representative utility's optimal level of emissions, E^*, equates its marginal abatement cost, $-\frac{\partial C(\bar{Q},E^*)}{\partial E}$, to the marginal cost of emissions, p.

2.2. Emissions Trading in Autarky

We now consider the specific design of the emissions trading system and the resulting outcomes—allowance prices and emissions—that occur in this market in autarky. Although an emissions trading policy has many design parameters through which the system can be adjusted, this analytical model focuses on two, and arguably the most important, of these policy parameters the level of the cap and the cost containment mechanism.

[5] Our simulation results show small changes in the retail price of electricity and the quantity of electricity consumed, which is roughly consistent with this assumption that a fixed quantity of electricity is demanded.

The intended emission cap initially distributes $\bar{A}$ allowances, each of which authorizes the holder to emit one unit of CO_2. These allowances are auctioned in a multi-unit, uniform-price auction. [6] This auction has a reserve price of p^F, known as a price floor, below which no allowances will be auctioned. Additional allowances beyond the intended cap are also available for purchase at a price of p^C, which is known as the price ceiling. These parameters specify the supply of allowances in the market. This supply correspondence is given by:

$$S(p) = \begin{cases} [0, \bar{A}] & \text{if} \quad p = p^F \\ \bar{A} & \text{if} \quad p \in (p^F, p^C) \\ [\bar{A}, \infty) & \text{if} \quad p = p^C \end{cases}$$

In words, up to $\bar{A}$ allowances are supplied at the auction reserve price of p^F; exactly $\bar{A}$ allowances are supplied if the market-clearing price is between p^F and p^C; and an unlimited number of allowances are available at a price of p^C. The presence of the auction reserve price and the availability of additional allowances constrain the resulting market prices and are collectively known as a price collar. To model a policy that does not include such a cost containment mechanism, the price floor and ceiling can be set at zero[7] and an arbitrarily large number, respectively. To model a CO_2 tax, the price floor and ceiling can be set at the same level, equal to the tax.

Emitters of CO_2 comply with the trading program by obtaining one allowance for each unit of CO_2 emitted. The first-order condition from the utility's cost minimization yields the equilibrium relationship between CO_2 emissions and the opportunity cost of each unit of emissions. Because one allowance is required for each unit emitted, inverting this relationship gives the demand for allowances as a function of the allowance price, $D(p)$, with $\frac{dD(p)}{dp} < 0$ over the relevant range of prices and allowances considered here.

The emissions market clears at the allowance price that equates the supply and demand of allowances and the total level of CO_2 emissions is equal to the market-clearing quantity of

---

[6] In many emissions trading systems, some allowances are freely allocated to emitters or other agents for political or economic reasons, such as building political support for the trading system or to compensate firms for their cost of compliance, which can have important implications for firm entry and exit and emissions leakage. In this analytical model, however, freely allocated allowances will affect market outcomes only if the market is sufficiently oversupplied through free allocation and no allowances are purchased in the auction, which will yield an allowance price below the auction reserve price. This is an extreme case that has not been observed in any allowance markets to date, although bilateral (spot) market prices have been observed to fall below auction reserve prices during periods between auctions.

[7] We assume free disposal of allowances, so the allowance price will never be negative. In the absence of an explicit price floor, zero is the effective price floor.

allowances. The allowance price, p^A, and level of emissions, E^A, that result from this market in autarky are:

$$(p^A, E^A) = \begin{cases} (p^F, D(p^F)) & \text{if} \quad p = p^F \\ (D^{-1}(\bar{A}), \bar{A}) & \text{if} \quad p \in (p^F, p^C) \\ (p^C, D(p^C)) & \text{if} \quad p = p^C \end{cases}$$

2.3. One-for-One Linking

We now consider two independent emissions trading systems, denoted by subscripts i and j, that link through the trade of emissions allowances. All characteristics of the representative utility—such as the cost function and the required quantity of generation—and characteristics of the policy—such as allowances issued and price collars—can vary across the different systems. Emitters in each system can comply with the emissions trading system by holding allowances issued by either system, and the linking is at a "one-for-one" rate, meaning one allowance covers one unit of emissions in either system.

This linking of the trading systems has important implications for the resulting allowance prices, emissions, and costs of compliance. We assume there is no arbitrage across the systems, so all allowances have the same market price, p, regardless of which system issued them; that is, $p = p_i = p_j$.[8] We first show how this linked market and equilibrated allowance price affects the effective price collar faced by the combined market.9 (All proofs appear in the appendix.)

Proposition 1. With one-for-one linking, the price collar of the linked market is given by:

i. The linked price ceiling is the minimum of the two price ceilings in autarky, $p^C = \min\{p_i^C, p_j^C\}$.

ii. If allowances from both systems are used, then the linked price floor is the maximum of the two price floors in autarky, $p^F = max\{p_i^F, p_j^F\}$. If only allowances from the system with a lower price floor are used, then the price falls below that linked price floor.

---

[8] If this were not the case, then any emitter holding the higher-priced allowance could arbitrage the allowance price difference by selling the higher-priced allowance and buying a lower-price allowance.

[9] We assume there is overlap in the two independent price collars; that is, $\min\{p_i^C, p_j^C\} > \max\{p_i^F, p_j^F\}$. If this is not true, it is unlikely the two systems would link due to differences in what allowance prices would be politically acceptable in each of these systems.

In other words, except in the case of an extreme oversupply of allowances, the effective price collar on the linked market will be the "tightest" combination of the individual price floors and price ceilings and the observed allowance price will be in the intersection of the sets of possible prices that would be observed in the systems under autarky.

We next show how the equilibrium allowance price in the linked market relates to the allowance prices in autarky. For this and the remainder of this section, suppose that the allowance price under autarky is weakly lower in system i than in system j, $p_i^A \leq p_j^A$.[10]

Proposition 2. With one-for-one linking, the allowance price is (weakly) between the allowance prices in autarky, $p_i^A \leq p \leq p_j^A$.

This change in allowance prices resulting from linking implies how emissions in each system are affected.

Corollary 2.1. With one-for-one linking, in the system with the lower allowance price in autarky, emissions are (weakly) less than in autarky ($E_i \leq E_i^A$); similarly, in the system with the higher allowance price in autarky, emissions are (weakly) greater than under autarky ($E_j \geq E_j^A$).

In other words, if allowance prices in the two markets differ in autarky, the prices in a linked system will converge to a single allowance price that is between (or equal to one of) the two allowance prices in autarky. As a result, emissions will shift from the system with the lower autarkic price to the system with the higher autarkic price, when comparing emissions under trading to those under autarky. However, the difference in emissions in the two jurisdictions may not perfectly offset and aggregate emissions when linked may differ from that in autarky. For simplicity, we first consider the linking of two emissions trading systems without price collars and show that linking causes allowances prices and emissions to move in a predictable way.

Proposition 3. With no price collars, one-for-one linking yields greater aggregate emissions than autarky ($E_i + E_j > E_i^A + E_j^A$) if and only if exactly one of the systems is non-binding in autarky ($p_i^A = 0$). Otherwise, aggregate emissions when linked are equal to total emissions in autarky ($E_i + E_j = E_i^A + E_j^A$). In either case, aggregate emissions are never greater than the sum of the emissions caps ($E_i + E_j \leq \overline{A}_i + \overline{A}_j$).

[10] If this is not the case, system i can be re-indexed as system j and vice-versa, yielding $p_i^A \leq p_j^A$.

When the trading systems have price collars, however, it is difficult to make general statements about allowance prices and emissions.[11] We can, however, make some general statements about aggregate emissions when both systems are binding in autarky.

Proposition 4. When both systems are binding in autarky, the effect of one-for-one linking is:

i. If both allowance prices in autarky are inside the linked price collar ($p^F < p_i^A \leq p_j^A < p^C$), then aggregate emissions when linked are equal to the sum of the emissions caps ($E_i + E_j = E_i^A + E_j^A = \overline{A}_i + \overline{A}_j$).

ii. If one allowance price is at or below the linked price floor and the other is inside the linked price collar ($p_i^A \leq p^F < p_j^A < p^C$), then aggregate emissions when linked are (weakly) less than the sum of the emissions caps ($E_i + E_j \leq E_i^A + E_j^A = \overline{A}_i + \overline{A}_j$).

iii. If one allowance price is at or above the linked price ceiling and the other is inside the linked price collar ($p^F < p_i^A < p^C \leq p_j^A$), then aggregate emissions when linked are (weakly) greater than the sum of the emissions caps ($E_i + E_j \geq E_i^A + E_j^A = \overline{A}_i + \overline{A}_j$).

In addition to potentially affecting aggregate emissions, the linking of trading systems also has an important implication for the efficiency of the emissions reductions. This effect is easiest to see when comparing the cost of a linked system with that of systems in autarky that yield the same aggregate emissions, so the change in emissions does not confound the change in costs.

Proposition 5. One-for-one linking yields (weakly) lower total costs than the combined costs of the systems in autarky with the same aggregate emissions ($C_i + C_j \leq C_i^A + C_j^A$). However, the system with the lower allowance price in autarky has (weakly) higher costs when linked ($C_i \geq C_i^A$), and the system with the higher allowance price in autarky has (weakly) lower costs when linked ($C_j \leq C_j^A$).

2.4. Linking with an Exchange Rate

We finally consider two systems that link with an allowance exchange rate other than one-for-one. This exchange rate, which we denote r, is the number of allowances from system j that are equivalent for compliance purposes to one allowance from system i. In other words, for

[11] For example, if one system is at its price floor and the other at its price ceiling in autarky, the linked allowance price could fall anywhere in between, yielding aggregate emissions in the linked system that may be less than, greater than, or equal to aggregate emissions in autarky, which itself may be less than, greater than, or equal to the sum of the two emissions caps.

each unit of CO_2 emitted by the utility in system i, it must have either one allowance from system i or r allowances from system j. Similarly, for each unit of CO_2 emitted by the utility in system j, it must have either one allowance from system j or $1/r$ allowances from system i.

We showed above that linking at a one-for-one rate has implications for the resulting allowance prices, emissions, and costs of compliance. When linking at a rate other than one-for-one, the exchange rate used can importantly affect these outcomes. We again assume there is no arbitrage across the systems; this means the price of an allowance from system i is r times the price of an allowance from system j, $p_i = rp_j$.[12] We first show how this linked market affects the effective price collar faced by the combined market.[13] We express the price collar in terms of system i allowances; the price collar for system j allowances is $1/r$ times these prices. Similarly, quantities of allowances from system j are converted to the equivalent number of allowances from system i. We also introduce the term "exchange-adjusted prices" to refer to the set of prices relevant to a particular system once the exchange rate is taken into account. For example, the relevant prices in system i are p_i and rp_j, while the relevant prices in system j are $\frac{1}{r}p_i$ and p_j.

