

UC Berkeley

Energy Institute Working Papers

Title

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

Permalink

<https://escholarship.org/uc/item/3bb694qz>

Author

Larson, Thor

Publication Date

2026-10-30



Energy Institute WP 367

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

Thor Larson

October 2026

Energy Institute at Haas working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to review by any editorial board. The Energy Institute acknowledges the generous support it has received from the organizations and individuals listed at <https://haas.berkeley.edu/energy-institute/about/funders/>.

© 2026 by Thor Larson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit is given to the source.

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

Thor Larson

October 22, 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 17.5 million tons CO₂e and reduced welfare by \$331 million annually – an outcome that was ex-ante foreseeable. A reformed linked market could have remained emissions-neutral while instead increasing welfare by up to \$32 million annually. I argue that political and institutional frictions prevented such reforms from being implemented and propose 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.

*Larson: thor_larson@berkeley.edu. I owe special thanks to Jim Saltee, Meredith Fowlie, Severin Borenstein, and Aaron Smith for their feedback, advice, and wisdom throughout. I am also grateful for valuable feedback at key stages of this project from Alan Auerbach, Matt Backus, James Bushnell, Lucas Davis, Jo Ellery, Shachar Kariv, Larry Karp, Ethan Ligon, Aprajit Mahajan, Charles Séguin, Joseph Shapiro, Quitzé Valenzuela-Stookey, and Paige Weber; from seminar participants at UC Berkeley, Stanford, the University of Toronto, UCSB, the Montréal Environmental and Natural Resource Economics Workshop, and the World Congress of Environmental and Resource Economists. ChatGPT-5.5 and Claude Sonnet 4.6 provided valuable proofreading and coding assistance throughout this process.

I Introduction

Policy makers regularly set the terms of market integration, which determine if and how market participants may trade with one another. When trade occurs, it can affect welfare through two channels. First, gains from trade reallocate production toward lower-cost producers (Ricardo, 1817). Second, by changing prices, trade may change the aggregate level of economic activity. If economic activity generates an externality, market integration can then have an ambiguous effect on overall welfare.

This insight has significant implications for environmental regulation, where cap-and-trade programs, also known as emissions trading systems (ETS), are the most widely used market-based approaches to reducing greenhouse gas emissions.¹ Today, more than three dozen regional, national, and supranational cap-and-trade markets operate independently from one another. I refer to this fragmented structure as *carbon autarky*. Scholars have long argued that integrating (“*linking*”) these autarkic markets should generate meaningful gains from trade by allowing emissions reductions to occur where they are least costly (Stavins, 1998).²

This conclusion assumes that cap-and-trade linking cannot affect overall emissions levels. Because today’s operating cap-and-trade programs are better classified as “hybrid” markets, in which prices are *collared* by price floors and price ceilings (Goulder and Schein, 2013),³ market linking affects not just where emissions occur, but also overall emissions levels. Linking hybrid cap-and-trade markets may therefore increase emissions and reduce welfare if the resulting increase in emissions externality damages exceeds the gains from trade.

This paper studies how linking hybrid cap-and-trade markets can adversely affect emissions and welfare, and which market design provisions can mitigate these effects. I develop a framework that quantifies the emissions and welfare consequences of linking, evaluates alternative market designs, and characterizes the political constraints that shape feasible reforms.

¹The United Nations Environment Programme (2025) predicts climate change externalities will reduce annual global GDP by 4% by 2050 (\$4 trillion / yr). World Bank (2025) reports that 28% of global greenhouse gas emissions are covered by 80 market-based instruments. A patchwork of 37 distinct cap-and-trade programs collectively covers 23% of global emissions, while 43 regional carbon taxes collectively cover just over 5% of global emissions.

²When regions A and B link their cap-and-trade programs, they agree that A’s carbon allowances may be used to satisfy compliance obligations in B and vice versa. Beyond cost-efficiency, linking markets may also generate other benefits, including reduced price volatility and increased liquidity (Jaffe et al., 2009).

³Most major cap-and-trade programs incorporate mechanisms that constrain allowance price volatility. Some employ explicit price collars (CAQC, RGGI, and New Zealand), while others adjust allowance supply or otherwise adjust market rules at pre-defined price levels (Korea, the UK). Even when markets lack explicit price collars, they face a natural near-zero price floor and some *de facto* price ceiling, above which policy makers may intervene and revise policy.

I first develop a theoretical framework that explains how binding price collars lead to ambiguous welfare effects. For instance, consider a global market that efficiently prices emissions and a tiny jurisdiction T that zealously over-issues allowances. While unlinked, T attains an inefficiently low level of abatement, but this inefficiency has a small impact on global welfare because T’s emissions footprint is small. When T links with larger markets, however, its surplus allowances can substitute for foreign abatement, lowering carbon prices, reducing global abatement, and increasing total emissions. Linking with T therefore worsens rather than corrects the underlying inefficiency. In a general model, I derive the set of primitives needed to characterize the welfare consequences of linking.

I then apply this framework to 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), the CAQC market faces minimal supranational confounders.⁴ I focus on the program’s first authorization period, which encompasses the eight years from 2013 to 2020.

I model the CAQC market by constructing each region’s allowance demand as the difference between business-as-usual (BAU) emissions and the supply of emissions abatement. This approach builds on the framework of [Borenstein et al. \(2019\)](#), who estimate allowance demand from the gap between baseline emissions and abatement opportunities. However, their empirical strategy is designed for a single market and cannot recover the joint allowance demand curves needed to evaluate how linking affects equilibrium emissions, prices, and welfare. I therefore extend this framework by forecasting BAU emissions using a dynamic factor model (DFM) estimated on historical emissions trends ([Stock and Watson, 2016](#)). The DFM pools information across macroeconomic and sectoral series to capture common shocks and provides a flexible framework for estimating allowance demand across multiple linked markets.⁵ I construct abatement supply as the sum of price-responsive fuel demand reductions across four emissions-intensive fuels and price-non-responsive components such as comple-

⁴Relevant supranational confounders would include policies adopted by higher-level governing authorities that affect emissions throughout the linked market. Examples include EU-wide climate and energy policies across EU member states and U.S. federal policies affecting RGGI states.

⁵The factor structure provides a parsimonious approach that can extend to settings with more than two linked markets by incorporating additional regional emissions series.

mentary regulations, carbon offsets, and emissions reshuffling.⁶ I find that the linked CAQC market is at the price floor with high probability (96.6% in the preferred specification), a result consistent with existing literature.

Estimating regional allowance demand curves allows me to construct counterfactual autarkic and linked-market equilibria. 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 than the linked market (99.4% of the time), whereas the autarkic Québec market is rarely at its price floor. In fact, the autarkic Québec market is constrained by a price *ceiling* 73.2% of the time. When markets link, expected emissions increase by 17.5 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 \$331 million per year.⁷

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 sufficiently increase expected abatement in the linked market to restore expected aggregate emissions to their *ceteris paribus* autarky level. 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.

The results above present a puzzle. Allowance prices remained at the price floor for much of the linked market’s operation, yet policy makers did not adopt readily available reforms that could have eliminated the welfare losses from linking. More strikingly, the welfare losses, emissions increases, and distributional consequences of CAQC linking were broadly predictable *ex ante* from information available to policy makers. Why, then, was the market designed as it was?

Resolving this puzzle requires distinguishing between socially efficient policies and politically incentive-compatible ones. Market provisions that maximize global welfare need not maximize the objectives of the policy makers responsible for adopting them. To understand

⁶Emissions reshuffling describes the reallocation of low-emissions electricity claims toward regulated entities while leaving physical generation unchanged. It arises from incomplete regulatory coverage and can undermine the effectiveness of emissions policies (Bushnell et al., 2008).

⁷Unless otherwise stipulated, all values are reported in 2013 USD. \$1 2013 USD is approximately equivalent to \$1.39 USD in 2026. <https://www.in2013dollars.com/us/inflation/2013?amount=1>

observed market design, therefore, we must first analyze the incentives and constraints that shape policy choices.

I conduct this analysis using two complementary approaches. First, I perform a qualitative review of the policy record, including legislative history, regulatory documents, and contemporaneous statements by policy makers. 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. This suggests that policy makers place significant weight on the cost of reaching their declared emissions targets when evaluating climate policy – and comparatively less weight on emissions damages.

I 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 2.5), emissions exceed autarky levels, though the increase is smaller than under default linking. Welfare gains under unilateral Québec-led trading ratios are \$222 million per year, just over 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.⁸ A related theoretical literature examines the role of price collars and allowance reserves (Fell et al., 2012; Murray et al., 2009), though this work generally focuses on unlinked cap-and-trade systems. The interaction between such price-containment mechanisms and linking has been noted informally (Burtraw et al., 2013), but has not been formally characterized. Furthermore, this paper is the first to analyze the California-Québec system bilaterally; existing work on this mar-

⁸Green (2021) writes: “[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 *delinking*, but do so in a purely theoretical setting.

ket 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 long tradition of research 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 sits within a growing literature that incorporates uncertainty into models of climate policy (Stock and Watson, 2016; Bushnell and Smith, 2025; Borenstein et al., 2019).

The paper also contributes to work on incentive-compatible second-best climate policy. The analysis emphasizes that feasible climate reforms are constrained both by the internal incentives of policy makers, who may pursue distributional, electoral, or institutional objectives beyond aggregate welfare maximization (Grossman and Helpman, 1994), and by the external decentralization of the climate externality, which places emissions damages beyond the jurisdictional reach of any individual regulator (Helm, 2003; Agrawal et al., 2024; Ostrom, 1990). This connects the paper to work on commitment, delegation, and bureaucratic discretion (Besley and Coate, 1998; Acemoglu, 2003; Alesina and Tabellini, 2007). The partial identification exercise extends tools from structural inference to the study of policy maker behavior (McFadden, 1974; Manski, 2007; Bombardini et al., 2023).

II Theoretical model

This section develops a theoretical model of linked cap-and-trade markets with price collars. The model shows that price collars fundamentally alter the welfare effects of linking: while linking without binding collars generates the standard gains from trade by reallocating abatement across regions, linking with collars can also change aggregate emissions and associated externality damages. The overall effect of linking on global welfare therefore depends on the tradeoff between cost savings and emissions changes. I then provide a set of primitives needed to determine the welfare consequences of linking.

II.a Setup

Consider two regions i and j that operate cap-and-trade programs. For an allowance price p , define the abatement supply in region i as $A_i(p)$ and analogously for j . For expositional ease, assume price-responsive abatement is linear in price: $A_i(p) = \alpha_i \cdot p$.⁹ Denote i 's business-as-usual emissions as BAU_i and i 's stock of allowances as q_i .

