Rethinking Linking: Effects of Cap-and-Trade Integration on Equilibrium Emissions and Welfare

DRAFT: DO NOT CITE OR CIRCULATE WITHOUT PERMISSION

Thor Larson

March 26, 2026

Abstract

Conventional economic reasoning suggests that linking cap-and-trade programs creates efficiency gains by lowering costs while holding emissions constant. When cap-and-trade programs have price floors or ceilings, however, linking may increase emissions and decrease welfare. I investigate this phenomenon within California and Québec, where cap-and-trade linkage increased emissions by 18.3 million tons CO2e and reduced welfare by $347 million annually – an outcome that was ex-ante foreseeable. A reformed linked market could have remained emissions-neutral while instead increasing welfare by $2 to $30 million annually. I argue that political and institutional frictions prevented such reforms from being implemented, proposing that policy makers prioritized cutting emissions cheaply over equating marginal damages and benefits. Given these constraints, linking with trading ratios remains incentive-compatible and recovers over three-fifths of potential emissions-neutral welfare gains.

I Introduction

Public-goods externalities have long been addressed through command-and-control regulation. The landmark insights of Pigou (1920) and Coase (1960) – that market-based instruments can also efficiently internalize externalities – marked a fundamental shift in economic thinking. This perspective has been especially influential in the context of environmental externalities, where a large literature shows that both price-based and quantity-based regulation can efficiently price pollution.[1] In practice, quantity-based regulation of emissions is far more popular: worldwide, roughly 80% of all regulated emissions are covered by 37 regional, national, and supranational markets called *cap-and-trade* programs or emissions trading systems (ETS).[2]

These 37 cap-and-trade markets generally operate independently from one another, a condition I refer to as *autarky*. Economic theory predicts that linking these autarkic markets produces Kaldor-Hicks improvements since linked markets experience cost-effective abatement, reduced volatility, and increased liquidity for the same quantity of emissions (Stavins, 1998; Jaffe et al., 2009).[3] This prediction, however, relies on the assumption that cap-and-trade operates as a pure quantity instrument. In fact, most operational cap-and-trade programs are better classified as "hybrid" markets where market prices are collared by price floors and ceilings (Goulder and Schein, 2013).[4]

Linking hybrid cap-and-trade markets may increase emissions and, moreover, may reduce welfare if the increase in external damages outweighs the cost reductions from trade. For an extreme example, consider a tiny jurisdiction T whose regional policy maker issues so many excess allowances that all regional emissions are covered and allowance prices in T fall to a floor of zero.[5] No abatement occurs in T, but T's small emissions footprint means that its contribution to global emissions and global welfare is limited. If T links with other

[1] On price-based regulation, see standard treatments of Pigouvian taxation; on quantity regulation, see Montgomery (1972) and Baumol and Oates (1988). Weitzman (1974) compares the relative performance of price and quantity instruments under uncertainty.

[2] World Bank (2025) reports that 28% of global greenhouse gas emissions are covered by market-based instruments. Of this, cap-and-trade programs collectively cover 23% of global emissions, while carbon taxes collectively cover just over 5% of global emissions.

[3] When regions A and B link their cap-and-trade programs, they agree that A's carbon allowance may be used to satisfy compliance obligations in B and vice versa.

[4] Even when markets lack explicit price collars, prices are constrained by a natural near-zero price floor of ϵ and some *de facto* price ceiling (above which policy makers will ostensibly intervene and revise policy). For instance, prices in the EU ETS collapsed to near-zero prices in 2007-08 due to issued allowances exceeding emissions demand; and recent policy discussions suggest that a price ceiling may be introduced in response to the ongoing Iran crisis (Foy, 2026).

[5] Of course, the first-best policy is for T to set its allowance quantity at the globally efficient level. When all regional policy makers (including T) are either unwilling or unable to change T's policies, this first-best outcome is unobtainable.

regions, though, T's surplus of allowances allow other regions to increase emissions as well. This effect depresses global carbon prices, reduces abatement globally, and increases total externality damages. From a global welfare perspective, then, linking with T exacerbates inefficiency rather than correcting it. Furthermore, nothing about this result hinges on T's price floor being fixed at zero, and the theoretical model presented in the remainder of the paper shows that the welfare effects of linking are generally ambiguous whenever price floors or ceilings bind.

This paper studies when linking hybrid cap-and-trade markets adversely affects emissions and welfare and which market provisions can mitigate these effects. Social welfare depends on emissions damages and abatement costs. If emissions increases are too large, the increase in damages may outweigh the cost savings from linking. By this same logic, market provisions that limit emissions increases may be welfare-improving even if they reduce cost efficiency. Moreover, focusing only on the efficiency of these provisions risks overlooking political and institutional frictions that determine which provisions can be feasibly implemented. By explicitly analyzing incentive-compatibility, we can better explain why inefficient policies persist and identify second-best provisions that are plausibly able to correct them. In this paper, I analyze, test, and quantify the efficiency and incentive-compatibility of many such provisions.

I study the California-Québec (CAQC) joint cap-and-trade program, the world's fourth largest ETS and the largest and longest-running linked cap-and-trade market. Covering over 310 million tons of greenhouse gas emissions in 2025, the CAQC market tracks emissions in major polluting sectors, including electricity generation, natural gas combustion, industrial production, and transportation. California-Québec provides an ideal context for studying the effects of linking cap-and-trade programs because, unlike other integrated markets (such as the European Union's EU ETS or the American Northeast's RGGI), there are minimal supranational confounders affecting emissions in the CAQC market. I focus on the program's first authorization period, which encompasses the eight-year period between 2013 and 2020.

I model the CAQC market by building on the approach of Borenstein et al. (2019). The core idea is to construct each region's emissions demand as the difference between business-as-usual (BAU) emissions and the supply of emissions abatement. BAU emissions forecasts are generated from a dynamic factor model estimated on historical emissions trends (Stock and Watson, 2016), which pools information from multiple macroeconomic and sectoral series to capture common shocks. Abatement supply is constructed as the sum of price-responsive demand reductions across four emissions-intensive fuels, together with price non-responsive components such as complementary regulations, offsets, and emissions reshuffling. Consistent

with prior work, the linked CAQC market is at the price floor with high probability (99.5% in the preferred specification).

Separately constructing regional emissions demand enables me to investigate the effects of linking. I first consider *ceteris paribus* autarky, where each region's stock of allowances remains unchanged in unlinked markets. The autarkic California market attains its price floor with even higher probability (99.9% of the time), whereas the autarkic Québec market is rarely at a price floor. In fact, the autarkic Québec market is constrained by a price ceiling 74.5% of the time. When markets link, expected emissions increase by 18.3 million tons of CO2 equivalent (MM T CO2e) in the linked market relative to *ceteris paribus* autarky. After accounting for gains from trade, and assuming a social cost of carbon (SCC) of $190 per ton, I estimate welfare losses from linking at $347 million per year.[6]

These welfare losses are not inevitable, and policy makers could introduce various provisions to counteract the adverse consequences of linking with collars. I focus on reforms that keep expected total emissions equal to their *ceteris paribus* autarky levels. This approach allows me to compare the relative cost-effectiveness of various provisions without taking a position on the social cost of carbon. Among these policies, price-floor increases are the most efficient, followed by allowance-quantity reductions. Tariff-like instruments – allowance border adjustment mechanisms and trading ratios – are the least efficient since they differentially raise prices across regions. Québec is a distributional loser across counterfactuals because all market provisions require the allowance-importing region to bear higher abatement costs.

These results challenge conventional expectations. Linking the CAQC market reduced global welfare and, at the same time, lowered Québec's regional welfare – and these outcomes were broadly predictable ex-ante given available information. Why, then, did Québec's policy makers choose to link? I address this question using two complementary approaches. First, I conduct a qualitative analysis of the policy record, including legislative history, regulatory documents, and contemporaneous statements by policymakers. Second, I develop a novel partial-identification framework that uses observed policy choices to infer the objective functions consistent with those choices. Both approaches support the same conclusion.

Policy makers do not appear to design climate policy in an economically efficient manner; rather, their actions can be rationalized as *declared-target cost minimization*. This means they aim to *minimize the costs* of reducing emissions below a *pre-declared emissions target*. This pre-declared target is generally set at an inefficient level (typically priced below the SCC), with policy inertia making revision of this target costly and unlikely. Within these constraints, regulators still attempt to meet mandated goals at the lowest possible cost.

[6]Unless otherwise stipulated, all values are reported in 2013 USD. One 2013 USD is approximately equivalent to 1.39 USD in 2026. https://www.in2013dollars.com/us/inflation/2013?amount=1

This suggests that policy makers place significant weight on declared emission targets when evaluating climate policy – and comparatively underweight actual emissions.

We can then analyze which market provisions are incentive-compatible for this type of policy maker. If bilateral negotiation and transfer payments are possible, increasing the price floor is the only jointly incentive-compatible reform. Emissions-neutral trading ratios are incentive-compatible only as a unilateral policy, though a unilaterally acting Québec would choose a lower trading ratio than under cooperative design. Under this lower trading ratio (approximately 2.8), emissions exceed autarky levels, though the increase is smaller than under default linking. Welfare gains under unilateral Québec-led trading ratios are $238 million per year, roughly three-fifths of the potential welfare gains from the most efficient emissions-neutral provision.

Related Literature To my knowledge, this is the first scholarly study to empirically evaluate the effects of linking two cap-and-trade programs.[7] It is also the first paper to analyze the California-Québec system bilaterally; existing work on this market has focused almost exclusively on California (Borenstein et al., 2019; Schmalensee and Stavins, 2017; Rabe, 2016). More broadly, the paper relates to research on the design and performance of emissions trading systems (Stavins, 1998; Rabe, 2016; Cullenward and Victor, 2020).

This work builds upon a long tradition of studying climate policy design, particularly the comparison of price- and quantity-based instruments (Weitzman, 1974; Pizer, 2002; Goulder and Schein, 2013). It also engages with newer proposed instruments, including carbon border adjustment mechanisms (Clausing and Wolfram, 2023; Clausing et al., 2025) and linking with trading ratios (Woerman, 2023; Holland and Yates, 2015; Fowlie and Muller, 2019; Muller and Mendelsohn, 2009). The core modeling approach builds on a growing literature that incorporates uncertainty into models of climate policy (Bushnell and Smith, 2025; Borenstein et al., 2019).

The discussion of incentive-compatible second-best reforms is informed by perspectives from political economy, where policy makers have been found to pursue objectives that extend beyond simple emissions minimization (Gaarder et al., 2026), including work examining noncooperative climate policy in decentralized settings (Helm, 2003; Carbone et al., 2009). More generally, the framework relates to questions of decentralized public-goods provision when the scope of externalities exceeds the jurisdictional reach of policy makers (Agrawal et al., 2024; Ostrom, 1990). The identification framework extends work in partial identifica-

[7] Green (2021) writes that, "[T]here is little research on the effects of linking carbon markets. California and Québec linked their markets in 2014, yet I was unable to find any scholarly assessments of their joint performance." I am not aware of work published since then that contradicts this assessment. Pizer and Yates (2015) examine cap-and-trade *de*linking, but do so in a purely theoretical setting.

tion and structural inference by applying it to the study of policy maker behavior (McFadden, 1974; Manski, 2007; Molinari, 2020).

II Theoretical model

Consider two regions i and j that operate cap-and-trade programs. For expositional ease, define abatement supply in region i according to $A_i(p) \equiv \max\{A_{0i}, A_{0i} + A_{1i}p\}$ and analogously for j.[8] Denote i's business-as-usual emissions as BAU_i and i's stock of allowances as q_i. Price collars – price floor $\underline{\kappa}$ and price ceiling $\bar{\kappa}$ – are the same in both markets.

II.a Unlinked markets

When price collars do not bind, all uncovered emissions must be abated:

$$\text{BAU}_i - q_i = A_i(p)$$

Define the *slack* in region i as $s_i \equiv \text{BAU}_i - A_{0i} - q_i$, a measures the excess or shortfall of allowances relative to baseline emissions net of free abatement at a carbon price of zero. The uncollared price of allowances in autarky is:

$$p_i^\circ(q_i) \equiv \frac{s_i}{A_{1i}} \tag{1}$$

I henceforth suppress the argument q_i for brevity. Price collars constrain the range of values that prices can take.

$$p_i = \begin{cases} \underline{\kappa} & \text{if } p_i^\circ < \underline{\kappa} \\ p_i^\circ & \text{if } \underline{\kappa} \leq p_i^\circ \leq \bar{\kappa} \\ \bar{\kappa} & \text{if } p_i^\circ > \bar{\kappa} \end{cases}$$

Emissions $e_i \equiv \text{BAU}_i - A_i(p_i)$ are similarly constrained:

$$e_i = \begin{cases} \bar{e}_i & \text{if } p_i^\circ < \underline{\kappa} \\ q_i & \text{if } \underline{\kappa} \leq p_i^\circ \leq \bar{\kappa} \\ \underline{e}_i & \text{if } p_i^\circ > \bar{\kappa} \end{cases}$$

[8] A more general treatment in Appendix A.2 shows that this section's results hold for any abatement supply curve with $A_i' > 0$ and $A_i'' > 0$.

with *maximum emissions* $\bar{e}_i \equiv \text{BAU}_i - A_i(\underline{\kappa})$ occurring at the price floor and with *minimum emissions* $\underline{e}_i \equiv \text{BAU}_i - A_i(\bar{\kappa})$ occurring at the price ceiling.

The total cost of abatement in i is then:

$$C_i(q_i) = \frac{A_{1i}}{2} p_i^2 \tag{2}$$

II.b Linked markets

The uncollared linked market price p^* is then written:

$$p^* \equiv \frac{s_i + s_j}{A_{1i} + A_{1j}} \tag{3}$$

Prices and emissions in the linked market are also collared:

$$p = \begin{cases} \underline{\kappa} & \text{if } p_i^* < \underline{\kappa} \\ p_i^* & \text{if } \underline{\kappa} \leq p_i^* \leq \bar{\kappa} \\ \bar{\kappa} & \text{if } p_i^* > \bar{\kappa} \end{cases}$$

$$e = \begin{cases} \bar{e}_i + \bar{e}_i & \text{if } p^* < \underline{\kappa} \\ q_i + q_j & \text{if } \underline{\kappa} \leq p^* \leq \bar{\kappa} \\ \underline{e}_i + \underline{e}_j & \text{if } p^* > \bar{\kappa} \end{cases}$$

Total abatement costs in the linked market are $C_i(p) + C_j(p)$. As a result, the change in abatement costs $\Delta C \equiv \Big(C_i(p) - C_i(p_i)\Big) + \Big(C_j(p) - C_j(p_j)\Big)$ is then:

$$\Delta C = \sum_{s \in \{i,j\}} \frac{A_{1s}}{2} (p^2 - p_s^2) \tag{4a}$$

When the linked market attains an interior price ($p = p^\circ$), (4a) can be reorganized as follows:

$$\Delta C = -\frac{A_{1i} A_{1j}}{2(A_{1i} + A_{1j})} (p_i^\circ - p_j^\circ)^2 + \sum_{s \in \{i,j\}} \frac{A_{1s}}{2} (p_s^{\circ 2} - p_s^2) \tag{4b}$$

(4b) decomposes the cost change into two terms. The first term is strictly negative and captures gains from trade: equalizing prices across regions leads to cost savings whenever $p_i^\circ \neq p_j^\circ$. The second term captures the cost of any aggregate abatement change induced by binding price collars. This term takes value 0 when neither collar binds, but is more generally of ambiguous sign.[9]

[9] A third term, $\sum_{s \in \{i,j\}} \frac{A_{1s}}{2} (p^{\circ 2} - p^2)$, vanishes from this expression because of the assumption $p = p^\circ$.

The other object of interest is the change in emissions:

$$\Delta e \equiv e - (e_i + e_j) \tag{5}$$

The total change in welfare, ΔWF, is a function of cost changes ΔC and emission changes Δe. Assume emissions are uniformly mixing and incur constant marginal damages ϕ:

$$\Delta WF \equiv -\Delta C - \phi \cdot \Delta e$$

ΔC and Δe vary depending on initial conditions, so I summarize the effects of linking on abatement costs, emissions, and welfare in Table 1. Derivations of all cases are detailed in Appendix A.1.

i \ j	Floor	Interior	Ceiling
Floor	C $-$ e $-$	C $\downarrow$ e $\uparrow$	C ? e ?
Interior	C $\downarrow$ e $\uparrow$	C $\downarrow$ e $-$	C ? e $\downarrow$
Ceiling	C ? e ?	C ? e $\downarrow$	C $-$ e $-$

(a) Costs and emissions

i \ j	Floor	Interior	Ceiling
Floor	$-$	?	?
Interior	?	$\uparrow$	?
Ceiling	?	?	$-$

(b) Aggregate welfare

Table 1: Effects of linking on (a) costs C and emissions e and (b) aggregate welfare

II.c Slack and differential additionality

Figure 1 illustrates the emissions effects of linking with j across the range of q_i, under three cases for j's autarkic market: j is at its price floor, interior, or price ceiling. As q_i increases, the price ceiling binds for sufficiently low q_i and the price floor binds for sufficiently high q_i; the corresponding cutoffs are shown above and below the axis respectively, for the unlinked and linked markets.

