Gradual Linking of Carbon Markets*

Simon Quemin Jan Abrell Dallas Burtraw Michael Pahle Matt Woerman

This version: March 2026

Abstract

The efficiency of carbon pricing hinges on markets broad enough to encompass heterogeneous economies and sectors. Linking carbon markets is gaining increasing attention, in particular for advanced systems in which caps are rapidly decreasing and prices are rising. Paradoxically, realizing linking surfaces economic and political threats that tend to be more severe the more advanced a system is. One is an immediate convergence of prices due to intertemporal arbitrage opportunities that emerge already when firms anticipate linking in the future. This paper shows that ensuring gradual price convergence across linked markets requires two instruments—one to establish the exchange rate at which permits trade in each period, another to align intertemporal arbitrage incentives with the targeted sequence of exchange rates—and characterizes how they should be calibrated jointly. These tools may enable policymakers to control price convergence while initiating linking, and offer more flexibility to manage the distributional and environmental impacts of linking, thus overcoming a main barrier to higher stringency. A calibrated numerical application to the EU ETS1 and ETS2 illustrates the different price paths and emission levels under various gradual linking policies.

Keywords Linking, banking, EU ETS.

JEL classification H23, Q58.

*Quemin: Potsdam Institute for Climate Impact Research (PIK), Telegrafenberg A31, D-14473 Potsdam; Climate Economics Chair, Paris-Dauphine University PSL, Place du Maréchal de Lattre de Tassigny, F-75016 Paris; Électricité de France R&D, SYSTEME Department, EDF Lab Paris-Saclay, 7 Boulevard Gaspard Monge, F-91120 Palaiseau; simon.quemin@pik-potsdam.de. Abrell: Chair of Energy Economics, University of Basel, Faculty of Business and Economics, Peter-Merian-Weg 6, CH-4002 Basel, jan.abrell@unibas.ch Butraw: Resources for the Future (RFF), 1616 P Street NW, Washington DC 20036; burtraw@rff.org Pahle: Potsdam Institute for Climate Impact Research (PIK), Telegrafenberg A31, D-14473 Potsdam; michael.pahle@pik-potsdam.de Woerman: Colorado State University, Department of Agricultural and Resource Economics, 301 University Avenue, Fort Collins, CO 80523-1172; matt.woerman@colostate.edu.

1 Introduction

Economic costs in carbon pricing programs pose a crucial barrier to realizing international or multilateral commitments to reduce greenhouse gas emissions. There are currently 80 different carbon pricing programs in place around the world covering 28% of global emissions, mostly through carbon markets (19% of emissions), but with individual program characteristics and widely varying prices (ICAP, 2025; World Bank, 2025). The discrepancy in abatement costs and prices across markets raises overall policy costs, setting back progress globally. Economic advice is to harmonize and link markets to broadly identify least-cost abatement opportunities; however, this brings an associated harmonization (equalization) of marginal costs, presenting severe distributional and political economy challenges that may explain why linkages between markets between remain few and far between.[1] Paradoxically, the greater the potential efficiency gains from linking, the more severe these challenges may be.

Linking is a form of scope expansion or policy defragmentation, which has the potential to improve market efficiency and abatement cost effectiveness. Numerous other benefits of linking include administrative advantages for regulators and firms, greater price stability and resilience in a larger market, and enhanced credibility of program goals leading to greater policy durability (ICAP, 2026; Pahle et al., 2026). Challenges to linking have also been recognized in the literature. Regulatory disadvantages may arise from the linked system's inconsistency with original domestic policy objectives and from the partial *de facto* cession of discretionary control over domestic programs (Flachsland et al., 2009; Burtraw et al., 2013).[2]

A primary challenge, which is our focus, is the distributional outcomes that may arise through price movement, affecting consumers and the fiscal revenues resulting from linking. In principle, these implications could be dealt with by suitable emission permit allocations or financial transfers so that all jurisdictions gain from linking, but negotiating this outcome has been difficult to achieve. [3] The distributional challenge may be aggravated in the case of linking an advanced system featuring a relatively high price with a newly developed system featuring a relatively low price (Roy et al., 2025). This is because linking may lead to sizable price

[1]Examples are the links between the EU and Switzerland and expected inclusion of the UK, the eleven eastern US states linked in the Regional Greenhouse Gas Initiative, and between California and Québec with the anticipated inclusion of Washington.

[2]Conversely, linking also may insulate policy from capture by special interests within each jurisdiction. See Section 2 for a detailed literature review.

[3]When such a redistribution proves to be infeasible, uniform pricing may be inefficient (Abrell et al., 2018; Bataille et al., 2018; Stiglitz, 2019).

shifts in both systems, which might incur additional adjustment costs or erode the value of investments. Further, in advanced systems, deeply entrenched design paths may exist that further complicate a necessary alignment of the rules between systems (Metcalf & Weisbach, 2012; Mehling et al., 2018, 2024). All this leads to the paradoxical situation in which larger efficiency gains from linking come along with greater challenges too.

The tradeoff between costs and distributional outcomes is especially salient in the EU, which maintains the world's oldest carbon emissions trading system (ETS1) covering the electricity and industrial sectors, currently with a price of approximately 75 EUR per ton of carbon. In 2027, the EU will launch a second emissions trading system (ETS2) covering the transportation and buildings sectors, with prices that remain highly uncertain for now and dependent on the reach of companion policies (Günther et al., 2025). These programs will initially be maintained separately, notably because of distributional concerns about impacts on households and businesses and because the ETS2 needs time to find equilibrium first. In the next decade, linking the two systems may be envisaged, notably because of the inefficiencies otherwise associated with price discrepancy in the long run.

The anticipation of such linking, however, may lead to early and politically problematic price convergence through arbitrage.[4] A recent example of such anticipatory "pricing in" is the reduced price spread between the EU and UK ETS prices following the announcement of the intention to link the two systems. This illustrates that the effects of linking can materialize well before the link is in place and indicates that a successful phase-in of the link requires addressing the temporal dimension.

Globally one observes a similar challenge is replicated across sectors and countries, with carbon prices that are even more disparate. Rapidly equating carbon prices by linking markets would result in distributional outcomes that are politically untenable, and possibly even sub-optimal (Lang, 2026). Decarbonization policies should be not only cost efficient but also feasible, fair, effective, and credible in the long term. Fabra & Reguant (2024) provide an overview of the various dimensions of this broader notion of efficiency and of the various associated tradeoffs. Gradual linking policies as considered in this paper should be viewed in this context, as a way to balance and deal with cost efficiency, distributional, environmental, and credibility impacts. A pressing policy problem is how to design carbon markets to capture gains from trade while governing the distributional impacts and incremental convergence in

[4] As we will see in detail, this effect occurs through the banking of permits for future compliance. While our model considers spot markets for permits only, we note that the intertemporal effects we characterize in this paper would be reinforced by the existence of markets for futures contracts.

prices over the long run.

This work develops a theoretical and numerical model to identify policy levers for regulators to gradually link carbon markets—that is, to allow the trade of permits without equating permit prices in the short run, while creating a pathway toward full linkage and price convergence in the long run. Such an outcome can manage the short-run distributional impacts that would otherwise make the energy transition politically unsustainable, while still maintaining long-run cost effectiveness. We show that a linking policy ensuring gradual price convergence across markets over time needs two instruments—one to establish the exchange rate at which permits trade in each period, another to align permit banking and intertemporal arbitrage incentives with the targeted sequence of exchange rates.

We consider a gradual linking policy "successful" if it results in a Pareto-improvement along three dimensions—allocative efficiency, environmental efficacy, and meeting distributional goals. In other words, policy "success" means three goals are met: some price convergence across markets occurs yielding near-term economic efficiency gains and shaping expectations that will guide investments affecting future market outcomes. Second, that price convergence is controlled and gradually increasing over time to address distributional aspects and political economy constraints. Third, that gradual convergence should occur without environmental harm (total emissions cannot increase).