II.b Unlinked markets

The demand for allowances is the difference between business-as-usual emissions and abatement supply: $D_i(p) \equiv \text{BAU}_i - A_i(p)$. When price collars do not bind, allowance demand equals the stock of emissions allowances:

$$D_i(p) = q_i$$

Define the *slack* in region i as $s_i \equiv \text{BAU}_i - q_i$. Slack measures the excess or shortfall of allowances relative to baseline emissions when carbon prices are zero and price-non-responsive abatement occurs. The uncollared market-clearing price of allowances in autarky, p_i° , is:

$$p_i^\circ(q_i) \equiv \frac{s_i}{\alpha_i} \quad (1)$$

I henceforth suppress the argument q_i for brevity. Price collars – a price floor $\underline{\kappa}$ and a price ceiling $\bar{\kappa}$ – are the same in both markets and constrain the range of values that prices can take. Define the collaring function $\kappa(\cdot)$:

$$\kappa(p) \equiv \begin{cases} \underline{\kappa} & \text{if } p < \underline{\kappa} \\ p & \text{if } \underline{\kappa} \leq p \leq \bar{\kappa} \\ \bar{\kappa} & \text{if } p > \bar{\kappa} \end{cases}$$

Collared prices in the autarkic market are

$$p_i \equiv \kappa(p_i^\circ)$$

Collars similarly constrain emissions $e_i \equiv \text{BAU}_i - A_i(p_i)$.¹⁰

⁹A more general treatment in Appendix A.3 shows that this section's results hold for any abatement supply curve with $A'_i > 0$ and $A''_i > 0$.

¹⁰This model assumes that any unsold allowances are freely disposed of. In many markets, however, unsold allowances are instead moved to a reserve account where they may be sold in future periods. Linking, then, may not increase cumulative emissions in such cases since it also reallocates emissions across time. Whether this intertemporal reshuffling ultimately increases total emissions depends on future slack: so long as future

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

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

In general, the total variable cost of abatement in i reflects the amount of price-responsive abatement needed to absorb market slack:

$$C_i(q_i) \equiv \int_0^{s_i(q_i)} A_i^{-1}(a) da \quad (2)$$

When price-responsive abatement is linear, these costs simplify to:

$$C_i(q_i) = \frac{\alpha_i}{2} p_i^2$$

II.c Linked markets

In a linked market, firms may freely trade allowances across jurisdictions. The uncollared linked market price p^* satisfies:

$$p^* \equiv \frac{s_i + s_j}{\alpha_i + \alpha_j} \quad (3)$$

Prices and emissions in the linked market remain collared:

$$p = \kappa(p^*)$$

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

Total variable abatement costs in the linked market are $C_i(p) + C_j(p)$. The change in abatement costs from linking $\Delta C \equiv (C_i(p) - C_i(p_i)) + (C_j(p) - C_j(p_j))$ simplifies to:

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

allowance supply is perfectly elastic (at either a floor or a ceiling), an unsold allowance reserve acts identically to immediate retirement.

When the linked market attains an interior price ($p = p^*$), (4a) can be reorganized to:

$$\Delta C = -\frac{\alpha_i \alpha_j}{2(\alpha_i + \alpha_j)}(p_i^\circ - p_j^\circ)^2 + \sum_{s \in \{i,j\}} \frac{\alpha_s}{2}(p_s^{\circ 2} - p_s^2) \quad (4b)$$

(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.¹¹

The other object of interest is the change in emissions:

$$\Delta e \equiv e - (e_i + e_j) \quad (5)$$

The change in global welfare, $\Delta W F$, depends on cost changes ΔC and emission changes Δe . Assume emissions are uniformly mixing and incur constant marginal damages ϕ :

$$\Delta W F \equiv -\Delta C - \phi \cdot \Delta e \quad (6)$$

Both ΔC and Δe vary when price collars bind. Figure 1 demonstrates this phenomenon in a simplified example where only region i 's floor has the potential to bind. In panel (a), the floor in region i is nonbinding. Linking generates gains from trade without changing aggregate emissions. In panels (b) and (c), by contrast, region i is at its price floor in autarky and has unsold allowances. Linking brings some of these allowances into circulation, increasing aggregate emissions. Panel (c) presents an extreme case where i 's slack can fully satisfy j 's excess demand, driving the linked price down to region i 's price floor.

More broadly, Figure 2 illustrates the generally ambiguous welfare effects of linking by expressing the welfare effect of linking as a function of the uncollared, unlinked allowance prices p_i° and p_j° . I use these market outcomes, rather than the underlying primitives, because they are the objects that directly determine whether price collars bind. The thick black curves of Figure 2 trace the loci along which linking leaves either emissions or total abatement costs unchanged. The white, “bird-shaped” region denotes the set of autarkic markets in which linking weakly decreases both emissions and abatement costs. In this region, linking is welfare-improving for any positive SCC value. The central white square – the “body” of the bird – is where both autarkic markets attain interior prices. In this region, collars have no effect on market outcomes, and the canonical result of [Stavins \(1998\)](#) applies.

¹¹ A third term, $\sum_{s \in \{i,j\}} \frac{A_s}{2}(p^{*2} - p^2)$, vanishes from this expression because of the assumption $p = p^*$.

Outside the bird, the welfare effects of linking depend on the SCC ϕ . In regions where linking raises aggregate emissions but lowers total abatement costs (light gray), linking is welfare-improving only when ϕ is sufficiently small. In regions where linking lowers aggregate emissions but raises total abatement costs (medium gray), linking is welfare-improving only when ϕ is sufficiently large. In the top-right and bottom-left regions (dark gray), both jurisdictions are constrained by the same price collar regardless of the linking decision. Linking then has no effect on welfare because both jurisdictions effectively face the same carbon tax, leaving no scope for gains from trade. Appendix A.1 provides the detailed construction of Figure 2 and derives the welfare conditions associated with each region.

The figure also illustrates why the expected welfare effect of linking depends on the joint distribution of p_i° and p_j° . This distribution reflects the joint behavior of all variables that determine uncollared prices, including BAU emissions, price-responsive abatement, price-nonresponsive abatement, and allowance quantities. Correlation in any one primitive is not sufficient. What matters is how these primitives jointly map into uncollared prices and how often those prices fall into each welfare region of the figure.

II.d Sufficient primitives for welfare analysis

Marginal emissions damages ϕ and allowance demand $D_i(\cdot)$, and characterize welfare gains given collars $\underline{\kappa}, \bar{\kappa}$ and allowance quantities q_i . It follows that ΔC and Δe are recoverable from $\{D_i(\cdot), \underline{\kappa}, \bar{\kappa}, q_i\}$, and the full welfare change ΔW is recoverable from $\{\phi, D_i(\cdot), \underline{\kappa}, \bar{\kappa}, q_i\}$.

This result follows directly. Equation (2) expresses costs C entirely in terms of these primitives. Emissions e equal either allowance quantities (q_i) or collared emissions ($\bar{e}_i, \underline{e}_i$), and depend only on these primitives either way. Equation (6) expresses welfare in terms of C , e , and ϕ .

This stylized model isolates the core mechanisms through which collars affect welfare by abstracting away many institutional features of real-world cap-and-trade programs. In practice, these institutional features produce nonlinear abatement supply curves $A_i(\cdot)$ whose shape determines welfare gains. My empirical analysis accounts for the institutional features, using observed policy rules and data to construct $A_i(\cdot)$ and evaluate the welfare effects of linking.




Figure 1: Effects of a binding price floor on trade and welfare. Region i , with allowance demand D_i , issues q_i allowances and potentially triggers price floor $\underline{\kappa}$. After i links with region j , markets clear where i 's residual supply of allowances RS_i equals j 's residual allowance demand RD_j ; v_i^* allowances flowing from j to i . Welfare effects are ambiguous and depend on whether the floor binds before and after linking. In (a), the floor does not bind at any point; linking generates gains from trade with no change in aggregate emissions. In (b), the floor binds under autarky but not after linking; aggregate emissions rise, generating additional emissions damages. In (c), the floor binds both before and after linking; emissions increase, and only region j realizes gains from trade.




Figure 2: Effects of linking on global welfare as a function of uncollared, unlinked allowance prices. The “bird-shaped” white region depicts the set of autarkic markets where linking is welfare-improving for any positive SCC value (linking weakly decreases both emissions and total abatement costs).

III History and background

My empirical setting is the California-Québec (CAQC) emissions market. CAQC is the largest linked cap-and-trade market, the longest-operating linked cap-and-trade market, and the world’s fourth-largest cap-and-trade market overall. The CAQC market is also a desirable setting to study linking since fewer confounding factors obscure the effects of integration. 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 the members of other multiregional programs (such as RGGI and the EU ETS), California and Québec engage in comparatively little goods trade.¹² 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.

¹²In 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.

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

Policy history Below, I briefly review the evolution of climate policy in California and Québec (summarized in Figure 3). Appendix B presents a more detailed history.

California authorized a cap-and-trade program in 2006 as part of Assembly Bill 32 (AB 32), 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 (the state’s 1990 emissions had not previously been measured), 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 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 only other 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).¹³ 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 2010. By late 2011, California and

¹³This quote suggests that, had California and Québec not planned to link in 2009, Québec might instead have developed a cap-and-trade program with a less binding cap. In this section, I set this possibility aside. By the time the markets formally linked in 2014, there was little scope for policy amendments, making *ceteris paribus* autarky a plausible counterfactual. I return to this concern in Section VI, where policy makers are given greater latitude in designing policy.

Québec were the only jurisdictions that had adopted cap-and-trade regulations consistent with the WCI framework.

After a period of joint rulemaking, California and Québec launched their cap-and-trade programs in 2013. Joint auctions 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 CAQC linkage agreement was subsequently renewed by both regions through 2030, with further extensions through 2045 and a potential linkage with Washington State currently under consideration.




Figure 3: Timeline of cap-and-trade policy in California and Québec.

Prices, price collars, and allowance stocks Over the 2013-2020 time horizon, a total of 2,793.1 million metric tonnes (MM T) of allowances were issued, with California issuing 2,386.6 MM T and Québec issuing 406.5 MM T. Allowances were freely bankable across time and space with minimal restrictions.¹⁴

Allowance auctions occur quarterly in an auction jointly conducted by CARB and Québec’s Ministère de l’Environnement (MELCCFP). 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

¹⁴Covered entities settle their emissions obligations over multi-year compliance periods, rather than annually. The first three compliance periods are 2013-2014, 2015-2017, and 2018-2020. I follow [Borenstein et al. \(2019\)](#) in treating these periods as economically integrated, given “the relatively generous allowance quantities made available in the early year auctions, and the ability of ARB to shift some additional allowances from later years.” Similarly, I treat the California-Québec market as geographically integrated over this interval: although allowances were initially sold separately in 2013-2014, linking in 2014 made them fully fungible prior to the first compliance surrender deadline and allowed *ex-post* trade.