Shaded regions indicate where $\Delta e \neq 0$, reproducing the results of Table 1(a). Dark blue regions denote *differential additionality*: ranges of q_i where a collar is close to binding in one market but not the other, so that the linked and unlinked regimes differ in whether the marginal allowance is additional. Specifically, this arises when the linked market has slack while the unlinked market does not, or vice versa. These regions matter because a marginal allowance may be additional under linking but nonadditional in autarky – or vice versa – potentially distorting abatement incentives at the margin.

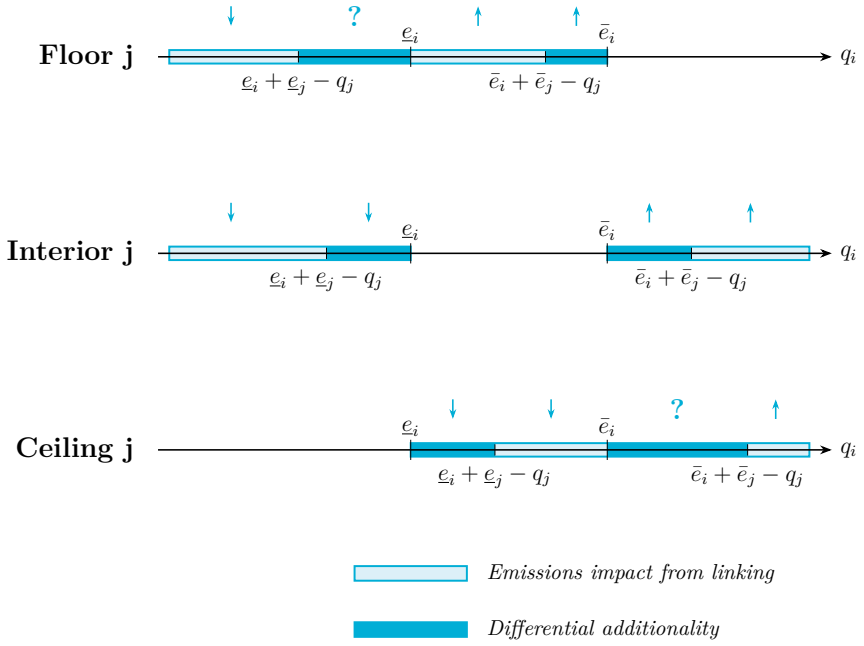


Figure 1: Emissions impacts of linking as a function of i's quantity and j's initial conditions. Values of q_i where emissions change are shaded, with effects on emissions notated with arrows above. Regions of differential additionality are in dark blue.

II.d Sufficient statistics for welfare analysis

Marginal emissions damages ϕ, abatement supply responsiveness A_{1i}, and "slack" $s_i \equiv \text{BAU}_i - A_{0i} - q_i$ are the sufficient statistics for welfare analysis when collars $\underline{\kappa}, \bar{\kappa}$ and allowance quantities q_i are known. .

ΔC can be written entirely in terms of these statistics, since $p_i^\circ = s_i/A_{1i}$ (1). Similarly, the emissions collars $\bar{e}_i = s_i - A_{1i}\underline{\kappa} + q_i$ and $\underline{e}_i = s_i - A_{1i}\bar{\kappa} + q_i$ are fully pinned down by the same statistics, so Δe is recoverable as well.. It follows that the full welfare change ΔW is recoverable from $\{\phi,\ A_{1i},\ s_i,\ \underline{\kappa},\ \bar{\kappa},\ q_i\}$ alone.

This stylized model isolates the core mechanism by which collars affect welfare by abstracting away many institutional features of real-world cap-and-trade programs. In practice, abatement supply functions are nonlinear, so the full functions $A_i(\cdot)$ must be known rather than just the slope parameters A_{1i}. My empirical analysis incorporates these institutional features and evaluates the welfare effects of linkage using observed policy rules and data.

III History and Background

My empirical setting is the California-Québec (CAQC) emissions market. CAQC is the largest and longest-operating linked cap-and-trade market, and is also the world's fourth-largest cap-and-trade market overall. The CAQC market also faces a reduced risk of confounding factors such as supranational policies, partial emissions coverage, or trade exposure when assessing the impacts of linking. Unlike the European Union's larger EU ETS, the CAQC market operates without supranational institutions enforcing compliance. Additionally, the CAQC market covers over three-quarters of total regional emissions, a much higher percentage than in many other markets (International Carbon Action Partnership, 2025). Finally, unlike other multiregional cap-and-trade programs (such as RGGI and the EU ETS), California and Québec engage in comparatively little goods trade.[10] Yet despite these attractive features, the CAQC market has drawn far less research attention than national systems such as the EU ETS, likely because it is framed as a regional rather than national program.

Within the CAQC market, California is the dominant actor. Its economy has historically been roughly ten times the size of Québec's while its annual emissions have been roughly five

[10] In 2023, total exports from California to Canada in 2023 totaled USD$0.8 billion, and imports to California from Québec totaled USD$1.9 billion.(Statistics Canada). Relative to the overall sizes of California's (USD$3.87 trillion, Bureau of Economic Analysis) and Québec's (USD$405 billion, Statistics Canada) GDP in 2023, net exports comprise only 0.25% of Québec's GDP and 0.027% of California's GDP.

times larger. A simple calculation then suggests that California's average emissions intensity is about half that of Québec's.

Policy History I provide a brief overview of the evolution of climate policy in California and Québec here; Appendix B presents a more detailed history.

California authorized a cap-and-trade program as part of Assembly Bill 32 (AB 32) in 2006, a sweeping piece of state climate legislation that required the state to return statewide emissions to their 1990 levels by 2020. AB 32 directed the California Air Resources Board (CARB) to develop a Scoping Plan to achieve this target, including the adoption of regulations for mandatory emissions reporting, the construction of a 1990 baseline emissions inventory (which was previously unknown), and the implementation of complementary policies like the state's Renewable Portfolio Standard. In parallel, California Governor Arnold Schwarzenegger sought to position California as an international climate leader by launching the Western Climate Initiative (WCI), a coalition of subnational jurisdictions that ultimately included eleven U.S. states and Canadian provinces by the end of 2008. WCI members were expected to design and legislatively authorize their own cap-and-trade programs, with the long-term objective of establishing an integrated regional carbon market.

Québec was the sole WCI member to fully implement the initiative's cap-and-trade vision. Québec Premier Jean Charest's Decree 1187-2009 set a 2020 target of 20% below 1990 emissions and explicitly framed this ambition as contingent on international cooperation, noting as early as 2009 that "it was possible to reduce [domestic] greenhouse gas emissions in the province to approximately 14% below the 1990 baseline, meaning that Québec had to achieve international reductions if it wanted to set a more ambitious target."(Government of Québec, 2021). The Québec parliament unanimously passed a cap-and-trade bill later in 2009.

Meanwhile, the WCI coalition began unraveling. As Rabe (2016) notes, "Much of the WCI alliance hinged on agreements between like-minded state governors and provincial premiers and often lacked legislative buy-in or supportive legislation." Cap-and-trade efforts stalled as these governors left office, and American enthusiasm waned further after the U.S. Senate failed to pass a federal cap-and-trade bill in 2009. By late 2011, California and Québec were the only jurisdictions that had adopted cap-and-trade regulations consistent with the WCI framework.

After a period of joint rule-making, California and Québec launched their cap-and-trade programs in 2013. Formal linkage began in late 2014, initially covering electricity and large industrial sources. The program expanded in 2015 to include transportation fuels and smaller industrial sources. The initial linked period, which is the focus of this paper, runs from

2013 through 2020. The linked CAQC agreement was subsequently renewed by both regions through 2030, with further extensions through 2040 and a potential linkage with Washington State currently under consideration.

Prices, Price Collars, and Allowance stocks Within the 2013-2020 time horizon, a total of 2,793.1 million metric tonnes (MM T) of allowances were issued, with California issuing 2386.6 MM T and Québec issuing 206.5 MM T. Allowances were freely bankable across time and space with minimal restrictions.

Allowance auctions occur quarterly in an auction jointly conducted by the CARB and Québec Ministère de l'Environnement (MELCCFP). However, auction prices are collared so as to keep prices within a fixed range. All sales must be at or above an auction reserve price (ARP), which was originally set at $10.50 USD. Additionally, allowance price containment reserves (APCR) release additional allowances into the market if prices exceed predetermined thresholds; A total of 161.7 MM T of allowances (121.8 MM T California; 39.9 MM T Québec) were supplied into these reserves, to be released in thirds if market prices reached pre-specified levels of $40, $45, and $50 USD.[11] Both ARP and APCR prices increase annually at a rate of 5% plus the rate of inflation. Figure 2 plots auction settlement prices, secondary market prices, and ARPs in the joint market for the 2014-2020 period. Prices in the joint market remained at the ARP at or near these price floors for nearly the entire duration of the program's first renewal period.

Contemporaneous SCC estimates from this time period fall within these price collars, though it is worth noting that all are noticeably greater than the ARP even after accounting for its annual 5% increases. For instance, the Canadian government's 2010 national valuation of the social cost of carbon of $27 CAD ($26.20 USD) was more than double the ARP (Environment and Climate Change Canada, 2016).

[11] California and Québec use different currencies, meaning price floors are potentially vary across regions as exchange rates fluctuate (Coutu, 2016). This is not a major concern during the period studied, where exchange rates remain near-parity (0.995). Moreover, Québec requires that units sold inside its jurisdiction be sold at the highest of the prices of both jurisdictions according to the daily average exchange rate of the Bank of Canada published on its website on the day preceding the sale (International Carbon Action Partnership, 2025)

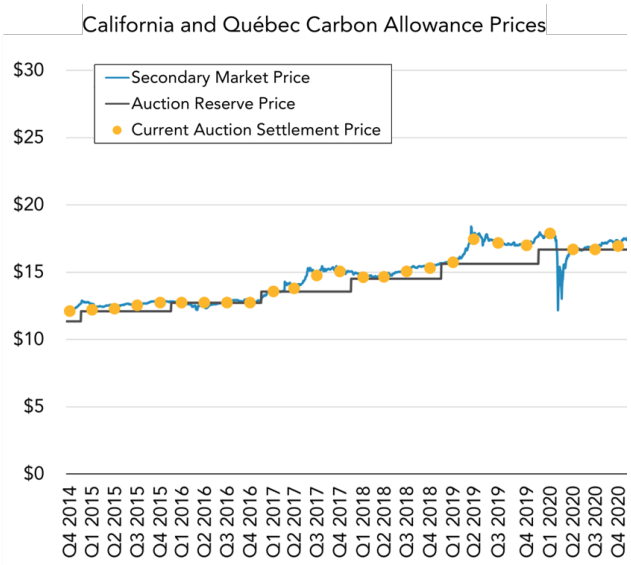


Figure 2: California and Québec average allowance prices, 2014-2020.
Image credit: California Air Resources Board

IV Empirical Analysis

This section studies the impact of linking on emissions and welfare by constructing regional emissions demand curves for California and Québec. To evaluate policy design from the perspective of regulators in 2013, I construct demand curves, $D_i(p)$, for the 2013-2020 compliance period using only ex-ante information that would have been available at the time the joint market was established. The ex-ante restriction matters because allowance supply decisions were made before the relevant compliance outcomes were realized; using ex-post information would overstate what policy makers could have anticipated.

One conventional approach to modeling emissions demand would be to directly estimate how emissions respond to variation in allowance prices using observed market data. That approach is insufficient here for practical and conceptual reasons. Practically, there is limited quasi-experimental allowance price variation; observed allowance prices were regularly at price floors during the 2013-2020 period, and there is no pre-period auction data before 2013. Conceptually, projecting future emissions would require forecasting future demand shifters such as macroeconomic conditions, technological change, preferences, and relative fuel prices, all of which are highly uncertain. In the language of Bushnell and Smith (2025), any economist attempting to model emissions demand must confront three sources of uncertainty: *model uncertainty* about how well her data-generating process (DGP) captures real-world behavior; *policy uncertainty* about any future policy changes that influence the

DGP; and *data uncertainty* about any other future input data realizations that will influence future emissions outcomes.

I account for policy and data uncertainty by combining the approach of Borenstein et al. (2019) with modern forecasting methods. This approach constructs business-as-usual (BAU) emissions forecasts and abatement supply independently and then computes emissions demand as the difference between BAU emissions and abatement supply. For BAU emissions, I estimate historical trends and forecast forward using a dynamic factor model (Stock and Watson, 2016), which pools information across multiple macroeconomic and sectoral series to capture common shocks. For abatement supply, I decompose abatement into price-responsive and price non-responsive components. Price-responsive abatement is constructed from the price response of fuel demand in four emissions-intensive fuels. Price non-responsive abatement reflects policies and activities that reduce emissions independent of allowance prices, including complementary regulations, offsets, and emissions reshuffling, which I implement following Borenstein et al. (2019).

I address model uncertainty in Appendix C, where sensitivity analyses show that these results are robust to alternative specifications and assumptions about BAU emissions and abatement supply.

Data The analysis draws on three data sources: emissions data, historical prices and quantities of energy-intensive fuels, and broader macroeconomic drivers of emissions.

Emissions Data The core datasets used in this analysis are the regional emissions inventories published by EIA and MELCCFP. These data are published annually at the sector-by-use-case level (e.g. "industrial emissions from the combustion of non-ferrous metal"). Data from 1990-2010 are used for ex-ante estimation; data post-2010 are used for forecast evaluation.[12] However, California emissions data pre-2006 were constructed retroactively as part of the state's 2008 Scoping Plan, with the 1990-2000 period being notably unreliable. I address this measurement problem in multiple ways. In the preferred specification, I use data from the EIA's State Energy Data System (SEDS) to track state-level emissions for the entire sample period. An alternate approach follows the cleaning process detailed in Borenstein et al. (2019) to clean CARB's published emissions inventories and is detailed in Appendix C.

From 2013-2014, the CAQC market covered only a "narrow" scope of emissions, comprising a portion of industrial natural gas emissions and all emissions from transportation and electricity generation.[13] Market coverage expanded to encompass a "broad" scope in

[12] Prior to the signing of the Kyoto Protocol, emissions were not tracked; thus pre-1990 data are inaccessible.

[13] Approximately 59% of all industrial emissions and natural gas consumptions were covered under the narrow scope.

2015 (and continuing to the present), which covers all narrow scope emissions, all industrial emissions, and all natural gas usage. This scope expansion was planned from the program's onset, making it necessary to forecast both narrow and broad scope emissions over time. Other emissions-intensive sectors, including waste management, agriculture, and forestry, remain uncovered "due to challenges associated with precise measurement, tracking or sector structure." (CARB 2008)

Other energy inputs Californian energy prices and quantity data are sourced from the U.S. Energy Information Administration (EIA) and the California Department of Tax and Fee Administration (CDTFA).[14]. Prices of electricity, natural gas, retail gasoline, and diesel are all taken from the EIA, as are electricity demand (MWh), electricity generation (MWh), and natural gas demand (MMBtu). Annual retail gasoline and diesel demand volumes are from the CDTFA.

Energy data for Québec come from multiple sources. Electricity prices are from HydroQuébec, while electricity generation and demand (MWh) are from Statistics Canada. Gasoline and diesel prices are calculated as the unweighted average of monthly prices across Montréal, Québec City, and Ottawa-Gatineau, using Statistics Canada data reported in liters.[15] These are converted to per-gallon prices using a factor of 3.785 liters per gallon. Natural gas price and volume data come from Statistics Canada as well,[16] with volumes reported in cubic meters and converted to MMBtu using a factor of 0.0353 m^3 per MMBtu. The average retail price is computed as the volume-weighted average across residential, commercial, and industrial users.

Macroeconomic Trends California GDP data are from the Bureau of Economic Analysis, and Québec GDP data are from Statistics Canada.[17] Annual CAD-USD exchange rates are taken from the Federal Reserve Bank of St. Louis.[18] Vehicle miles traveled (VMT) data are from the U.S. Department of Transportation and Statistics Canada.[19] Global oil prices are measured using Brent and WTI benchmarks from the EIA.