Proposition 6. With linking at an exchange rate of r, the price collar (in system i) is given by:

i. The linked price ceiling is the minimum of the two exchange-adjusted price ceilings in autarky, $p^C = \min\{p_i^C, rp_j^C\}$.

ii. If allowances from both systems are used, then the linked price floor is the maximum of the two exchange-adjusted price floors in autarky, $p^F = \max\{p_i^F, rp_j^F\}$; if only allowances from the system with a lower price floor are used, then the price falls below the linked price floor.

As a result, we show that the choice of the exchange rate affects which price ceiling and floor will make up the effective cost containment mechanism of the linked system.

Corollary 6.1. When linking with an exchange rate, which price ceiling and price floor make up the linked system's price collar depends on the choice of the exchange rate, r, relative to the ratio of price ceilings ($\frac{p_i^C}{p_j^C}$) and price floors ($\frac{p_i^F}{p_j^F}$), respectively:

---

[12] If this were not the case, then any emitter holding the higher-valued allowance could arbitrage the allowance price difference by selling the higher-valued allowance and buying the comparable number of lower-valued allowances.

[13] We assume there is overlap in the two independent price collars when the exchange rate is taken into consideration; that is, $\min\{p_i^C, rp_j^C\} > \max\{p_i^F, rp_j^F\}$. If this is not true, it is unlikely the two systems would link with an exchange rate of r due to differences in what allowance prices would be politically acceptable in each of these systems.

i. If the exchange rate is (weakly) greater than the ratio of price ceilings ($r \geq \frac{p_i^C}{p_j^C}$), the exchange-adjusted price ceiling of system i is the linked price ceiling ($p^C = p_i^C$). If the exchange rate is (weakly) less than the ratio of price ceilings ($r \leq \frac{p_i^C}{p_j^C}$), the exchange-adjusted price ceiling of system j is the linked price ceiling ($p^C = rp_j^C$).

ii. If the exchange rate is (weakly) less than the ratio of price floors ($r \leq \frac{p_i^F}{p_j^F}$), the exchange-adjusted price floor of system i is the linked price floor ($p^F = p_i^F$). If the exchange rate is (weakly) greater than the ratio of price ceilings ($r \geq \frac{p_i^F}{p_j^F}$), the exchange-adjusted price floor of system j is the linked price floor ($p^F = rp_j^F$).

This differs from one-for-one linking because it is not the "tightest" price collar in absolute terms that binds as we observed in Proposition 1, but rather it is the "tightest" exchange-adjusted price collar that binds. As a result, the higher of the two price ceilings may be the effective linked price ceiling if the exchange rate is set at a sufficient level, and similarly for the lower of the two price floors. We next show how the equilibrium allowance prices in the linked market relate to the allowance prices in autarky and to the choice of an exchange rate.

Proposition 7. With linking at an exchange rate of r, the allowance price is (weakly) between the exchange-adjusted allowance prices in autarky.

This change in allowance prices from autarky implies how linking affects emissions in each system.

Corollary 7.1. With linking at an exchange rate of r, in the system with the lower exchange-adjusted allowance price in autarky, emissions are (weakly) less than in autarky; similarly, in the system with the higher exchange-adjusted allowance price in autarky, emissions are (weakly) greater than under autarky.

We further show that the choice of the exchange rate, r, determines how prices and emissions are affected under trading as compared to under autarky.

Corollary 7.2. When linking with an exchange rate, the relationship between trade and autarky for both allowance prices and emissions in each system depends on the choice of the exchange rate, r, relative to the ratio of allowance prices under autarky, $\frac{p_i^A}{p_j^A}$.

i. If the exchange rate equals the ratio ($r = \frac{p_i^A}{p_j^A}$), allowance prices and emissions are equal under trade and autarky for each system ($p_i = p_i^A$, $E_i = E_i^A$, $p_j = p_j^A$, and $E_j = E_j^A$).

ii. If the exchange rate is greater than the ratio ($r > \frac{p_i^A}{p_j^A}$), the allowance price is (weakly) greater and emissions are (weakly) less than under autarky in system i ($p_i \geq p_i^A$ and $E_i \leq E_i^A$), and the allowance price is (weakly) less and emissions are (weakly) greater than under autarky in system j ($p_j \leq p_j^A$ and $E_j \geq E_j^A$).

iii. If the exchange rate is less than the ratio ($r < \frac{p_i^A}{p_j^A}$), the allowance price is (weakly) less and emissions are (weakly) greater than under autarky in system i ($p_i \leq p_i^A$ and $E_i \geq E_i^A$), and the allowance price is (weakly) greater and emissions are (weakly) less than under autarky in system j ($p_j \geq p_j^A$ and $E_j \leq E_j^A$).

Again this outcome differs from one-for-one linking because the linked allowance price is not necessarily between the allowance prices in autarky, but rather between the exchange-adjusted allowance prices. As a result, if the exchange rate is set at a sufficient level, the system with the greater allowance price under autarky may have an even greater price with trading, which would lead to fewer emissions in that system under trading than under autarky. In other words, the choice of the exchange rate affects which system is a net importer or exporter of allowances and emissions.

As with one-for-one trading, trade of allowances may not perfectly offset, however, and aggregate emissions when linked may differ from that in autarky. This outcome is particularly likely when trading with an exchange rate because one allowance covers a different amount of emissions in each system. For simplicity, consider the linking of two trading systems for which price collars do not bind, which includes the case of systems without price collars. [14]

Proposition 8. When linking at an exchange rate and when price collars do not bind ($p_i^F < p_i^A < p_i^C$, $p_j^F < p_j^A < p_j^C$, $p^F < p_i^A < p^C$, and $\frac{1}{r}p^F < p_j^A < \frac{1}{r}p^C$), aggregate emissions depend on the choice of the exchange rate, r:

i. Aggregate emissions are less than under autarky if and only if r is between $\frac{p_i^A}{p_j^A}$ and 1 ($\frac{p_i^A}{p_j^A} < r < 1$ if $p_i^A < p_j^A$; $1 < r < \frac{p_i^A}{p_j^A}$ if $p_i^A > p_j^A$).

ii. Aggregate emissions are equal to that under autarky if and only if $r = \frac{p_i^A}{p_j^A}$ or $r = 1$.

[14] Lazarus et al. (2015) discuss the impact of varied exchange rates of economic effectiveness and emissions outcomes. We formalize and generalize this discussion through propositions 8, 9, and 10 by considering general rather than linear marginal abatement cost curves.

iii. Aggregate emissions are greater than under autarky if and only if r is outside the range defined by $\frac{p_i^A}{p_j^A}$ and 1 ($r < \frac{p_i^A}{p_j^A}$ or $r > 1$ if $p_i^A \leq p_j^A$; $r < 1$ or $r > \frac{p_i^A}{p_j^A}$ if $p_i^A \geq p_j^A$).

In addition to potentially affecting aggregate emissions, the linking of trading systems with an exchange rate also has an important implication for the costs incurred by the utilities in each system.

Proposition 9. When linking at an exchange rate of r, the system with the lower exchange-adjusted allowance price in autarky has (weakly) higher costs when linked, and the system with the higher exchange-adjusted allowance price in autarky has (weakly) lower costs when linked.

As a result, the choice of the exchange rate, r, determines which system will incur greater total costs, and which will incur less total costs, under trading as compared to autarky.

Corollary 9.1. When linking with an exchange rate, the cost incurred by each system, as compared to autarky, depends on the choice of the exchange rate, r, relative to the ratio of allowance prices under autarky, $\frac{p_i^A}{p_j^A}$.

i. If the exchange rate is (weakly) greater than the ratio ($r \geq \frac{p_i^A}{p_j^A}$), the cost incurred by system i is (weakly) greater than under autarky ($C_i(\overline{Q_i}, E_i) \geq C_i\left(\overline{Q_i}, E_i^A\right)$) and the cost incurred by system j is (weakly) less than under autarky ($C_j\left(\overline{Q_j}, E_j\right) \leq C_j\left(\overline{Q_j}, E_j^A\right)$).

ii. If the exchange rate is (weakly) less than the ratio ($r \leq \frac{p_i^A}{p_j^A}$), the cost incurred by system i is (weakly) less than under autarky ($C_i(\overline{Q_i}, E_i) \leq C_i\left(\overline{Q_i}, E_i^A\right)$) and the cost incurred by system j is (weakly) greater than under autarky ($C_j\left(\overline{Q_j}, E_j\right) \geq C_j\left(\overline{Q_j}, E_j^A\right)$).

Finally, when linking with an exchange rate, the choice of the exchange rate can affect the efficiency of the emissions reductions. This effect is easiest to see when comparing the costs of two different linked systems, each with a different exchange rate, but each yielding the same level of aggregate emissions, and this level is less than aggregate emissions under autarky.

Proposition 10. When linking with an exchange rate, if two exchange rates yield the same level of emissions, and this level is less than aggregate emissions under autarky, then the exchange rate closer to 1 has a lower total cost.

3. Evaluating Readiness to Link—The Case of California and the Regional Greenhouse Gas Initiative

In this section, we evaluate the readiness of the California and RGGI emissions markets for linking and set the stage for an application of the analytical model.[15] Many studies point to the significant obstacles in linking two trading systems that are designed separately and the potential costs of linking without close harmonization of specific design features (Haites and Wang 2009, Zyla 2010). Burtraw et al. (2013) conduct an extensive evaluation of the design features of the California and RGGI systems, finding that the designs are already quite closely harmonized because of a long history of cooperation, information sharing, mutual learning, and replication of each other's designs. The authors conclude—based on criteria including the degree to which design features are aligned, and whether any misalignments of design features would be important for the functioning of a new enlarged allowance market or for political reasons stemming from economic or environmental preferences—that the California and RGGI trading systems are nearly ready to link. The discussion below focuses on four design features identified by Burtraw et al. (2013). We find that misalignments regarding these four features are either tolerable or relatively straightforward to align, and therefore that the programs could link quite easily in the future.

3.1. Comparability of the Emissions Cap

The main determinant of the stringency of the program and of allowance prices is the choice of how many allowances to issue (the emissions cap). RGGI allowances are denoted in short tons while California allowances are denoted in metric tons. This distinction is not a barrier to linking the markets but it implies that, unless one of the programs changes its unit of measurement, linking will require a conversion factor between the programs to achieve equivalent tons. Current allowance prices vary widely between RGGI (near $5 per short ton) and California (about $13 per metric ton), indicating large potential gains in the efficiency of overall emissions reductions. We consider all calculations in this paper in short tons.

Linking also implies flows of allowances between regions. Generally, the region with lower stringency is expected to export allowances, and therefore import revenue from the region with higher stringency (Corollary 2.1). Such transfers potentially present a political challenge but do not present a challenge to the functioning of the enlarged allowance market. Nonetheless, the

[15] This task is made somewhat easier because the RGGI program is limited to the electricity sector, and many design features of both programs are similar in this sector.

choice of stringency in a linked market must be balanced against the distributional outcome across and within systems.