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.¹⁵ Both ARP and APCR prices increase annually at a rate of 5% plus the rate of inflation. Figure 4 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 or near the ARP for nearly the entire duration of the program’s initial authorization period.

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](#)).




Figure 4: California and Québec average allowance prices, 2014-2020.
Image credit: California Air Resources Board

IV Empirical analysis

This section quantifies the effects of linking on emissions and welfare in the CAQC market using an empirical model. Section II.d identifies the five primitives required for this analysis. Price collars $\bar{\kappa}$, $\underline{\kappa}$ and allowance quantities q_i are set by law and treated as given. I construct

¹⁵California and Québec use different currencies, meaning price floors may vary across regions as exchange rates fluctuate ([Coutu, 2016](#)). This is not a major concern during the period studied, when exchange rates remain near parity (0.995). Moreover, Québec requires that units sold inside its jurisdiction be sold at the higher of the two jurisdictions’ prices 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](#)).

the three components of allowance demand D_i – business-as-usual emissions BAU_i , price-non-responsive abatement $A_{i,NPR}$, and price-responsive abatement $A_{i,PR}$ – using three separate procedures.¹⁶

I construct the components of allowance demand for the 2013-2020 compliance period using *only* information available to regulators at the time of market linking. This *ex-ante* perspective lets me study both the effects of linking and whether or not these effects were foreseeable. Linked allowance prices ultimately remained near the price floor throughout the 2013-2020 period (Figure 4), an outcome I show was highly probable based on information available when the market was established in 2013.

Markets clear at an equilibrium price where allowance supply equals allowance demand: $D_i(p) \equiv \text{BAU}_i - A_{i,NPR} - A_{i,PR}(p)$. We might be tempted to estimate D_i directly from observed market data by studying how emissions respond to variation in allowance prices, but this approach is insufficient 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 allowance 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.

To account for policy and data uncertainty, I adapt the approach of Borenstein et al. (2019), who construct business-as-usual (BAU) emissions forecasts and abatement supply independently and then recover allowance demand as the difference between the two. This decomposition separates the sources of uncertainty that matter most for market outcomes. Their implementation, however, relies on a vector error-correcting model (VECM) whose joint dynamics cannot be credibly estimated from the short annual history available once

¹⁶Abatement $A_i(p)$ satisfies $A_i(p) \equiv A_{i,NPR} + A_{i,PR}(p)$, where $A_{i,NPR}$ represents price-non-responsive abatement and $A_{i,PR}$ represents price-responsive abatement. Price-non-responsive abatement refers to emissions reductions induced by policies or actions that do not respond directly to the allowance price. For example, policy makers often implement complementary policies or technology mandates that require more abatement than firms would choose at the prevailing cap-and-trade price. Distinguishing BAU emissions from price-non-responsive abatement is not needed to obtain the theoretical results of Section II, nor can the two components be separately identified from the model alone. Nevertheless, this distinction proves useful when empirically studying the CAQC context.

multiple linked jurisdictions enter the system. I therefore depart from their BAU estimation strategy while constructing abatement supply along the lines they propose.

I instead forecast BAU emissions with a dynamic factor model ([Stock and Watson, 2016](#)), which pools information across many macroeconomic and sectoral series and summarizes their comovement with a small number of common factors, keeping the estimated system small as regions are added. For abatement supply, I distinguish price-responsive from price-non-responsive abatement. The former is constructed from the price response of fuel demand across four emissions-intensive fuels; the latter reflects reductions that occur independent of allowance prices, including complementary regulations, offsets, and emissions reshuffling.

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. Welfare estimates are highly sensitive to the choice of social cost of carbon ϕ . I use a value of \$190 per ton ([Environmental Protection Agency, 2023](#)) in the preferred specification and show how estimated welfare losses scale across a wide range of plausible SCC values.

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.¹⁷ 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.¹⁸ Market coverage expanded to encompass a “broad” scope in 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

¹⁷Prior to the signing of the Kyoto Protocol, emissions were not tracked; thus pre-1990 data are inaccessible.

¹⁸Approximately 59% of all industrial emissions and natural gas consumptions were covered under the narrow scope.

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).¹⁹ 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.²⁰ 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,²¹ with volumes reported in cubic meters and converted to MMBtu using a factor of 0.0353 m³ 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.²² Annual CAD-USD exchange rates are taken from the Federal Reserve Bank of St. Louis.²³ Vehicle miles traveled (VMT) data are from the U.S. Department of Transportation and Statistics Canada.²⁴ Global oil prices are measured using Brent and WTI benchmarks from the EIA.

Business-as-usual emissions Business-as-usual (BAU) emissions depend on a set of factors that evolve dynamically over time, reflecting changes in economic activity, technology, and energy markets. A credible model of BAU emissions ought to estimate how these factors evolve rather than imposing their dynamics *ex ante*. Conventional time-series approaches are appropriate when all factors are observable. The econometrician’s challenge, then, is

¹⁹<https://www.cdtfa.ca.gov/taxes-and-fees/spftrpts.htm>

²⁰Statistics Canada. Table 18-10-0001-01

²¹Statistics Canada. Table 25-10-0033-01

²²Statistics Canada. Table 36-10-0402-01

²³<https://fred.stlouisfed.org/series/AEXCAUS>

²⁴California 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.

that some predictive factors may be unobservable (latent). For example, a region’s electricity generation and its electricity consumption both reveal information about the underlying activity of the electricity sector.

The behavior of latent factors correlates with the co-movement of many related, observable variables. A larger set of observables, therefore, provides more information to help recover these factors and reduces omitted-variable concerns. However, incorporating many variables quickly saturates conventional time-series models (such as VECMs). Thus, we need some form of dimension reduction in order to parsimoniously incorporate this additional information.

The dynamic factor model (DFM) is a time-series model well-suited to this task. The DFM’s core identifying assumption is that, even for a large number of observable variables, the number of latent factors remains small. This restriction reduces dimensionality while remaining computationally facile and consistent with standard dynamic equilibrium macroeconomic theories – it asserts that the movement of a small number of key variables suffices to characterize the behavior of a larger economic system (Stock and Watson, 2016). In Appendix C, I detail alternate forecasting approaches that relax the DFM’s 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 cross-sectional or dynamic assumptions embedded in the DFM.

Let X_t^{CA} denote the vector of N California-specific observables in year t , with X_t^{QC} defined analogously for Québec. The vector $X_t \equiv [X_t^{CA}, X_t^{QC}]'$ stacks these region-specific series into a single $2N$ -dimensional input to the DFM.²⁵ In the preferred specification, X_t includes 18 (2×9) variables that characterize the BAU environment: sectoral emissions, electricity generation and consumption, vehicle miles traveled, regional GDP, regional natural gas prices, and crude oil prices.²⁶ The equations for the DFM are then:

$$\begin{aligned} X_t &= \mu + \Lambda F_t + u_t \\ F_t &= \tau_t + f_t \end{aligned} \tag{7}$$

This model has two components. The first equation extracts a low-dimensional set of factors from the observables, while the second decomposes those factors into low- and high-frequency components.

²⁵A vector error correction model (VECM) or related multivariate time-series model could, in principle, include all $2N$ variables directly. In short panels, however, this is infeasible because the parameter count grows quickly with the dimension of the system. An unrestricted transition matrix for the $2N$ -variable system contains $4N^2$ parameters, making the dynamics difficult to identify with the available time series. I implement a VECM using a smaller number observed variables as part of the sensitivity analyses discussed in Appendix C.

²⁶For oil prices, the prices of WTI and Brent crude are used.




Figure 5: Latent factor extraction. After grouping variables of interest (depicted: vehicle miles traveled, gray), I extract the first principal component (brown) within each group. The principal component summarizes the common variation across series, providing a low-dimensional representation that preserves the dominant dynamics while remaining interpretable.

Factor extraction The first equation of (7) asserts that, in each period t , X_t equals the sum of a constant μ , a linear combination of r mean-zero factors F_t , and an idiosyncratic error u_t . The $2N \times r$ factor-loading matrix Λ reduces dimensionality by linking the full vector of observables to a smaller set of latent factors. For this factor structure to be appropriate, idiosyncratic errors u_t must be serially uncorrelated (so that all systematic temporal dependence is captured by the factors) and weakly cross-sectionally correlated across regions (to avoid introducing additional common components that should instead appear as factors).

The econometrician has considerable discretion when choosing the factor-loading structure Λ and the number of factors r . There are multiple ways to group observables into factors, which introduce a tradeoff between model flexibility and factor intelligibility.²⁷ In the preferred specification, I group the 18 starting variables into 8 thematic factors. These factors summarize cross-regional variation within a single, economically-interpretable variable group. For example, grouping California and Québec vehicle miles traveled (VMT) and extracting their first principal component recovers their shared variation in travel demand. Figure 5 illustrates this procedure for the VMT group. I extract the first principal component from the California and Québec VMT series, then repeat this procedure across all other groups (natural gas prices, oil prices, GDP, electricity generation, and sectoral emissions) to recover the latent factors.

²⁷Equivalently, she may choose whether to impose a pre-defined block-diagonal structure on the factor-loading matrix Λ . See Appendix C.2 and Bushnell and Smith (2025) for further discussion.




Figure 6: Decomposition and forecasts of latent factors. Each factor (black) from Figure 5 is decomposed into a low-frequency trend (blue) and a high-frequency component (red). Forecasts are constructed by bootstrapping the low-frequency trend (100 draws) and, conditional on each, simulating forward 10 high-frequency paths via a VAR, yielding 1,000 total draws. Shaded areas denote the 5th and 95th percentiles for each component.

Factor decomposition The second equation of (7) asserts that the vector of latent factors F_t decomposes into two parts: a low-frequency term τ_t and a high-frequency term f_t . The τ_t terms capture long-run trends and are estimated as linear combinations of Fourier terms. Factor shocks f_t may be serially correlated, so I follow standard approaches and estimate the dynamics of these shocks using a vector autoregression (VAR) model. Appendix C.2 provides details. The first and third columns of Figure 6 depict how each estimated factor (black) is decomposed into a low-frequency component (blue) and a high-frequency component (red).

Forecasting with the DFM After estimating the DFM, I use the fitted model to generate bootstrap forecasts. Because the low-frequency trend may be imprecisely estimated in finite samples, I first generate 100 alternative trend realizations via bootstrap. Conditional on each bootstrapped trend, I then simulate 10 draws of future components: the low-frequency term by extending the estimated trend and the high-frequency term using a VAR. Combining these terms returns 1,000 overall forecasts for each latent factor. I show the distribution of forecasts for low- and high-frequency terms in the second and fourth columns of Figure 6.