Business-as-usual emissions I forecast business-as-usual (BAU) emissions using a dynamic factor model (DFM), a method that parsimoniously incorporates a large number of predictive variables. DFMs handle this large dimensionality by collapsing the full set of

[14] https://www.cdtfa.ca.gov/taxes-and-fees/spftrpts.htm
[15] Statistics Canada. Table 18-10-0001-01
[16] Statistics Canada. Table 25-10-0033-01
[17] Statistics Canada. Table 36-10-0402-01
[18] https://fred.stlouisfed.org/series/AEXCAUS
[19] California VMT: https://www.fhwa.dot.gov/policyinformation/statistics.cfm. Québec VMT are calculated through a combination of multiple datasets detailed in Appendix C. When necessary, kilometers are converted to miles using a factor of 1.609 km per mile.

predictive data into a small number of latent factors – a restriction that is both consistent with standard dynamic equilibrium macroeconomic theories and computationally facile (Stock and Watson, 2016). I detail alternate forecasting approaches in Appendix C that relax the low-dimensional latent factor restriction and impose weaker parametric structures. The similarity of BAU forecasts across specifications indicates that the core findings are not an artifact of the strong cross-sectional or dynamic assumptions embedded in the DFM.

Let $X_t = (X_{1t}^{CA}, X_{1t}^{QC} ... X_{Nt}^{CA}, X_{Nt}^{QC})'$ denote the vector of log-values of the $2N$ region-level variables used to characterize the BAU environment: sectoral emissions, electricity generation and consumption, vehicle miles traveled, regional GDP, regional natural gas prices, and crude oil prices. Regional values for both California (CA) and Québec (QC) are included for all variables.[20] The DFM is then written as:

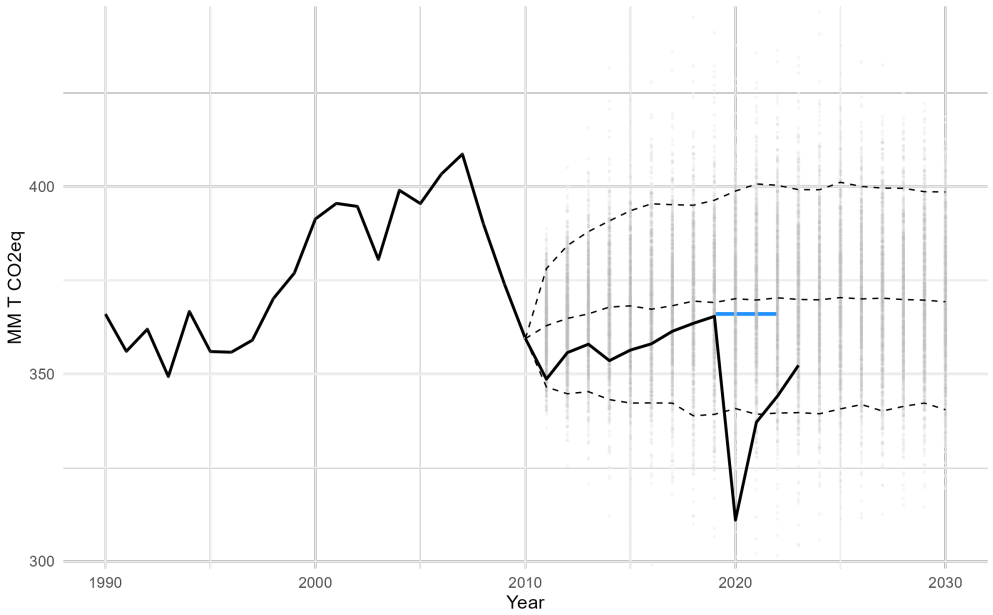
$$\begin{aligned} X_t &= \mu + \Lambda F_t + u_t \\ F_t &= \textstyle\sum_{\ell=1}^{L} A_\ell F_{t-\ell} + B_0 + B_1 t + \varepsilon_t \end{aligned} \tag{6}$$

Factors (F_t) evolve according to a vector autoregressive model (VAR) that includes a time trend t. The idiosyncratic component u_t represents the deviations of the observed variables from the factors. I also assume that the VAR includes sufficient lags (L) such that shocks ε_t are serially uncorrelated.

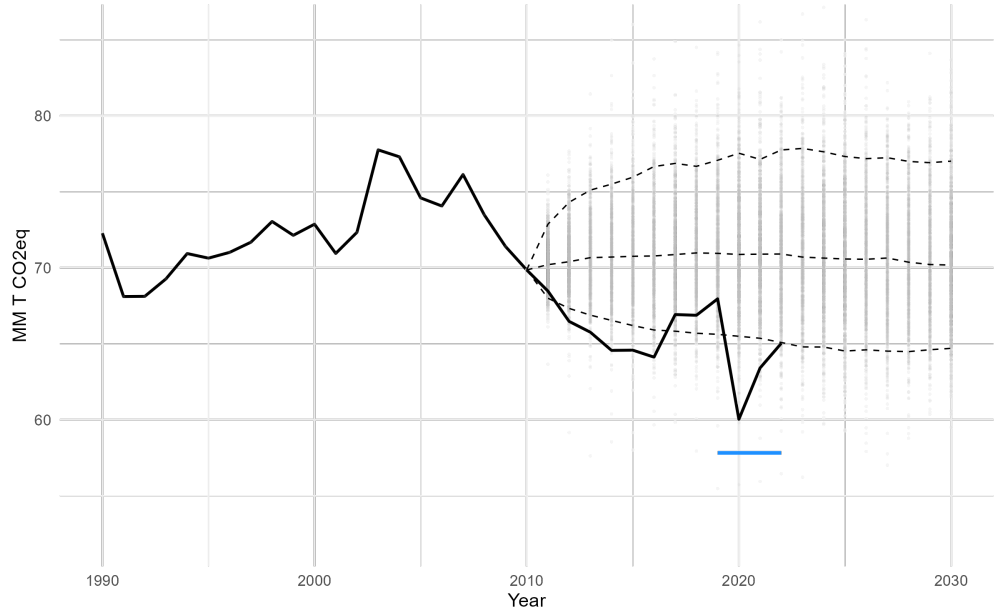
After estimating the underlying DFM from observed data, I simulate the empirical distribution of future emissions over 1,000 draws for California and Québec. I summarize these results in Table 2. Annual emissions from all covered sources are shown in the Broad scope column for each region. Cumulative emissions sum these values year-over-year, accounting for the exemption of transportation emissions and some natural gas/industrial emissions from the cap-and-trade program during the 2013-2014 period.

Figure 3 depicts the median, 5th, and 95th percentile of broad scope emissions by region; actual emissions are plotted in solid black. In addition, I plot each region's stated 2020 emissions goal in blue. These forecasts reveal a noticeable difference between California and Québec. Even without the emissions drop from the COVID-19 pandemic, California had a high likelihood of attaining 1990 emissions levels by 2020 absent any carbon-pricing policy; indeed, the state's emissions target lies at the 44th percentile of all emissions forecasts. In contrast, Québec was not on track to independently meet its 2020 target of 20% below 1990 emissions levels, which fall well below even the lowest-emission emissions forecasts. Indeed, Québec policy makers accurately assessed this circumstance as early as 2009: "Québec had to achieve international reductions if it wanted to set a more ambitious target." (Government of Québec, 2021).

[20] For oil prices, the prices of WTI and Brent crude are used.



(a) California



(b) Québec

Figure 3: Broad scope BAU forecasts for (a) California and (b) Québec. 1,000 bootstraps are drawn from a joint DFM trained on data from 1990-2010. Dashed lines depict the 5th, median, and 95th percentiles of each year's emissions draw. Solid lines depict observed emissions. The solid blue line represents each region's 2020 emissions target.

	California		Québec	
Year	Broad scope MMT	Cumulative MMT	Broad scope MMT	Cumulative MMT
2013	366.4	112.4	70.7	25.6
	(13.3)	(6.5)	(2.6)	(1.4)
2014	367.4	225.1	70.8	51.4
	(14.8)	(11.9)	(3.1)	(3.0)
2015	368.1	593.2	70.9	122.4
	(16.0)	(22.4)	(3.5)	(6.0)
2016	368.0	961.2	71.0	193.4
	(16.5)	(35.7)	(4.1)	(9.7)
2017	368.4	1329.7	71.2	264.5
	(16.8)	(49.3)	(4.3)	(13.6)
2018	368.9	1698.5	71.2	335.7
	(17.9)	(63.6)	(4.2)	(17.3)
2019	369.2	2067.7	71.4	407.2
	(18.4)	(77.7)	(8.7)	(24.8)
2020	369.8	2437.5	71.6	478.8
	(18.8)	(92.2)	(16.1)	(39.3)

Table 2: Summary of regional DFM emissions forecasts. Estimates are mean values of 1000 DFM draws. Values in parentheses are standard deviations of 1000 draws. All values are in millions of tons CO2e.

Abatement supply I infer regional abatement supply curves A_{QC}, A_{CA} by separating abatement supply into its price-responsive and price non-responsive components. Price-responsive abatement, which reflects the potential impacts of increased allowance prices on total abatement, encompasses demand response across four carbon-intensive fuels (electricity, gasoline, diesel, and natural gas). For each fuel, an increase in the price of carbon allowances ($ / T CO2e) translates into a proportional increase in the price of fuels ($ / fuel unit). This proportional relation is defined by each fuel's emissions intensity – a known physical constant.[21] Assuming this price increase is fully passed onto consumers, the corresponding decrease in the quantity of fuel demanded is governed solely by the fuel's demand elasticity.[22] I account for uncertainty in this demand elasticity by drawing fuel-specific demand elasticities from $\beta(2, 2)$ distributions over a range of plausible estimates.[23] Using these assumptions, I

[21] Emissions intensities are taken from the EPA for all fuels except electricity; electricity intensities are 2012 emissions divided by total demand in 2012. Natural gas: 0.0543 T CO2e / MMBtu. Electricity: California, 0.327 T/ kWh; Québec, 0.00279 T / kWh. Motor gasoline: California (E10 gasoline), 8.36441 kg CO2e/gal; Québec (E05 gasoline), 8.62164 kg CO2e/gal. Diesel fuel: 10.18 kg CO2e / gal, both regions.

[22] The assumption of 100 percent pass through is highly plausible for gasoline fuels, a likely overestimate for natural gas, and a potential underestimate for electricity – see Borenstein et al. (2019) for further discussion on this point. The effects of incomplete pass-through on overall abatement are small, so the simplifying assumptions made here are relatively innocuous.

[23] For natural gas, (-0.1,-0.3), as in Rubin and Auffhammer (2024). For gasoline and diesel, (-0.1, -0.2), as in Hughes et al. (2008). For electricity without any complementary policies in play, (-1.1, -1.9), as in

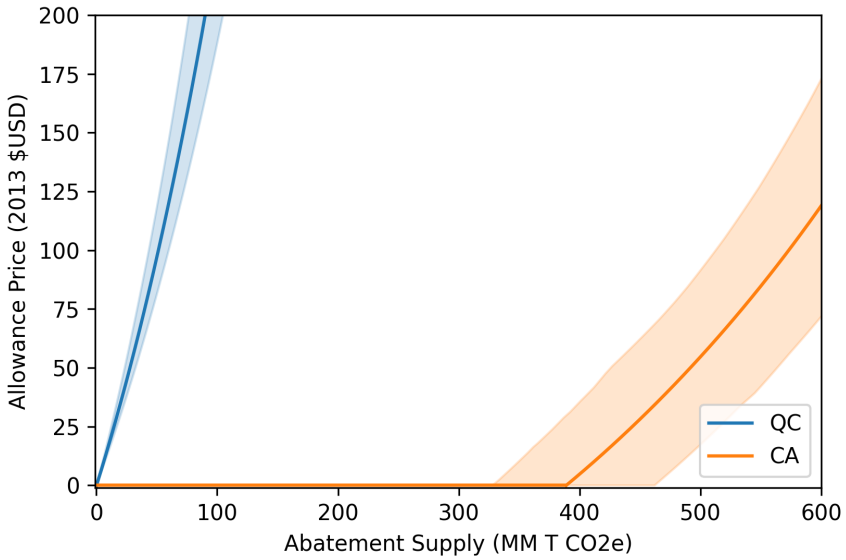


Figure 4: Abatement supply curves for Québec (blue) and California (orange), 1,000 draws. Solid lines denote mean regional abatement, with the shaded range bounding the 2.5th and 97.5th percentiles of abatement supply for each region.

map changes in allowance prices to changes in fuel volumes. Finally, I convert fuel volumes into emissions equivalents by once again using the known emissions intensities. Summing over all four fuel types yields the price-responsive abatement supply curve.

Price non-responsive abatement reflects the suite of energy policies and other activities expected to affect emissions independent of cap-and-trade prices. In California, these activities include complementary environmental policies (mandating a Renewable Portfolio Standard, increasing fuel efficiency standards), emissions offsets, and the reshuffling of electricity emissions into uncovered regions (by out-of-state generators). I implement these activities identically to Borenstein et al. (2019), drawing the impact of these abatement sources from corresponding $\beta(2, 2)$ ranges. Complementary policies are ignored in Québec since, during the 2013-2020 period, the province implemented few analogous large-scale complementary policies, approved few offsets projects, and was a net exporter of electricity in all years. By summing price non-responsive abatement and price-responsive abatement, I construct the overall distribution of regional abatement supply. I plot the mean, 2.5th percentile, and 97.5th percentile of the distribution of these regional abatement supply curves in Figure 4.

Buchsbaum (2022). For electricity with complementary policies in play, (-0.6, -0.9), ibid. These β draws can be drawn independently or with some degree of correlation. In practice, this has no meaningful effect on results. Uncertainty in price-responsive abatement is not a major determinant of allowance prices because the uncertainty of BAU emissions and of price non-responsive abatement are substantially larger.

Price Collars and Market Clearing Regional BAU forecasts and abatement supply curves combine to define regional emissions demand. Specifically, emissions demand in i is $D_i(p) \equiv BAU_i - A_i(p)$, where p is the allowance price. Because A_i increases monotonically, D_i is monotonically decreasing. In the linked market, the unique market-clearing price satisfies:

$$D_{CA}(p) + D_{QC}(p) = q_{CA} + q_{QC} \tag{7}$$

where q_i denotes regional allowance stocks. I then apply a price collaring function $\kappa \equiv \kappa(p)$ in accordance with the published ARP and APCR levels. As a consequence, regional emissions are a function of collared prices rather than market-clearing prices:

$$e_i = BAU_i - A_i(\kappa) \tag{8}$$

The CAQC joint market explicitly calculates net allowance imports v_i as proportional to the total number of allowances supplied. Allowance flow typically satisfies $v_i = q_i - e_i$, but if the auction reserve price (ARP) binds, not all allowances are sold and a more general formula is needed. Policy makers foresaw this possibility and developed a formula to calculate allowance flows in such cases. Denoting Québec's proportion of total supplied allowances as $\rho \equiv \frac{q_{QC}}{q_{CA}+q_{QC}}$(Government of Québec and California Air Resources Board, 2022), allowance flow is calculated as a function of actual emissions $e^{TOT} \equiv e_{QC} + e_{CA}$:

$$\begin{aligned} v_{QC} &= \rho \cdot e^{TOT} - e_{QC} \\ v_{CA} &= (1-\rho)e^{TOT} - e_{CA} \end{aligned} \tag{9}$$

When prices are not at the ARP floor, $e^{TOT} = q_{QC} + q_{CA}$ and (9) collapses to the conventional allowance flow definition. Observe also that $v_{QC} + v_{CA} = 0$ since markets clear.

I summarize the outcomes of market clearing in Figure 5 and Table 3 for 1,000 draws of BAU forecasts and 1,000 abatement supply curves. In Figure 5, I plot the estimated PDF of BAU emissions along with the joint market's collared abatement supply curve (bounded at the 2.5th and 97.5th percentiles). In expectation, the linked market is at its price floor

99.5% of the time and at its ceiling 0.1% of the time.[24] Expected allowance flows are 109 MM T from California to Québec.

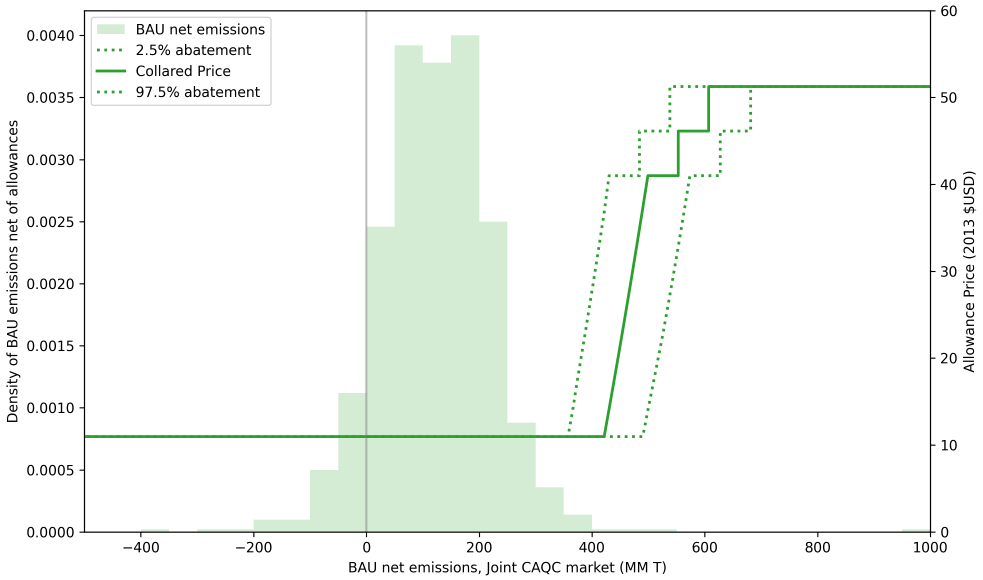


Figure 5: Net emissions and abatement supply for the joint California-Québec market, 2013-2020.