One possible solution for regulators sensitive to revenue transfers would be the implementation of an exchange rate in order to control allowance flows. For example, an exchange rate might specify that three RGGI allowances are equivalent to one California allowance—meaning a California entity could retire one California allowance or three RGGI allowances for one ton of emissions, and a RGGI entity could retire one California allowance or three RGGI allowances for three tons of emissions. This exchange rate is similar to the ratio of prices observed in the two programs. In principle, an exchange rate would allow a jurisdiction to balance the cost savings achieved by linking with political preferences, such as more localized control over allowances prices and wealth transfers.[16] However, as demonstrated in the analytical discussion in section 2 above and in the simulation model in section 4 below, the use of exchange rates would introduce uncertainty regarding overall emissions. To address this uncertainty the trading program might employ numerous other mechanisms to control allowance flows such as import quotas, unilateral linking, discount rates, and fees imposed on using allowances from other programs for compliance (Mehling and Haites 2009) (Lazarus et al. 2015). We therefore argue there are a variety of tools available that enable jurisdictions to control allowance flows and consequently a large different in allowance prices pre-link need not be an insurmountable barrier to linking.

3.2. Offsets

If linked jurisdictions have different restrictions placed on the use or eligibility of offset credits, the price of offset credits will be communicated between jurisdictions through the linked allowance market. This is described as the “free-up effect” and is expected to occur if offset rules are not aligned across jurisdictions (Sterk and Kruger 2009, p. 396) The free-up effect results in rules in one jurisdiction unilaterally increasing the supply of compliance instruments in the linked market; for example, if one program allowed the use of a particular type of offset while the other program intended to preclude its use. A jurisdiction may wish to preclude the use of specific offset types if it prefers a high carbon price or if it is risk averse with respect to the environmental integrity of the offset credits. These preferences are subverted if programs with varying standards are linked, leading some authors to identify misaligned offset rules as a key

[16] As programs evolve, the political acceptability of higher allowance prices may change, enabling an adjustment in the exchange rate through mandated, periodic reevaluation of the exchange rates, through automatic adjustments of the exchange rate via a pre-specified adjustment schedule, or indexed to an economic or environmental indicator.

barrier to linking (Tuerk et al. 2009, Sterk and Kruger 2009, Flachsland, Marschinski, and Edenhofer 2009). A discriminating program might impose import quotas, fees, or discount rates on offsets depending on their origin. This treatment would not solve the free-up effect, because the offsets would still be available in the other program, but it would ensure that they are not used for compliance in the discriminating program, which may help achieve political objectives. Because the free-up effect cannot be completely mitigated, regulators should place a high priority on aligning policies about offsets.

3.3. Price Collars

Marginal costs of compliance are determined primarily by the relative stringency (cap) of the individual programs and, as discussed above in section 2, price collars provide a method of managing costs when factors affecting the market are uncertain. However, different trigger prices for the floor and ceiling across linked systems could influence allowance flows and prices (Tuerk et al. 2009) and there also is a strong potential for differing floors to erode the environmental integrity of the linked programs as we discuss in section 2. If they are not aligned, linking could undermine the value of previous investments and thereby the confidence of investors going forward. Hence, the alignment of price floors and ceilings across programs poses a potential threat to the functioning of the market and is a focus of the modeling exercise in section 4.

One specific element of price collars poses a political and environmental challenge as well: whether additional allowances that might be available at a price ceiling come from "inside" or "outside" of the cap. In California, additional allowances come from under the cumulative cap through 2020. In RGGI allowances that are available at a price ceiling come from outside the cap. From a design standpoint, some advocates are likely to feel that environmental integrity, in the form of emissions reductions, can be guaranteed only if allowances come from under the cap (Harrison 2006).

3.4. Legal Contingencies

Provisions for changing the design of either program or for delinking are difficult to align and potentially important. Within RGGI, each state retains the ability to leave the program, leading to a strong emphasis on finding consensus on policy decisions (Pizer and Yates 2015). This process within RGGI places it on a different decision making schedule than that of California. Consequently, if formal linking were to occur, future changes to the combined program might be made unilaterally and on inconsistent time schedules.

The California Air Resources Board staff anticipates that if delinking were to occur, it would trigger a program review, as would be likely in RGGI as well. As predictable as the triggering of a review might be, the outcome is not. This element of uncertainty means compliance entities will recognize some risk associated with compliance instruments issued by the other jurisdiction. In particular, one is not likely to see banking of compliance instruments from the other jurisdiction. This failure to bank might imply a price difference in the market due to the different convenience yield that each instrument provides an investor, with some loss of market efficiency as a result, however, the technical issues associated with potential delinking are not likely to be fatal to the market. For example, on the date the decision to delink is announced, holdings of allowances from outside a given program are noted and those allowances assigned legitimacy for compliance (possibly within a limited period) or sold to the originating program (Haites and Wang 2009). This protocol was followed when New Jersey left RGGI in 2011; previously issued allowances from New Jersey that were banked were recognized as valid within RGGI. If that were not the protocol, one would not be likely to see banking of compliance instruments from the other jurisdiction. Alternatively, Newell et al. (2012) suggest a pegged currency system with separate currencies rather than a currency union. As long as linked trading systems maintain distinct units of account, which we interpret to include distinct registries, then they argue delinking should not be a problem.

4. Modeling Analysis of Linking the California and RGGI Markets

We now turn to our simulation of a link between the California and RGGI systems. We break this section into several parts: a short description of the model, a presentation of model results for both systems under autarky, and a discussion of our results under linked scenarios with different exchange ratios. We consider the effects of linking architecture (i.e., different exchange rates) and the unique designs of these two programs on several indicators of allowance market, electricity market, and emissions outcomes, as well as how different constituencies in the two regions are affected by the linking of the programs.

4.1. Model Description

We use the Haiku electricity market model to explore the implications of linking the California and RGGI trading systems. The model simulates investment and retirement decisions and system operation in 22 inter-connected regions spanning the continental United States over a

25-year horizon.[17] We focus on results for the year 2020, which gives program outcomes for a medium-run timeframe. Because these trading systems are represented within a national framework, changes in electricity generation and fuel use within these regions can have effects across the nation. However, because of geographic distance, there is no effective power flow between these regions so for general purposes it is sufficient to imagine that these electricity markets operate independently, except when we link their emissions trading systems. Our analysis focuses on the electricity sector, so we limit our modeling to the electricity portion of the California program; our model covers the continental US, so we assume no relationship between California and the Quebec or Ontario programs.[18]

4.2. Results for the Unlinked Programs

The California program, as modeled, results in emissions reductions from the electricity sector of roughly 10 percent below baseline levels in 2020 with an allowance price of $14.2 per ton,[19] about 12 percent above the price floor (in short tons). The program raises electricity prices by about 2 percent and lowers REC prices by about 16 percent compared with a baseline with no program.

In our model, the RGGI region has an allowance price in 2020 of $7.2 per ton. Emissions are 22 percent lower than a baseline with Phase 1 RGGI program specifications, under which allowances continue to be sold at the current price floor. The tightening of the RGGI cap results in only a minimal change in electricity price in the region, however, to the extent leakage occurs it will cause the overall emissions reduction to be less than we report for RGGI.

Figure 1 depicts estimated marginal abatement cost curves for the two regions and includes box points indicating the allowance price and level of reductions obtained in each of the unlinked programs in 2020 relative to the modeled baseline.[20]

---

[17] For more information about the RFF Haiku model, see Paul (Paul, Burtraw, and Palmer 2009).

[18] See the appendix for additional details on how we incorporate the California and RGGI programs into our simulation model.

[19] All prices are in 2009 dollars per short ton.

[20] For each region, the marginal abatement cost curve was constructed from the pairs of allowance price and level of reductions in that region obtained over several model scenarios, including ones not reported here. The depicted curve is the best linear fit of these price-reduction pairs; the flat portion of California's curve represents the price floor in that market. The resulting marginal abatement cost curve does not perfectly align with the results of each scenario, but rather it represents the average over all modeling scenarios. Changes in the costs of renewable technologies and expanded availability of natural gas has affected the marginal abatement costs significantly since the timeframe for this modeling excercise but the findings about linking remain fully relevant.

Figure 1. Marginal Abatement Cost Curves and Results for Unlinked Programs

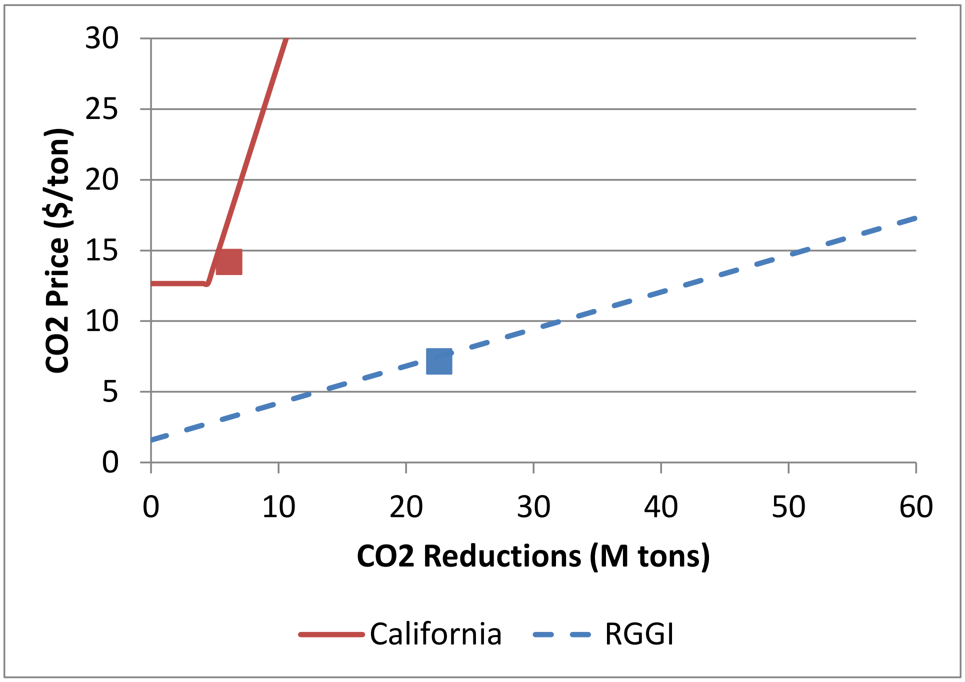

4.3. Results of One-for-One Linking

In order to explore exchange rates as a possible linking architecture, we simulate a scenario that allows for one-for-one trading of allowances between the two programs (an exchange rate of 1:1) and another that allows for three-for-one trading (an exchange rate of 3:1), which requires three RGGI allowances for each ton emitted in California and only one-third of a California allowance for each ton emitted in RGGI. The California allowance price is roughly double the RGGI allowance price in autarky, so these two scenarios span a range of exchange rates that are both below and above the ratio of prices under autarky. An allowance exchange rate can mediate the differences in marginal cost. However, as anticipated in section 2, the emissions outcome is not determined by the sum of the emissions caps in this context because an emissions allowance that is transferred between jurisdictions confers a license to emit different amounts in the two jurisdictions. In addition, although we do not explore this in the model, we expect the outcome also to be affected by the presence of companion policies that directly support technologies such as renewables or energy efficiency. The effect of these policies on the flow of allowance value between jurisdictions could introduce a strategic dimension to the policy architecture.