We show that, if exchange rates are used in isolation, the gradual linking policy fails at achieving these goals. Specifically, intertemporal arbitrage will cause future exchange rates to propagate back to earlier periods, making earlier exchange rates mute and the planned phase-in obsolete. Further, such a system can create perverse incentives for both economic and environmental outcomes, such as trade of permits in the economically inefficient direction—the lower-priced market buying permits from the higher-priced market in early periods—which effectively relaxes the stringency of the linked emissions cap.

Instead, combining exchange rates with intertemporal banking conversion rates can create a successful policy, and we define how the conversion rates for banked permits should be set to always sustain this successful policy in equilibrium. We further show that such a "two-dimensional" policy gives policymakers more flexibility to manage the distributional and environmental impacts of linking.

The paper then provides a calibrated numerical application to the EU ETS1 and ETS2 to demonstrate the real-world applicability of a two-dimensional gradual linking policy and to illustrate the mechanisms driving the results of different gradual linking policies. We also

quantify the price paths, emission levels, and welfare effects that may result under various gradual linking policies.

The paper is structured as follows. Section 2 reviews the literature and describes our contributions. Section 3 provides a simple theoretical model to study gradual linking policies and presents the main results. Section 4 illustrates the impacts of various gradual linking policies using a numerical model parametrized to the EU ETS1 and ETS2. Section 5 concludes. An Appendix collects the proofs, describes the implementation of the numerical model, and characterizes alternative instruments.

2 Literature

Our paper connects and contributes to two strands of literature that we review below.

2.1 Permit conversion rates

In a permit market with unrestricted banking and borrowing, intertemporal flexibility and arbitrage imply that the permit price grows at the private rate of interest (Rubin, 1996). The supply of permits is analogous to an exhaustible resource, which is optimally used according to the Hotelling depletion rule. When borrowing is restricted, the price grows at a rate at most as high as the interest rate, a rule that holds in expectation under uncertainty (Schennach, 2000). However, such a path may not be socially optimal as the private interest rate typically differs from the social interest rate.[5] More formally, the decentralized problem of minimizing abatement costs over time for firms is generally not congruent with the centralized problem of maximizing the consumer surplus net of abatement costs and environmental damages, both with flow and stock pollutants (Kling & Rubin, 1997; Leiby & Rubin, 2001).[6]

Against this background, the literature has essentially focused on how to address uncertainty about damage and abatement costs as well as asymmetric information about costs in permit markets with banking. Some flexibility in arbitraged price paths can be introduced with the use of non-unitary intertemporal trading ratios (ITRs) that adjust—typically discount—the emission or compliance value of banked permits in subsequent periods. For instance, Yates

[5] As a rule, capital market interest rates tend to be higher than social interest rates.

[6] With a flow pollutant, for instance, the former problem equates discounted marginal costs over time, whereas the latter equates marginal cost and marginal damage in every period (Kling & Rubin, 1997). In general, both solutions need not coincide. See also Leard (2013) for a discussion of the environmental impacts of banking and associated emission reshuffling over time.

& Cronshaw (2001) and Feng & Zhao (2006) compare the efficiency of banking regimes with different ITRs. In essence, ITRs enable a decoupling between discounted marginal abatement costs (and prices) across periods through adjusted intertemporal arbitrage. Such a decoupling is a prerequisite when regulators seek to adjust permit supply over time as new information arrives in order to achieve first best (Newell et al., 2005; Pizer & Prest, 2020).[7]

In the case of heterogeneous or localized pollution damages (e.g., non-uniformly mixed pollution), trading ratios across sources are generally warranted to operate a trade-off between minimizing abatement costs and minimizing damages (Mendelsohn, 1986; Montero, 2001; Muller & Mendelsohn, 2009; Fowlie & Muller, 2019; Kakeu, 2025). Trading ratios allow the flexible control of marginal costs across sources, which reduces cost effectiveness but enables one to discriminate emissions based on their location and damages caused. Because the focus of these papers is on spatial heterogeneity, granular but static frameworks are used. As a result, the interaction between trading ratios and ITRs, which are also aimed at controlling marginal costs across periods, is not considered.

The contribution of our paper lies in the combination of the spatial and temporal dimensions. For instance, Holland & Yates (2015) provide a general characterization of optimal trading ratios across space or time, but do not consider both dimensions at the same time. In this paper, we allow for both intra- and inter-period trade and show that both types of ratios—which for disambiguation we call linking and banking conversion rates—are needed to control the marginal costs and prices across markets and time periods. Notably, we characterize how they must be set in tandem to achieve a given objective.

Finally, note that linking and banking conversion rates are not just theoretical concepts, but are actually part of the existing policy toolkit. For instance, Pizer & Prest (2020) describe examples of ITRs used or proposed in US programs such as the Acid Rain Trading Program and the "Waxman-Markey Bill". In addition, the NOx Budget Program in the eastern US established a 2-for-1 surrender rate for the use of banked permits over a specified quantity (EPA, 2003). More recently, the US EPA's Good Neighbor Plan annually recalibrates the compliance value of the permit bank for emitting sources with higher emissions rates (EPA, 2023). Similarly, exchange rates for emission permits and discount rates for offsets have also been used or proposed in a variety of carbon markets, see notably Quemin & de Perthuis (2019) and Woerman (2023) for examples and discussions.

[7]Decoupling can also be achieved by the "conversion function" in Karp & Traeger (2025), which modifies the compliance value of permits over time without the direct use of ITRs.

2.2 Linking permit markets

The benefits and drawbacks of linking permit markets are well researched and long established (e.g., Flachsland et al., 2009; Tuerk et al., 2009). The main economic benefit is an increase in allocative efficiency that results from the equalization of (or a reduced dispersion in) permit prices between markets. Linking allows for a (more) cost-effective split of abatement efforts across markets as permits flow where they reduce compliance cost the most, and all markets are better off from linking in aggregate. Intuitively, the higher the initial price differential across markets, the higher the efficiency gains from such "effort sharing". Yet at the same time, the induced distributional and price impacts are larger, making the acceptability of linking less obvious and the decision to link more difficult to make politically (e.g., Burtraw et al., 2013; Quemin & de Perthuis, 2019). In addition, linking brings about efficiency gains from "risk sharing" as shocks can be pooled and spread over a larger market with larger buffer capacity (Doda et al., 2019). While price volatility decreases on average, it may increase in some linked markets when external shocks are "imported" through the link.

While in principle linking-induced efficiency gains could facilitate setting more ambitious emission targets to increase environmental benefits, the literature finds ambiguous results regarding the cap-setting dynamics under linking (e.g., Doda et al., 2019; Holtsmark & Weitzman, 2020). For instance, total emissions tend to increase as a result of linking when emission caps are set non-cooperatively (Helm, 2003), but this result is reversed in a dynamic setting where investment in abatement technology is endogenous (Holtsmark & Midttømme, 2021). Emissions may also increase as a result of linking when caps are intensity-based, as in tradable performance standards programs (Long et al., 2025) although well-chosen performance standards can reduce emissions in linked programs. In addition, there are economic reasons why permit prices could differ across markets, including differences in labor or capital markets and energy structures (Bataille et al., 2018), the public good nature of emission reductions (Chichilnisky & Heal, 1994), and pre-exisiting inequality or fairness issues (Abrell et al., 2018; Bierbrauer, 2025). Other drawbacks of linking include reduced autonomy in terms of governance, market design, and price developments.