Finally, I use the forecasted latent factors to construct forecasts of the underlying variables. I estimate an ARIMA model relating each variable to the latent factors in-sample, and then project these relationships forward using the simulated factor paths. Aggregating across emissions sectors yields empirical distributions of total emissions for California and Québec, which I report in Table 1. 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.

Year	California		Québec	
	Broad scope MM T	Cumulative MM T	Broad scope MM T	Cumulative MM T
2013	367.3 (13.3)	112.4 (6.3)	70.7 (2.3)	25.6 (1.2)
2014	368.2 (14.9)	225.1 (11.4)	70.8 (2.7)	51.4 (2.4)
2015	368.8 (16.8)	594.0 (22.5)	70.8 (3.1)	122.2 (5.0)
2016	369.0 (18.1)	963.0 (37.6)	70.9 (3.4)	193.0 (8.0)
2017	369.4 (19.7)	1332.3 (54.4)	71.0 (3.7)	264.1 (11.3)
2018	369.8 (22.0)	1702.2 (73.5)	71.1 (4.3)	335.2 (15.1)
2019	370.3 (24.1)	2072.5 (94.4)	71.7 (10.6)	406.9 (22.6)
2020	371.3 (26.8)	2443.8 (118.1)	71.8 (12.8)	478.7 (33.5)


Table 1: Summary of regional DFM emissions forecasts. Estimates are mean values of 250 DFM draws. Values in parentheses are standard deviations of 250 draws. All values are in millions of tons CO₂e.

Figure 7 depicts the median, 5th, and 95th percentile of forecasted broad scope emissions by region; actual emissions are plotted in solid black. I plot each region’s stated 2020 emissions goal in blue. The California and Québec forecasts differ considerably. Even without the emissions drop from the COVID-19 pandemic, California was highly likely to attain 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 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).

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,



(a) California



(b) Québec

Figure 7: 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 (black). Dashed lines depict the 5th, median, and 95th percentiles of each year's emissions draw. Solid lines depict observed emissions. Actual emissions are shown in blue, with the thick blue line representing each region's 2020 emissions target.

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.²⁸ 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.²⁹ 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.³⁰ Using these assumptions, I 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 8.

²⁸Emissions 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 / MWh; Québec, 0.00279 T / MWh. 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.

²⁹The 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.

³⁰These β draws can be drawn independently or with some degree of correlation. For electricity without any complementary policies in play, I use the medium-run elasticities of (-1.1, -1.9), as in [Buchsbaum \(2022\)](#). For electricity with complementary policies in play, (-0.6, -0.9), *ibid.* For natural gas and gasoline, medium-run elasticities are unavailable so I use short-run elasticities (natural gas: (-0.1,-0.3), as in [Rubin and Auffhammer \(2024\)](#); gasoline and diesel: (-0.1, -0.2), as in [Hughes et al. \(2008\)](#)). Using short-run elasticities will tend to underestimate price-responsive abatement supply, though the robustness exercises in Appendix C.3 suggest this will not meaningfully affect 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.




Figure 8: Uncollared abatement supply curves for California (blue) and Québec (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.

Price collars and market clearing Regional BAU forecasts and abatement supply curves combine to define regional allowance demand. Specifically, allowance demand in i is $D_i(p) \equiv \text{BAU}_i - A_i(p)$, 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:

$$D_{CA}(p) + D_{QC}(p) = q_{CA} + q_{QC} \quad (8)$$

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:

$$e_i = \text{BAU}_i - A_i(\kappa(p)) \quad (9)$$

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 $e^{TOT} \equiv e_{QC} + e_{CA}$:

$$\begin{aligned} v_{QC} &= \rho \cdot e^{TOT} - e_{QC} \\ v_{CA} &= (1 - \rho)e^{TOT} - e_{CA} \end{aligned} \quad (10)$$

When prices are not at the ARP floor, $e^{TOT} = q_{QC} + q_{CA}$ and (10) 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 9 and Table 2 for 1,000 draws of BAU forecasts and 1,000 abatement supply curves. In Figure 9, 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). The abatement supply curve's perfectly inelastic sections correspond to the ARP (floor) and the three APCR (ceiling) prices. The probability mass of BAU emissions is concentrated in the floor section of the abatement supply curve, depicting the finding that, in expectation, the linked market is at its price floor 99.5% of the time and at its ceiling 0.1% of the time.³¹ Expected allowance flows are 108 MM T from California to Québec.




Figure 9: Net emissions and (collared) abatement supply for the joint California-Québec market, 2013-2020.

³¹These 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 CARB data but smaller variance, features which reduce the probabilities of non-price floor outcomes.

Linked price (\$USD 2013)	QC Emissions (MM T)	CA Emissions (MM T)	Pr(Floor)	Pr(Ceiling)
11.42 (3.42)	471.8 (33.5)	2,027.8 (122.8)	0.995	0.001

Table 2: 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 allowance 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 uncollared market-clearing price satisfies $D_i(p_i^\circ) = q_i$, collared autarkic prices are $\kappa(p_i^\circ)$, and regional emissions are a function of collared regional prices; $e_i^\circ \equiv \text{BAU}_i - A_i(\kappa(p_i^\circ))$.

I summarize results from this exercise in Figure 10 and Table 3 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 73.2% 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	44.58 (0.35)	453.7 (34.0)	0.010	0.732
CA	11.12 (0.05)	2,028.4 (122.8)	0.999	0.001

Table 3: 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 17.5 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 Califor-



(a) California



(b) Québec

Figure 10: Net emissions and abatement supply for California and Québec, *ceteris paribus autarky*, 2013-2020.

nia price floor. This price decrease mechanically increases allowance 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. Figure 11, the empirical analog of Figure 2, illustrates this point graphically. The distribution of unlinked, uncollared prices lies mostly below the “wings” of the bird, where linking raises emissions without generating sufficient cost savings to offset the associated damages. As a result, linking reduces welfare for sufficiently high social costs of carbon.




Figure 11: Empirical distribution of uncollared, unlinked allowance prices between California and Québec. Blue points reflect the joint distribution of uncollared, unlinked, prices over 1000 draws. Dashed gray lines represent the ARP price floor (\$10.87) and the first APCR price ceiling (\$40.99). The black “bird”-shaped region illustrates the region where linking is welfare-improving for all social costs of carbon. Its regions of nonconvexity reflect the presence of additional APCR price tiers.

To quantify welfare losses, first recall that variable regional abatement costs are:

$$C_i(A_i) \equiv \int_0^{p_i^o(A_i)} A_i(s) ds \quad (11)$$

In Table 4, I report expected abatement costs by region in linked markets and *ceteris paribus* autarky. The total cost decrease of \$504 million is driven by cost decreases in Québec, a combination of decreased overall abatement and gains from trade.

The total welfare impact of linking depends on the social cost of carbon ϕ . For low SCCs, the reduction in abatement costs outweighs the increases in emissions damages; for large SCCs, this trend reverses. When marginal emissions are constant, the welfare change is precisely Equation (6). In Figure 12, I plot total welfare change as a function of SCC, assuming constant marginal emissions.³² The breakeven SCC value of \$28.87 / T (where emissions damages equal cost savings) is close to the SCC estimates available to regulators at the time of linking. Using the SCC reported by the United States in 2013 (\$35 / ton, using a 3% discount rate), average losses are \$182 million over the 2013-2020 period, or \$23 million per year [Interagency Working Group \(2013\)](#).³³ More recent SCC estimates are considerably higher, so the welfare cost of linking looks considerably worse *ex post* than policy makers might have anticipated. Using an SCC of \$190 ([Environmental Protection Agency, 2023](#)), welfare losses from linking are an order of magnitude greater: \$2.6 billion over the 2013-2020 period, or \$331 million per year.

	Autarky (\$MM)	Linked (\$MM)	Change (\$MM)
QC	559.19 (190.23)	41.44 (50.58)	-517.75 (191.91)
CA	142.93 (131.16)	156.79 (184.35)	13.86 (127.22)
Combined	702.12 (235.78)	198.24 (234.43)	-503.89 (241.38)

Table 4: Summary of expected abatement costs in *ceteris paribus* autarky and linked markets, 2013-2020, 1,000 draws. All values in 2013 USD.

Sensitivity analysis The results in this section come from a single model. Although the model is flexible, it may still miss important features of the real-world data-generating process. I therefore address this *model uncertainty* – that is, the robustness of the main findings to alternative assumptions about both BAU emissions and abatement behavior – through a series of sensitivity analyses.

For BAU forecasting, I consider alternative factor groupings and alternative input data while retaining the DFM structure; I also employ forecasting models that impose progressively less structure on the data-generating process. For abatement supply, I vary the

³²With constant marginal damages, the breakeven SCC satisfies $\sum_i \Delta C_i / \sum_i \Delta e_i$. If damages are convex, the breakeven SCC is even lower.

³³Other contemporaneous SCC estimates are [Interagency Working Group \(2013\)](#) at a 2.5% discount-rate (\$55 / T; \$71 million lost per year) and [Moore and Diaz \(2015\)](#)’s substantially higher \$220 / T (\$418 million lost per year). All welfare losses are reported in 2013 dollars using the [U.S. Inflation Calculator](#).




Figure 12: Total welfare change as a function of the assumed social cost of carbon, 2013-2020. I assume marginal emissions damages are constant.

role of price-nonresponsive policies and the demand elasticities that govern price-responsive abatement. Appendix C describes these exercises in detail, and Table C.1 summarizes the results.

The core conclusions of Section IV are stable across specifications. In all cases, the linked market is highly likely to bind at the price floor, which increases expected emissions relative to autarky and reduces welfare for sufficiently large social costs of carbon. Across the robustness checks, expected emissions increases range from 6.6 to 18.7 MM T CO₂e. California binds at its price floor in 89-99% of draws, Québec binds at its ceiling in 48-79% of draws, and the joint market binds at the price floor in 78-99% of draws. The implied breakeven SCC ranges from \$27 to \$44 per ton, while annual welfare losses range from \$113 million to \$355 million.