With one-for-one trading of allowances between the two programs, higher California allowance prices suggest that allowances would flow from RGGI to California.[21] The result would cause allowance prices to rise in RGGI and fall in California (Proposition 2). As allowance prices in RGGI rise as a result of linking, emissions in RGGI would be expected to exhibit a greater response than in California because the supply of emissions reductions is more elastic in RGGI than it is in California, as illustrated in Figure 2. The extent to which the equilibrium allowance price in California can fall as a result of imports from RGGI is constrained by the California price floor.[22] Linking the programs with one-for-one trading imposes the California price floor on both markets (Proposition 1). As a result, allowance prices in RGGI rise by nearly 80 percent, while allowance prices in California fall by only about 10 percent before they reach the floor.

Figure 2. Marginal Abatement Cost Curves and Results under One-for-One Trading

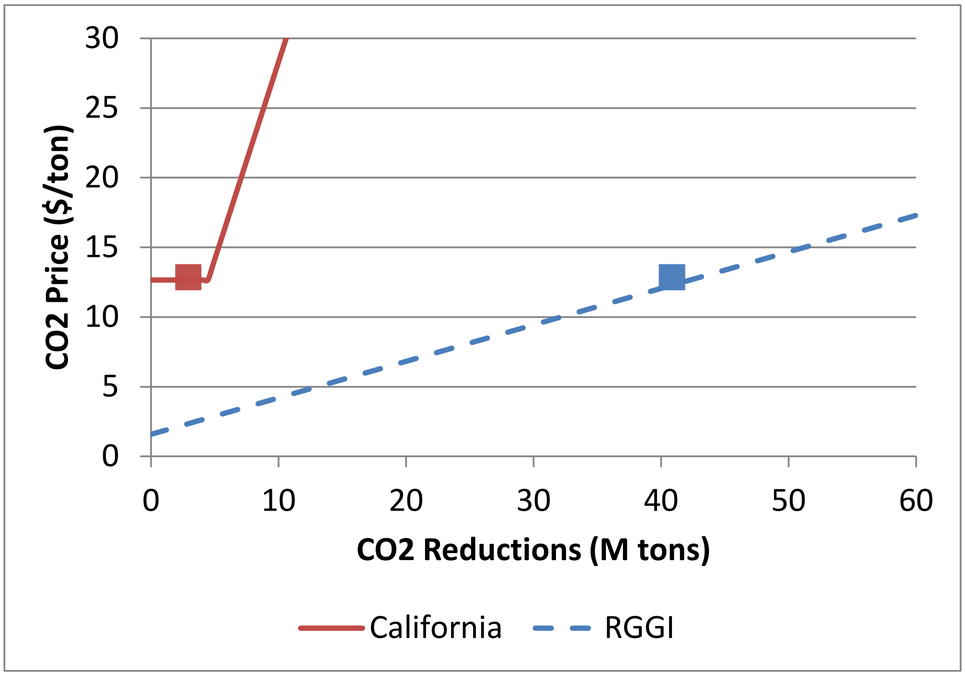

[21] One-for-one trading actually would involve a conversion factor as the programs are currently organized because a California allowance is denominated in metric tons and a RGGI allowance is denominated in short tons. One-for-one trading corresponds to equivalent tons.

[22] If RGGI supplied enough allowances to satisfy demand in both markets, the price would fall below the California price floor (Proposition 1), but otherwise the market price of allowances in RGGI will be bid up to the California floor as we find in our modeling.

One-for-one linking has three other important effects:

- *Emissions*. Linking shifts the location of CO_2 emissions from RGGI to California (Corrolary 2.1). Emissions from generators covered by the California program rise by 5 percent, while emissions in RGGI fall by 23 percent compared with emissions when unlinked. As a result of the price floor being spread across the two programs, total emissions from the two programs combined are lower than when they are not linked; combined emissions in the two regions are 26 percent below baseline levels with one-for-one linking, compared with 17 percent below baseline when the programs operate separately (Proposition 4).
- *Retail electricity prices*. Linking of the two programs has virtually no effect on electricity price in California because of the allocation of allowance revenues to local distribution companies.[23] The average electricity price in RGGI is roughly 1 percent higher in 2020 as a result of linking. In RGGI, most of the allowance revenues go to energy efficiency programs, which reduce electricity demand and price.[24]
- *Potential leakage*. As a result of the higher allowance prices in RGGI due to linking, power imports into the region increase by roughly 15 percent (the increase is equivalent to 5 percent of total consumption), suggesting that linking at one-for-one may contribute to emissions leakage in the RGGI region. Incentives for leakage in California would presumably be reduced because emissions prices fall with linking.

A comparison of total costs across scenarios in the model is not straightforward because the model has a detailed representation of regulatory structure. RGGI is modeled as a competitive power market. However, California resembles a cost-of-service territory with average cost pricing in the model, so changes in electricity price can have unintuitive outcomes on the cost measure (e.g., welfare can increase when electricity prices rise). Consequently, we focus on the distribution of costs within each system.

The distributional effects of linking clearly differ across geography and constituencies as displayed in Table 1. The effects of linking are reported in dollars per megawatt-hour (MWh),

[23] We assume this revenue is used to offset changes in electricity price. In practice, the majority of revenue associated with the auction of allowances to the electricity sector is returned as per customer account dividends received biannually, so customers see prices rise for five months before seeing a credit in the sixth month, thereby mostly preserving the perception that electricity prices are higher due to the program.

[24] The Haiku model has endogenous representation of the reduction in demand resulting from investments in energy efficiency. We adopt conservative assumptions about the effectiveness of those expenditures in reducing demand.

with positive values representing net benefits and negative values representing net costs, and are disaggregated into the effects on allowance value, resource cost, and electricity price. Table 1 shows that one-for-one trading leads to a small electricity price increase in RGGI, which hurts consumers. Because allowance prices rise in RGGI, the government collects more revenue from the allowance auction. This revenue is used to pay for energy efficiency and thus contributes to the low impact on electricity price. Fossil generators in RGGI benefit from the higher electricity price, but this benefit is outweighed by the combination of higher allowance costs and higher operating costs[25] (Proposition 5).[26]

In California, one-for-one trading has net positive effects for both consumers and fossil generators. Lower wholesale electricity prices affect consumers positively, but much of that effect is wiped out by the lower allowance revenues going to local distribution companies. Fossil producers are hurt by the lower electricity prices, but reductions in allowance costs and overall resource costs (Proposition 5) more than compensate.

[25] Resource cost/flow for fossil generators is the cost category that most closely matches the cost function of the analytic model, $C(Q,E)$.

[26] Note that the net effect indicated in the table is not strictly additive across interest groups, because the use of revenues to the government influences the outcome for consumers and generators.

Table 1. Incidence of Benefits and Costs of One-for-One Trading

$/MWh	RGGI			California		
	Consumers	Government	Fossil Generators	Consumers	Government	Fossil Generators
Allowance Value		1.4	-2.8	-1.1		0.6
Resource Cost/Flow			-6.9			3.8
Electricity Price	-1.0		1.0	1.2		-1.2
Net Effect	**-1.0**	**1.4**	**-8.6**	**0.1**	**n/a**	**3.2**

4.4. Results for Three-for-One Trading

Three-for-one trading provides a rough adjustment for the relative stringencies of the two programs but reduces the opportunities for costs savings from shifting CO_2 emissions from RGGI to California because of the requirement that three RGGI allowances be surrendered for every ton of emissions from sources regulated under the California program. Conversely, regulated sources in RGGI require only one-third of a California allowance to cover one ton of emissions in RGGI. The trading ratio also means that the effective minimum price in RGGI is by construction one-third of the price floor in California. This trading ratio lowers demand for RGGI allowances in California and increases demand for California allowances in RGGI, compared to one-for-one linking. Figure 3 shows the resulting allowance prices and emissions reductions from three-for-one trading. Note that the RGGI allowance price is read off the right-hand axis, which is one-third of the California allowance price on the left-hand axis. The box point is away from the line because it indicates the outcome from the specific modeling scenario and the line is the linear prediction over this range. As anticipated by Proposition 7, the resulting allowance price in RGGI of $5/*ton* is between the exchange-adjusted prices in autarky, where the price for compliance in RGGI using a RGGI allowance is $7.2/*ton* and the exchange-adjusted price using a California allowance is $4.73/*ton*.[27]

[27] Similarly, the resulting allowance price in California of $15/*ton* is between the exchange-adjusted prices in autarky, where the price for compliance in California using a California allowance is $14.2/*ton* and the exchange-adjusted price of using a RGGI allowance is $21.6/*ton*.

Figure 3. Marginal Abatement Cost Curves and Results of Three-for-One Trading

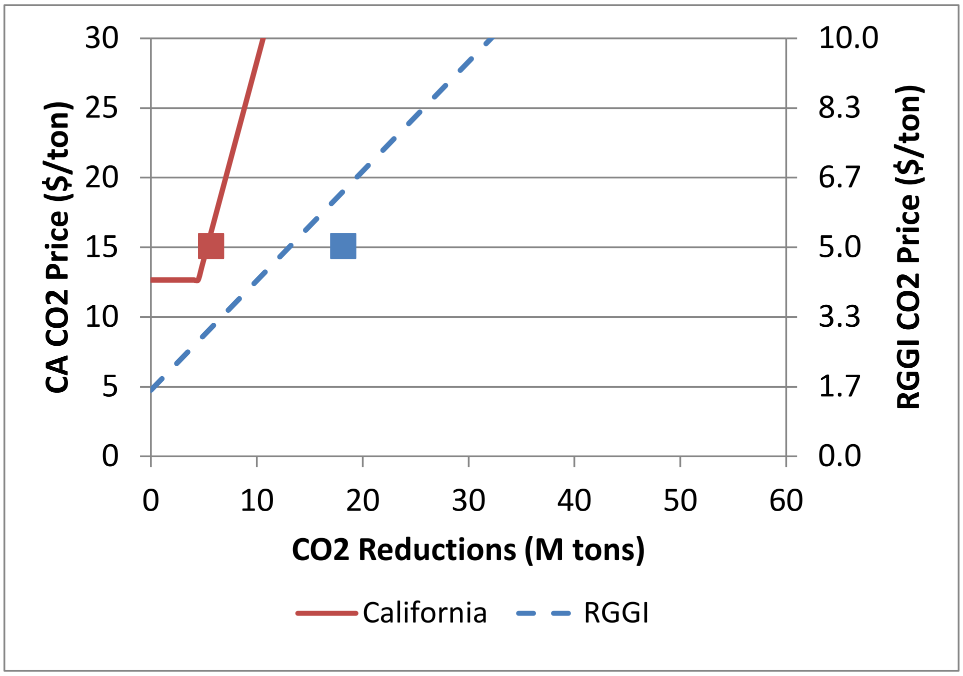

Linking at three-for-one compared with an unlinked regime has several other consequences:

- *Emissions*. This program leads to a 6 percent increase in emissions in RGGI relative to the unlinked case and only a small change in emissions in California (Corollaries 7.1 and 7.2). As a result, total emissions of CO_2 in the two regions increase to 14 percent below baseline levels, compared with 17 percent below baseline levels when the programs are unlinked (Proposition 8).
- *Retail electricity prices*. With three-for-one trading, linking has only a small effect on retail electricity price in California, but the average retail electricity price in RGGI increases by about 1 percent relative to the unlinked program.
- *Potential leakage*. Power imports into RGGI fall and total generation in RGGI rises as a result of the reduction in allowance cost associated with producing power in the region, suggesting leakage is less of a concern in RGGI under three-for-one trading.