Against this background, linking with restrictions or "controls" that may be gradually lifted over time has the potential to capture efficiency gains all the while managing distributional impacts and initial differences across markets if expectations about future prices are supported (Dolphin et al., 2023). The literature has studied the effects of such restrictions to trade, such as exchange rates or quantitative trading limits, but only in static frameworks

(Quemin & de Perthuis, 2019; Woerman, 2023). Conversely, dynamic models of linking assume unrestricted banking and borrowing, and do not feature exchange rates, trade limits, or intertemporal trading ratios (Doda et al., 2019; Holtsmark & Midttømme, 2021).

Our paper closes this gap by characterizing how effective gradual linking policies should be designed when banking is allowed. Specifically, we show how a policy that performs well in a static setting (e.g., linking conversion rates to allow for the flexible control of the price ratios between markets) may be insufficient and lead to detrimental outcomes in a dynamic setting. Intuitively, because current permit prices reflect their expected value in the future, intertemporal arbitrage entails that the future linking conversion rate (provided that it is known) is effective before it is in place while the current linking conversion rate is mute. Thus, implementing gradual linking requires driving a wedge between markets' Hotelling price paths, which can only occur with the use of intertemporal trading ratios that are set to sustain the targeted sequence of linking conversion rates. In other words, the joint use of linking and banking conversion rates allow policymakers to accommodate the need to communicate early about future policy to avoid credibility issues (e.g., Salant, 2016; Dubois et al., 2025) with a smooth transition to the future policy regime that would otherwise materialize and be priced in "too early" through arbitrage.

Our premise in this paper is that the policy focus is on ensuring gradual price convergence across linked markets. Such endeavor may result from various political economy constraints that reach beyond economic efficiency (e.g., Fabra & Reguant, 2024), which we take as given and abstract from. These include adjustment costs and some form of systemic inertia that may call for a phase-in of the price and distributional implications of full linking in the long run.[8] In other words, we do not consider the socially optimal gradual linking policy (if there is one), but we aim to assess the joint impacts of a targeted sequence of linking and banking conversion rates on cost and environmental effectiveness. This allows us to delineate first-pass policy recommendations for policy selection—e.g., we identify value ranges where rates are welfare improving. Because linking and banking conversion rates may have opposite effects on emissions depending on how they are set, we utilize calibrated numerical simulations to illustrate which effect is more likely to dominate.

[8] See notably Williams III (2010) and Vogt-Schilb et al. (2018) for related analyses.

3 Model

In this section, we describe the modeling framework and characterize the market equilibrium under various linking scenarios and gradual linking policies. The model is kept as simple as possible to capture and study the key gradual linking dynamics.

3.1 Modeling framework

The model borrows from and adjusts the model developed in Pizer & Yates (2015) to study gradual linking in a simple framework.[9] The main novelty is that the model features both exchange rates and intertemporal trading ratios. In comparison, existing models only feature one or the other (see Section 2). As we will show, both instruments are necessary to guarantee that a linking policy successfully achieves the policy goal of gradual price convergence.

Specifically, we consider a two-period, two-market model in which firms are regulated under separate or linked emission permit markets. Periods are indexed by $t \in \{1, 2\}$ and markets are indexed by $m \in \{H, L\}$. H denotes the high abatement cost or high permit price market under autarchy, whereas L denotes the low cost or low permit price market. Permits are denominated by their market of origin and vintage, and ω_{mt} denotes the emission cap or permit endowment in market m and period t. Without loss of generality, we assume that permits are freely allocated.[10] Firms are required to surrender (the equivalent of) one period-t m-permit for each unit of emissions generated in market m and period t.

Permit conversion rates. Linking and banking give firms some flexibility to shift emissions across markets and time periods, respectively. We consider linking and banking policies that use permit conversion rates to adjust the compliance value of permits (i.e., the number of emission permits that must be surrendered for each unit of emissions) when they are traded across markets or are banked for future compliance. Importantly, the use of non-unitary conversion rates makes the volume of aggregate emissions variable, i.e., the total number of emission units is no longer equal to the total number of permits $\sum_{mt} \omega_{mt}$. As a result, environmental effectiveness joins gradual price convergence, cost effectiveness, and distributional impacts as integral considerations of linking policy.

[9] By contrast, the model in Pizer & Yates (2015) is designed to study "delinking".

[10] This is for ease of exposition and inconsequential for the results under the assumption of perfect competition. As the numerical implementation of the model shows in Appendix B, our results hold for any mix of auctioning and free allocation.

Linking allows permits originating from one market to be used for compliance in the other market. The linking policy we consider adjusts and governs the compliance value of traded permits through a *linking conversion rate* ρ_t that is specific to period t, whereby ρ_t L-permits are equivalent to (i.e., fungible with) one H-permit in period t, and vice versa.[11] Let Δ_t denote the net quantity of permits—measured in H-permits or units of emissions in market H—traded from market L to market H and used for compliance in period t.[12]

Banking allows permits allocated in period 1 to be saved and used for compliance in period 2. The banking policy we consider adjusts and governs the compliance value of banked permits through a *banking conversion rate* δ_m that is specific to market m, whereby one period-1 m-permit banked in market m is equivalent to (i.e., fungible with) δ_m period-2 m-permits. Let B_m denote the amount of banked period-1 permits carried over into period 2 in market m under the assumption that borrowing is not allowed, i.e., $B_m \geq 0$.[13]

Equilibrium. For ease of exposition, we characterize the market equilibrium under different linking and banking policies with an optimization problem. We consider a representative agent seeking to minimize aggregate abatement costs across markets and time periods. Let e_{mt} and $C_{mt}(e_{mt})$ denote aggregate emissions and abatement cost functions in market m and period t. The functions C_{mt} are defined with respect to emissions (i.e., $-C'_{mt}$ denote marginal abatement costs) and satisfy standard assumptions (i.e., $C'_{mt} < 0$, $C''_{mt} > 0$). For simplicity, we further assume the discount factor is one and ignore positivity constraints on emissions.[14]

[11] Two broad mechanisms for exchange of permits when they trade across markets can be conceived. One is exchange at a conversion rate, analogous to currency exchange, with retirement of the original permit and re-issuance of a new one. In this approach, when a firm in H buys a permit from L, that L-permit is retired in L and ρ_t^{-1} H-permits are re-issued in H in its place. In the second approach, the value of a permit is recognized at the time of its use for compliance. That is, there is no retirement and re-issuance and every L-permit traded can directly be surrendered against ρ_t^{-1} units of emissions in H, and vice versa. We believe the second approach is likely to be preferred by policymakers, notably as it ensures that the denomination of permits is immutable and regulators can always keep track of permits by their origin and vintage.

[12] In some policy configurations, permits may flow in the "unnatural" direction from H to L (i.e., $\Delta_t < 0$).

[13] A few permit markets, such as the EU ETS, allow for limited (implicit) borrowing, but changing the constraint from $B_m \geq 0$ to $B_m \geq -\overline{B}_m$ with $\overline{B}_m > 0$ the borrowing limit does not affect our results.

[14] Alternatively, assuming $-C'_{mt}(e_{mt}) \to \infty$ in the limit as $e_{mt} \to 0$ ensures $e_{mt} > 0$ in equilibrium. This holds in the numerical simulations where marginal abatement cost functions are quadratic (see Section 4).