Linked price ($USD 2013)	QC Emissions (MM T)	CA Emissions (MM T)	Pr(Floor)	Pr(Ceiling)
11.07 (1.54)	362.99 (52.26)	2,131.32 (101.71)	0.995	0.001

Table 3: Summary of linked market outcomes, 2013-2020, 1,000 draws. Pr(ceiling) sums the probabilities of the market being at any of the three APCR price ceilings.

Ceteris paribus autarky Because I have constructed regional emissions demand curves, it is easy to study the counterfactual effects of unlinking the California and Québec cap-and-trade markets. In *ceteris paribus* autarky, each region's stock of allowances remains unchanged. Thus, each region's market-clearing price satisfies $D_i(p_i^{AUT}) = q_i$ and regional emissions are a function of collared regional prices; $e_i^{AUT} \equiv BAU_i - A_i(\kappa(p_i^{AUT}))$.

[24]These percentages are somewhat higher than Borenstein et al. (2019), who estimate the market as being at its floor 94.3% of the time. This difference stems predominantly from different underlying datasets for Californian emissions in the 1990-2000 period. Borenstein et al. (2019) apply a manual haircut to CARB-reported emissions during this period, which are in turn faulty as CARB estimates were all collected retroactively beginning in 2004. I instead utilize the federally-published SEDS data, which more precisely measure state-level trends each year. Using SEDS data produces a BAU distribution with similar mean to the CARB data but smaller variance, features which reduce the probabilities of non-price-floor outcomes.

I summarize results from this exercise in Figure 6 and Table 4 over the same sample of BAU forecasts and abatement supply curves used in the linked market analysis. *Ceteris paribus* autarky reveals stark asymmetries in market clearing outcomes across the two regions. The autarkic California market is at the price floor with even higher probability (99.9% of the time), while the autarkic Québec market is rarely at a price floor. In fact, the autarkic Québec market is constrained by a price ceiling 74.5% of the time. This result suggests that the allowance supply exceeded allowance demand only in California.

	Price ($USD 2013)	Emissions (MM T)	Pr(Floor)	Pr(Ceiling)
QC	44.88 (0.45)	453.57 (39.56)	0.000	0.745
CA	11.00 (0.01)	2,022.42 (98.23)	0.999	0.000

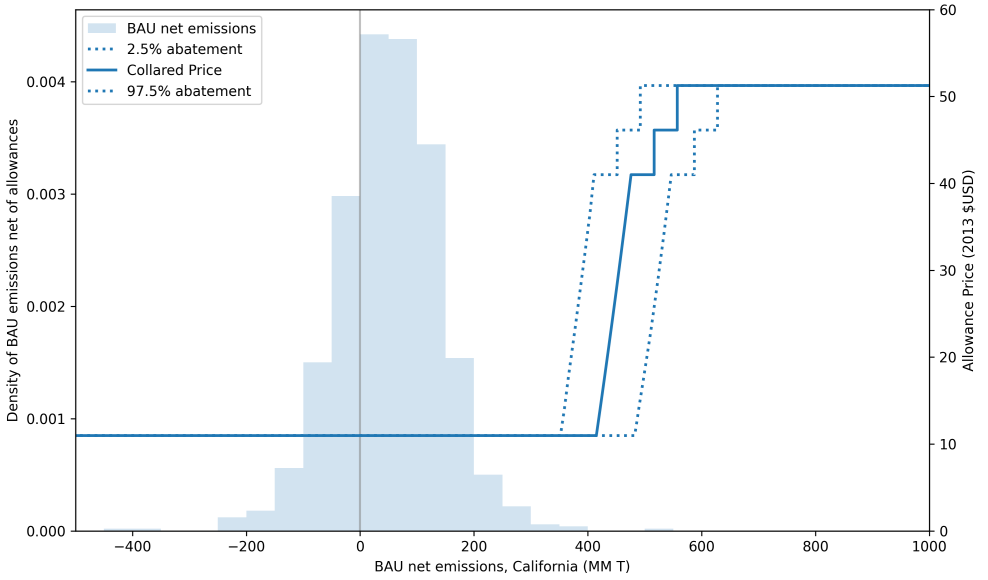
Table 4: Summary of *ceteris paribus* autarky market outcomes, 2013-2020, 1,000 draws. Pr(ceiling) sums the probabilities of the market being at any of the three APCR price ceilings.

The consequences of one-sided allowance oversupply have a significant impact on emissions and welfare. Differencing total emissions between the linked market and *ceteris paribus* autarky reveals an expected increase of 18.3 MM T CO2e over the 2013-2020 period. This finding – linking cap-and-trade programs increased emissions – is driven precisely by California's one-sided allowance oversupply. As the preceding analysis demonstrates, the California market was going to be at a price floor with high probability regardless of its linking decision. By linking with California, the price of carbon in Québec was dragged down to the California price floor. This price decrease mechanically increases emissions demand in Québec and therefore global emissions.

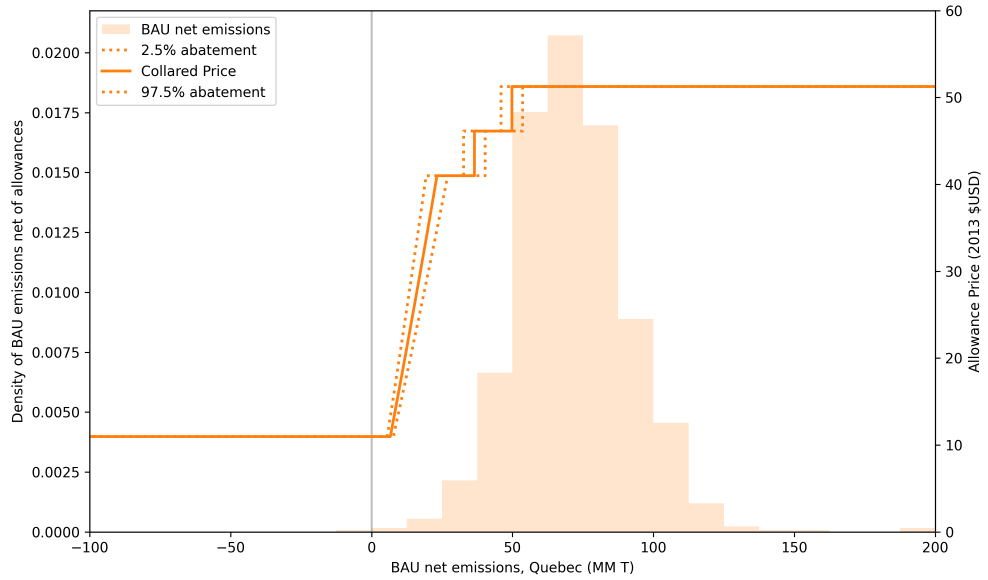
These emissions increases were also welfare-decreasing. In principle, emissions increases need not decrease welfare if the gains from trade and market integration are sufficiently large. This does not appear to be the case in the context of California-Québec. Variable regional abatement costs are:

$$C_i(A_i) \equiv \int_0^{A_i} p_i(s)\, ds \tag{10}$$

In Table 5, I report expected abatement costs by region in linked markets and *ceteris paribus* autarky. The total cost decrease of $523 million is driven by cost decreases in Québec, a combination of decreased overall abatement and gains from trade. Given the expected emissions increase of 18.3 MM T CO2e, linking would have been welfare-improving only if the social cost of carbon were less than or equal to $28.53 / T – well below even



(a) California



(b) Québec

Figure 6: Net emissions and abatement supply for California and Québec, *ceteris paribus autarky*, 2013-2020.

contemporaneous estimates of the SCC (Interagency Working Group, 2013).[25] At an SCC of $190 (Ricke et al., 2018), welfare losses from linking are $348 million per year.[26]

	Autarky ($MM)	Linked ($MM)	Change ($MM)
QC	562.54 (179.11)	36.69 (22.12)	-525.84 (178.90)
CA	136.26 (42.29)	139.38 (80.97)	3.11 (55.95)
Combined	698.80 (185.12)	176.07 (102.58)	-522.73 (191.14)

Table 5: Summary of expected abatement costs in *ceteris paribus* autarky and linked markets, 2013-2020, 1,000 draws. All values in 2013 USD.

V Counterfactual Policy

The preceding section finds that, relative to *ceteris paribus* autarky, linking the California and Québec emissions markets increased expected aggregate emissions by 18.3 MM T CO2e. At the same time, linking these markets generated sizable gains from trade. In principle, any market provision that reduces emissions damages while preserving gains from trade would be welfare-improving. Evaluating such reforms is difficult, however, because the precise welfare effects of emissions reductions depend on assumptions about the social cost of carbon. I circumvent this issue by studying a set of *emissions-neutral market provisions*, defined as modifications to the joint market that restore expected aggregate emissions to their *ceteris paribus* autarky level. By holding aggregate emissions fixed, I can compare the relative cost-effectiveness of various provisions without taking a position on the social cost of carbon.

I consider three classes of provisions: increases in the joint price floor, reductions in the total allowance supply, and tariff-like instruments. Adjustments to the prices and quantities of allowances relate to the longstanding debate over the relative performance of price-based versus quantity-based regulation (Weitzman, 1974; Goulder and Schein, 2013). Tariff-like instruments, whereby one region unilaterally introduces policy, study climate policy under partial coordination, a growing literature that includes carbon border adjustment mechanisms (Clausing et al., 2025; Clausing and Wolfram, 2023) and linking with trading ratios

[25] This assumes constant marginal damages, in which case the breakeven SCC satisfies $\sum_i \Delta C_i \; / \; \sum_i \Delta e_i$. If damages are convex, the breakeven SCC is even lower.

[26] 190$USD 2017 are, once adjusted for inflation, equivalent to 180.57$USD 2013. https://www.in2013dollars.com/us/inflation/2017?endYear=2013&amount=190

(Woerman, 2023). I first describe the mechanics of each policy and then compare their performance. These provisions differ both in their cost-effectiveness and in their distributional consequences.

Price floor increases An increase in the joint market's auction reserve price raises the minimum allowance price and therefore increases total expected abatement. At the same time, a higher floor reduces the probability of an interior equilibrium because the range of prices over which the market clears without binding constraints becomes smaller. The emissions-neutral price floor increase is a new collaring function $\tilde{\kappa}$ such that total expected emissions equal their *ceteris paribus* autarky level:

$$\mathbb{E}[\sum_i e_i^{AUT}] = \mathbb{E}[\sum_i BAU_i - \sum A_i(\tilde{\kappa}(p))]$$

In the preferred specification, the emissions-neutral $\tilde{\kappa}$ has a price floor of \$17.78.

Allowance quantity reduction Reducing allowance quantities lowers the likelihood that the market hits its price floor and raises the likelihood it hits its ceiling. Total expected abatement increases, though the effect on the likelihood of an interior solution is ambiguous.

The new market-clearing price $\tilde{p}$ for some quantity reduction ΔQ satisfies:

$$\sum_i D_i(\tilde{p}) = \sum_i q_i - \Delta Q$$

The emissions-neutral ΔQ^* is such that total expected emissions equal their *ceteris paribus* autarky level:

$$\mathbb{E}[\sum_i e_i^{AUT}] = \mathbb{E}[\sum_i BAU_i - \sum A_i(\kappa(\tilde{p}^*)]$$

In the preferred specification, the emissions-neutral Q^* is 276 MM T.

The distributional effects of a quantity reduction depend on which region retires the allowances. Recall from (9) that net allowance flows are proportional to the stock of supplied allowances. A region that reduces its stock of allowances also increases its net allowance inflow. Because these flows are revenue neutral, they do not affect total welfare, but they matter for incidence. I consider three ways that quantity reduction might be accomplished: with California unilaterally retiring all allowances; with Québec retiring all allowances, and with the allowances being retired proportionally between both regions so that ρ does not change.

Tariff-like instruments Tariff-like instruments impose a tax on allowances entering Québec, raising the effective allowance price faced by Québec emitters. This increases abatement in Québec and raises total expected abatement, but it also induces an inefficient reallocation of emissions toward California. As a result, we should expect emissions-neutral tariff-like instruments to be less cost-effective than the bilateral price floors increases or quantity reductions discussed above. I consider both specific and ad valorem tariffs.

Specific tariffs: Under a specific (per-unit) tariff schedule, allowances imported into Québec face a flat surcharge τ. This policy is analogous to the European Union's carbon border adjustment mechanism (CBAM), which imposes tariffs on emissions-intensive imports. Conceptually, we might imagine a specific tariff on allowances as a further "upstream" CBAM since emissions are a production input for goods like steel and electricity. Following the logic of Shapiro (2021), moving the tariff upstream targets emissions more directly and may reduce downstream distortions.

For a specific tariff τ, market-clearing prices satisfy: $D_{CA}(\tilde{p}) + D_{QC}(\tilde{p} + \tau) = \sum_i q_i$

The emissions-neutral τ^* solves:

$$\mathbb{E}[\sum_i e_i^{AUT}] = \mathbb{E}[\sum_i BAU_i - A_{CA}(\kappa(\tilde{p}^*)) - A_{QC}(\kappa(\tilde{p}^*) + \tau^*)]$$

In the preferred specification, the emissions-neutral specific tariff τ^* is \$33.47.

Trading ratios and ad-valorem tariffs: Trading ratios have received limited but growing attention in the cap-and-trade literature (Woerman, 2023; Holland and Yates, 2015). Operating similarly to currency market exchange rates, a trading ratio of r between Québec and California indicates that r California allowances are equivalent, for compliance purposes, to one Québec allowance. No-arbitrage conditions then imply that the price of Québec allowances under a trading ratio are r times that of California allowances. In this way, the trading ratio is isomorphic to an ad valorem tariff of $r - 1$ at the level of aggregate welfare. Unlike an ad valorem tariff, however, trading ratios also mechanically alter net allowance flows by favoring the region with the higher effective allowance price: $\tilde{\rho} = \frac{r \cdot q_{QC}}{r \cdot q_{QC} + q_{CA}}$. I adopt the trading ratio notation for this manuscript.

For a trading ratio r, market-clearing prices satisfy:

$$D_{CA}(\tilde{p}) + D_{QC}(r \cdot \tilde{p}) = \sum_i q_i$$

The emissions-neutral r^* solves:

$$\mathbb{E}[\sum_i e_i^{AUT}] = \mathbb{E}[\sum_i BAU_i - A_{CA}(\kappa(\tilde{p}^*)) - A_{QC}(r^* \cdot \kappa(\tilde{p}^*))]$$

In the preferred specification, the emissions-neutral trading ratio r^* is 3.7.

Effects on costs and welfare Since all of the preceding provisions are emissions-neutral, the welfare increase of each provision versus autarky equals the decrease in total abatement costs. In the top panel of Figure 7, I show the aggregate cost impacts of each provision relative to autarky. Total costs decrease for all provisions, reflecting gains from trade. Price floor increases are most cost-effective ($260 MM over the eight-year period), followed by quantity reductions ($92 MM). Tariff-like instruments are the least cost-effective provisions (trading ratios, $19.6 MM; specific tariffs, $19.5 MM), since tariffs distort prices across regions and induce an inefficient division of abatement.[27] The dashed line represents the decrease in total abatement costs relative to autarky in the default (no-provision) linked market ($523 MM). Total abatement costs necessarily increase relative to the default linked market because total abatement increases under all provisions.

In the bottom panel of Figure 7, I show the distribution of abatement costs by region and provision. Québec is a distributional loser for all emissions-neutral provisions, bearing more cost than both autarky and the default linked market. Since Québec is a net importer of allowances, any provision that raises the effective price of allowances also increases Québec's allowance trade deficit. From the Québec perspective, then, the most favorable emissions-neutral provisions are unilateral allowance reductions by California and the introduction of trading ratios. Both of these provisions assign Québec a greater proportion of total supplied allowances ρ, thus decreasing Québec's net imports and shifting most of the additional abatement burden onto California.