With three-for-one trading, the benefits of linking in the RGGI region accrue primarily to fossil generators, which face lower allowance and resource costs (Proposition 9 and Corollary 9.1). Consumers in the region also see slight benefits from lower electricity prices, while government revenues from allowance sales are lower because of lower allowance prices in RGGI. In California, fossil generators are negatively affected by the higher allowance costs and the lower electricity price.[28] For consumers, higher allowance value results in a direct benefit in the form of allowance revenue rebates, which complement the reduction in wholesale electricity costs relative to the unlinked scenario. Although electricity prices would be expected to increase with an increase in allowance prices, the assignment of allowance value to local distribution companies and the dynamic nature of capacity investments and electricity consumption in Haiku result in lower electricity prices in California under this scenario.

Table 2. Incidence of Benefits and Costs of Three-for-One Trading

$/MWh	RGGI			California		
	Consumers	Government	Fossil Generators	Consumers	Government	Fossil Generators
Allowance Value		-0.5	1.2	0.3		-0.4
Resource Cost/Flow			1.2			0.5
Electricity Price	0.2		-0.2	0.5		-0.5
Net Effect	**0.2**	**-0.5**	**2.2**	**0.7**		**-0.4**

4.5. Sensitivity Analysis

The consequences of linking the RGGI and California CO_2 markets depend importantly on other factors that affect the electricity markets in those two regions. The central case scenarios discussed above adopt assumptions that are consistent with the Annual Energy Outlook (AEO) 2011, but several of those assumptions are highly uncertain. In the sensitivity analysis discussed in this section, we analyze the effects of alternative forecasts of natural gas prices and of electricity demand growth on the outcomes in these two regional CO_2 allowance markets and, in particular, on the effects of linking the two markets. In the first sensitivity case, we consider the effects of higher natural gas price projections. In the second sensitivity case, we explore the

[28] Proposition 9 and Corollary 9.1 predict fossil generators in California should experience an increase in resource cost/flow as a result of three-for-one linking. However, due to the complex dynamics of investment and retirment in Haiku, which are not captured by our analytic model, our simulation results show these generators seeing a slight reduction in these costs.

effects of combined higher natural gas price projections and higher electricity demand growth. The sensitivity analysis is relevant not only because it allows for consideration of different projections for gas prices and electricity demand, but it illustrates how the effects of linking might change if these future paths shift due to unanticipated exogenous changes over the course of the trading programs.

4.5.1. High Natural Gas Price Sensitivity

Higher projected gas prices[29] have different effects on the generation mix in the two regions. In RGGI, coal capacity is greater, and thus coal claims a larger share of the generation market under high gas prices than it does with the low gas prices assumed in the central case. Higher gas prices also raise the cost of reducing emissions from the sector, as reflected by the marginal CO_2 abatement cost curve for RGGI shown in Figure 4, which is steeper than the comparable curve in Figure 2. In California, the main CO_2-emitting electricity capacity is fired by natural gas, although emissions associated with out-of-state coal generation serving California customers are also included in the trading system. High gas prices discourage the use of these local existing generators and encourage investment in non-emitting technologies, thereby lowering the marginal cost of reducing emissions in California as also seen by comparing Figure 4 to Figure 2. Total CO_2 emissions in the baseline with high gas prices are much greater than with lower gas prices, so the programs that target a particular cap yield greater reductions relative to the baseline with high gas prices than they do with lower gas prices.

[29] In this sensitivity case, natural gas forecasts track those in AEO 2009, which projected total natural gas consumption in 2020 of 21.53 trillion cubic feet (TCF) at an average wellhead price of $6.84/MMBtu, whereas the AEO 2011 projected total natural gas consumption in 2020 of 25.34 TCF at an average wellhead price of $4.47/MMBtu.

Figure 4. Marginal Abatement Cost Curves and Results without Linking under High Gas Prices

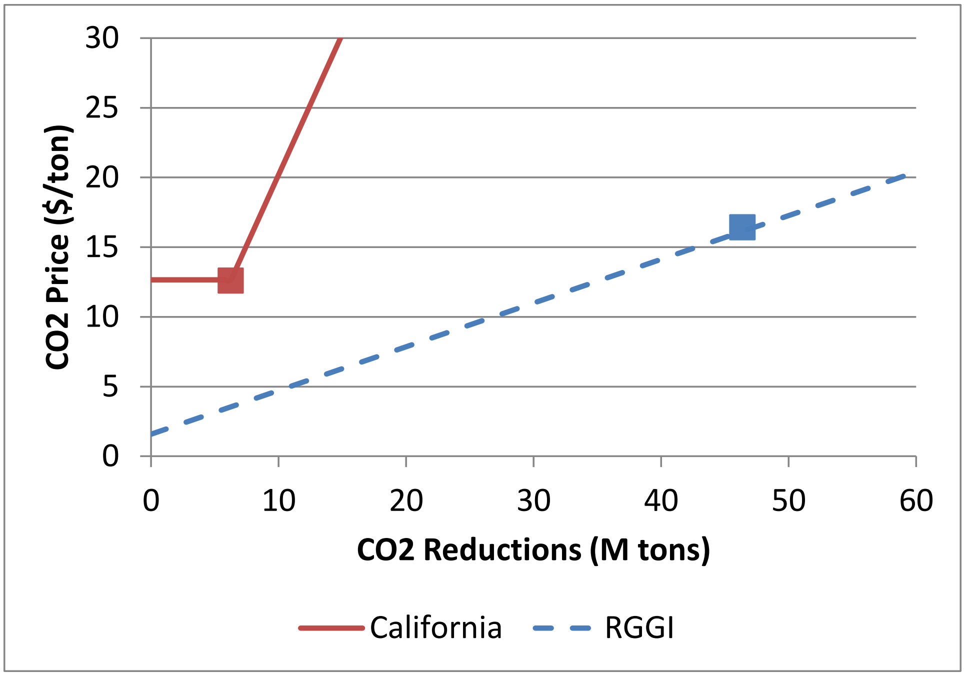

In the high gas price sensitivity case, when the two regional cap-and-trade programs are not linked, the price of CO_2 allowances in California falls to the floor, while the allowance price in RGGI is much higher than in the reference case presented earlier and actually exceeds the allowance price in California. As a consequence, linking the two programs at one-for-one pulls the allowance price in California up off its floor (Proposition 2), and thus emissions in California (Corollary 2.1) and across the two regions (Proposition 3) increase with linking (not depicted). However, linking results in very little allowance trading between RGGI and California.

The switch in relative allowance prices across the two unlinked programs, with RGGI now having higher prices than California, suggests that trading allowances at three-for-one will not result in aggregate emissions reductions. Indeed, three-for-one trading leads to allowances flowing from California to RGGI (not depicted). Allowance prices in California are twice as high as the floor, while the ability to use one-third of a California allowance to cover a ton of emissions in RGGI results in a substantial reduction in the RGGI allowance prices (Proposition 7 and Corollary 7.2). Total emissions in the two regions increase by roughly 10 percent compared with the unlinked case (Proposition 8).[30] On the other hand, linking at three-for-one reduces

[30] Proposition 8 requires that no price collars bind. In this case, the California price floor is marginally binding under autarky. Our results conform to the prediction of Proposition 8, even though not all requirements of the proposition are met.

power imports into RGGI by about 4.5 percent relative to the unlinked case and thus would lower the amount of leakage out of RGGI.

4.5.2. Combined High Natural Gas Price and High Electricity Demand Sensitivity

When electricity prices and electricity demand growth are both aligned with assumptions in AEO 2009, overall emissions in the absence of the two programs are even higher, and thus the abatement required to reach the emissions caps established in each of the two separate programs is substantial. When the two regional programs are in place but not linked, the prices of emissions allowances in the two regions are close, as shown in Figure 5. In both cases, allowance prices are substantially above the price floor in California.

Figure 5. Marginal Abatement Cost Curves and Results without Linking under High Gas Prices and High Electricity Demand

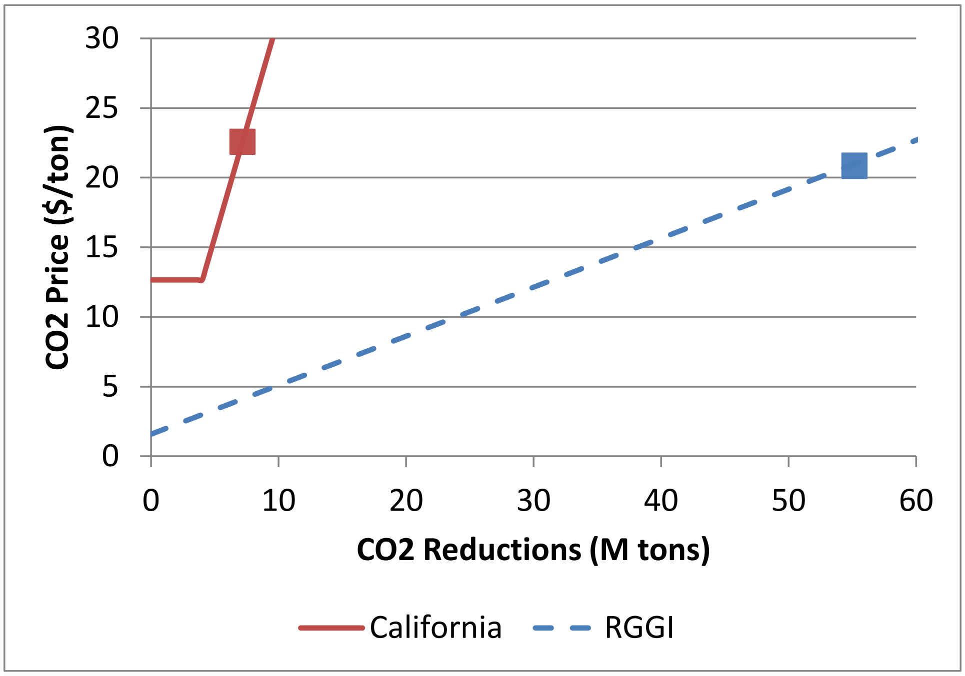

Because the prices in the unlinked programs are similar, linking the two programs at one-for-one has little effect in either market (not depicted). Little allowance trade occurs and total emissions are basically the same (Proposition 3), equal to the sum of the caps (Proposition 4). Under three-for-one trading, allowances flow from California to RGGI, and the allowance price in California is bid up to the California cost containment reserve (soft price ceiling) of $50.6 (Propositions 6 and 7 and Corollary 6.1). As a result, total CO_2 emissions in the two programs rise by roughly 5 percent relative to the unlinked case (proposition 8),[31] both because of RGGI buying allowances at the exchange rate (Corollaries 7.1 and 7.2) and because more allowances are offered in California to help support the allowance price ceiling. However, linking at three-for-one does reduce the amount of power imports into RGGI relative to the unlinked case, thereby potentially reducing emissions leakage.