The resulting optimization problem is

$$\min_{e_{mt},\Delta_t,B_m} \sum_{t\in\{1,2\}} \sum_{m\in\{H,L\}} C_{mt}(e_{mt}) \tag{1a}$$

$$\text{subject to} \quad e_{H1} = \omega_{H1} + \Delta_1 - B_H \quad (p_{H1}) \tag{1b}$$

$$e_{L1} = \omega_{L1} - \rho_1\Delta_1 - B_L \quad (p_{L1}) \tag{1c}$$

$$e_{H2} = \omega_{H2} + \Delta_2 + \delta_H B_H \quad (p_{H2}) \tag{1d}$$

$$e_{L2} = \omega_{L2} - \rho_2\Delta_2 + \delta_L B_L \quad (p_{L2}) \tag{1e}$$

$$\Delta_1 \leq e_{H1} \quad (\gamma_H) \tag{1f}$$

$$-\rho_1\Delta_1 \leq e_{L1} \quad (\gamma_L) \tag{1g}$$

$$\text{and} \quad -B_m \leq 0 \quad (\lambda_m) \quad \text{for } m \in \{H, L\} \tag{1h}$$

The first four constraints given by (1b–1e) reflect policy compliance in each market and time period, i.e., emissions equal the allocation of permits net of both trade between markets and banking between periods.[15] The next two constraints given by (1f–1g) reflect constraints on the trade of permits between markets given that firms will not over-comply with the policy, i.e, the level of emissions places an upper bound on the total number of permits traded into the market and surrendered for compliance in period 1.[16] The final set of constraints given by (1h) reflects the assumption that borrowing is not allowed in either market. The shadow prices or Lagrange multipliers associated with all constraints are given within parentheses.

The standard assumptions about abatement cost functions ensure that the solution to this cost-minimization problem is equivalent to the competitive market equilibrium outcome with many firms in each market. Notably, the competitive equilibrium permit price in market m

[15]For example, constraint (1b) states that emissions in market H and period 1 must equal the number of permits allocated in that market and period plus the inflow of permits from market L (measured in H-permit equivalents) and minus the permits banked for use in period 2. Note that this constraint explicitly disallows under-compliance and over-compliance with the policy (the latter would not be cost minimizing).

[16]Note that combining constraints (1f–1g) with (1b–1c) yields an alternative set of constraints on banking, rather than trade: $B_m \leq \omega_{m1}$. These banking constraints reflect the fact that the number of permits banked cannot exceed the number of permits that are allocated domestically. As in Pizer & Yates (2015), we use the trade constraints rather than the (mathematically equivalent) banking constraints for the interpretation of the resulting Lagrange multipliers. Notably, this specification allows us to carefully disentangle the effects of and interactions between linking and banking policies.

and period t is equal to p_{mt}. The equilibrium is characterized by the first-order conditions

$$-C'_{m1}(e_{m1}) = p_{m1} - \gamma_m \quad \text{for } m \in \{H, L\} \tag{2a}$$

$$-C'_{m2}(e_{m2}) = p_{m2} \quad \text{for } m \in \{H, L\} \tag{2b}$$

$$p_{H1} - \gamma_H = \rho_1(p_{L1} - \gamma_L) \tag{2c}$$

$$p_{H2} = \rho_2 p_{L2} \tag{2d}$$

$$\text{and } p_{m1} = \delta_m p_{m2} + \lambda_m \quad \text{for } m \in \{H, L\} \tag{2e}$$

Equation (2a) shows that period-1 permit prices can be driven above marginal abatement costs. This wedge exists when $\gamma_m > 0$, which occurs when market m uses only traded foreign permits—no domestic permits—for compliance in period 1. In this case, all of market m's permits in period 1 are instead banked, so the period-1 price is driven by the permit's future compliance value in period 2, rather than its contemporaneous compliance value. As we will show, some linking policies create a strong incentive for this trading and banking pattern, resulting in this wedge between permit price and marginal abatement cost. Equation (2b) shows the more intuitive result that period-2 permit prices will always equal marginal abatement costs, because permits have no additional future value in the final period.

Equation (2c) gives the linking arbitrage condition for period 1. When $\gamma_H = \gamma_L = 0$, this condition states that relative permit prices in period 1 will be governed by the period-1 linking conversion rate ρ_1. As above, however, if $\gamma_m > 0$ for at least one market, then the m-permit's future value causes relative prices in period 1 to deviate from this specified ratio.[17] We will show that some linking policies, despite the use of a sequence of linking conversion rates intended to control relative prices over time, will yield immediate price convergence in line with future linking conversion rates due to this influence of future permit values (i.e., ρ_2 will apply already in period 1 and ρ_1 will be obsolete). Equation (2d) gives the linking arbitrage condition for period 2, showing that relative permit prices in period 2 will always be governed by the period-2 linking conversion rate ρ_2.

Equation (2e) gives the banking arbitrage condition for each market. When market m banks permits in period 1 (i.e., the no-borrowing constraint does not bind and $\lambda_m = 0$), the relative permit prices in market m are governed by the banking conversion rate δ_m, meaning the price in period 1 is linked to its future value in period 2 adjusted by δ_m. As we describe above,

[17] Combining conditions (2a) and (2c), we see that the relative marginal abatement costs in period 1 will always equal ρ_1, even though relative permit prices may not. As we describe above, $\gamma_m > 0$ when only one market's permits are used for compliance across both markets in period 1. In that case, that market's permit price—adjusted for the linking conversion rate ρ_1—sets the marginal abatement cost in both markets.

it is this intertemporal dependency that can create price wedges across markets in period 1. Alternatively, when the no-borrowing constraint binds and $\lambda_m > 0$, this intertemporal dependency is broken.

Gradual linking policies. To discipline the analysis of the different combinations of linking conversion rates (ρ_1, ρ_2) and banking conversion rates (δ_H, δ_L), we consider four types of illustrative linking policies in turn, as shown in Table 1. By proceeding in steps through special cases of the general first-order conditions in (2), we build intuition for why both linking and banking conversion rates are needed to implement gradual linking successfully, and how they should be calibrated together to that end.

Table 1: Main linking policies considered

Policy description	ρ_1	ρ_2	δ_H	δ_L
Full 1:1 linking in both periods with 1:1 banking	1	1	1	1
Delayed 1:1 linking in period 2 with 1:1 banking	n.a.	1	1	1
One-dimensional gradual linking policy	var.	1	1	1
Two-dimensional gradual linking policy	var.	var.	var.	var.

Note: n.a.: not applicable. var.: variable.

The first type of policy features two fully linked markets with unitary linking and banking conversion rates. This is the reference case that achieves the least cost for a given level of aggregate emissions. The second type delays linking, which is only allowed in period 2 but is anticipated in period 1. This case illustrates that the least cost may still obtain through banking and intertemporal arbitrage. The third and fourth types seek to implement gradual linking, respectively with linking conversion rates only, and with both linking and banking conversion rates. As we will show, the latter two-dimensional policy can guarantee gradual linking implemented as desired while the former one-dimensional policy does not in general, with potential perverse side-effects.[18]

3.2 Full 1:1 linking in both periods with 1:1 banking

We start with the reference least-cost case for the level of aggregate units of emissions given by $\sum_{mt} \omega_{mt}$, which features 1:1 linking in both periods and 1:1 banking in both markets. The

[18]We take the linking and banking conversion rates—as well as the emission caps—as given and do not formally consider the normative criteria underlying their selection.

first-order conditions in (2) can be rewritten as

$$\begin{aligned} -C'_{rt}(e^{fl}_{mt}) &= p^{fl}_{mt} \text{ for } m \in \{H, L\} \text{ and } t \in \{1, 2\}, \\ p^{fl}_{Ht} &= p^{fl}_{Lt} \equiv p^{fl}_{t} \text{ for } t \in \{1, 2\}, \\ \text{and } \lambda^{fl}_{m} &= p^{fl}_{m1} - p^{fl}_{m2} = p^{fl}_{1} - p^{fl}_{2} \equiv \lambda^{fl} \text{ for } m \in \{H, L\}, \end{aligned} \tag{3}$$

which characterize the full linking equilibrium, denoted by the superscript fl.