These estimates assume that all unsold allowances are freely disposed. In practice, however, allowances can be banked for future use. If the future market clears at an interior price, linking with banking merely shifts emissions forward in time (rather than increasing cumulative emissions). Resulting damages are smaller, then, since a time-shifted ton of emissions causes only \$2–\$4 in damage.³⁴ In contrast, if future prices sit at a binding floor or ceiling, or

³⁴I compute the cost of shifting one ton of emissions ten years into the future using the formula $\Delta_{2013} = SCC_{2013} - SCC_{2023} \cdot (1 + r)^{-10}$. The \$2–\$4 range spans the inflation-adjusted 2.5% and 3% discount rate series in [Interagency Working Group \(2013\)](#).

if they fall to zero because the program expires, future allowance supply is already perfectly elastic and banking does not shift emissions. In these cases, linking raises emissions by the same amount as under free disposal.

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 17.5 MM T CO₂e and reduced global welfare. Had the joint market linked with additional provisions, these adverse outcomes could have been avoided.

I consider three classes of provisions: increases in the joint price floor, reductions in the total allowance supply, and tariff-like instruments. These provisions differ both in their cost-effectiveness and in their distributional consequences, though all draw upon a rich body of economic literature. 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 imposes a surcharge on imported allowances, relate to a growing literature studying climate policy under partial coordination, including carbon border adjustment mechanisms (Clausing et al., 2025; Clausing and Wolfram, 2023) and linking with trading ratios (Woerman, 2023). I describe the mechanics of each candidate policy and then compare their performance.

Because welfare changes are sensitive to the assumed social cost of carbon, I focus on a set of *emissions-neutral market provisions*. Emissions-neutral provisions are modifications to the joint market that restore expected aggregate emissions to their *ceteris paribus* autarky level ($\mathbb{E}[\sum_i e_i^\circ]$). By holding aggregate emissions fixed, I can compare the relative cost-effectiveness of various provisions (and thus the welfare impacts) without taking a position on the social cost of carbon.

Price floor increases An increase in the joint market’s auction reserve price raises the minimum allowance price and therefore increases total expected abatement.³⁵ 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 induces a new collaring function $\tilde{\kappa}$ such that total expected emissions equal their *ceteris paribus* autarky level:

³⁵I assume that price ceilings are unchanged in this counterfactual.

$$\mathbb{E}[\sum_i e_i^o] = \mathbb{E}[\sum_i \text{BAU}_i - \sum_i A_i(\tilde{\kappa}(p))]$$

In the preferred specification, the emissions-neutral $\tilde{\kappa}$ has a price floor of \$19.30.

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:

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

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

$$\mathbb{E}[\sum_i e_i^o] = \mathbb{E}[\sum_i \text{BAU}_i - \sum_i A_i(\kappa(\tilde{p}^*))]$$

In the preferred specification, the emissions-neutral ΔQ^* is 263 MM T.

The distributional effects of a quantity reduction depend on which region retires the allowances. Recall from (10) 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 provisions 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 “up-

stream” 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: $D_{CA}(\kappa(\tilde{p})) + D_{QC}(\kappa(\tilde{p}) + \tau) = \sum_i q_i$
The emissions-neutral τ^* solves:

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

In the preferred specification, the emissions-neutral specific tariff τ^* is \$32.13.

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: $\tilde{p} = \frac{r \cdot q_{QC}}{r \cdot q_{QC} + q_{CA}}$. I adopt the trading ratio notation for this manuscript.

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

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

The emissions-neutral r^* solves:

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

In the preferred specification, the emissions-neutral trading ratio r^* is 3.3.

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 13](#), I show the aggregate cost savings of each provision relative to autarky. Total costs decrease for all provisions relative to autarky, reflecting gains from trade. Price floor increases are most cost-effective (\$256 MM over the eight-year period), followed by quantity reductions (\$101 MM). Tariff-like instruments are the least cost-effective provisions (trading ratios, \$8.1 MM; specific tariffs, \$30.2 MM), since tariffs distort

prices across regions and reduce gains from trade. The dashed line represents expected cost savings experienced in the baseline (no-provision) linked market (\$504 MM). Cost savings necessarily increase under no-provision linking because the total level of abatement decreases.

In the bottom panel of Figure 13, 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 \$1.0 MM (trading ratios) and \$32 MM (price floor increases) versus autarky. Combining these gains with the estimated annual welfare losses of \$331 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 \$330 to \$360 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.

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?

The idea that real-world policy can be inefficient is not new. Representative democracies struggle to commit to future policy outcomes (Besley and Coate, 1998), and elected officials regularly pursue goals other than aggregate welfare maximization (Grossman and Helpman, 1994). When the interests of politicians and citizens diverge, Coasean bargaining alone cannot restore efficiency; what is needed instead is an “outside agency with the coercive capacity to enforce such arrangements” (Acemoglu, 2003).³⁶ Delegation can, however,

³⁶In California, CARB’s institutional structure places it close to Acemoglu (2003)’s ideal since the agency operates with considerable regulatory independence from the state legislature. Québec’s Ministry of Envi-




Figure 13: Expected relative abatement cost savings (versus autarky) across emissions-neutral market provisions. Expected abatement cost savings in the default (no-provision) linked market are shown in dashed lines. All values in 2013 USD.

introduce new inefficiencies depending on the discretion afforded to the regulator and the objectives it pursues (Alesina and Tabellini, 2007, 2008). The CAQC market illustrates this point.

I argue that climate policy delegation in California-Québec – and, likely, in other regions as well – is defined by three intertwined frictions. They are as follows:

- (i) Lawmakers set domestic emissions targets at inefficient levels.
- (ii) Policy inertia makes target revision very costly (if not impossible) for regulators.
- (iii) Regulators still aim to achieve these targets as cost-effectively as possible.

Together, these three frictions produce a testable hypothesis: regulators aim to minimize the cost of achieving a pre-declared emissions target, which is generally set at an inefficient level. I term this objective *declared-target cost minimization* (DTCM).

Of course, just because a friction can explain observed outcomes does not mean that it must. Other frictions may also be at play, and I do not attempt to claim the DTCM hypothesis as the sole explanation. Instead, I provide qualitative and quantitative evidence that this hypothesis plays an important role in the CAQC market. I then return to the market provisions from Section V and assess which provisions are feasible for a regulator operating under this hypothesis.

The adverse outcomes documented here do not arise simply because regulators pursue inefficient objectives. As Helm et al. (2003) argue, the optimal second-best climate regulator may place a higher weight on emissions damages than a standard cost-benefit calculation implies.³⁷ Rather, these adverse outcomes result from regulators overemphasizing cost minimization while under-weighting the damages from emissions increases.

Qualitative evidence Declared-target cost minimization reflects the division of responsibilities between lawmakers and regulators. Lawmakers, 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. This division of responsibility produces three frictions which impede efficient climate policy.

First, 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. Motivated by Gaarder

ronment operates with considerably less autonomy, as it is a standard cabinet ministry directly accountable to the National Assembly.

³⁷Second-best policy often deliberately distorts behavior away from the first-best allocation to maximize welfare (Lipsey and Lancaster, 1956). The classic example from monetary policy is Rogoff (1985)’s “conservative” central banker, who places greater weight on inflation stabilization than society does on average.

et al. (2026)’s caution that the “effectiveness of discretionary climate policies can be severely undermined by decision-makers’ non-emission motives,” I outline four such incentives: (1) distributional concerns, i.e., an aversion to imposing concentrated costs on particular industries, regions, or constituencies (Hughes and Urpelainen, 2015; Mildemberger, 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.

Second, policy inertia makes target revision costly and unlikely. The division of governmental responsibility means that lawmakers and regulators must agree to enact revisions together.³⁸ 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. In the CAQC market, the costs of revision appear to have outweighed the benefits on at least two occasions. Allowance prices sat transparently at the floor for nearly eight years, visible to regulators and widely discussed, yet the caps were never revised. Moreover, Ontario’s rapid entry and exit from the WCI over just six months in 2016 represented a significant unforeseen shock, yet it too prompted no major policy response.

Third, 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

³⁸Even 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.

“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 l’énergie, 2015).

Partial identification and quantitative evidence I also provide suggestive quantitative evidence in support of the DTCM hypothesis. Ideally, one would estimate policy maker objectives directly from observed political choices, as in Bombardini et al. (2023). That approach is not feasible here. The empirical record contains too little cross-sectional or within-legislator variation to uniquely identify regional preferences. There is, fundamentally, only one realized institutional choice from which to infer preferences (California and Québec chose to link their cap-and-trade markets). Moreover, Québec’s cap-and-trade legislation passed unanimously, leaving no observed legislative disagreement with which to distinguish competing objective functions.

This identification problem is not specific to this setting. Even with richer voting or bargaining data, point identification would generally remain impossible unless the econometrician observed the full feasible choice set. Some alternative policies were publicly discussed, but many feasible bargains, amendments, implementation details, and institutional constraints were private, informal, or never recorded. Appendix D formalizes this point: when the full choice set is unobserved, observed policy choices cannot uniquely identify policy maker objectives.

I therefore use a novel partial identification strategy. Rather than asking which objective function policy makers truly maximized, I ask which candidate objective functions can rationalize the observed decision to link. The revealed-preference restriction is simple: a region should agree to link only if linking makes that region weakly better off than a feasible outside option. Under the assumptions in Appendix D, this restriction yields an incentive-compatibility test whenever at least one outside option is known. In this paper, I treat autarky as the known outside option. Whether the relevant autarky counterfactual is *ceteris paribus* autarky or some other autarkic market depends on policy inertia, which I test for directly.

For a candidate policy maker type θ_i , let $V_i^{AUT}(\theta_i)$ denote region i ’s expected payoff under autarky, and let $V_i^{LINK}(\theta_i)$ denote region i ’s expected payoff under linking.

Region i ’s *perceived gain* from linking is then

$$\Delta V_i(\theta_i) \equiv V_i^{LINK}(\theta_i) - V_i^{AUT}(\theta_i).$$

A candidate type θ_i is inconsistent with the observed linking decision if either region has negative perceived gains from linking.

I test seven candidate types of policy makers, where each type is defined by the policy maker’s aims (his objectives) and inertia (the breadth of outside options against which linking is evaluated). A *regional-welfare maximizer* aims to maximize welfare in his own jurisdiction, while a *global-welfare maximizer* aims to maximize welfare across both jurisdictions. Finally, *cost minimizers* aim to minimize the cost of attaining a specified emissions target. These targets need not correspond to the efficient emissions level, so the relevant target must be specified separately. I consider two fixed-target versions. An *expected-target* cost minimizer uses the region’s expected emissions under constrained autarky as his target, while a *declared-target* cost minimizer uses his region’s official 2020 emissions goal. If a policy maker is *unconstrained*, he may set allowance quantities at different level under linking and under autarky in order to maximize his objective. If he is instead *constrained*, policy inertia prevents him from adjusting allowance quantities in response to the linking decision. The exact functional forms, first-order conditions, and implementation details for each objective are provided in Appendix D.