5. Conclusion

This paper provides a framework for analysis of the linking of emissions trading systems. We develop an analytical framework for linking of programs with different features, including stringency, as measured by allowance prices, and a cost containment mechanism. We then apply that framework to the potential linking of the California and RGGI CO_2 emissions allowance markets. In a qualitative evaluation, we conclude these markets are almost ready for linking when evaluated based on administrative measures and the expected functioning of a common market. We then simulate the linking of these programs using a model of the US electricity sector; we analyze the programs in a stylized way, incorporating only the electricity sector portion of the now economy-wide cap-and-trade program in California. Despite the near readiness for linking, our simulation exercise suggests the difference in stringency and the different program designs introduce potentially difficult outcomes under linked market scenarios, some of which are predicted by our analytical model and others that are unexpected.

Our analytical and simulation models demonstrate that formal linking of emissions trading systems may lead to aggregate emissions that differ from the sum of the caps of the two programs when they operate independently. We find two-way uncertainty to the emissions outcome of linking; that is, emissions can be either lower or higher under a linked market. One reason this uncertainty could result is the presence of cost containment measures, either price floors or ceilings, that adjust the number of emissions allowances introduced in one program in response to allowance prices but which have effects that propagate across both programs when they are linked. The use of an exchange rate to reconcile differences in stringency between the programs also could have the effect of changing aggregate emissions. This consequence of linking might become increasingly apparent if relative marginal abatement costs change over time, for example, due to changes in fuel prices or electricity demand. In addition, other aspects

[31] Proposition 8 requires no price collars to bind. In this case, the link allowance price reaches the California price ceiling. Our results conform to the prediction of Proposition 8, even though not all requirements of the proposition are met.

of program design that could lead to this outcome include the treatment of offsets or efforts to contain leakage, some of which have been anticipated previously in the literature.

Linking also has important implications for the economic costs of the trading systems. Our analytical model finds that one-for-one linking improves the cost-effectiveness of emissions reductions, although the use of an exchange rate undoes some of these cost reductions and may even yield a linked system that is costlier that the combination of the independent systems. Additionally, a variety of subtle distributional effects emerge even when aggregate emissions are equal to the sum of the two independent caps, but which might be exacerbated when total emissions change. We consider the effects on three constituencies—consumers, producers and government. We find that whenever linking occurs, at least one of these groups suffers negative effects. Increasing attention is also being given to the distribution of emissions reductions that result from carbon trading programs, despite the global nature of climate change. This attention is focused on the concern that not all communities see reductions in conventional air pollutants or receive other environmental benefits in equal measure, and some may be made worse due to the flexible implementation of emissions trading and other carbon pricing schemes. Economic approaches to environmental policy typically separate these effects from the central goal of carbon pricing, which is to achieve greenhouse gas reductions at the least cost. In general, linking programs and expanding the coverage of programs is expected to contribute to this central goal. However, our research highlights other issues that should be anticipated, including changes in the total emissions of the regulated pollutant and potentially uneven distributional outcomes among the affected constituencies, but more generally would also include changes in conventional air pollutants. Policymakers may need to consider and compensate for these distributional effects if linking occurs.

The path forward for linking would appear significantly easier if programs initially have comparable stringency before linking is pursued—a criterion that is mandated but somewhat imperfectly defined by California state law under Senate Bill 1018. When comparable stringency is in place, then the expanded carbon market offers economic benefits as well as resilience to external factors, such as changes in weather or economic activity. In the meantime, until formal linking is achieved, incremental alignment of institutions, program design, and stringency represents an important but informal linking by degrees that points toward eventual broad-based carbon policy.

References

Bodansky, Daniel M., Seth A. Hoedl, Gilbert E. Metcalf, and and Robert N. Stavins. 2015. "Facilitating Linkage of Climate Policies Through the Paris Outcome." *Climate Policy* (July): 1–17.

Bohm, P. 1992. "Distributional Impacts of Allowing International Trade in CO2 Emissions Quotas." *The World Economy* 15 (1): 107–114.

Burtraw, Dallas, Karen L. Palmer, Clayton Munnings, Paige Weber, and Matt Woerman. 2013. "Title." Resources for the Future Discussion Paper 13–04, Washington DC.

Burtraw, Dallas, Karen Palmer, Anthony Paul, and Matt Woerman. 2012. "Secular Trends, Environmental Regulations, and Electricity Markets." *The Electricity Journal* 25 (6): 35–47.

Carbon Pulse. 2015. China Plans Tiered Discounting System for Carrying Over CO_2 Units Into National ETS.

Doda, Baran, and Luca Taschini. 2016. "Carbon Dating: When Is It Beneficial to Link ETSs?" *Grantham Research Institute*.

Fischer, Carolyn. 2003. "Combining rate-based and cap-and-trade emissions policies." *Climate Policy* 3 (Supplement 2): S89–S103.

Flachsland, Christian, Robert Marschinski, and Ottmar Edenhofer. 2009. "To link or not to link: benefits and disadvantages of linking cap-and-trade systems." *Climate Policy (Earthscan)* 9 (4): 358–372. doi: 10.3763/cpol.2009.0626.

Haites, Erik, and X. Wang. 2009. "Ensuring the Environmental Effectiveness of Linked Emissions Trading Schemes Over Time." *Mitigation and Adaptation Strategies for Global Change* 14: 465–476.

Harrison, David. 2006. Interactions of Cost-Containment Measures and Linking of Greenhouse Gas Cap-and-Trade Programs. In *Technical Update 1013315*. Palo Alto, CA: Electric Power Research Institute.

Helm, C. . 2003. "International Emissions Trading with Endongenous Allowance Choices." *Journal of Public Economics* 87: 2732–2747.

Holland, Stephen P. , and A. J. Yates. 2015. "Optimal Trading Ratios for Pollution Permit Markets." *Journal of Public Economics* 125: 16–27.

Hung, Ming-Feng, and Daigee Shaw. 2005. "A Trading-Ratio System for Trading Water Pollution Discharge Permits." *Journal of Environmental Economics and Management* 49 (1): 83–102.

Jaffe, Judson, Matthew Ranson, and Robert N. Stavins. 2009. "Linking Tradable Permit Systems: A Key Element of Emerging International Climate Policy Architecture." *Ecology Law Quarterly* 36: 789–808.

Lazarus, Michael, Lambert Schneide, Carrie Lee, and Harro van Asselt. 2015. Options and Issues for Restricted Linking of Emissions Trading Systems. Berlin, Germany: International Carbon Action Partnership.

Macinante, Justin 2016. "Networking Carbon Markets—Key Elements of the Process." *World Bank Group Discussion Paper*.

Mehling, Michael , and Erik Haites. 2009. "Mechanisms for Linking Emissions Trading Schemes." *Climate Policy* 9: 169–184.

Metcalf, Gilbert E., and D. Weisbach. 2012. "Linking Policies When Tastes Differ: Global Climate Policy in a Heterogeneous World." *Review of Environmental Economics and Policy (Winter 2016)* 6 (1): 110–129.

Newell, G. Richard, A. William Pizer, and Daniel Raimi. 2012. "Title." RFF Discussion Paper, Washington, DC.

Newell, Richard G. , William A. Pizer, and Daniel Raimi. 2013. "Carbon Markets 15 Years After Kyoto: Lessons Learned, New Challenges." *Journal of Economic Perspectives* 27 (1): 123–146.

Paltsev, Sergey V. . 2001. "The Kyoto Protocol: Regional and Sectoral Contributions to the Carbon Leakage." *The Energy Journal* 22 (4): 53–79.

Paul, Anthony, Dallas Burtraw, and Karen Palmer. 2009. Haiku Documentation: Electricity Market Model version 2.0. Washington, D.C.: Resources for the Future.

Pizer, William. A. , and Andrew J. Yates. 2015. "Terminating Links Between Emissions Trading Programs." *Journal of Environmental Economics and Management* 71: 142–159.

Ranson, Matthew, and Robert N. Stavins. 2016. "Linkage of greenhouse gas emissions trading systems: learning from experience." *Climate Policy* 16 (3): 284–300. doi: 10.1080/14693062.2014.997658.

Sterk, Wolfgang, and Joseph Kruger. 2009. "Establishing a transatlantic carbon market." *Climate Policy (Earthscan)* 9 (4): 389–401. doi: 10.3763/cpol.2009.0623.

Tuerk, Andreas, Michael Mehling, Christian Flachsland, and Wolfgang Sterk. 2009. "Linking carbon markets: concepts, case studies and pathways." *Climate Policy (Earthscan)* 9 (4): 341–357. doi: 10.3763/cpol.2009.0621.

World Bank, and Ecofys. 2016. States and Trends of Carbon Pricing. In *Washington, D.C.: World Bank Group*.

Zyla, K. A. . 2010. "Linking Regional Cap-and-Trade Programs: Issues and Recommendations." *Washington, DC: Georgetown Climate Center*.

Appendix

Proofs of Propositions and Corollaries

Proof of Proposition 1.

i. Suppose, without loss of generality, that $p_i^C \leq p_j^C$. Also suppose $p > p_i^C$. Then there exists an arbitrage opportunity to buy an allowance from system i at p_i^C and sell at the market price of $p > p_i^C$, which violates the no arbitrage condition. Thus, $p \leq p_i^C$, so p_i^C is the price ceiling.

ii. Suppose, without loss of generality, that $p_i^F \geq p_j^F$ and $D_i(p_i^F) + D_j(p_i^F) > \overline{A}_J$. Also suppose $p < p_i^F$. Then $D_i(p) + D_j(p) > \overline{A}_J \geq S_j(p)$, so the market does not clear. Thus, $p \geq p_i^F$, so p_i^F is the price floor.
Suppose, without loss of generality, that $p_i^F > p_j^F$ and $D_i(p_i^F) + D_j(p_i^F) < \overline{A}_J$. Also suppose $p \geq p_i^F$. Then $D_i(p) + D_j(p) < \overline{A}_J \leq S_j(p)$, so the market does not clear. Thus, $p < p_i^F$.