Result 1. *With 1:1 linking in both periods, prices and marginal abatement costs are equated across markets in each period. In addition, with 1:1 banking the single price is either constant, or it is decreasing across periods when the no-borrowing constraint is binding.*

Proof. See Appendix A.1. □

Result 1 is the analog of the Rubin-Schennach result for a single market with 1:1 banking (since 1:1 linking is effective in both periods, the two markets are de facto merged into one).[19] The possible price differential between periods reflects the shadow price of the no-borrowing constraint (i.e., the cost saving from borrowing one permit if this were allowed). Whether the no-borrowing constraint is binding or not depends on the permit allocation between periods in each market (i.e., the larger the allocation in period 1 is relative to period 2, the less likely the constraint will be binding).

Importantly, there typically are multiple least-cost solutions, i.e., different combinations of the flexibility levers provided by linking and banking can sustain the least-cost paths $\{e^{fl}_{mt}\}$ and $\{p^{fl}_{t}\}$.[20] For instance, given a least-cost solution, the least cost may also obtain by (i) trading one permit less in period 1 and one permit more in period 2; (ii) banking one permit less in H and one permit more in L. In fact, for a stock pollutant, the two flexibility levers are substitutable to some extent in that there are different equivalent allocations of the overall least-cost volume of trade across periods so long as banking can be adjusted accordingly.

[19] There is no need to shift "extra" period-1 L-permits to H for compliance in period 2 through increased trading in period 1, increased banking in L, and reduced banking in H, because trading can also take place in period 2. Formally, $\Delta^{fl}_{1} < e^{fl}_{H1}$ (or equivalently $B^{fl}_{H} < \omega_{H1}$) and $\gamma^{fl}_{H} = 0$ always holds.

[20] For interior solutions (i.e., when boundary trading or banking conditions are not binding), the set of first-order conditions and emission constraints does not uniquely determine trading and banking levels (only prices are uniquely pinned down). This is a standard feature in models with perfect arbitrage, linear constraints, and no convex trading costs or banking adjustment costs.

3.3 Delayed 1:1 linking in period 2 with 1:1 banking

Pushing the above reasoning one step further, the least cost may obtain even in the absence of linking in period 1 (i.e., only with 1:1 linking in period 2 and 1:1 banking across periods) provided that 1:1 linking in period 2 is anticipated in period 1 and sufficient banking occurs.

Result 2. *If 1:1 linking delayed to period 2 but is anticipated in period 1, the least-cost emission and price paths obtain when banking is positive in both markets. In this case, 1:1 price convergence occurs already in period 1 even though the 1:1 linkage is not yet effective.*

Proof. See Appendix A.2. □

Result 2 shows that linking in period 1 can be dispensed with so long as banking is positive in both markets and establishes a full connection between markets in both periods. Specifically, because $p_{H2} = p_{L2}$ holds through 1:1 linking in period 2, when $p_{m1} = p_{m2}$ also holds through banking and arbitrage in each market, then $p_{H1} = p_{L1}$ even in the absence of 1:1 linking in period 1.[21] This occurs when banking provides sufficient flexibility to shift the overall least-cost volume of trade to period 2. Even though permits are not directly fungible in period 1, they may indirectly be traded by adjusting banking and subsequent trade in period 2.

Result 2 also shows that delaying or announcing 1:1 linking at a future date can cause price paths to converge through banking and arbitrage before the link is effective.[22] Importantly, this backpropagation mechanism that sustains "early" price convergence is the same as the one that is at play in the case of gradual linking considered in the following sections. Of course, early convergence may be partial or may not occur at all when banking and arbitrage are constrained or when the implementation of the link is deemed uncertain.[23]

[21] Formally, this holds when the least-cost paths in (3) are the solution to (1) with $\rho_2 = 1$, $\delta_H = \delta_L = 1$, and (1f) replaced by $\Delta_1 = 0$. This depends on the feasibility of the least-cost emission paths relative to the permit allocation profiles. The more frontloaded permit allocation is, the more likely this feasibility is.

[22] Because 1:1 banking in both markets causes prices to equalize immediately, policies using only banking conversion rates to achieve gradual price alignment could be devised (without trading in period 1 and with 1:1 trading in period 2). This type of linking policy is arguably unlikely, as it lacks a linking spirit, and we do not derive its properties analytically. However, we consider it in numerical illustrations in Section 4.2.4, notably because it showcases the main effects of banking conversion rates.

[23] Formally, if ϕ denotes the subjective probability of linking in period 2 (as viewed by market participants in period 1) and banking is positive in market m, then $p_{m1} = p_{m2} = \phi p^{2l}_{m2} + (1-\phi) p^{a}_{m2}$, where the superscripts $2l$ and a refer to period-2 linking and autarchy. In this case, early price convergence occurs only partially.

3.4 One-dimensional gradual linking policy

We now turn to the case where policymakers seek to implement linking gradually using only a sequence of linking conversion rates to govern the ratios of prices across markets over time. We call this a "one-dimensional" gradual linking policy, as banking conversion rates are not used. Specifically, we consider that linking is effective in both periods, with policymakers seeking to have permits trading at parity no sooner than in period 2 (i.e., $\rho_1 > \rho_2 = 1$) but without companion banking policy being implemented in parallel (i.e., $\delta_H = \delta_L = 1$).

Result 3. *Let $\rho_1 > \rho_2 = 1$ and $\delta_H = \delta_L = 1$. When banking is positive in both markets, permits trade at parity already in period 1, i.e., ρ_1 is mute and $\rho_2 = 1$ applies in both periods. When banking is zero in at least one market, permits trade at parity only in period 2 and the price ratio in period 1 is in general set by ρ_1.*

Proof. See Appendix A.3. □

Result 3 shows that full 1:1 price convergence already occurs in period 1 when both markets have positive banking, even though the current linking conversion rate is above one. Because prices are expected to have the same price in period 2 due to 1:1 linking, prices immediately converge on identical Hotelling paths by arbitrage, which makes ρ_1 obsolete for governing the price ratio in period 1. Only in "corner" cases where banking is zero in one or both markets may ρ_1 be effective in period 1, as the intertemporal connection between markets breaks down.[24] This connection is more likely to hold in full when both markets have strong incentives to bank, which is more likely to happen when permit allocation is frontloaded relative to the least-cost emission paths. It follows that gradual linking requires market prices to remain on separate Hotelling paths, an objective which requires an additional instrument to adjust banking incentives accordingly. Such "two-dimensional" gradual linking policies, using both linking and banking conversion rates, are studied in the next section.

A striking feature of the leveling of the terms of trade over time (i.e., $\rho_1 > \rho_2 = 1$) is that the quantity of emissions allowed may increase, despite the fixed number of permits. Firms will seek to shift as much trade as possible from L to H from period 1 to period 2 to minimize the trade-induced contraction in overall emissions and associated compliance costs. This is because trading ρ_1 L-permits to H in period 1 allows emissions in H to increase by one unit

[24] It suffices that the connection breaks down in one market for the price ratio in period 1 to be governed by ρ_1, as the period-1 price in that market reflects the current terms of trade. Moreover, in some edge cases, the price ratio in period 1 is ambiguous and comprised strictly between 1 and ρ_1 (see Appendix A.3).

and constrains emissions in L to decrease by ρ_1 units, whereas trading in period 2 is emission-preserving. Crucially, it may even be profitable for firms to trade from H to L in period 1, if the gain from relaxing the overall emission constraint outweighs the cost inefficiency of that trade. For each H-permit traded in this way, firms in L can bank the equivalent ρ_1 L-permits and then sell only one L-permit to effectively return that H-permit to firms in H in period 2, which increases emissions by $\rho_1 - 1$. Overall, all firms will prefer to have L-permits banked, and firms in H will only bank once firms in L have banked all of their period-1 permits.