Overall, this section shows that the increase in emissions under the default linked market was not inevitable. Emissions-neutral market reforms would have increased global welfare by between $19.5 MM (specific tariffs) and $260 MM (price floor increases) versus autarky. Combining these gains with the estimated annual welfare losses of $347 MM from Section IV implies that, relative to the default linked market, the emissions-neutral provisions considered here would have raised total annual welfare by approximately $350 to $380 MM. By extension, many emissions-reducing reforms would also have been cost-neutral or even cost-saving relative to the default linked market, and therefore welfare-improving versus autarky for sufficiently large SCC estimates.

[27] The comparative advantage of ad-valorem tariffs over specific tariffs parallels classic results by Suits and Musgrave (1953) and Bishop (1968) documenting the advantages of ad-valorem taxation over specific taxation.

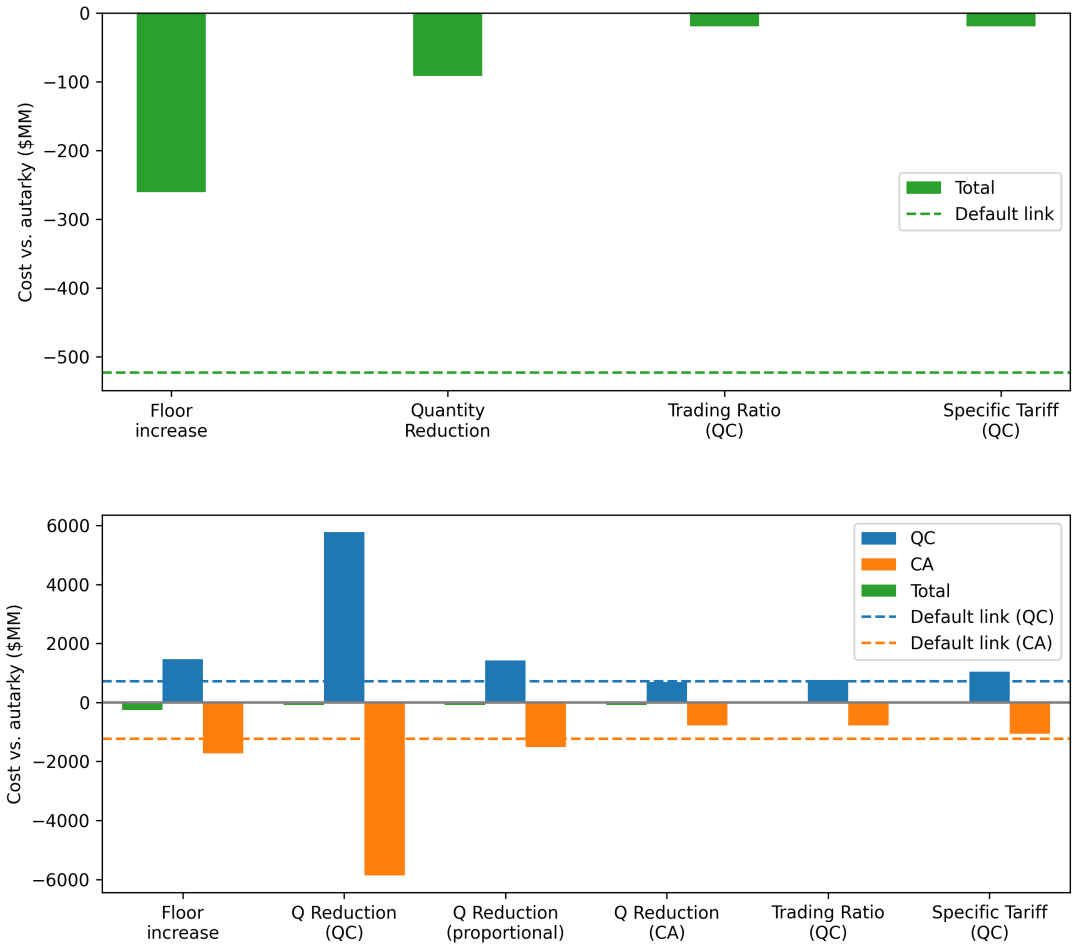


Figure 7: Expected relative abatement costs (versus autarky) across emissions-neutral market provisions. Expected abatement costs in the default (no-provision) linked market are shown in dashed lines. All values in 2013 USD.

VI Feasible Counterfactuals: Policy Maker Incentives and Frictions

The preceding analysis shows that many alternative designs would have improved welfare relative to the observed linked market. An immediate follow-up question is then: if better policies were possible, why were they not adopted?

One explanation argues that policy inertia made welfare-improving provisions impossible. Revising program rules may have been politically costly, time-consuming, or institutionally complex. This interpretation is hard to square with the evidence: prices sat transparently at the floor for nearly eight years, and this was both visible to regulators and widely discussed. Moreover, since these binding collars were ex-ante identifiable, it is unclear why policy makers would have chosen default linking when designing the program.

A second explanation argues that policy makers lacked the ambition to implement more efficient policy. This interpretation emphasizes policy makers' motivations as the drivers of policy, yet canonical economic motivations struggle to explain the observed outcomes. Section IV shows that linking increased global emissions and reduced global welfare for most plausible social cost of carbon values; a Québec policy maker maximizing global welfare would not have chosen to link under these conditions. At the same time, Section V shows that Québec was a distributional loser from linking, paying higher abatement costs while also experiencing higher emissions. A Québec policy maker maximizing regional welfare would likewise have declined to link. Taken seriously, the ambition explanation therefore implies that Québec legislators were maximizing neither global welfare nor regional welfare.

A broader conceptual approach is necessary. I develop a formal framework that models policy makers as distinct *types*, defined by their objectives and the set of policies available to them. I then show how these types can be partially identified from observable outcomes under additional assumptions, and argue that policy makers in CAQC are best understood as *declared-target cost-minimizers*. Finally, I return to the market provisions in Section V and assess which provisions are feasible for this type of policy maker.

Model Policy maker i faces a finite choice set of actions $\alpha_i \subset \mathcal{A}_i$; the total space of actions is $\alpha = \prod_i \alpha_i$. A policy maker earns payoffs according to an **objective function** $f_i : \alpha \to \mathbb{R}$; $f_i \in \mathcal{F}_i$. Define the *type* θ_i of policy maker i as his corresponding pair (f_i, α_i). Thus, policy maker i's type encodes both what he wants to do (f_i) and what he is able to do (α_i).

This means that each policy maker has a best-response action $\alpha_i^* : \alpha_{-i} \to \alpha_i$:

$$\alpha_i^*(\alpha_{-i}|\theta_i) = \arg\max_{a \in \alpha_i} f_i(\alpha_i, \alpha_{-i})$$

Assume best responses are single-valued so that the Nash equilibrium is unique. Across all policy makers, the Nash equilibrium $\alpha^* : \prod_i (\mathcal{F}_i \times \alpha_i) \to \alpha$ takes the form:

$$\alpha^* \equiv \left(\alpha_i^*(\alpha_{-i}^* | \theta_i) \right)$$

Partial identification of types Suppose the econometrician does not observe f_i. Her ability to identify f_i – and, by extension, the type θ_i – depends on her knowledge of the choice set α_i. She observes that action α_i^* belongs to α_i, but she may not know which other actions were actually available.

If the econometrician's only information about α_i is that it contains α_i^*, the econometrician cannot uniquely identify policy makers' types. For instance, both the types $(f_i, \alpha_i = \{\alpha_i^*\})$ and $(f_i = \mathbb{I}(\alpha_i = \alpha_i^*), \alpha_i)$ can rationalize outcome α^*. $(f_i, \alpha_i = \{\alpha_i^*\})$ represents an arbitrary objective function when the policy maker's only available choice is $\{\alpha^*\}$. $(f_i = \mathbb{I}(\alpha_i = \alpha_i^*), \alpha_i)$ represents an arbitrary choice set containing α_i^* and the objective function "choose α_i^*".

At another extreme, full knowledge of α_i can point-identify f_i with sufficient restrictions on the set of admissible objective functions. This is the logic behind much of the discrete choice literature (McFadden, 1974; Berry et al., 1995). When the researcher observes the full menu of alternatives and specifies a structured preference system, observed data can be used to estimate the underlying parameters that govern behavior.

This paper charts a middle course. Econometricians seldom observe the complete menu of options facing policy makers in practice. Political bargaining occurs largely in private, institutional constraints are only partly visible, and many proposals never advance to any official stage of consideration. Without visibility into these alternatives, point identification is impossible. Nonetheless, institutional processes usually reveal a subset of feasible options. Public debate, competing bills, and plausible policy variants often reveal at least some credible alternatives.

The approach in this paper uses these credible alternatives to develop a partial identification strategy that relies on weaker assumptions. I now describe those assumptions.

A1 **Infallibility**: Policy makers do not make mistakes. They select, from all available policies, the one that maximizes their expected objective function.

A2 **Symmetric Information**: Types are common knowledge among all policy makers.

A3 **Immutable Types**: Types remain fixed over time.

A4 **Known Outside Option (KOO)**: Alongside the observed outcome α^*, at least one unselected outside option $\alpha_i' \neq \alpha_i^*, \alpha_i' \in \alpha_i$ is well-defined and known to the econometrician.

A1 (infallibility) is a standard assumption of hedonic economic models. Infallibility crucially does not rule out inefficient objective functions or systematic mistake-making (both of which can be encoded within f_i); rather, it rules out idiosyncratic processes where a policy maker does not implement his most-preferred feasible policy.[28] In the words of Cullenward and Victor (2020), infallibility asserts that "Troubles with creating effective market-based strategies...do not reflect policy makers' failure to understand how markets work." A2 (symmetry) restricts solutions to pure-strategy equilibria, avoiding the need to model mixed strategies. A3 (immutability) similarly reduces the problem to a static problem, avoiding the need to solve a dynamic problem. Symmetry and immutability could be relaxed within a richer model, but this is not the focus of this paper.

Assumption A4, the known outside option, is essential for partial identification.[29] Specifically, the econometrician's knowledge that i does *not* choose the KOO α_i' provides her with partial information regarding f_i and a way to test various f_i for feasibility.

Both objective functions f_i and choice sets α_i can be tested using this framework. To do so, potential choice sets must be tested jointly alongside objective functions using the Incentive Compatibility Hypothesis Test (ICHT):

Proposition 1 (Incentive Compatibility Hypothesis Test)**.** *The econometrician ought to reject any hypothesized type (f_i, α_i') where, for some KOO $\alpha_i' \in \alpha_i'$,*

$$f_i(\alpha_i', \alpha_{-i}^*(\alpha_i')) > f_i(\alpha_i^*, \alpha_{-i}^*)$$

Proposition 1 (The ICHT) states that the econometrician should reject the hypothesis "i's objective is f_i" if the KOO α_i' is a feasible and profitable deviation for the f_i-minded policy maker. This logic readily applies to the preceding results: the core finding of Section IV encourages the econometrician to reject the hypothesis, "Québec was a global welfare-maximizing policy maker with the option to enforce *ceteris paribus* autarky" because such a planner could have enforced *ceteris paribus* autarky and raised global welfare.

At this point, two brief notes on the limitations of the ICHT are warranted. First, the ICHT is ad hoc. It relies on the econometrician to propose candidate types, creating the risk

[28] The version of A1 used here is strong due to data sparsity; with richer observed choice data, even unobserved, idiosyncratic preference shocks can be added. Such assumptions allow models like McFadden (1974) and Berry et al. (1995) to recover consumer preferences from market-level shares.

[29] Work by Manski (2007) considers a similar question in a consumer choice setting. Molinari (2020) provides an overview of scholarship in this area.

that the true type is never tested. Second, the ICHT provides only partial identification. It allows us to reject certain hypotheses about policy maker behavior, but other incorrect types will remain unrejected.[30] For these reasons, the ICHT should not be viewed as a substitute for stronger identification strategies when richer data environments are available. Rather, the strength of the ICHT is its ability to rule out otherwise plausible interpretations of policy maker behavior in settings where data are comparatively sparse.

Hypothesis tests of policy maker types I proceed now by testing seven candidate policy maker types using the ICHT. In all cases, the region's KOO is a variant of autarkic cap-and-trade. I differentiate this autarky option across candidate types as either *constrained* or *unconstrained*. In unconstrained autarky, the policy maker is free to choose any emissions cap q_i, potentially selecting different caps under autarky and linking.[31] For this type of planner, observed choices in the linked market reveal information about the value of q_i that maximizes the policy maker's objective. In constrained autarky, the regional emissions cap q_i cannot change between autarky and the linked market. Such constraints reflect policy inertia or institutional frictions that make revising targets costly. In general, constrained types are assumed to design policy to satisfy domestic objectives under autarky, not in the linked market.

I simultaneously consider three classes of policy maker objectives. A first possibility is that policy makers *maximize regional welfare*, weighting only outcomes within their own jurisdiction. A second possibility is that they *maximize global welfare*, weighting welfare in all regions equally. A third possibility is that policy makers *minimize costs* of reducing emissions below an emissions target $\bar{e}_i$, If $\bar{e}_i$ coincides with the socially efficient emissions level, the cost-minimizing planner also maximizes global welfare. However, cost-minimizing planners are permitted to set $\bar{e}_i$ at inefficient levels.

With two autarky environments (constrained and unconstrained) and three objective functions, one might expect six policy maker types. The count increases to seven because constrained cost minimization can be understood two ways. One version interprets $\bar{e}_i$ as the region's *expected* emissions under an autarky counterfactual when q_i allowances are issued. A second version interprets $\bar{e}_i$ as each region's officially-*declared* emissions goals, which are looser than expected autarky emissions in California and stricter than expected autarky

[30] The type mentioned earlier, $(f_i = \mathbb{I}(\alpha_i = \alpha_i^*), \alpha_i)$, is one such unfalsifiable example.

[31] Although one could allow complementary policies to adjust alongside allowance quantities, that richer environment is not the focus here.

emissions in Québec.[32] I term these types *expected-target cost minimization* and *declared-target cost minimization*, respectively.

Each planner type θ is characterized by a structural parameter, denoted $\psi(\theta)$. $\psi(\theta)$ is a sufficient statistic in the sense of Chetty (2009); concretely, ψ encapsulates both (i) how regions choose allowance quantities and (ii) how regions calculate their objective function. Consider an unconstrained global welfare-maximizing planner, whose objective is equivalently stated as minimizing the sum of abatement costs and emissions damages ($\sum_i C_i(A_i) + \psi \cdot e^{TOT}$), where ψ represent his perceived global marginal damages. This planner chooses quantities so that global marginal damages equal global marginal abatement costs. When markets are linked, he sets quantities so that the expected linked market price equals marginal damages: $\mathbb{E}[p] = \psi$. In autarky, meanwhile, this planner chooses allowance quantities so that global marginal damages equal regional marginal abatement costs: $C_i'(A_i^{AUT}) = \psi$. Using these two equilibrium conditions, ψ can be calibrated from the observed linked market and then used to estimate quantities in the autarky counterfactual.[33] I follow this process for all seven candidate types. The specific moment conditions needed to identify $\psi(\theta)$ for each type of policy maker are detailed in Appendix D.

Table 6 reports Wilcoxon signed-rank tests of the joint hypothesis that the objectives of California (CA) and Québec (QC) are greater under default linking than under autarky. Of the seven candidate types considered, only the cost minimization hypotheses remain unrejected. Of these hypotheses, only *declared-target cost minimization* is consistently unrejected across specifications (Appendix D). The Wilcoxon test nonparametrically compares median differences between autarky and default linked outcomes. The test assumes that differences are symmetrically distributed around the median and that observations are independent across simulation draws; rejection of the null implies that both regions' objectives are systematically higher under linking. Failure to reject suggests that the observed differences could be attributed to chance variation, and therefore corresponds with a rejection of the underlying type assumption. Test statistics take values of 0 or 1 when the empirical medians are sufficiently far apart.

Motivating declared-target cost minimization One may worry that the declared-target cost-minimizing type was effectively selected ex-post to rationalize observed outcomes. This concern is valid and merits consideration. Here, I argue that declared-target cost minimization reflects the division of responsibilities between lawmakers and regulators. Law-

[32] This occurs since California's autarkic market usually prices at a price floor whereas Québec's autarkic market prices at its ceiling.

[33] In a richer data context, for instance Bombardini et al. (2023), ψ can be estimated from richer choice data rather than calibrated.

Type	Objective	Wilcoxon Test Statistic
Unconstrained	*Global Welfare*	1.0 *Reject*
Constrained		1.0 *Reject*
Unconstrained	*Regional Welfare*	1.0 *Reject*
Constrained		1.0 *Reject*
Unconstrained	*Regional Abatement Cost*	1.0 *Reject*
Constrained, expected-target		0.04 *Do not reject*
Constrained, declared-target		0.00 *Do not reject*

Table 6: Results of Wilcoxon sign-rank tests for seven types of policymakers. The joint hypothesis tested posits that the objectives of both California and Québec attain greater values under linking than under the KOO autarky option.

makers, typically led by an executive, set emissions targets and define the broad policy framework, while regulators are charged with implementing those targets using the tools authorized by statute. Within this division of responsibilities, three entwined frictions impede efficient climate policy: (i) Lawmakers set domestic emissions targets at inefficient levels; (ii) Policy inertia makes target revision very costly; and (iii) Regulators still aim to achieve these targets as cost-effectively as possible.