Results In Table 5, I report normalized perceived gains for each candidate type:

$$\Delta V_i^N(\theta_i) \equiv \frac{\Delta V_i(\theta_i)}{\text{sd}(\Delta V_i(\theta_i))}.$$

This statistic expresses the expected perceived gain from linking in standard-deviation units, making magnitudes comparable across objectives measured in different units. I then use Wilcoxon signed-rank tests to compare paired objective values under linking and autarky without imposing normality. A candidate type is rejected when either region’s objective is systematically higher under autarky than under linking. Normalization does not affect these tests because it rescales all paired differences for a given region and type by the same positive constant, leaving their signs and ranks unchanged.

The only candidate type that cleanly rationalizes the observed decision is declared-target cost minimization. Under this objective, both California and Québec have positive perceived gains from linking. The expected-target version is much more weakly rationalized: Québec clearly gains, but California’s perceived gain is close to zero. All other candidate types fail the incentive compatibility test because, under each of them, at least one region would have preferred autarky to linking. This does not prove that policy makers uniquely or literally maximize declared-target cost minimization, but it does show that the observed link-

ing decision is consistent with the DTCM hypothesis and inconsistent with many plausible alternatives.

Type	Objective	QC perceived gains (normalized)	CA perceived gains (normalized)	Wilcoxon Test Statistic
Unconstrained	<i>Maximize Global Welfare</i>	-3.545	-3.545	1.000
Constrained		-2.147	-2.147	1.000
Unconstrained	<i>Maximize Regional Welfare</i>	-17.266	-17.583	1.000
Constrained		-24.043	19.511	1.000
Unconstrained	<i>Minimize Regional Abatement Cost</i>	8.397	-4.473	1.000
Constrained, expected-target		6.880	0.157	0.000
Constrained, declared-target		5.998	19.012	0.000

Table 5: Results of Wilcoxon sign-rank tests for seven types of policy makers, $N = 1,000$. QC and CA margins report perceived gains from linking relative to the known autarky option, normalized by the within-specification standard deviations of the corresponding objective functions. Unconstrained policy makers may set allowance quantities at different level under linking and under autarky; constrained types may not. Policy makers aim to either maximize global welfare, maximize regional welfare, or minimize the cost of attaining a certain emissions goal. Positive margins favor linking. The Wilcoxon test statistic is the larger of the one-sided signed-rank p -values for Québec and California, so the joint hypothesis is rejected at level α when the reported value is less than α . Test statistics take values of 0 or 1 when empirical medians are sufficiently far apart.

Feasible counterfactuals I re-run the counterfactual analysis of emissions-neutral provisions for declared-target, cost-minimizing policy makers and present results in Figure 14. The policy makers’ expected payoffs (“perceived costs”) are lower under default linking (dashed lines) than in autarky in both regions, 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.

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 15, I show that Québec’s regional cost-minimizing trading ratio is 2.5, considerably lower than the emissions-neutral trading ratio of 3.3. 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 \$222 MM relative to the default linked market, roughly three-fifths of total potential welfare gains.³⁹ 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 16.

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 shaped by legislators and regulators through a delegated process. When this process prioritizes achieving a pre-declared emissions targets rather than equating marginal benefits and marginal damages, policy is not made “on the margin” and adverse outcomes can easily occur. Even superficially similar climate goals – “reduce emissions by

³⁹Using an SCC of \$35 / T increases welfare by \$22 million relative to default linking (two-fifths of total potential gains); an SCC of \$55 / T increases welfare by \$54 million per year (one-half of total potential gains); and an SCC of \$220 / T increase welfare by of \$279 million per year (three-quarters of total potential gains).

X,” “emissions must be X,” and “reduce emissions by X%” – may imply very different policy choices and outcomes. An open question, then, is whether the policy making process can be amended to implement efficient climate policy. Investments in information, analytical tools, or institutional design may help legislators to enact more efficient policy. If these interventions are unworkable, the logic of [Rogoff \(1985\)](#) suggests that empowering climate regulators with greater independence and a mandate to reduce emissions may attain a second-best outcome.




Figure 14: Expected relative perceived abatement cost savings (versus autarky) across emissions-neutral market provisions for declared-target cost-minimizing policy makers. Expected perceived abatement cost savings 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 15: 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 2.5, which is less than the emissions-neutral trading ratio of 3.3




Figure 16: Expected joint welfare (versus ceteris paribus autarky) for a range of market provisions at SCC \$190.

References

- Acemoglu, D. (2003). Why not a political coase theorem? social conflict, commitment, and politics. Journal of comparative economics, 31(4):620–652.
- Agrawal, D. R., Brueckner, J. K., and Brühlhart, M. (2024). Fiscal federalism in the twenty-first century. Annual Review of Economics, 16.
- Alesina, A. and Tabellini, G. (2007). Bureaucrats or politicians? part i: a single policy task. American Economic Review, 97(1):169–179.
- Alesina, A. and Tabellini, G. (2008). Bureaucrats or politicians? part ii: Multiple policy tasks. Journal of Public Economics, 92(3-4):426–447.
- Assemblée nationale du Québec (2024). Projet de loi 69, loi assurant la gouvernance responsable des ressources énergétiques, 43^e législature.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. Econometrica, 63(4):841–890.
- Besley, T. and Coate, S. (1998). Sources of inefficiency in a representative democracy: a dynamic analysis. American Economic Review, pages 139–156.
- 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.
- Burtraw, D., Palmer, K., Munnings, C., Weber, P., and Woerman, M. (2013). Linking by degrees: Incremental alignment of cap-and-trade markets. RFF Discussion Paper 13-04, Resources for the Future.
- Bushnell, J., Peterman, C., and Wolfram, C. (2008). Local solutions to global problems: Climate change policies and regulatory jurisdiction. Review of Environmental Economics and Policy.
- 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].
- 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.
- 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.
- Environmental Protection Agency (2023). Report on the social cost of greenhouse gases: estimates incorporating recent scientific advances. Environmental Protection Agency.
- Fell, H., Burtraw, D., Morgenstern, R. D., and Palmer, K. L. (2012). Soft and hard price collars in a cap-and-trade system: A comparative analysis. Journal of Environmental Economics and Management, 64(2):183–198.
- 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.
- 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.
- Grossman, G. M. and Helpman, E. (1994). Protection for sale. American Economic Review, 84(4):833–850.
- Helm, C. (2003). International emissions trading with endogenous allowance choices. Journal of Public Economics, 87(12):2737–2747.

- Helm, D., Hepburn, C., and Mash, R. (2003). Credible carbon policy. Oxford Review of Economic Policy, 19(3):438–450.
- 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.
- Lipsey, R. G. and Lancaster, K. (1956). The general theory of second best. The review of economic studies, 24(1):11–32.
- 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.
- Moore, F. C. and Diaz, D. B. (2015). Temperature impacts on economic growth warrant stringent mitigation policy. Nature climate change, 5(2):127–131.
- Muller, N. Z. and Mendelsohn, R. (2009). Efficient pollution regulation: getting the prices right. American Economic Review, 99(5):1714–1739.
- Müller, U. K. and Watson, M. W. (2016). Measuring uncertainty about long-run predictions. Review of Economic Studies, 83(4):1711–1740.

- Murray, B. C., Newell, R. G., and Pizer, W. A. (2009). Balancing cost and emissions certainty: An allowance reserve for cap-and-trade. Review of Environmental Economics and Policy, 3(1):84–103.
- Ostrom, E. (1990). Governing the commons: The evolution of institutions for collective action. Cambridge university press.
- 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 l'énergie (2015). Demande – r-3924-2015. April 17, 2015.
- Ricardo, D. (1817). On the Principles of Political Economy and Taxation. John Murray, London.
- Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. The quarterly journal of economics, 100(4):1169–1189.
- 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.
- United Nations Environment Programme (2025). Global Environment Outlook 7: A Future We Choose.
- 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 Conditional welfare

I leverage Equations (4a) and (4b) to summarize the effects of linking on abatement costs, emissions, and welfare conditional on each region's unlinked autarky outcome. There are three possible outcomes for each region's autarky price (an interior solution, a price ceiling, or a price floor), leading to a total of six possible cases after accounting for symmetry. I report these conditional results in Table A.1. and provide derivations below.

$i \backslash j$	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

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

(b) Aggregate welfare

Table A.1: Effects of linking on (a) costs C and emissions e and (b) aggregate welfare, conditional on autarkic market conditions.

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:

$$\Delta C = -\frac{A_i A_j}{2(A_i + A_j)}(p_i^\circ - p_j^\circ)^2 \leq 0 \quad (\text{A1})$$

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):

$$\Delta C = \frac{A_j}{2}(\underline{\kappa}^2 - p_j^2) \quad (\text{A2})$$

$\Delta C < 0$ 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:

$$\Delta C = -\frac{A_i A_j}{2(A_i + A_j)}(p_i^\circ - p_j^\circ)^2 + \frac{A_i}{2}(p_i^{\circ 2} - \underline{\kappa}^2) \quad (\text{A3})$$

Since $q_i > \bar{e}_i$, $p_i^\circ < \underline{\kappa}$. Thus, all three terms in (A3) 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):

$$\Delta C = \frac{A_j}{2}(\bar{\kappa}^2 - p_j^2) \quad (\text{A4})$$

$\Delta C > 0$ 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:

$$\Delta C = -\frac{A_i A_j}{2(A_i + A_j)}(p_i^\circ - p_j^\circ)^2 + \frac{A_i}{2}(p_i^{\circ 2} - \bar{\kappa}^2) \quad (\text{A5})$$

Since $q_i < \underline{e}_i$, $p_i^\circ > \bar{\kappa}$. The first term of (A5) 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 Welfare analysis and “bird” construction

This section constructs Figure 2 using the case results from Section A.1. It focuses on the $\Delta e = 0$ and $\Delta C = 0$ loci, their shapes, and their relative positions in (p_i°, p_j°) space.

The $\Delta e = 0$ loci identify the points where linking leaves aggregate emissions, and hence aggregate abatement, unchanged. Such points can only arise when one region is constrained by a price floor and the other by a price ceiling. For example, if region i is at the floor and region j is at the ceiling, total price-responsive abatement is $A_i(\underline{\kappa}) + A_j(\bar{\kappa})$. The upper $\Delta e = 0$ locus therefore begins at $(\underline{\kappa}, \bar{\kappa})$ and traces out the combinations of uncollared autarkic prices that preserve this same level of total abatement after linking. The lower locus is defined analogously and begins at $(\bar{\kappa}, \underline{\kappa})$.