Result 4. *Let $\rho_1 > \rho_2 = 1$ and $\delta_H = \delta_L = 1$. Linking leads to increased total emissions when some trade occurs in the cost-inefficient direction in period 1, which is more likely to occur the higher ρ_1 is and the more firms have room to adjust banking and shift trade to period 2.*

Proof. See Appendix A.4. □

Result 4 extends the result in static linking models that, despite the fixed number of permits, total emissions vary as the compliance value of permits is modified when they trade across markets due to linking conversion rates. Specifically, in a static setting, linking conversion rates may lead to increased emissions, but for that to happen, they must be set in an arguably poor manner ex ante—namely, when ρ is above the ratio of autarchy prices p_H/p_L or below 1 (Quemin & de Perthuis, 2019; Woerman, 2023).[25] In a dynamic setting with banking, total emissions may increase even when ρ_1 is set adequately from a static viewpoint—namely, when ρ_1 is set between 1 and the ratio of autarchy prices p_{H1}/p_{L1}. Total emissions are more likely to increase relative to no linking the more the terms of trade are distorted across periods (i.e., the higher ρ_1 is) and the more L (resp. H) can increase (resp. decrease) banking.[26] Notably, because this happens through some trade from the high-cost market to the low-cost market, the cost effectiveness of meeting the relaxed emission constraint is also reduced.

Together, Results 3 and 4 show that introducing a linking policy without a companion banking policy in markets that allow banking is likely to result in a failed attempt at implementing gradual linking. Notably, equilibrium price ratios under a one-dimensional gradual linking policy are state-dependent in the sense that they depend on the demand for banking, and ultimately, on the allocation of permits across the two periods. In addition, when prices gradually converge as intended, a one-dimensional policy typically leads to higher total emissions

[25] When $\rho > p_H/p_L$, the dispersion of marginal abatement costs and prices is greater than under autarchy, abatement occurs where it is more expensive, and permits flow in the "unnatural" direction ($H \to L$). When $\rho < 1$, the initial ranking of marginal abatement costs and prices between markets is reversed, and although permits flow in the "natural" direction ($L \to H$), $\rho < 1$ means each traded permit increases total emissions.

[26] That is, the more permit allocations are frontloaded relative to least-cost emission paths, as L has more room to bank permits and H has more room to sell permits to L in period 1.

than provided for by emission caps, driven by some cost-inefficient trading. That is, gradual linking may only be implemented—if at all—at the cost of increased emissions and reduced cost effectiveness. Finally, a one-dimensional policy gives no flexibility to policymakers in shaping the price paths to 1:1 linking when it happens to be successful at controlling price ratios over time. In that case, H's price path can only be decreasing over time to reach L's (flat) price path from above.[27] This provides limited room to address the impacts of direct 1:1 linking that gradual linking is meant to smooth out and cushion.

3.5 Two-dimensional gradual linking policy

We finally consider the general case where policymakers use linking and banking conversion rates jointly, which we call a "two-dimensional" gradual linking policy. This has the potential to address the different limitations of a one-dimensional policy identified above.

Result 5. *The rate of price convergence across linked markets is governed either by the ratio of linking conversion rates or the ratio of banking conversion rates. In general,*

$$\frac{p_{H1}}{p_{L1}} \Big/ \frac{p_{H2}}{p_{L2}} = \begin{cases} \delta_H/\delta_L \text{ when banking is positive in both markets,} \\ \rho_1/\rho_2 \text{ when banking is zero in one or both markets.} \end{cases}$$

Therefore, the targeted price ratios of ρ_1 in period 1 and ρ_2 in period 2 are always sustained in equilibrium provided that banking conversion rates satisfy $\delta_H/\delta_L = \rho_1/\rho_2$.

Proof. See Appendix A.5. □

Result 5 shows that the price ratio across markets in period 1 may effectively be governed by the targeted linking conversion rate ρ_1 only in the "corner" cases where banking is zero in one or both markets (spatial arbitrage prevails).[28] In the "interior" cases where banking is positive in both markets, banking conversion rates supersede linking conversion rates and govern the rate of price convergence across markets (intertemporal arbitrage prevails). Therefore, for gradual linking to occur as intended in all cases—notably if one expects interior cases with positive banking to be most relevant in practice—banking conversion rates must be set in line with the targeted linking conversion rates as per the "rule" $\delta_H/\delta_L = \rho_1/\rho_2$. In comparison, recall that equilibrium price ratios under a one-dimensional policy are state-dependent. More

[27] This occurs because H does not bank in that case. Using a one-dimensional policy, L's price path cannot be increasing over time to reach H's price path from below.

[28] In some edge cases, the rate of price convergence is not unambiguously determined (see Appendix A.5).

generally, Result 5 reflects that gradual price convergence across linked markets requires the use of banking conversion rates to drive a wedge between price paths that are otherwise only governed by time discounting and shadow prices associated with banking constraints.

"Well-set" banking conversion rates adjust banking incentives for the distortion in trading incentives across periods arising from the use of different linking conversion rates. Recall that setting $\rho_1 > \rho_2$ incentivizes firms to maximize trade from L to H in period 2 and possibly even trade from H to L in period 1, with banking as large as possible in L and as low as possible in H. These strategies can be countered by making banking in L relatively less attractive than banking in H, which occurs with banking conversion rates set according to $\delta_H/\delta_L = \rho_1/\rho_2 > 1$.[29] More specifically, when this "rule" is satisfied, banking conversion rates preserve the economic value of permits along the entire time-market chain, accounting for the changes in traded permits' compliance value as governed by linking conversion rates. At the margin of an interior equilibrium, the incentive for firms in L to {trade less in period 1, bank more, trade more in period 2} is on par with the incentive to {trade more in period 1, bank less, trade less in period 2}—and vice versa for firms in H. It follows that there are multiple equivalent trading-banking "routes" to move permits between any two time-markets.[30]

An important policy aspect of Result 5 is that multiple combinations of banking conversion rates can implement the same targeted sequence of price ratios. This gives policymakers some flexibility to deal with the price, economic, and environmental impacts of gradual linking. Regarding the price paths to 1:1 linking in the future, recall that δ_m governs the slope of the price path in market m. For instance, with $\delta_H > 1$ and $\delta_L < 1$, price convergence over time may unfold with a gradually decreasing price path in H and a gradually increasing price path in L. Regarding environmental and cost effectiveness, one has the following results.

Result 6. *Let $\rho_1 > \rho_2 = 1$ and $\delta_H/\delta_L = \rho_1$. There is at least one pair of banking conversion rates that is emission-preserving with respect to overall emission caps. Relative to 1:1 linking from period 1 with 1:1 banking, the economic cost of gradual linking at constant overall emissions is increasing with ρ_1. Relative to autarchy in both periods with 1:1 banking, cost effectiveness is improved provided that both periodic prices are reduced in H and increased in L as a result of gradual linking, respectively, which can occur only if $\delta_H > 1$ and $\delta_L < 1$.*

Proof. See Appendix A.6. □

[29] For instance, if $\rho_1 = 2\rho_2$, setting $\delta_H = 2\delta_L$ neutralizes these strategies: An L-permit is twice as valuable for compliance in H in period 2 as in period 1, but because the compliance value in L of a banked L-permit is halved in period 2, the two effects cancel each other out.

[30] The multiplicity of least-cost solutions in Sections 3.2 and 3.3 is a special case of the model indeterminacy.

Result 6 demonstrates that gradual price alignment can be implemented as intended without any loss of environmental effectiveness. However, achieving both of these policy objectives comes at the cost of reduced cost effectiveness, as compared to 1:1 linking from the start. Intuitively, the economic cost of gradual policy is inversely related to the speed of the transition: The slower the convergence to 1:1 linking (i.e., the larger ρ_1), the larger the resulting marginal abatement cost distortions and, consequently, the overall economic burden. Importantly, however, cost effectiveness is improved relative to autarchy when permit prices are reduced in H and increased in L in both periods, respectively, which is the gradual linking policy configuration that can be expected. Overall, Result 6 extends the result in static linking models that, when ρ is set between 1 and the ratio of autarchy prices p_H/p_L, total emissions are reduced, while cost effectiveness is improved relative to autarchy but always lower than under 1:1 linking (Quemin & de Perthuis, 2019; Woerman, 2023).