Lawmakers set domestic emissions targets that are often economically inefficient. As shown above, prices in the linked CAQC market were well below standard estimates of the social cost of carbon, implying substantial underpricing of damages. There are at least six reasons why this could have occured. Gaarder et al. (2026) caution that the "effectiveness of discretionary climate policies can be severely undermined by decision-makers' non-emission motives." These motives include: (1) distributional concerns, i.e., an aversion to imposing concentrated costs on particular industries, regions, or constituencies (Hughes and Urpelainen, 2015; Mildenberger, 2020); (2) ambition signaling via simple and symbolically ambitious goals (Collins et al., 2005); (3) increased policy durability by building difficult-to-reverse institutions and commitments (Rosenbloom et al., 2019); and (4) short-term political incentives, which create pressure to adopt initiatives that generate visible, immediate benefits while postponing most of the costs and difficulties (Cullenward and Victor, 2020). Beyond

non-emissions motives, institutional and information frictions can further hamper efficient target setting. These include: (5) Limited administrative and institutional capacity, which may constrain the kinds of policies that can be realistically designed and enforced (Meckling and Nahm, 2018); and (6) early legislative target setting, since emissions targets are typically chosen before reliable emissions data are available. In an extreme example, California legislators did not even know the state's 1990 emissions when they proclaimed that "statewide greenhouse gas emissions shall be reduced to the 1990 level by 2020." (AB 32) It is difficult to imagine that targets chosen under such conditions would consistently coincide with global efficiency.

Policy inertia then makes target revision costly and unlikely. The division of governmental responsibility means that lawmakers and regulators must agree to enact revisions together.[34] For lawmakers, climate policy is often driven by the efforts of highly motivated executives. These executives may be out of power by the time revisions are needed (indeed, neither Governor Schwarzenegger nor Premier Charest led their respective regions by the time cap-and-trade markets opened in 2013); revising legislation may not be politically desirable or beneficial to successive governments. Furthermore, revisiting earlier decisions requires reopening statutes, re-engaging stakeholders, and re-exposing the policy to political conflict, which are all risky and costly actions.

Within these constraints, regulators still try to achieve mandated goals at the lowest possible cost. Statutes in both California and Québec explicitly direct regulators to pursue "cost-effective" or "least-cost" implementation (California Air Resources Board, 2018; Assemblée nationale du Québec, 2024). Agencies regularly evaluate complementary programs using cost-benefit metrics and reject or scale back proposals that fail cost-effectiveness criteria (California Air Resources Board, 2018; Régie de lénergie, 2015).

Feasible counterfactuals I re-run the counterfactual analysis of emissions-neutral provisions for declared-target, cost-minimizing policy makers and present results in Figure 8. I plot expected *perceived* costs, that is, costs according to a declared-target cost minimizer's objective function. Regional expected perceived costs are lower under default linking (dashed lines) than in autarky, consistent with the result that these policy makers prefer linking markets to autarky. Provisions are *feasible* for a given region if the region's perceived cost under the provision is lower than both autarky and default linking.

[34]Even when granted broad latitude, regulators are able to unilaterally revise the caps of cap-and-trade programs. For instance, Senate Bill X1-2 granted CARB broad authority to design accounting methods, compliance mechanisms, and allowance distribution procedures for the cap-and-trade program – but it did not empower regulators to adjust the state's 2020 emissions goal.

If transfer payments are possible, increasing the price floor is jointly feasible; in Section V I showed that price floor increases are also the most cost-effective provision. When transfer payments are not possible, however, distributional concerns prevent price floor increases and all other emissions-neutral provisions from being jointly feasible. Transfer payments require bilateral negotiation, which may be difficult, costly, and slow in practice. This makes unilateral approaches appealing insofar as they only need to be acceptable to one jurisdiction to implement.

For this reason, Québec-led trading ratios merit further consideration. Trading ratios expand the supply of California allowances available at the price floor; although Québec imports more allowances once trading ratios are introduced, those imports tend to remain priced at the floor, reducing Québec's perceived costs relative to default linking. A Québec regulator acting unilaterally might then find trading ratios attractive more attractive than other market provisions. Despite this promise, unilateral policies carry risks. First, unilateral provisions may be inherently less durable. California is worse off under trading ratios than under default linking, making the policy vulnerable to reversal. In dynamic contexts, unilateral trading ratios can lead the entire market to unravel. Second, unilateral provisions fail to internalize foreign benefits from linking. In Figure 9, I show that Québec's regional cost-minimizing trading ratio is 2.8, considerably lower than the emissions-neutral trading ratio of 3.7. Emissions fall compared with default linking but still exceed autarky, and welfare remains below the cooperative benchmark. Such a trading ratio would increase annual welfare by $238 MM relative to the default linked market, roughly three-fifths of total potential welfare gains. The relative welfare gains from emissions-neutral price floor increases, emissions-neutral trading ratios, Québec-led unilateral trading ratios, and default linking are illustrated in Figure 10.

VII Discussion

This paper documents the adverse consequences of linking hybrid cap-and-trade markets. Adverse outcomes are not inevitable: a range of market provisions can both avoid emissions increases while preserving gains from trade, and many of these provisions are politically feasible through either bilateral or unilateral action. The results are especially relevant as California and Québec consider further linking cap-and-trade markets with Washington State, as the European Union contemplates re-linking with the United Kingdom, and more broadly as policy makers explore coordination of decentralized climate policy.

Climate policy is often designed by policy makers who prioritize achieving emissions targets rather than equating marginal benefits and marginal damages of actual emissions.

Adverse outcomes result in such contexts because policy is not made "on the margin". An open question, then, is whether policy makers can be motivated to implement efficient policy. If so, investments in information, analytical tools, or institutional design may prove worthwhile. If not, a careful study of the institutional and political constraints that shape climate policy is essential. Even superficially similar goals – "reduce emissions by X," "emissions must be X," and "reduce emissions by X%" – may imply very different policy choices and outcomes.

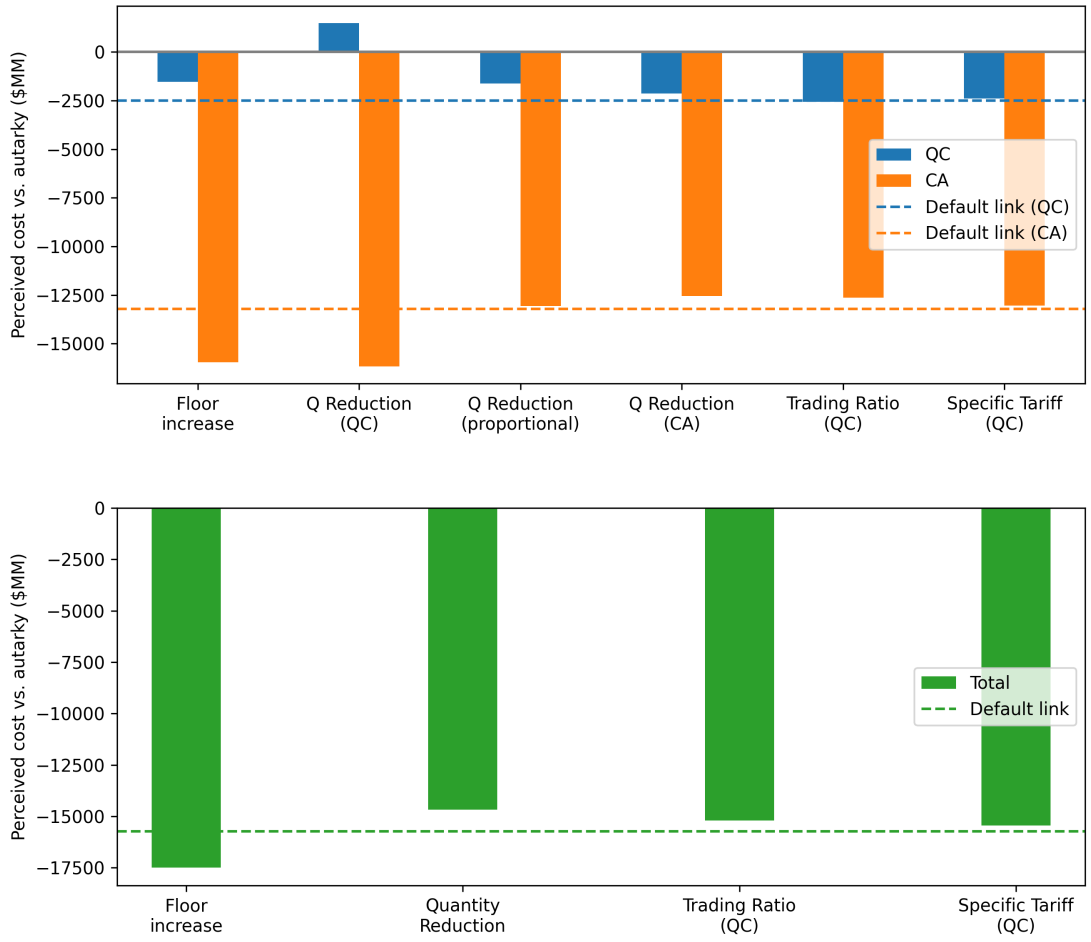


Figure 8: Expected relative perceived abatement costs (versus autarky) across emissions-neutral market provisions for declared-target cost-minimizing policy makers. Expected perceived abatement costs in the default (emissions-increasing) linked market are shown in dashed lines. For a provision to be feasible, it must be lower-cost than both ceteris paribus autarky and default linking. All values in 2013 USD.

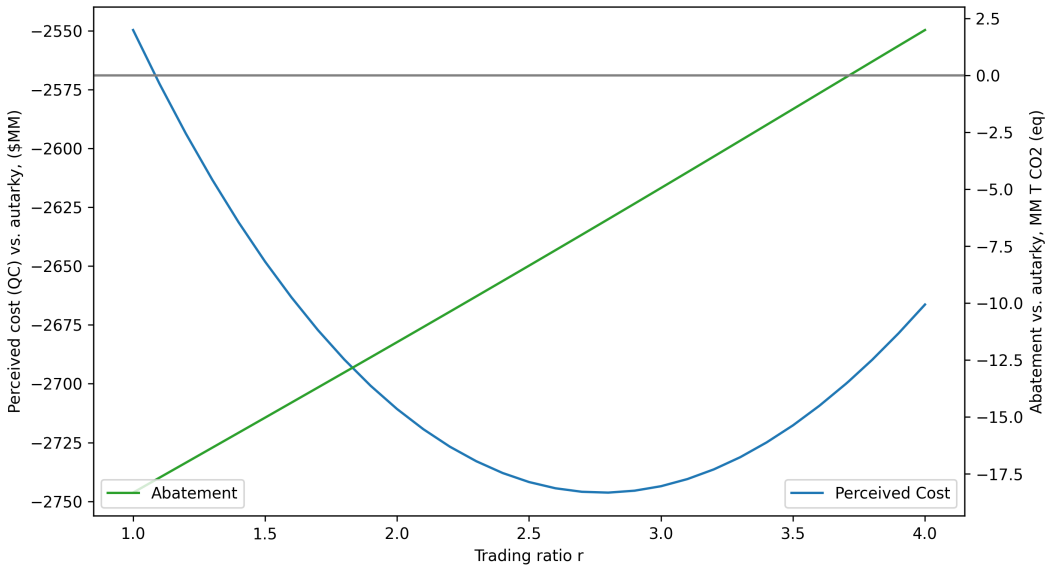


Figure 9: Expected relative perceived abatement costs in Québec and global abatement (versus autarky) for a range of trading ratios. Québec's perceived cost is minimized at a trading ratio of 2.8, which is less than the emissions-neutral trading ratio of 3.7

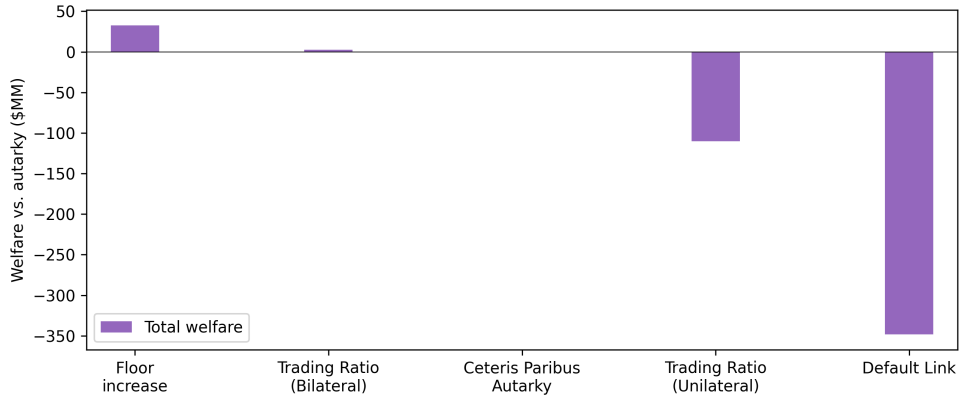


Figure 10: Expected joint welfare (versus ceteris paribus autarky) for a range of market provisions at SCC $190.

References

Agrawal, D. R., Brueckner, J. K., and Brülhart, M. (2024). Fiscal federalism in the twenty-first century. Annual Review of Economics, 16.

Assemblée nationale du Québec (2024). Projet de loi 69, loi assurant la gouvernance responsable des ressources énergétiques, 43^{e} législature.

Baumol, W. J. and Oates, W. E. (1988). The theory of environmental policy. Cambridge University Press.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. Econometrica, 63(4):841–890.

Bishop, R. L. (1968). The effects of specific and ad valorem taxes. The Quarterly Journal of Economics, 82(2):198–218.

Bombardini, M., Li, B., and Trebbi, F. (2023). Did us politicians expect the china shock? American Economic Review, 113(1):174–209.

Borenstein, S., Bushnell, J., Wolak, F. A., and Zaragoza-Watkins, M. (2019). Expecting the unexpected: Emissions uncertainty and environmental market design. American Economic Review, 109(11):3953–3977.

Buchsbaum, J. (2022). Long-run price elasticities and mechanisms: Empirical evidence from residential electricity consumers. Energy Institute Working Paper, 331.

Bushnell, J. and Smith, A. (2025). Modeling uncertainty in climate policy: An application to the us inflation reduction act. Environmental and Energy Policy and the Economy, 6(1):4–41.

California Air Resources Board (2018). Ab 32 global warming solutions act of 2006 ,[fact sheet].

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

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. Annu. Rev. Econ., 1(1):451–488.

Clausing, K. A., Colmer, J. M., Hsiao, A., and Wolfram, C. (2025). The global effects of carbon border adjustment mechanisms.

Clausing, K. A. and Wolfram, C. (2023). Carbon border adjustments, climate clubs, and subsidy races when climate policies vary. Journal of Economic Perspectives, 37(3):137–162.

Coase, R. (1960). The problem of social cost. The Journal of Law and Economics, 3:1–44.

Collins, J. C., Collins, J., Porras, J. I., and Porras, J. (2005). Built to last: Successful habits of visionary companies. Random House.

Coutu, E. (2016). Incertitude des taux de change dans le système de plafonnement et déchange avec bornes de prix: le cas du marché du carbone entre la californie et le québec. Master's thesis, Université du Québec à Montréal. English translation by Charles Seguin: Exchange rate uncertainty in the cap-and-trade system with price collars: the case of the carbon market between California and Québec.

Cullenward, D. and Victor, D. G. (2020). Making climate policy work. John Wiley & Sons.

Environment and Climate Change Canada (2016). Technical update to environment and climate change canadas social cost of greenhouse gas estimates. Cat. No.: En14-202/2016E-PDF.

Fowlie, M. and Muller, N. (2019). Market-based emissions regulation when damages vary across sources: What are the gains from differentiation? Journal of the Association of Environmental and Resource Economists, 6(3):593–632.

Foy, H. (2026). Eu agrees on review of emissions trading system as energy price concerns grow. Financial Times. Accessed: 2026-03-25.

Gaarder, I., Grindaker, M., Meling, T. G., and Mogstad, M. (2026). Green waste. Working Paper 34649, National Bureau of Economic Research.

Goulder, L. H. and Schein, A. R. (2013). Carbon taxes versus cap and trade: a critical review. Climate Change Economics, 4(03):1350010.

Government of Québec (2021). Report on the achievement of Québec's greenhouse gas emissions reduction target for 2020.