These $\Delta e = 0$ loci are downward-sloping and convex, a result that follows from the curvature of the abatement functions. Along a locus, total abatement is constant and implicitly defines function g such that

$$A_i(p_i^\circ) + A_j(g(p_i^\circ)) = \bar{A}.$$

Differentiating once gives

$$g'(p_i^\circ) = -\frac{A'_i(p_i^\circ)}{A'_j(g(p_i^\circ))} < 0,$$

so the locus is downward sloping. Differentiating again gives

$$g''(p_i^\circ) = -\frac{A''_i(p_i^\circ) + A''_j(g(p_i^\circ)) (g'(p_i^\circ))^2}{A'_j(g(p_i^\circ))}.$$

Since abatement is increasing and concave in prices, $A'_i, A'_j > 0$ and $A''_i, A''_j < 0$, implying $g''(p_i^\circ) > 0$. Thus each $\Delta e = 0$ locus is convex in (p_i°, p_j°) space.

The $\Delta C = 0$ locus is constructed in the same spirit, but holds the change in total abatement costs fixed rather than total abatement. Because each region's cost is increasing and convex in abatement, the same isoquant logic implies that the $\Delta C = 0$ locus is convex and downward-sloping. The $\Delta C = 0$ locus passes through the regions where at least one jurisdiction is at its autarkic price ceiling and connects through $(\bar{\kappa}, \bar{\kappa})$. It lies strictly above the corresponding $\Delta e = 0$ locus because, in that region, linking increases emissions. The additional emissions reduce abatement costs, while gains from trade reduce costs further; the $\Delta C = 0$ boundary marks where these cost reductions are just offset by the increase in abatement required elsewhere.

A.3 General model

Suppose there are $N \geq 2$ regions. Define region i 's general abatement supply function $A_i(p)$, with $A_i(0) = 0$, $A'_i(p) > 0$, and $A''_i(p) < 0$. Call the set of regions $\mathcal{R}$.

The total cost of achieving emissions q_i is

$$C_i(q_i) = \int_0^{s_i} A_i^{-1}(a) da = \int_0^{p_i} t A'_i(t) dt.$$

The change in total cost from linking is then

$$\Delta C = \sum_{r \in \mathcal{R}} \int_{p_r}^p t A'_r(t) dt. \quad (\text{A6})$$

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

$$\Delta C = \sum_{r \in \mathcal{R}} \int_{p_r}^{p_r^\circ} t A'_r(t) dt + \sum_{r \in \mathcal{R}} \int_{p_r^\circ}^{p^\circ} t A'_r(t) dt + \sum_{r \in \mathcal{R}} \int_{p^\circ}^p t A'_r(t) dt.$$

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

$$\Delta C = \sum_{r \in \mathcal{R}} \int_{p_r}^{p_r^\circ} t A'_r(t) dt + \sum_{r \in \mathcal{R}} \int_{p_r^\circ}^{p^\circ} t A'_r(t) dt. \quad (\text{A7})$$

The first term has ambiguous sign. For two regions, it is positive at a ceiling, negative at a floor, and zero in the interior; with more than two regions, it is similarly ambiguous. The second term reflects gains from trade and is negative so long as $p \neq p_r^\circ$.⁴⁰ The results

⁴⁰To prove this, fix total abatement at $\sum s_r$. The interior linked price p° equalizes marginal costs across regions and therefore minimizes $\sum C_r$ subject to $\sum A_r(p_r) = \sum s_r$, with strict convexity of A implying a

of Section A.1 continue to hold for these more general functional forms because $A_i(p)$ is monotonically increasing.

A.4 Slack and differential additionality


Figure A.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 A.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. Differential additionality arises whenever a linked market has slack while an unlinked market does not, or vice versa. When this phenomenon occurs, linking distorts abatement incentives at the margin.

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

unique minimum at p° . Since p_r° also satisfies this constraint but does not equalize marginal costs when $p \neq p_r^\circ$, it must yield weakly higher total cost.



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 MM T CO₂e (CARB, 2018). CARB’s first Scoping Plan, published at the end of 2008, outlined preliminary business-as-usual 2020 emissions forecasts (approximately 600 MM T 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 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). 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 pub-

lic 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).⁴¹ 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 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.⁴²

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)

⁴¹Waste management, agriculture, and forestry remained uncovered "due to challenges associated with precise measurement, tracking or sector structure." (CARB 2008b)

⁴²California 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.

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 allowance 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.

C DFMs and alternate specifications

C.1 Data cleaning

California emissions The baseline model uses SEDS as the source of California sectoral emissions because it provides a standardized, consistently constructed, and widely used measure of state-level energy-related emissions. These data are a top-down measure of emissions, however, and therefore do not exactly match the bottom-up accounting used to determine compliance obligations under cap-and-trade. In principle, the preferred emissions measure would be the state’s Mandatory Reporting Regulation (MRR) data, which assign emissions obligations to covered entities and are therefore closest to the object regulated by the program. The limitation is that MRR data were first assembled in 2013 (as part of the cap-and-trade program).

One potential alternative is California’s Greenhouse Gas Emissions Inventory, which reports annual sectoral emissions back to 1990 and has been used in prior work. The Inventory is also a top-down measure of emissions, however, and it differs from MRR data in the 2013-2015 period. As [Borenstein et al. \(2019\)](#) note, “The Inventory’s top-down approach to quantifying emissions differs importantly from the bottom-up method of accounting for facility-specific emissions under the cap-and-trade program. In particular, the Inventory

likely overstates emissions from industrial activity relative to those covered in the first compliance period of the cap-and-trade program.” [Borenstein et al. \(2019\)](#) therefore construct an adjusted Inventory series that uses the overlap between the Inventory and MRR data in 2013-2014 to align the historical Inventory series with emissions measured under the compliance accounting system.

I use the [Borenstein et al. \(2019\)](#) adjusted Inventory series as an alternative input to the DFM. This data estimates an even lower California emissions level in the 1990’s, leading to lower BAU forecasts overall. The lower BAU forecast results in a greater expected surplus of California allowances and raises the probability that price floors binds.

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.⁴³ 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 1987-2023.⁴⁴ 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 2008.⁴⁵ For 2009-2010, I use U.S. average light-duty fuel economy.⁴⁶ 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 $X_t = (X_{1t}^{CA}, X_{1t}^{QC} \dots X_{Nt}^{CA}, X_{Nt}^{QC})'$ within a comparatively few

⁴³Statistics Canada Table [23-10-0097-01](#).

⁴⁴Statistics Canada Table [23-10-0066-01](#).

⁴⁵[TransportPolicy.net](#). Canada: Light-duty fuel consumption and GHG

⁴⁶Department of Transportation: [NHTSA CAFE PIC Report: Fleet Performance Dashboard](#)

number of factors. The specific DFM used in this paper is:

$$\begin{aligned}
X_t &= \mu + \Lambda F_t + u_t \\
F_t &= \tau_t + f_t \\
f_t &= \sum_{\ell=1}^L A_\ell f_{t-\ell} + \varepsilon_t \\
\Delta \tau_t &= \Psi(\Psi' \Psi)^{-1} \Psi' \Delta F_t
\end{aligned} \tag{C1}$$

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
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

Within each group, I demean and standardize each series prior to PCA. Let $X_t \in \mathbb{R}^N$ denote the standardized data vector at times $t = 1, 2, \dots, T$. Latent factors are using principal component analysis (PCA),

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

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

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. Additionally, this approach requires estimating all $r \cdot N$ elements of $\hat{\Phi}$, increasing the risk of misspecification in short time series.

An alternative is named factor normalization. This requires some portion of $\hat{\Phi}$ to be diagonal, which improves interpretability by forcing particular factors correspond to particular variables. 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 imposes block-diagonal restrictions on $\hat{\Phi}$. This method groups variables by type and extracts the first principle component from each group. 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.

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 $t = 1, 2, \dots, T$,⁴⁷ 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.

⁴⁷ $\Psi_j(s) = \sqrt{2} \cos(j\pi s)$ applied to $s = \frac{t-0.5}{T}$. See [Müller and Watson \(2016\)](#) for further details on spectral decomposition.

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

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

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 $\dot{z} \sim t(0, \widehat{\text{Var}}(\Psi^* \Delta \hat{F}))$. 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:

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

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

As before, $\ddot{\varepsilon}_t$ are randomly drawn with replacement from empirical VAR residuals, and trend noise is distributed $\ddot{z} \sim t(0, \widehat{\text{Var}}(\Psi^* \Delta \hat{F}))$. The combined factor forecast is then $\ddot{F}_{T+k} \equiv \ddot{f}_{T+k} + \ddot{\tau}_{T+k}$. [TO ADD: Writeup on conditional adjustment]

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.

C.3 Alternate models and sensitivity assessment

Table C.1 summarizes results from the alternative specifications discussed in this section.

Alternate DFM groupings I also consider relaxing the block-diagonal restrictions on Φ . Scree plot analysis suggests that a five-component specification is appropriate. This specification performs well for mean and median forecasts, but generates a small number of extreme upper-tail predictions. These tails arise because relaxing the grouping restrictions increases the residual variation in u_t . When these larger residuals are resampled and the forecasts are exponentiated, a small number of draws can produce very large emissions forecasts.

Alternate BAU forecasting – VECM A more flexible model imposes no latent factor structure on potential drivers of growth. Instead, it considers the four variables of interest directly:

$$D_t = \log[Broad_t^{CA}, Narrow_t^{CA}, Broad_t^{QC}, Narrow_t^{QC}]'$$

In a Vector Error Correction Model (VECM), these variables are assumed to be cointegrated, so that despite each series being nonstationary in levels, linear combinations of them may be stationary. The VECM is:

$$D_t = \mu + \Xi D_{t-1} + \Gamma \Delta D_{t-1} + \epsilon_t$$

where Ξ captures long-run level dynamics and Γ captures short-run adjustment. The cointegration rank (in this case, $r = 1$) is selected via Johansen’s trace test.

The VECM offers greater flexibility than the DFM by placing fewer restrictions on how the four emissions series interact. Its chief disadvantage is that, with the estimation sample spanning only 1990–2010, the unrestricted transition matrix Ξ introduces 16 free parameters alongside the 16 in Γ and 4 intercepts, leaving little degrees of freedom. This parameter proliferation raises the risk of overfitting in-sample and producing unstable out-of-sample forecasts, particularly for Québec, where the shorter history of reliable emissions data further limits identification. Sensitivity analysis in Table C.1 confirms that the VECM produces median forecasts close to the DFM baseline but exhibits substantially larger forecast variance, especially for Québec.

Alternate BAU forecasting – sampling with replacement An even more flexible model imposes no parametric structure on the distribution of emissions growth rates, instead drawing directly from their empirical distribution.