Finally, note that policymakers may wish to deviate from the emission-preserving banking conversion rates (e.g., to gain some flexibility over the slope of the price path in each market) or they may be politically constrained to do so. In this case, increasing (resp. decreasing) both banking conversion rates along the isoline $\delta_H/\delta_L = \rho_1$ will tend to increase (resp. decrease) emissions.[31] In Section 4, we will numerically explore the market and distributional outcomes for the universe of possible price ratios in period 1 and banking conversion rates.

3.6 Extensions

For expositional clarity, we elaborated the model for two markets and two time periods only. This suffices to establish the general mechanisms and dynamics at work. However, the basic framework can easily be extended to multiple periods, multiple markets, and uncertainty.[32] Below, we sketch how our results generalize and discuss additional considerations that would apply to such extensions. This might provide helpful guidance for future, more policy-oriented work in this direction.

Multiple markets. Gradual linking may involve more than two markets. The necessary relation between linking and banking conversion rates to ensure gradual price convergence in Result 5 naturally extends to an arbitrary number of markets. Let $\rho_t^{m\tilde{m}}$ denote the linking

[31] Policymakers may also deviate from the "rule". In that case, setting $\delta_H/\delta_L >$(resp. $<$)ρ_1/ρ_2 entails lower or slower (resp. greater or faster) price convergence than targeted because banking incentives are "over" (resp. "under") adjusted w.r.t. trading incentives, so the price ratio is higher (resp. lower) than ρ_1 in period 1.

[32] The numerical model in Appendix B also allows for multiple firms within each market.

conversion rate when permits flow from market m to market $\tilde{m}$ in period t. In a multilateral linkage, linking conversion rates must be set consistently, i.e., $\rho_t^{m\tilde{m}} = 1/\rho_t^{\tilde{m}m}$ and $\rho_t^{m\tilde{m}} = \prod_{k=m}^{\tilde{m}-1} \rho_t^{kk+1}$ for all pair of markets $(m, \tilde{m})$. Assuming positive banking in all markets, for any market pair $(m, \tilde{m})$ one has

$$\frac{\delta_m}{\delta_{\tilde{m}}} = \frac{\rho_1^{m\tilde{m}}}{\rho_2^{m\tilde{m}}}, \tag{5}$$

showing that any combinations of targeted price ratios $(\rho_1^{m\tilde{m}}, \rho_2^{m\tilde{m}})$ can be sustained in equilibrium with a combination of banking conversion rates calibrated to the associated changes in linking conversion rates across periods for each market pair. In other words, multilateral gradual linking can be seen and implemented as a combination of bilateral gradual linkages, provided linking conversion rates are defined consistently across all markets.

Multiple periods. Gradual linking may unfold over more than two periods. The necessary relation between linking and banking conversion rates to ensure gradual price convergence in Result 5 naturally extends to an arbitrary number T of periods. Let δ_m^t denote the banking conversion rate between periods t and $t+1$ in market m. Assuming positive banking in both H and L until period T, the intertemporal arbitrage condition yields $\prod_{t=1}^{T-1} \delta_H^t/\delta_L^t = \rho_1/\rho_T$. Importantly, this can be unpacked into

$$\frac{\delta_H^1}{\delta_L^1} \times \cdots \times \frac{\delta_H^{T-1}}{\delta_L^{T-1}} = \frac{\rho_1}{\rho_2} \times \cdots \times \frac{\rho_{T-1}}{\rho_T}, \tag{6}$$

showing that a targeted sequence of price ratios $(\rho_1, \ldots, \rho_T)$ can be sustained in equilibrium with a sequence of banking conversion rates calibrated to the associated changes in linking conversion rates across periods, i.e., $\delta_H^t/\delta_L^t = \rho_t/\rho_{t+1}$. In other words, dynamic gradual linking can be seen and implemented as a sequence of one-on-one gradual linkages. Notably, gradual convergence to 1:1 linking may unfold in any chosen number of steps. As a rule, the more steps there are, the smaller the ρ_t/ρ_{t+1} ratios need to be, and hence the smaller the deviations from 1:1 banking implied by the associated δ_H^t/δ_L^t ratios are.

Uncertainty and shocks. So long as policymakers commit to a given sequence of price ratios $(\rho_1, \ldots, \rho_T)$, the introduction of uncertainty or shocks in the model does not change the definition of the companion sequence of banking conversion rates, as the same intertemporal conditions hold in expectations. However, with the realization of shocks or the arrival of new information, policymakers may want to adjust the pace of price convergence. Because our gradual linking policy rule can be broken down period-on-period and because it is independent

of history,[33] it can accommodate such policy updates. For instance, if in period $\tilde{t}$ policymakers revise future price ratios to $(\tilde{\rho}_{\tilde{t}+1}, \ldots, \tilde{\rho}_T)$, then it suffices to adjust the banking conversion rates such that $\delta_H^{\tilde{t}}/\delta_L^{\tilde{t}} = \rho_{\tilde{t}}/\tilde{\rho}_{\tilde{t}+1}$ and $\delta_H^t/\delta_L^t = \tilde{\rho}_t/\tilde{\rho}_{t+1}$ for $t \geq \tilde{t}+1$.[34] In particular, price convergence may be accelerated, slowed down, or even reverted.

3.7 Alternative policies

Finally, we consider alternative policies to implement gradual linking that may have traction with policymakers, namely quantitative limits. For instance, trade volume may be limited to an increasing share of the domestic cap over time. Arguably, the appeal of these policies lies in their apparent simplicity, the perception that they offer direct control over the extent of linking, and their ability to preserve aggregate emissions. However, we show that they can fail to achieve gradual price alignment and may even produce unintended outcomes.

Studying these policies requires a different problem formulation, as laid out in Appendix A.7. In principle, in a deterministic framework, for each pair of linking conversion rates (ρ_1, ρ_2) and banking conversion rates (δ_H, δ_L), one may find a "dual" set of quantitative constraints on trade in both periods and on banking in both markets that leads to the same equilibrium outcomes. Intuitively, just like linking and banking conversion rates are needed to address the spatial and temporal dimensions of gradual alignment of prices, quantitative limits on trade and banking are needed in the dual problem.

Interestingly, however, conversion rates and quantitative limits are not equivalent. Generally, a quantitative limit must be binding to have an effect on the price path in equilibrium. If the constraint is not binding, perhaps due to regulator error or secular market changes, then price convergence will have occurred sooner than intended. In contrast, a conversion rate is always effective at the margin, affecting the permit price ratio in each period and intertemporal arbitrage. The complex interaction between banking and trading—as demonstrated before in the case of one-dimensional policies using only linking conversion rates—may even result in price divergence over time with quantitative limits on trade.

Various practical issues also arise with quantity-based policies. First, binding limits generate economic rents whose distribution among firms and markets is not clear ex ante. Second, by

[33] To see this, note that the sequence of banking conversion rates that sustains $(\rho_1, \ldots, \rho_T)$ in equilibrium can also be defined recursively starting backwards from period T, and that our results hold independently of the initial stock of permits held by firms.

[34] As in the deterministic case, we do not consider the normative criteria underlying the selection of or the adjustment in the linking conversion rates.

analogy to quantitative limits on offset use, the most likely approach to implementation of quantitative limits is to impose a constraint on the surrender ratio of permits by each firm. Limits on banking would likely require holding limits on foreign permits by each firm. In contrast, conversion rates apply at the market rather than firm level, facilitating implementation and calibration. Finally, a quantity-based approach also has efficiency implications because quantitative limits blur the information contained in prices, which is the informational basis for regulatory oversight, and the reduction in the informational efficiency of prices is exacerbated by banking and intertemporal arbitrage.