Government of Québec and California Air Resources Board (2022). Accounting mechanism for article 8 of the 2017 linkage agreement. Developed pursuant to Article 8 of the Agreement on the Harmonization and Integration of Cap-and-Trade Programs.

Green, J. F. (2021). Does carbon pricing reduce emissions? a review of ex-post analyses. Environmental Research Letters, 16(4):043004.

Helm, C. (2003). International emissions trading with endogenous allowance choices. Journal of Public Economics, 87(12):2737–2747.

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

Hughes, J. E., Knittel, C. R., and Sperling, D. (2008). Evidence of a shift in the short-run price elasticity of gasoline demand. The Energy Journal, 29(1):113–134.

Hughes, L. and Urpelainen, J. (2015). Interests, institutions, and climate policy: Explaining the choice of policy instruments for the energy sector. Environmental Science & Policy, 54:52–63.

Interagency Working Group (2013). Technical support document:-technical update of the social cost of carbon for regulatory impact analysis-under executive order 12866. Environmental Protection Agency.

International Carbon Action Partnership (2025). Emissions trading worldwide: Status report 2025.

Jaffe, J., Ranson, M., and Stavins, R. N. (2009). Linking tradable permit systems: A key element of emerging international climate policy architecture. Ecology LQ, 36:789.

Manski, C. F. (2007). Partial identification of counterfactual choice probabilities. International Economic Review, 48(4):1393–1410.

McFadden, D. (1974). The measurement of urban travel demand. Journal of public economics, 3(4):303–328.

Meckling, J. and Nahm, J. (2018). The power of process: State capacity and climate policy. Governance, 31(4):741–757.

Mildenberger, M. (2020). Carbon captured: How business and labor control climate politics. MiT Press.

Molinari, F. (2020). Microeconometrics with partial identification. Handbook of econometrics, 7:355–486.

Montgomery, W. D. (1972). Markets in licenses and efficient pollution control programs. Journal of economic theory, 5(3):395–418.

Muller, N. Z. and Mendelsohn, R. (2009). Efficient pollution regulation: getting the prices right. American Economic Review, 99(5):1714–1739.

Ostrom, E. (1990). Governing the commons: The evolution of institutions for collective action. Cambridge university press.

Pigou, A. C. (1920). The Economics of Welfare. Macmillan, London.

Pizer, W. A. (2002). Combining price and quantity controls to mitigate global climate change. Journal of public economics, 85(3):409–434.

Pizer, W. A. and Yates, A. J. (2015). Terminating links between emission trading programs. Journal of Environmental Economics and Management, 71:142–159.

Rabe, B. G. (2016). The durability of carbon cap-and-trade policy. Governance, 29(1):103–119.

Régie de lénergie (2015). Demande – r-3924-2015. April 17, 2015.

Ricke, K., Drouet, L., Caldeira, K., and Tavoni, M. (2018). Country-level social cost of carbon. Nature Climate Change, 8:895–900.

Rosenbloom, D., Meadowcroft, J., and Cashore, B. (2019). Stability and climate policy? harnessing insights on path dependence, policy feedback, and transition pathways. Energy Research & Social Science, 50:168–178.

Rubin, E. and Auffhammer, M. (2024). Quantifying heterogeneity in the price elasticity of residential natural gas. Journal of the Association of Environmental and Resource Economists, 11(2):319–357.

Schmalensee, R. and Stavins, R. N. (2017). Lessons learned from three decades of experience with cap and trade. Review of Environmental Economics and Policy.

Shapiro, J. S. (2021). The environmental bias of trade policy. The Quarterly Journal of Economics, 136(2):831–886.

Stavins, R. N. (1998). What can we learn from the grand policy experiment? lessons from so2 allowance trading. Journal of Economic perspectives, 12(3):69–88.

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In Handbook of macroeconomics, volume 2, pages 415–525. Elsevier.

Suits, D. B. and Musgrave, R. A. (1953). Ad valorem and unit taxes compared. The Quarterly Journal of Economics, 67(4):598–604.

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

Woerman, M. (2023). Linking carbon markets with different initial conditions. Journal of Environmental Economics and Management, 119:102820.

World Bank (2025). State and Trends of Carbon Pricing 2025. World Bank, Washington, DC. License: Creative Commons Attribution CC BY 3.0 IGO.

A Theoretical Model: Proofs and Extensions

A.1 Cases

I. Two interior solutions When both regions achieve interior market clearing in autarky, $p_i = p_i^\circ$ and $p_j = p_j^\circ$. The linked market must also achieve an interior solution. As a result, linking has no effect on emissions ($\Delta e = 0$). At the same time, the expression for ΔC simplifies:

$$\Delta C = -\frac{A_{1i}A_{1j}}{2(A_{1i} + A_{1j})}(p_i^\circ - p_j^\circ)^2 \leq 0 \tag{A11}$$

This reproduces two familiar results: first, costs are weakly decreasing when markets link due to gains from trade; second, the slopes A_1 are sufficient information to estimate these gains from trade.

Since $\Delta e = 0$ and $\Delta C \leq 0$, $\Delta WF \geq 0$

II. Two floors When both regions are at a price floor in autarky($p_i = p_j = \underline{\kappa}$), the linked market's price is also at its floor ($p = \underline{\kappa}$). Total emissions are $\bar{e}_i + \bar{e}_i$ whether linked or unlinked; thus, $\Delta e = 0$, $\Delta C = 0$ and $\Delta WF = 0$.

Intuitively, two markets at their price floors are indistinguishable from an emissions tax of $\underline{\kappa}$. There is no aggregate welfare gain from these regions conducting the same emissions tax jointly rather than independently.

III. Two ceilings By analogous logic to case II, neither costs nor emissions change when both linking markets are already at their price ceilings. $\Delta e = 0$, $\Delta C = 0$ and $\Delta WF = 0$.

IV. One floor, one interior solution Suppose $p_i = \underline{\kappa}$ for some allowance quantity $q_i > \bar{e}_i$ while $p_j = p_j^\circ$, implying $q_j < \bar{e}_j$. Emissions in autarky are $\bar{e}_i + q_j$. The linked market either achieves an interior solution or attains the price floor; I consider both cases here:

If the linked market is at its floor ($p = \underline{\kappa}$), the price $p = \underline{\kappa}$ and linked emissions e are $\bar{e}_i + \bar{e}_j$. Since $q_j < \bar{e}_j$, $\Delta e > 0$. The change in costs is given by (4a):

$$\Delta C = \frac{A_{1j}}{2}(\underline{\kappa}^2 - p_j^2) \tag{A12}$$

$\Delta C < 0$ since $p_j > \underline{\kappa}$.

If the linked market is at an interior price p^*, linked emissions are $e = q_i + q_j$. Since $q_i > \bar{e}_i$, $\Delta e > 0$. Equation (4b) simplifies to:

$$\Delta C = -\frac{A_{1i}A_{1j}}{2(A_{1i} + A_{1j})}(p_i^\circ - p_j^\circ)^2 + \frac{A_{1i}}{2}(p_i^{\circ 2} - \underline{\kappa}^2) \tag{A13}$$

Since $q_i > \bar{e}_i$, $p_i^\circ < \underline{\kappa}$. Thus, all three terms in (A13) are negative and $\Delta C < 0$.

Thus, regardless of the ultimate price of the linked market, linking causes an increase in total emissions and a decrease in total abatement costs. The effect on welfare is ambiguous.

V. One ceiling, one interior solution Suppose i attains the ceiling of $\bar{\kappa}$ for some allowance quantity $q_i < \underline{e}_i$. j is at an interior price, implying $q_j > \underline{e}_j$. Emissions in autarky are $\underline{e}_i + q_j$. The linked market either achieves an interior solution or attains the price floor; I consider both cases just as before:

If the linked market is at its ceiling, the price $p = \bar{\kappa}$ and linked emissions e are $\underline{e}_i + \underline{e}_j$. Since $q_j > \underline{e}_j$, $\Delta e < 0$. The change in costs is given by (4a):

$$\Delta C = \frac{A_{1j}}{2}(\bar{\kappa}^2 - p_j^2) \tag{A14}$$

$\Delta C > 0$ since $p_j > \underline{\kappa}$.

If the linked market is at an interior price, linked emissions are $e = q_i + q_j$. Since $q_i < \underline{e}_i$, $\Delta e < 0$. Simplifying Equation (4b) as before:

$$\Delta C = -\frac{A_{1i}A_{1j}}{2(A_{1i} + A_{1j})}(p_i^\circ - p_j^\circ)^2 + \frac{A_{1i}}{2}(p_i^{\circ 2} - \bar{\kappa}^2) \tag{A15}$$

Since $q_i < \underline{e}_i$, $p_i^\circ > \bar{\kappa}$. The first term of (A15) is negative while the latter two terms are positive; thus the total effect on ΔC is ambiguous.

These cases reveal a difference relative to the one floor, one interior solution case. While total emissions increase during linking, the effect on total costs is ambiguous. When the linked market is at a price ceiling, costs surely increase; at interior prices, these cost increases are potentially offset by gains from trade.

VI. One floor, one ceiling When one market is at its floor and the other is at its ceiling, linking has ambiguous effects on emissions, costs, and welfare. This is easily demonstrated by two examples, in which I assume i to be at the floor and j to be at a ceiling. If the linked market falls to the price floor, the only effect is a reduction of allowance prices in j. This causes total emissions to increase and total costs to decrease. Analogously, if the linked market attains its price ceiling, its only effect is an increase in allowance prices in i, which decrease emissions and increase total costs.

A.2 General Model

Define region i's general abatement supply function

$$A_i(p) = \begin{cases} A_{0i}, & \text{if } p < 0, \\ A_{0i} + A_{1i}(p), & \text{if } p \geq 0, \end{cases}$$

with $A_{1i}(0) = 0$, $A'_{1i}(p) > 0$, and $A''_{1i}(p) > 0$.

The total cost of achieving emissions q_i is

$$C_i(q_i) = \int_0^{s_i} A_{1i}^{-1}(a)\, da = \int_0^{p_i} t\, A'_{1i}(t)\, dt.$$

The change in total cost from linking is then

$$\Delta C = \sum_{s \in \{i,j\}} \int_{p_s}^{p} t A'_{1s}(t)\, dt. \tag{A16}$$

Expand this equation around the interior price points p, p_s°:

$$\Delta C = \sum_{s \in \{i,j\}} \int_{p_s}^{p_s^\circ} t A'_{1s}(t)\, dt + \sum_{s \in \{i,j\}} \int_{p_s^\circ}^{p^\circ} t A'_{1s}(t)\, dt + \sum_{s \in \{i,j\}} \int_{p^\circ}^{p} t A'_{1s}(t)\, dt.$$

For an interior equilibrium $p = p^\circ$, the last term vanishes and I recover an expression analogous to (4b):

$$\Delta C = \sum_{s \in \{i,j\}} \int_{p_s}^{p_s^\circ} t A'_{1s}(t)\, dt + \sum_{s \in \{i,j\}} \int_{p_s^\circ}^{p^\circ} t A'_{1s}(t)\, dt. \tag{A17}$$

The first term has ambiguous sign: it is positive at a ceiling, negative at a floor, and zero in the interior. The second term reflects gains from trade and is negative whenever $p_i^\circ \neq p_j^\circ$.[35] The results of Section A.1 continue to hold for these more general functional forms because $A_i(p)$ is monotonically increasing.

[35] To prove this, fix total abatement at $s_i + s_j$. The interior linked price p° equalizes marginal costs across regions and therefore minimizes $C_i + C_j$ subject to $A_i(p_i) + A_j(p_j) = s_i + s_j$, with strict convexity of A implying a unique minimum at (p°, p°). Since (p_i°, p_j°) also satisfies this constraint but does not equalize marginal costs when $p_i^\circ \neq p_j^\circ$, it must yield weakly higher total cost.

B A History of California and Québec Climate Policy

In 2006, California lawmakers, led by Republican governor Arnold Schwarzenegger, passed Assembly Bill 32 (AB 32) and set a goal of achieving 1990 levels by 2020. The state's main environmental regulatory agency, the California Air Resources Board (CARB), was tasked with developing a cap-and-trade program as part of this bill. Later that year, Governor Schwarzenegger attempted to position California as an international climate leader by helping to launch the Western Climate Initiative (WCI). The WCI envisioned widespread interstate climate cooperation. In February 2007, a bipartisan coalition of governors from California and four other western U.S States (Arizona, New Mexico, Oregon, and Washington) signed a nonbinding agreement "directing their respective states to develop a regional target for reducing greenhouse gas emissions, participate in a multi-state registry to track and manage greenhouse gas emissions in the region, and develop a market-based program to reach the target." (WCI, 2007). This idea proved politically popular, and by the end of 2008 eleven U.S. states and Canadian provinces had joined the WCI. WCI members were expected to build their own cap-and-trade systems by quantifying the cap, designing the emissions-regulatory framework, and securing legislative authorization for regulators to administer and enforce the market.

In California, policy makers first had to construct historical emissions and future emissions forecasts because neither were known when AB 32 passed.Voluntary emissions reporting was available through the California Climate Action Registry (SB 1771, 2001), though this was known by lawmakers to be incomplete (CARB, 2008b). CARB published a 1990-2004 greenhouse gas emissions inventory in 2007 and set the 2020 limit at 427 MMT CO2e (CARB, 2018). CARB's first Scoping Plan, published at the end of 2008, outlined preliminary business-as-usual 2020 emissions forecasts (approximately 600 MMT CO2e), emphasizing that "Reducing greenhouse gas emissions to 1990 levels means cutting approximately 30 percent from business-as-usual emission levels projected for 2020." (CARB, 2008a)

Unlike California, Québec policy makers had access to more than 15 years of emissions history. The Canadian federal government began tracking province-level sectoral emissions after signing the Kyoto Protocol in 1990. Premier Jean Charest and his governing Liberal Party had previously used this data when developing Québec's First Climate Plan in 2006. Premier Charest's 2009 Decree 1187-2009 set Québec's 2020 emissions target at 20% below 1990 emissions levels. This goal was predicated on international cooperation: Québec policy makers noted, as early as 2009, that "it was possible to reduce [domestic] greenhouse gas emissions in the province to approximately % below the 1990 baseline, meaning that Québec had to achieve international reductions if it wanted to set a more ambitious target." (Gov-

ernment of Québec, 2021). With these targets in hand, the Québec parliament unanimously passed an amendment to the Environment Quality Act in June 2009 empowering regulators to create a cap-and-trade market.

California and Québec were uniquely able to pass cap-and-trade legislation. Initial enthusiasm in other states waned quickly. As Rabe (2016) notes, "Much of the WCI alliance hinged on agreements between like-minded state governors and provincial premiers and often lacked legislative buy-in or supportive legislation." As governors left office, and as lawmakers failed to pass regional carbon market bills, the WCI eroded. Even in California, opponents sought to suspend AB 32 through Proposition 23 in 2010 - a measure voters decisively rejected. In 2011, all other US states and the province of Manitoba exited the WCI, leaving only California, Québec, and a few other Canadian provinces (S&P Global, 2011). Though the provinces of Ontario (exited 2018), British Columbia (exited 2020), and Nova Scotia (exited 2024) remained members of the WCI, none developed consequential cap-and-trade programs at any point.

This exodus meant that most policy design was performed by Californian and Québec regulators, who collaborated to draft rule language, held industry workshops, and held public hearings related to all emissions reductions initiatives. California defined covered and uncovered emissions sources for the cap-and-trade program. All major emissions sectors were covered, including electricity, natural gas, and transportation - a property rare among operating emissions markets and a property Québec found especially attractive (Decree No. 378-2008).[36] Complete regulations governing cap-and-trade markets passed through each region's legislature in 2011. However, debates over the rules for carbon offsets delayed full cap-and-trade deployment until these rules were finalized in 2012, and the first carbon offset credits were issued in California in September 2013 (California Governor's Office, 2013).