For $D_t = \log[\text{Broad}_t^{CA}, \text{Narrow}_t^{CA}, \text{Broad}_t^{QC}, \text{Narrow}_t^{QC}]'$ define the vector of log growth rates $G_t \equiv \Delta D_t$. Let $\hat{\rho}_k$ denote the estimated first-order autocorrelation of the k th element of G_t , computed over the estimation sample.

For each bootstrap replication $b = 1, \dots, B$, the procedure draws a sequence $\{\tilde{G}_1^b, \dots, \tilde{G}_T^b\}$ by sampling T rows with replacement from the empirical distribution of $\{G_t\}$, preserving the joint cross-series distribution in each draw. To incorporate persistence, a smoothed growth rate is constructed for each series k and period $j \geq 2$:

$$\hat{G}_{k,j}^b = \hat{\rho}_k \tilde{G}_{k,j-1}^b + (1 - \hat{\rho}_k) \tilde{G}_{k,j}^b, \quad (\text{C2})$$

with $\hat{G}_{k,1}^b = \tilde{G}_{k,1}^b$. Forecasted log-levels are then recovered recursively as $\hat{D}_{k,j}^b = \hat{D}_{k,j-1}^b + \hat{G}_{k,j}^b$, initialized at the observed start value of $D_{k,2010}$.

Price-responsive abatement [Buchsbaum \(2022\)](#) estimates medium-run electricity demand elasticities for California between -0.6 and -0.9 , considerably larger in magnitude than comparable estimates for transportation fuels (-0.1 to -0.2) or natural gas (-0.1 to -0.3). [Borenstein et al. \(2019\)](#) instead propose drawing electricity elasticities from the same ($-0.1, -0.2$) range used for transportation fuels. While this likely understates the true responsiveness of electricity demand, it preserves relatively similar elasticities across fuel types and avoids assigning electricity an outsized role in aggregate abatement supply.

Price-non-responsive abatement As an additional robustness exercise, I shut down all price-non-responsive abatement and instead allow substantially larger price elasticities for all fuel types, drawing elasticities from $(-1.1, -1.9)$ in line with the higher estimates in [Buchsbaum \(2022\)](#). This specification generates much stronger price-responsive abatement supply and substantially increases the role of allowance prices in determining equilibrium emissions. Despite these large changes in the abatement mechanism, the core equilibrium results remain largely unchanged: the linked market continues to spend most realizations at the price floor, implying higher equilibrium emissions under linking.

	<i>Baseline DFM (1)</i>	<i>DFM (Alt. elasticity) (2)</i>	<i>DFM (no CP) (3)</i>	<i>DFM (BBWZ data) (4)</i>	<i>DFM (alt. groupings) (5)</i>	<i>VECM (6)</i>	<i>Bootstrap (7)</i>
N Draws	1000	250	250	250	250	250	250
Cumulative BAU (MM T)							
CA	2443.8 (118.1)	2443.8 (118.1)	2443.8 (118.1)	2182.3 (92.8)	2461.4 (200.5)	2390.8 (272.9)	2433.9 (230.1)
QC	478.7 (33.5)	478.7 (33.5)	478.7 (33.5)	483.3 (24.5)	495.8 (54.2)	460.6 (22.2)	461.4 (32.4)
Autarky – CA							
Price (\$USD 2013)	11.12 (0.05)	11.28 (0.17)	11.93 (0.18)	10.98 (0.01)	12.14 (0.41)	13.05 (0.37)	12.22 (0.52)
CA Emissions (MM T)	2028.4 (122.8)	2084.8 (134.4)	2236.3 (133.8)	1763.7 (114.5)	2034.1 (211.9)	1972.8 (265.5)	2017.5 (227.4)
P(Floor)	0.999	0.985	0.896	1.000	0.944	0.925	0.934
P(Ceiling)	0.001	0.004	0.004	0.000	0.012	0.037	0.010
Autarky – QC							
Price (\$)	44.58 (0.35)	45.04 (0.35)	45.04 (0.35)	45.87 (0.50)	45.29 (0.17)	38.88 (0.45)	37.65 (0.20)
QC Emissions (MM T)	453.7 (34.0)	454.7 (29.0)	454.7 (29.0)	457.5 (29.4)	470.0 (60.9)	438.6 (18.1)	440.2 (26.8)
P(Floor)	0.010	0.005	0.005	0.012	0.004	0.010	0.072
P(Ceiling)	0.731	0.754	0.754	0.788	0.754	0.481	0.491
Linked							
Price (\$)	11.42 (3.42)	11.88 (4.63)	13.13 (6.03)	11.06 (1.39)	13.50 (7.98)	13.51 (8.15)	12.94 (6.61)
CA Emissions (MM T)	2027.8 (122.8)	2084.1 (134.2)	2225.6 (139.3)	1763.5 (114.4)	2031.3 (212.0)	1971.8 (264.3)	2016.0 (225.1)
QC Emissions (MM T)	471.8 (33.5)	472.7 (28.3)	472.0 (28.4)	476.4 (28.3)	487.3 (60.7)	452.5 (21.8)	453.7 (32.0)
P(Floor)	0.995	0.949	0.776	0.996	0.887	0.892	0.895
P(Ceiling)	0.001	0.007	0.005	0.000	0.030	0.043	0.018
Change in Abatement (MM T)	-18.3	-17.3	-6.6	-18.7	-14.5	-12.9	-12.0
Breakeven SCC (\$)	28.52	29.18	43.72	29.09	30.14	27.40	27.82
Annual ΔWF (MM \$)	-348.0	-327.6	-112.7	-354.7	-272.2	-246.8	-229.2

Table C.1: Effects of linking on emissions and welfare across specifications. (1) is the baseline model; all other model modify this model along one dimension. (2) calibrates price-responsive abatement using elasticities from [Borenstein et al. \(2019\)](#) rather than [Buchsbaum \(2022\)](#); (3) omits all sources of price non-responsive abatement; (4) uses an alternate California emissions data as inputs to the DFM; (5) imposes no groups on the DFM and instead inputs the first 8 factors from PCA; (6) forecasts BAU emissions using a 4-variate VECM; (7) forecasts BAU emissions using a resampled bootstrap. All models besides (1) use 250 random draws to improve computational efficiency.

D A model of policy maker types

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** $V_i : \alpha \rightarrow \mathbb{R}$; $V_i \in \mathcal{V}_i$. Define the *type* θ_i of policy maker i as his corresponding pair (V_i, α_i) . Thus, policy maker i 's type encodes both what he wants to do (V_i) and what he is able to do (α_i).

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

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

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

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

Partial identification of types Suppose the econometrician does not observe V_i . Her ability to identify V_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 the complete set of available other actions.

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 $(V_i, \{\alpha_i^*\})$ and $(\mathbb{I}(\alpha_i = \alpha_i^*), \alpha_i)$ can rationalize outcome α^* . $(V_i, \{\alpha_i^*\})$ represents an arbitrary objective function when the policy maker's only available choice is $\{\alpha_i^*\}$. $(\mathbb{I}(\alpha_i = \alpha_i^*), \alpha_i)$ 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 V_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 V_i); rather, it rules out idiosyncratic processes where a policy maker does not implement his most-preferred feasible policy.⁴⁸ 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.⁴⁹ Specifically, the econometrician’s knowledge that i does *not* choose the KOO α'_i provides her with partial information regarding V_i and a way to test various V_i for feasibility.

Both objective functions V_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 (V_i, α'_i) where, for some KOO $\alpha'_i \in \alpha'_i$,*

$$V_i(\alpha'_i, \alpha_{-i}^*(\alpha'_i)) > V_i(\alpha_i^*, \alpha_{-i}^*)$$

⁴⁸The 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.

⁴⁹Work by [Manski \(2007\)](#) considers a similar question in a consumer choice setting. [Molinari \(2020\)](#) provides an overview of scholarship in this area.

Proposition 1 (The ICHT) states that the econometrician should reject the hypothesis “ i ’s objective is V_i ” if the KOO α'_i is a feasible and profitable deviation for the V_i -minded policy maker. In the context of the CAQC market, this proposition becomes the testable hypothesis: $V_i^{LINK}(\theta_i) > V_i^{AUT}(\theta_i)$

Applying Proposition 1 to the findings of Section IV, for instance, 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 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.⁵⁰ 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.⁵¹ 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 $\tilde{e}_i$. If $\tilde{e}_i$ coincides with the socially efficient emissions

⁵⁰The type mentioned earlier, $(V_i = \mathbb{I}(\alpha_i = \alpha_i^*), \alpha_i)$, is one such unfalsifiable example.

⁵¹Although one could allow complementary policies to adjust alongside allowance quantities, that richer environment is not the focus here.

level, the cost-minimizing planner also maximizes global welfare. However, cost-minimizing planners are permitted to set $\tilde{e}_i$ 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 $\tilde{e}_i$ as the region's *expected* emissions under an autarky counterfactual when q_i allowances are issued. A second version interprets $\tilde{e}_i$ as each region's officially-*declared* emissions goals, which are looser than expected autarky emissions in California and stricter than expected autarky emissions in Québec.⁵² I term these types *expected-target cost minimization* and *declared-target cost minimization*, respectively.

I calibrate the relevant parameters of each planner type θ from the relevant equilibrium FOCs corresponding to V_i . With richer data, θ can be estimated instead using a variety of approaches, e.g. [Bombardini et al. \(2023\)](#).

D.1 ICHT moment conditions

Each candidate policy maker type maximizes his objective functions V_i over a choice set α_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.

In general, region i 's welfare is function of the domestic cost of abatement $C_i(BAU_i - e_i)$, 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:

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

Unconstrained regional welfare maximizers Unconstrained policy makers 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 $C'_i(BAU_i - e_i) = p$, the market-clearing price:

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

⁵²This occurs since California's autarkic market usually prices at a price floor whereas Québec's autarkic market prices at its ceiling.

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

Constrained regional welfare maximizers A constrained policy maker 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, $WF_i = -C_i(BAU_i - q_i) - \phi_i \cdot e^{TOT}$. 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, $WF^{AGG} \equiv \sum -C_i(BAU_i - e_i) - \phi^G \cdot e^{TOT}$, 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 5, 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:

$$C_i(A'_i) + (p + \frac{dp}{dq}(v'_i - v_i)) \cdot v'_i \quad (\text{D1})$$

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: $\tilde{e}_i \equiv \mathbb{E}[\text{BAU}_i - A_i(p) + v_i]$. Intuitively, if the linked market delivered emissions above or below the target in expectation, the policy maker 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 $\tilde{e}_i = q_i$.

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 policy makers 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: $\tilde{e}_i = \mathbb{E}[\text{BAU}_i - A_i(p_i^\circ)]$, 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.