Proof of Proposition 2. Suppose, without loss of generality, that $p_i^A \leq p_j^A$. Consider two cases, (i) $p < p_i^A \leq p_j^A$, and (ii) $p_i^A \leq p_j^A < p$.

i. If $p < p_i^A \leq p_j^A$, then $D_i(p) > D_i(p_i^A)$ and $D_j(p) > D_j(p_j^A)$, so $D_i(p) + D_j(p) > D_i(p_i^A) + D_j(p_j^A)$. Additionally, $S_i(p) \leq S_i(p_i^A)$ and $S_j(p) \leq S_j(p_j^A)$, so $S_i(p) + S_j(p) \leq S_i(p_i^A) + S_j(p_j^A)$. Thus, $D_i(p) + D_j(p) > S_i(p) + S_j(p)$, so the market does not clear.

ii. By a similar argument, if $p_i^A \leq p_j^A < p$, then $D_i(p) + D_j(p) < S_i(p) + S_j(p)$, so the market does not clear.

Thus, it cannot be that $p < p_i^A \leq p_j^A$ or $p_i^A \leq p_j^A < p$, so $p_i^A \leq p \leq p_j^A$.

Proof of Corollary 2.1. Suppose, without loss of generality, that $p_i^A \leq p_j^A$. By Proposition 2, $p_i^A \leq p \leq p_j^A$. Then $D_i(p) \leq D_i(p_i^A)$, so $E_i \leq E_i^A$. Similarly, $D_j(p) \geq D_j(p_j^A)$, so $E_j \geq E_j^A$.

Proof of Proposition 3. There are three cases to consider: (i) neither system is binding, (ii) both systems are binding, and (iii) only one system is binding. Suppose, without loss of generality, that $p_i^A \leq p_j^A$.

i. Suppose $p_i^A = p_j^A = 0$. By Proposition 2, $p = 0$, so $D_i(p) + D_j(p) < \overline{A}_\iota + \overline{A}_J$. Then $D_i(p) = D_i(p_i^A)$ and $D_j(p) = D_j(p_j^A)$, so $E_i + E_j = E_i^A + E_j^A < \overline{A}_\iota + \overline{A}_J$.

ii. Suppose, without loss of generality, that $0 < p_i^A \leq p_j^A$. Then $D_i(p_i^A) = \overline{A}_\iota$ and $D_j(p_j^A) = \overline{A}_J$, so $E_i^A + E_j^A = \overline{A}_\iota + \overline{A}_J$. By Proposition 2, $p > 0$. Then $D_i(p) + D_j(p) = \overline{A}_\iota + \overline{A}_J$, so $E_i + E_j = E_i^A + E_j^A = \overline{A}_\iota + \overline{A}_J$.

iii. Suppose, without loss of generality, that $p_i^A = 0$ and $p_j^A > 0$. Then $D_i(p_i^A) < \overline{A_\iota}$ and $D_j(p_j^A) = \overline{A_J}$, so $E_i^A + E_j^A < \overline{A_\iota} + \overline{A_J}$. By Proposition 2, $0 \leq p \leq p_j^A$. If $p = 0$, then $D_i(p) + D_j(p) < \overline{A_\iota} + \overline{A_J}$, $D_i(p) = D_i(p_i^A)$, and $D_j(p) > D_j(p_j^A)$, so $E_i^A + E_j^A < E_i + E_j < \overline{A_\iota} + \overline{A_J}$. If $p > 0$, then $D_i(p) + D_j(p) = \overline{A_\iota} + \overline{A_J}$, so $E_i^A + E_j^A < E_i + E_j = \overline{A_\iota} + \overline{A_J}$.

Proof of Proposition 4. Suppose $D_i(p_i^A) = \overline{A_\iota}$ and $D_j(p_j^A) = \overline{A_J}$, so $E_i^A + E_j^A = \overline{A_\iota} + \overline{A_J}$. Also suppose, without loss of generality, that $p_i^A \leq p_j^A$.

i. Suppose $p^F < p_i^A \leq p_j^A < p^C$. By Proposition 2, $p^F < p < p^C$, so $D_i(p) + D_j(p) = \overline{A_\iota} + \overline{A_J}$. Thus, $E_i + E_j = E_i^A + E_j^A = \overline{A_\iota} + \overline{A_J}$.

ii. Suppose $p_i^A \leq p^F$ and $p^F < p_j^A < p^C$. By Proposition 2, $p^F \leq p < p^C$. If $p^F < p < p^C$, then $E_i + E_j = \overline{A_\iota} + \overline{A_J}$ as above. If $p = p^F$, then $D_i(p) + D_j(p) \leq \overline{A_\iota} + \overline{A_J}$. Thus, $E_i + E_j \leq E_i^A + E_j^A = \overline{A_\iota} + \overline{A_J}$.

iii. Suppose $p^F < p_i^A < p^C$ and $p_j^A \geq p^C$. By Proposition 2, $p^F < p \leq p^C$. If $p^F < p < p^C$, then $E_i + E_j = \overline{A_\iota} + \overline{A_J}$ as above. If $p = p^C$, then $D_i(p) + D_j(p) \geq \overline{A_\iota} + \overline{A_J}$. Thus, $E_i + E_j \geq E_i^A + E_j^A = \overline{A_\iota} + \overline{A_J}$.

Proof of Proposition 5. Suppose, without loss of generality, that $p_i^A \leq p_j^A$. Also suppose $E_i + E_j = E_i^A + E_j^A$. By Proposition 2, $p_i^A \leq p \leq p_j^A$, and by Corollary 2.1, $E_i \leq E_i^A$ and $E_j \geq E_j^A$. Hence, $C_i(\overline{Q_\iota}, E_i) \geq C_i(\overline{Q_\iota}, E_i^A)$ and $C_j(\overline{Q_J}, E_j) \leq C_j(\overline{Q_J}, E_j^A)$. Additionally, the difference in total cost, ΔC, is $\Delta C = \int_{E_i^A}^{E_i} \frac{\partial C_i(\overline{Q_\iota}, e)}{\partial E} de + \int_{E_j^A}^{E_j} \frac{\partial C_j(\overline{Q_J}, e)}{\partial E} de$. From each utility's first-order condition, $-\frac{\partial C_i(\overline{Q_\iota}, E_i)}{\partial E} = p \geq p_i^A = -\frac{\partial C_i(\overline{Q_\iota}, E_i^A)}{\partial E}$ and $-\frac{\partial C_j(\overline{Q_J}, E_j)}{\partial E} = p \leq p_j^A = -\frac{\partial C_j(\overline{Q_J}, E_j^A)}{\partial E}$. Thus, $\Delta C \leq (E_i^A - E_i)p - (E_j - E_j^A)p = 0$.

Proof of Proposition 6.

i. Suppose, without loss of generality, that $p_i^C \leq r p_j^C$. Also suppose $p_i > p_i^C$. Then there exists an arbitrage opportunity to buy an allowance from system i at p_i^C and sell at the market price of $p_i > p_i^C$, which violates the no arbitrage condition. Thus, $p_i \leq p_i^C$, so p_i^C is the price ceiling.

ii. Suppose, without loss of generality, that $p_i^F \geq r p_j^F$ and $D_i(p_i^F) + \frac{1}{r} D_j\left(\frac{1}{r} p_i^F\right) > \frac{1}{r} \overline{A_J}$. Also suppose $p_i < p_i^F$. Then $D_i(p_i) + \frac{1}{r} D_j(p_j) > \frac{1}{r} \overline{A_J} \geq \frac{1}{r} S_j(p_j)$, so the market does not clear. Thus, $p_i \geq p_i^F$, so p_i^F is the price floor.

Suppose, without loss of generality, that $p_i^F > r p_j^F$ and $D_i(p_i^F) + \frac{1}{r} D_j\left(\frac{1}{r} p_i^F\right) < \frac{1}{r} \overline{A_J}$. Also suppose $p_i \geq p_i^F$. Then $D_i(p_i) + \frac{1}{r} D_j(p_j) < \frac{1}{r} \overline{A_J} \leq \frac{1}{r} S_j(p_j)$, so the market does not clear. Thus, $p_i < p_i^F$.

Proof of Corollary 6.1.

i. If $r \geq \frac{p_i^C}{p_j^C}$, then $p_i^C \leq rp_j^C$, so $p^C = p_i^C$ by Proposition 6. Similarly, if $r \leq \frac{p_i^C}{p_j^C}$, then $p^C = rp_j^C$.

ii. If $r \leq \frac{p_i^F}{p_j^F}$, then $p_i^F \geq rp_j^F$, so $p^F = p_i^F$ by Proposition 6. Similarly, if $r \geq \frac{p_i^F}{p_j^F}$, then $p^F = rp_j^F$.

Proof of Proposition 7. Suppose, without loss of generality, that $p_i^A \leq rp_j^A$. Consider two cases, (i) $p_i < p_i^A \leq rp_j^A$, and (ii) $p_i^A \leq rp_j^A < p_i$.

i. If $p_i < p_i^A \leq rp_j^A$, then $D_i(p_i) > D_i(p_i^A)$ and $\frac{1}{r}D_j(p_j) > \frac{1}{r}D_j(p_j^A)$, so $D_i(p_i) + \frac{1}{r}D_j(p_j) > D_i(p_i^A) + \frac{1}{r}D_j(p_j^A)$. Additionally, $S_i(p_i) \leq S_i(p_i^A)$ and $\frac{1}{r}S_j(p_j) \leq \frac{1}{r}S_j(p_j^A)$, so $S_i(p_i) + \frac{1}{r}S_j(p_j) \leq S_i(p_i^A) + \frac{1}{r}S_j(p_j^A)$. Thus, $D_i(p_i) + \frac{1}{r}D_j(p_j) > S_i(p_i) + \frac{1}{r}S_j(p_j)$, so the market does not clear.

ii. By a similar argument, if $p_i^A \leq rp_j^A < p_i$, then $D_i(p_i) + \frac{1}{r}D_j(p_j) < S_i(p_i) + \frac{1}{r}S_j(p_j)$, so the market does not clear.

Thus, it cannot be that $p_i < p_i^A \leq rp_j^A$ or $p_i^A \leq rp_j^A < p_i$, so $p_i^A \leq p_i \leq rp_j^A$ and $\frac{1}{r}p_i^A \leq p_j \leq p_j^A$.

Proof of Corollary 7.1. Suppose, without loss of generality, that $p_i^A \leq rp_j^A$. By Proposition 2, $p_i^A \leq p_i \leq rp_j^A$. Then $D_i(p_i) \leq D_i(p_i^A)$, so $E_i \leq E_i^A$. Similarly, $\frac{1}{r}D_j(p_j) \geq \frac{1}{r}D_j(p_j^A)$, so $E_j \geq E_j^A$.