Price collars were the most impactful rule added during the policy design phase of development. The collars ensured that all allowances were sold above an Auction Reserve Price (a price floor) that increased each year; similarly, a stock of Auction Containment Reserves were set aside and would only be disbursed if prices exceeded pre-stipulated price thresholds. In effect, this made the CAQC cap-and-trade market effectively a hybrid of a cap-and-trade market and a carbon tax. The theoretical efficiency of such a hybrid market has been subject to much debate within the economics literature (Goulder and Schein, 2013; Weitzman, 2014). As mentioned above, though, the price floors in the joint market meant that Québec's cap-and-trade market was dragged down to the binding price floor of the

[36]Waste management, agriculture, and forestry remained uncovered "due to challenges associated with precise measurement, tracking or sector structure." (CARB 2008b)

California market. Until 2020, prices remained at or near the price floor. Policy makers did not react immediately to this phenomenon by reducing allowance stocks.[37]

Subsequent renewals of the cap-and-trade program have followed similar timing. In 2015, Québec set a 2030 emissions target of 37.5% below 1990 levels (Décret 1018-2015) and a 2050 target of 80-95% below 1990 levels (later revised into a net-zero 2050 goal). California set similarly ambitious targets in 2016, with a 2030 emissions target of 40% below 1990 levels (Executive Order B-30-15, Senate Bill 32[13]) and carbon neutrality by 2045 (Executive Order B-55-18). Lawmakers in both regions renewed the cap-and-trade program through 2030 (California: AB 398[14]; Québec: Order in Council 1126-2017). A second renewal effort is currently underway, and California stated its intent to renew the cap-and-trade program through 2040 earlier this year (AB 1207 and SB 840). Québec is expected to sign a similar intent to renew later this year (IETA, 2025). Washington State is also expected to link its cap-and-trade market with the joint CAQC market in the future, though negotiations are not yet finalized. (RFF, 2025) Complementary Environmental Policies Both California and Québec implemented many complementary environmental policies alongside their cap-and-trade programs. These programs led to large-scale emissions reductions and bear mentioning, since these additional sources of abatement further reduced emissions demand. California's first Scoping Plan also called for expanding complementary environmental policies, including California's Renewable Portfolio Standard, the Low Carbon Fuel Standard, Vehicle Emissions Standards, rooftop-solar subsidies (CARB, 2008b), and carbon offsets programs. CARB estimated that existing complementary policies would reduce emissions by 15% but that expansion could help the state better meet its 2020 goals. California's internal rulemaking concluded in 2011, meaning CARB was able to begin enforcing other policies (the Renewable Portfolio Standard, the Low Carbon Fuel Standard) in those years. Québec began pursuing complementary climate policies as early as the 1990s, when the province successfully transitioned almost all of the province's electricity sector to hydroelectric power. By the time Québec passed its First Climate Plan in 2006, the government outlined dozens of policy instruments, including adoption of California vehicle standards, subsidization of energy-efficient air heaters and water heaters, investment in wind and hydroelectric generation, and a CAD$720 million investment in public transit.

[37]California unilaterally retired approximately 12 MM tons worth of allowances from 2016-2017, a small number compared with the 2800 MM tons covered emissions. Prices remained at the price floor even after this retirement.

C DFMs and alternate specifications

C.1 Data Cleaning

California Emissions Forthcoming.

Québec VMT Official vehicle miles traveled (VMT) series for Québec are only available for the period 1999-2009 and are therefore insufficient for the full sample period.[38] I do not use these reported VMT data. Instead, I construct a consistent VMT series by combining fuel sales data with estimates of average vehicle fuel economy.

Specifically, I use annual gasoline and diesel sales volumes for Québec, measured in liters, from Statistics Canada for the period 1987-2023.[39] These fuel volumes are converted into implied vehicle kilometers traveled using average fuel consumption rates for light-duty vehicles. Canadian average fuel consumption for light-duty trucks is available through 2008[40]; for 2009-2010, I use U.S. average light-duty fuel economy.[41] This substitution is reasonable given the high degree of integration in North American vehicle markets and the similarity of vehicle fleets over this period.

I convert implied vehicle kilometers traveled to vehicle miles traveled using a conversion factor of 1.609 kilometers per mile. The resulting series provides a continuous and internally consistent measure of Québec VMT over the full sample period (1990-2010) used in the analysis.

C.2 DFM

The purpose of a dynamic factor model is to encapsulate the dominant trends and correlations of a large number of variables $X_t = (X_{1t}^{CA}, X_{1t}^{QC} ... X_{Nt}^{CA}, X_{Nt}^{QC})'$ within a comparatively few number of factors. The specific DFM used in this paper is:

$$\begin{aligned} X_t &= \mu + \Lambda F_t + u_t \\ F_t &= \textstyle\sum_{\ell=1}^{L} A_\ell F_{t-\ell} + B_0 + B_1 t + \varepsilon_t \end{aligned} \tag{C18}$$

I use this DFM to forecast future outcomes using the following procedure:

1. Obtain factors $\hat{F}_t$ from data X_t
2. Fit a "ground-truth" VAR ($\hat{A}_\ell$) on these factors

[38] Statistics Canada Table 23-10-0097-01.
[39] Statistics Canada Table 23-10-0066-01.
[40] TransportPolicy.net. Canada: Light-duty fuel consumption and GHG
[41] Department of Transportation: NHTSA CAFE PIC Report: Fleet Performance Dashboard

3. Generate bootstrapped factor histories $\dot{F}_t$ using the VAR
4. Create factor forecasts $\ddot{F}_{T+k}$ from each factor histories and the ground-truth VAR
5. Use factor forecasts to generate forecasts of $\ddot{X}_{T+k}$ in future periods

C.2.i Obtaining factors

I extract factors from variables using a principal component analysis (PCA). As Bushnell and Smith (2025) note, there are multiple valid normalizations and grouping choices the econometrician might employ when using PCA to summarize a high-dimensional dataset, each with potential strengths and drawbacks.

One approach is principal components normalization, which constructs factors as linear combinations of the full set of observables in X_t (using eigenvectors of the sample covariance matrix, or alternatively the correlation matrix). This choice produces factors that are orthogonal to one another and explain as much of the variation in X_t as possible. However, because each factor typically loads on many series, this normalization often yields factors that are not easily interpretable and therefore difficult to use for structural analysis.

An alternative is named factor normalization. This approach fixes the factor loading matrix so that particular factors correspond to particular variables (or blocks of variables), which improves interpretability. A drawback is that the resulting factors need not be orthogonal, and may exhibit nontrivial correlation with one another. Importantly, this choice affects the normalization of the factors but not the span of the estimated factor space.

I adopt Bushnell and Smith (2025)'s hybrid approach that groups variables by type. For example, I group all motor gasoline price series across regions into a single block and extract the first principal component of that block. This yields factors that are both empirically dominant within each block and economically interpretable.

Within each group, I demean and standardize each series prior to PCA. Let $X_t \in \mathbb{R}^N$ denote the standardized data vector for a given group at time t, for $t = 1, 2, \ldots, T$. The group-specific factor is obtained as the first principal component,

$$\hat{F}_t = \hat{\Phi}' X_t,$$

where $\hat{\Phi} \in \mathbb{R}^N$ is the associated loading vector.

After extracting factors, I estimate the coefficients $(\hat{\mu}, \hat{\Lambda})$ and compute empirical residuals $\{\hat{u}_t\}_{t=1}^T$ from the ARIMA model:

$$X_t = \mu + \Lambda F_t + u_t.$$

C.2.ii Constructing the "Ground Truth" VAR

I model the VAR as a trend-stationary process using a low-frequency trend function. To estimate this, I first construct the array of factor first differences $\Delta\hat{F} = [\Delta\hat{F}_t]'$. Letting Ψ denote the matrix of low-frequency cosine basis functions evaluated at $t = 2, \ldots, T$, the estimated low-frequency trend component of $\Delta\hat{F}_t$ is the least-squares projection

$$\Delta\hat{\tau} = \Psi(\Psi'\Psi)^{-1}\Psi'\Delta\hat{F}$$

The levels trend $\{\hat{\tau}_t\}_{t=1}^T$ is recovered by cumulating differences: $\hat{\tau}_t = \sum_{s=1}^t \Delta\hat{\tau}_s$, with normalization $\hat{\tau}_1 = 0$.

Then, I define the detrended factor series:

$$\hat{f}_t \equiv \hat{F}_t - \hat{\tau}_t$$

The detrended factors $\{\hat{f}_t\}_{t=1}^T$ are also mean-zero since Ψ includes a constant (frequency-0) term. This makes them suitable inputs to a VAR(L), which takes the form

$$\hat{f}_t = \sum_{\ell=1}^{L} \hat{A}_\ell \hat{f}_{t-\ell} + \hat{\varepsilon}_t$$

where VAR coefficients ($\hat{A}_\ell$) summarize how each factor responds over time to shocks in itself and in the other factors, capturing both persistence and cross-factor interactions in the common component of X_t. In the preferred specification, a lag of $L = 1$ is used.

C.2.iii Generating bootstrapped factor histories

Using the "ground-truth" VAR coefficients $\hat{A}_\ell$, I generate a bootstrapped series of detrended factor histories $\{\dot{f}_t\}_{t=1}^T$ according to

$$\dot{f}_t = \sum_{\ell=1}^{L} \hat{A}_\ell \dot{f}_{t-\ell} + \dot{\varepsilon}_t$$

with initial condition $\dot{f}_1 = \hat{f}_1$ and where $\{\dot{\varepsilon}_t\}_{t=1}^T$ are randomly sampled with replacement from VAR residuals $\{\hat{\varepsilon}_t\}$.

At the same time, I generate bootstrapped historical trend components $\{\dot{\tau}_t\}_{t=1}^T$ recursively as

$$\dot{\tau}_t = \dot{\tau}_{t-1} + \hat{\tau}_0 + \Psi^*_{t-1}\dot{z}, \qquad \dot{\tau}_1 = 0$$

where Ψ^*_{t-1} denotes the $(t-1)$th row of low-frequency basis Ψ, excluding the frequency-0 (constant) term. Random draws are distributed $\dot{z} \sim t\Big(0,\ \widehat{\text{Var}}(\Psi^{*'}\Delta\hat{F})\Big)$. The combined factor history is then $\dot{F}_t \equiv \dot{f}_t + \dot{\tau}_t$.

C.2.iv Forecasting future factors and variables

For a given bootstrapped factor history $\{\dot{f}_t\}_1^T$ and $\{\dot{\tau}_t\}_1^T$, I forecast factors and variables over a time horizon h by iterating the "ground truth" VAR forward in time. This process is identical to the preceding section, though I crucially allow for multiple forecasts to be constructed for each factor history. Consequently, the distribution of forecasts reflects two distinct sources of uncertainty: uncertainty about the underlying pre-period factor history $\{\ddot{F}_t\}_{t=1}^T$, and uncertainty about the subsequent evolution of factors after time T.

Factor forecasts are explicitly written as:

$$\ddot{f}_{T+k} = \sum_{\ell=1}^{L} \hat{A}_\ell \ddot{f}_{T+k-\ell} + \ddot{\varepsilon}_{T+k}, \qquad \ddot{f}_T = \dot{f}_T$$

$$\ddot{\tau}_{T+k} = \hat{\tau}_0 + \ddot{\tau}_{T+k-1} + \Psi^*_k \ddot{z} \qquad \ddot{\tau}_T = \dot{\tau}_T$$

As before, $\ddot{\varepsilon}_t$ are randomly drawn with replacement from empirical VAR residueals, and trend noise is distributed $\ddot{z} \sim t\Big(0,\ \widehat{\text{Var}}(\Psi^{*'}\Delta\hat{F})\Big)$. The combined factor forecast is then $\ddot{F}_{T+k} \equiv \ddot{f}_{T+k} + \ddot{\tau}_{T+k}$.

Finally, I recover underlying variables $\ddot{X}_t$ from factor forecasts. Using estimated coefficients $(\hat{\mu}, \hat{\Lambda})$ and drawing idiosyncratic residuals $\ddot{u}_{T+k}$ with replacement from $\{\hat{u}_t\}_{t=1}^T$:

$$\ddot{X}_{T+k} = \hat{\mu} + \hat{\Lambda}\ddot{F}_{T+k} + \ddot{u}_{T+k}$$

ε_t and u_t are drawn separately because the factor structure makes these terms uncorrelated within the model.

D ICHT Moment Conditions

Each candidate policy maker type maximizes his objective functions f_i over a choice set A_i. The first order conditions of these optimization problems are characterized by a deep parameter $\psi(\theta)$. In this section, I detail the specific moment conditions needed to identify $\psi(\theta)$ for each type of policy maker.

Let Π_i represent the regional benefits of emissions. In general, region i's welfare is function of the domestic benefits of emissions $\Pi_i(e_i)$, total actual emissions e^{TOT}, and the

region's net trade balance v_i. Specifically, for regional marginal damages ϕ_i and allowance price p, regional welfare is:

$$WF_i \equiv \Pi_i(e_i) - \phi_i \cdot e^{TOT} + p \cdot v_i$$

Unconstrained regional welfare maximizers Unconstrained policymakers are free to set allowance quantities q_i strategically within linked and unlinked markets. Therefore, observed allowance quantities in the linked market must satisfy policy makers' first-order conditions. If an unconstrained policy maker aims to maximize regional welfare, he sets q_i to maximize WF_i. The first-order condition simplifies because $\Pi_i'(e_i) = p$, the market-clearing price:

$$\phi_i \cdot \frac{de^{TOT}}{dq_i} = p + \frac{dp}{dq_i} v_i$$

The structural parameter $\psi = (\phi_{CA}, \phi_{QC})$ is the vector of all regional marginal damages.

Constrained regional welfare maximizers A constrained policymaker cannot adjust allowance quantities in anticipation of linking. Thus, if region i's policy maker sets allowance allocations q_i° to maximize regional welfare, WF_i, he does so based on expected autarky outcomes. In autarky, there is no allowance flow, so $q_i = e_i$. Thus, $WF_i = \Pi_i(q_i) - \phi_i \cdot e^{TOT}$. Letting p_i° denote each region's autarkic allowance price, the first order condition is then:

$$\phi_i \cdot \frac{de^{TOT}}{dq_i} = p_i^\circ$$

.

The structural parameter $\psi = (\phi_{CA}, \phi_{QC})$ is once again the vector of all regional marginal damages.

Unconstrained global welfare maximizers When policy makers cooperate, they act to jointly maximize global welfare, $WF^{AGG} \equiv \sum \Pi_i(e_i) - \phi^G \cdot e^{TOT}$, where ϕ^G denotes the global marginal damages from emissions. Because policy makers are unconstrained, q_i is chosen to maximize WF^{AGG}, and observed allowance quantities must satisfy policy makers' first-order conditions:

$$\phi^G = p.$$

The structural parameter $\psi = \phi^G$ is the scalar representing global marginal damages.

Constrained global welfare maximizers Similarly to the preceding paragraph, the structural parameter $\psi = \phi^G$ is sufficient to estimate welfare changes for constrained welfare maximizers. indeed, this is the exercise carried out in Section IV. This type assumption predicts policy makers in all regions would set autarkic quantities so that $p_i = \phi^G$ – an assumption refuted by the empirical price distributions of Québec and California autarky allowances. This hypothesis is therefore rejected absent additional assumptions about the source of price differences. Even so, I report welfare estimates for a SCC of \$190 in Table 6, which repeat results from Section IV.

Unconstrained cost minimizers Cost-minimizing policy makers fix an emissions target $\bar{e}_i$ and seek to minimize the cost of achieving it. If $\bar{e}_i$ is set at the globally efficient level, this objective is indistinguishable from global welfare maximization; in general, $\bar{e}_i$ is not required to be set efficiently. Regions count both domestic abatement $A_i(p)$ and net allowance purchases v_i as sources of abatement, and i's total cost of abatement in a linked market is $C_i(A_i) - p \cdot v_i$.

If Region i's emissions in a linked market equilibrium do not equal its emissions target, it must pursue additional abatement through a combination of domestic abatement A_i' and foreign abatement v_i'. To a first-order approximation, the total expected cost of reaching $\bar{e}_i$ is then:

$$C_i(A_i') + (p + \frac{dp}{dq}(v_i' - v_i)) \cdot v_i' \tag{D19}$$

Because policy makers are unconstrained, observed allowance quantities in the linked market must satisfy policy makers' first-order conditions. The expected cost of attaining $\bar{e}_i$ is minimized when the emissions target equals expected linked-market emissions: $\bar{e}_i \equiv \mathbb{E}[BAU_i - A_i(p) + v_i]$. Intuitively, if the linked market delivered emissions above or below the target in expectation, the policymaker could reduce expected compliance costs by adjusting the allowance quantity, since the same emissions outcome could be achieved more cheaply through a different mix of domestic abatement and allowance purchases. When price collars never bind, this condition further implies $\bar{e}_i = q_i$.

The structural parameter $\psi = (\bar{e}_{CA}, \bar{e}_{QC})$ is the vector of all regional emissions targets.

Constrained cost minimizers For constrained cost minimizers, expected abatement costs are calculated identically to the unconstrained cost minimization case. Because policymakers are constrained, they cannot adjust allowance quantities in anticipation of linking, and must instead minimize costs conditional on a fixed target.

For expected-target cost minimizers, this emissions target is expected emissions in autarky: $\bar{e}_i = \mathbb{E}[BAU_i - A_i(p_i^\circ)]$, i.e.

For declared-target cost minimizers, $\bar{e}_i$ is the emissions target stated by policy makers, such as "20% below 1990 emissions levels by 2020". I assume a linear emissions trajectory from 2013 to 2020 and evaluate expected cumulative emissions over this period.

In both cases, the structural parameter $\psi = (\bar{e}_{CA}, \bar{e}_{QC})$ is the vector of all regional emissions targets.