Proof of Corollary 7.2.

i. If $r = \frac{p_i^A}{p_j^A}$, then $p_i^A = rp_j^A$. By Proposition 7, $p_i = p_i^A = rp_j^A$ and $p_j = \frac{1}{r}p_i = p_j^A$. Then, by Corollary 7.1, $E_i = E_i^A$ and $E_j = E_j^A$.

ii. If $r > \frac{p_i^A}{p_j^A}$, then $p_i^A < rp_j^A$. By Proposition 7, $p_i \geq p_i^A$ and $p_j \leq p_j^A$. Then, by Corollary 7.1, $E_i \leq E_i^A$ and $E_j \geq E_j^A$.

iii. If $r < \frac{p_i^A}{p_j^A}$, then $p_i^A > rp_j^A$. By Proposition 7, $p_i \leq p_i^A$ and $p_j \geq p_j^A$. Then, by Corollary 7.1, $E_i \geq E_i^A$ and $E_j \leq E_j^A$.

Proof of Proposition 8. Suppose $p^F < p_i^A < p^C$ and $\frac{1}{r}p^F < p_j^A < \frac{1}{r}p^C$. By proposition 7, $p^F < p_i = rp_j < p^C$, so the linked system binds. Hence, $D_i(p_i) + \frac{1}{r}D_j(p_j) = \bar{A}_i + \frac{1}{r}\bar{A}_j = D_i(p_i^A) + \frac{1}{r}D_j(p_j^A)$, so $D_j(p_j) - D_j(p_j^A) = -r[D_i(p_i) - D_i(p_i^A)]$. The difference in aggregate emissions is given by $\Delta E = (E_i + E_j) - (E_i^A + E_j^A) = [D_i(p_i) - D_i(p_i^A)] + [D_j(p_j) - D_j(p_j^A)]$. Substituting from above, $\Delta E = [D_i(p_i) - D_i(p_i^A)] - r[D_i(p_i) - D_i(p_i^A)] = (1-r)[D_i(p_i) - D_i(p_i^A)]$. There are three cases to consider:

i. This difference is negative if and only if $1 - r < 0$ or $D_i(p_i) - D_i(p_i^A) < 0$ but not both. When $r < 1$, this is true only when $p_i > p_i^A$, which corresponds to $r > \frac{p_i^A}{p_j^A}$ by Corollary 7.2. Similarly, when $r > 1$, this is true only when $r < \frac{p_i^A}{p_j^A}$. Hence, $\Delta E < 0$ if and only if $\frac{p_i^A}{p_j^A} < r < 1$ or $1 < r < \frac{p_i^A}{p_j^A}$. Note that the former can only occur when $p_i^A < p_j^A$, and the latter when $p_i^A > p_j^A$.

ii. This difference is zero if and only if $1 - r = 0$ or $D_i(p_i) - D_i(p_i^A) = 0$. The former is true only when $r = 1$. By Corollary 7.2, the latter is true only when $r = \frac{p_i^A}{p_j^A}$. Hence, $\Delta E = 0$ if and only if $r = 1$ or $r = \frac{p_i^A}{p_j^A}$.

iii. This difference is positive if and only if $1 - r$ and $D_i(p_i) - D_i(p_i^A)$ have the same sign. Both are negative only when $r > 1$ and $p_i > p_i^A$, which corresponds to $r > \frac{p_i^A}{p_j^A}$ by Corollary 7.2. Similarly, both are negative only when $r < 1$ and $r < \frac{p_i^A}{p_j^A}$. Hence, $\Delta E > 0$ if and only if $r > \max\left\{1, \frac{p_i^A}{p_j^A}\right\}$ or $r < \min\left\{1, \frac{p_i^A}{p_j^A}\right\}$.

Proof of Proposition 9. Suppose, without loss of generality, that $p_i^A \le r p_j^A$. By Proposition 2, $p_i^A \le p_i = r p_j \le r p_j^A$, and by Corollary 2.1, $E_i \le E_i^A$ and $E_j \ge E_j^A$. Hence, $C_i(\overline{Q_\iota}, E_i) \ge C_i(\overline{Q_\iota}, E_i^A)$ and $C_j(\overline{Q_J}, E_j) \le C_j(\overline{Q_J}, E_j^A)$.

Proof of Corollary 9.1.

i. If $r \ge \frac{p_i^A}{p_j^A}$, then $p_i^A \le r p_j^A$. By Proposition 9, $C_i(\overline{Q_\iota}, E_i) \ge C_i(\overline{Q_\iota}, E_i^A)$ and $C_j(\overline{Q_J}, E_j) \le C_j(\overline{Q_J}, E_j^A)$.

ii. If $r \le \frac{p_i^A}{p_j^A}$, then $p_i^A \ge r p_j^A$. By Proposition 9, $C_i(\overline{Q_\iota}, E_i) \le C_i(\overline{Q_\iota}, E_i^A)$ and $C_j(\overline{Q_J}, E_j) \ge C_j(\overline{Q_J}, E_j^A)$.

Proof of Proposition 10. Consider two exchange rates, $r^1$and r^2. The resulting prices and quantities of these exchange rates are also denoted by a superscript 1 and 2, respectively. Suppose $E_i^1 + E_j^1 = E_i^2 + E_j^2 \le E_i^A + E_j^A$. There are two cases to consider, $p_i^A \le p_j^A$ and $p_i^A \ge p_j^A$:

i. Suppose $p_i^A \le p_j^A$. By Proposition 8, $\frac{p_i^A}{p_j^A} \le r^1 \le 1$ and $\frac{p_i^A}{p_j^A} \le r^2 \le 1$. Suppose, without loss of generality, that $\frac{p_i^A}{p_j^A} \le r^1 \le r^2 \le 1$. Then, by Corollary 7.2, $p_i^A \le p_i^1 \le p_i^2$ and $p_j^A \ge p_j^1 \ge p_j^2$; also by Corollary 7.2, $E_i^A \ge E_i^1 \ge E_i^2$ and $E_j^A \le E_j^1 \le E_j^2$. Now define $\Delta C_i = C_i(\overline{Q_i}, E_i^2) - C_i(\overline{Q_i}, E_i^1)$ and, similarly, $\Delta C_j = C_j(\overline{Q_j}, E_j^2) - C_j(\overline{Q_j}, E_j^1)$. From each utility's first-order condition, $p_i^1(E_i^1 - E_i^2) \le \Delta C_i \le p_i^2(E_i^1 - E_i^2)$ and $-p_j^1(E_j^2 - E_j^1) \le \Delta C_j \le -p_j^2(E_j^2 - E_j^1)$. Because $E_i^1 - E_i^2 = E_j^2 - E_j^1$, these expressions can be summed to give an expression for the difference in the total cost, $\Delta C \le (p_i^2 - p_j^2)(E_i^1 - E_i^2) \le 0$. The last inequality comes from $r^2 \le 1$, so $p_i^2 = r^2 p_j^2 \le p_j^2$. Hence, the total cost is less under exchange rate r^2.

ii. Suppose $p_i^A \ge p_j^A$. By Proposition 8, $1 \le r^1 \le \frac{p_i^A}{p_j^A}$ and $1 \le r^2 \le \frac{p_i^A}{p_j^A}$. Suppose, without loss of generality, that $1 \le r^1 \le r^2 \le \frac{p_i^A}{p_j^A}$. Then, by Corollary 7.2, $p_i^1 \le p_i^2 \le p_i^A$ and $p_j^1 \ge p_j^2 \ge p_j^A$; also by Corollary 7.2, $E_i^1 \ge E_i^2 \ge E_i^A$ and $E_j^1 \le E_j^2 \le E_j^A$.From each utility's first-order condition, $p_i^1(E_i^1 - E_i^2) \le \Delta C_i \le p_i^2(E_i^1 - E_i^2)$ and $-p_j^1(E_j^2 - E_j^1) \le \Delta C_j \le -p_j^2(E_j^2 - E_j^1)$. Because $E_i^1 - E_i^2 = E_j^2 - E_j^1$, these expressions can be summed to give an expression for the difference in the total cost, $\Delta C \ge (p_i^1 - p_j^1)(E_i^1 - E_i^2) \ge 0$. The last inequality comes from $r^1 \ge 1$, so $p_i^1 = r^1 p_j^1 \ge p_j^1$. Hence, the total cost is less under exchange rate r^1.

Thus, in either case, the total cost of the linked ETS system is less when the exchange rate is closer to 1.

Modeling California and RGGI

The modeling analysis of linking involves comparing the results of linked programs with those from unlinked programs. The first step in modeling the effects of linking is to specify the requirements imposed by the two trading systems on electricity generators within each region. In the case of California, the program extends beyond the state border, as those who deliver power to the California market that is generated outside the state also must surrender allowances to cover the associated CO_2 emissions. Throughout this modeling analysis, the central case assumptions regarding fuel prices and underlying electricity demand growth projections are based on now outdated assumptions in the US Energy Information Administration's (EIA) 2011 Annual Energy Outlook.[32] However, the main results are robust to different underlying parameters. We

[32] For more information on those assumptions, see the description of the baseline scenario in (Burtraw et al. 2012).

investigate in sensitivity analysis how different fuel costs and assumed rates of electricity demand growth would affect the trading and emissions outcomes.

The Haiku model solves for selected simulation years through 2035. In this analysis, we select 2015, 2017, and 2020 as the primary simulation years covering the time period for California's cap-and-trade program. The phase-in of emissions caps in California is coincident with a dramatic ramp up in the requirements of the renewable portfolio standard and thus the rapid introduction of renewables, which has important implications for allowance flows between California and RGGI. To capture these effects, we focus on 2020.

We model an emissions cap in California's electricity sector in order to achieve allowance prices roughly comparable with those anticipated by futures prices in the summer of 2012, about $18 per short ton in 2020 (in 2009 dollars). We use the resulting cumulative emissions across all years at these prices to create a trajectory of cap levels that decrease linearly each year, and we solve the model over the entire horizon through 2035. We assume that the emissions levels must not exceed the cap in each year, meaning no banking of allowances for future use occurs. The price floor in California rises at 5 percent per year in real terms, reflecting the program design. There is no explicit offset market or description of companion (technology) policies other than the renewable portfolio standard, but the electricity sector contribution to the cap is calculated taking these policies into account.

We include emissions associated with electricity imports into California under the cap to reflect regulators' intent to control emissions leakage. We assume that no contract shuffling in the imported power market will take place in response to the requirement to surrender allowances on imported power. In our model, the decision at the margin about whether to import power uses the marginal emissions rate for each neighboring region that exports power to California. The volume of allowances required for imported power is based on the average emissions rate for each neighboring region.

We model the RGGI cap by simulating a $6 per short ton allowance price (in 2009 dollars) on CO_2 emissions in 2015 that rises at 5 percent per year in real terms.[33] We use the resulting cumulative emissions across all years at these prices to create a trajectory of cap levels, which start at baseline levels at the beginning of the time horizon and decrease linearly each year. We assume that the cap is binding (emissions levels will hit the cap) and that no banking of allowances for future use occurs.

[33] We assume the price floor is unchanged.