4 Numerical approach

In this section, we use a numerical model to illustrate the effects of and trade-offs inherent to different policy choices to implement gradual linking. The EU is considered as a case study because linking between two separate carbon markets is an imminent issue (although direct price convergence may be a political obstacle) and the initial price discrepancy is substantial (at least in our calibration). This makes the case for gradual linking especially salient and allows our simulations to span a wide range of policy parameter values (and in turn a large array of policy choices and possible outcomes).[35] We proceed in two steps: First, we describe the model calibration. Second, we present the numerical results and identify the policy effects and trade-offs in the different policy scenarios under consideration.

4.1 Model calibration

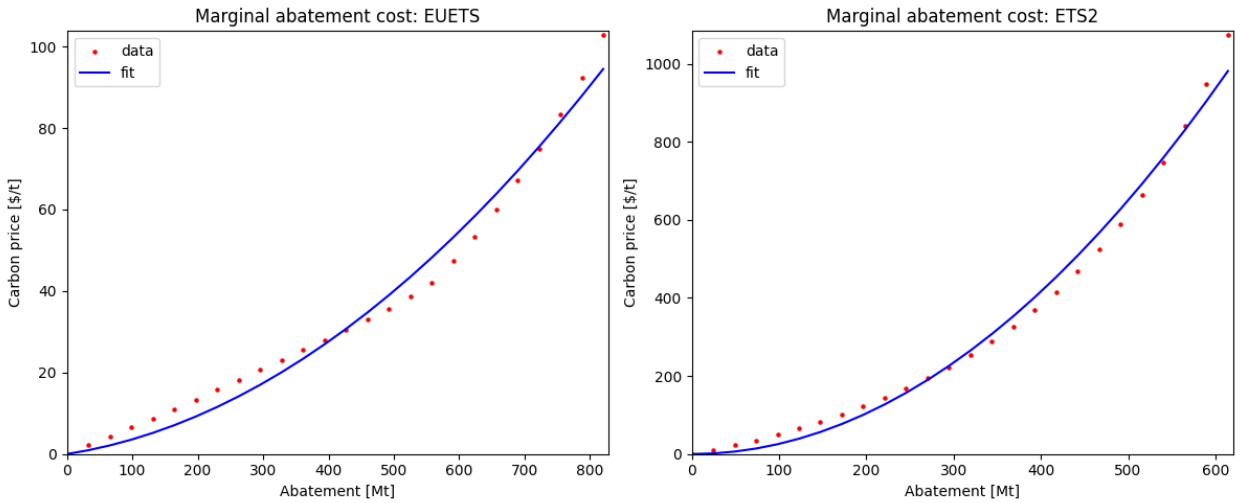
Following our theoretical approach, the numerical model is formulated as a partial equilibrium model. A rigorous model formulation is provided in Appendix B.1. The numerical model is parameterized by calibrating the coefficients of the marginal abatement cost (MAC) curves and baseline emissions for two markets—the EU ETS for large installations in the electricity, heat and industrial sectors launched in 2005 (ETS1) and that for road transport, buildings and smaller installations in other sectors to be launched in 2028 (ETS2). Linking between

[35] Arguably, this large discrepancy is also a drawback in terms of realism of the considered gradual linking policy using only one intermediary step to 1:1 linking (i.e., the two-period model may require unrealistically large one-period adjustments in policy parameter values). In practice, the transition to 1:1 linking between heterogeneous markets would likely happen in a smoother way (i.e., using more steps), which would require smaller periodic adjustments to the policy parameter values. For instance, (6) shows that with more steps, the ρ_t/ρ_{t+1} ratios can be closer to 1, hence the δ_H^t/δ_L^t ratios can be closer to 1 also, which can be achieved with smaller deviations from 1:1 banking in both markets.

the two markets is not planned currently, but may be considered from 2030 on.

We follow the literature deriving the MAC coefficients based on a global computable general equilibrium (CGE) model (see e.g., Böhringer et al., 2006; Klepper & Peterson, 2006). The methodology is described in detail in Abrell & Rausch (2017). The static multi-sector multi-region model is based on the GTAP 10 database (Aguiar et al., 2019) using 2014 as base year. The CGE resolves all EU countries with sufficient sectoral details to represent the coverage of the two ETSs. For each ETS, we impose a sequence of reduction requirements ranging from 0 to 50% in steps of 2%. For each sector and reduction requirement, the CGE model produces the price that clears the permit market and the associated abatement level. To determine the coefficients of the MAC curves, we fit a quadratic function under the restriction that the intercept is zero (Figure 1). Baseline emissions are determined for the base year 2014 and taken from the CGE model (Table 2). All $ values are denominated in USD_{2014}.

Figure 1: Numerical marginal abatement cost curves



Note: Red dots are the price-abatement pairs obtained from the CGE model and the blue lines are the MAC curves fitted under the assumption of a quadratic function with zero intercept.

The model parameters are assumed to be constant over time, i.e., we assume that no technological progress takes place and that baseline emissions stay constant over time.[36] We make these assumptions for lack of good or uniformly agreed parameters characterizing technological progress for abatement technologies or the permanence of emission reductions.

[36]This is without loss of generality for the purpose of illustrating our theoretical results, which hold under the assumption that abatement cost functions are exogenously given and independent of the permit price.

Table 2: Numerical model parameters

	Marginal Abatement Cost Function Terms			Baseline Emissions
	Constant [$]	Linear [$/t]	Quadratic [$/t^2]	[10^6t]
ETS1	0	0.025025	0.00011	1640
ETS2	0	0	0.002601	1280

4.2 Numerical results

Following our theoretical approach, the numerical approach also considers two periods, 2027 and 2030.[37] The absolute emission reduction targets are derived from the annual emission caps provided in current legislation in the ETS1 and ETS2 perimeters (Table 3). In the main scenarios considered, we assume that all permits are frontloaded in each market, i.e., permits associated with the emission caps for both periods are allocated in the first period.[38]

Table 3: Emission reduction targets

	2014	2027		2030	
	Observed [Mt]	Target [Mt]	Reduction	Target [Mt]	Reduction
ETS1	1630	1095	33%	825	49%
ETS2	1350	1036	23%	849	37%
Total	2980	2131	28%	1674	44%

Note: Observed ETS1 emissions are taken from the European Environmental Agengy's Emissions Trading viewer. We use the value for the year 2014 including the correction for scope expansion and excluding non-EU countries. Observed ETS2 emissions are the 2015 emissions estimated by Graichen & Ludig (2024), who do not provide a value for 2014. Note that historical emissions and baseline emissions from the CGE model are line (compare Tables 2 and 3). Emission targets (caps) are taken from Günther et al. (2025) and Pahle et al. (2025) based on current legislation. The reduction column provides the reduction relative to the year 2014.

In the following, we consider different types of linking policies through which policymakers seek to implement gradual price alignment, with ETS1 and ETS2 permits trading at parity no sooner than in the second period. Before proceeding, we present the reference scenarios

[37] Since the original calibration of the model, the start of ETS2 has been postponed by one year, to 2028. For simplicity, we assume a zero continuation value for banked permits at the end of the second period. A positive continuation value would only affect our results quantitatively. Moreover, as banking volumes would increase, the need to use both linking and banking conversion rates would become even more salient.

[38] Technically, full frontloading is considered as a way to relax the constraint that permit banking is capped by the amount of permits allocated in the first period. This allows us to focus on interior equilibria and to span a wider range of relevant outcomes in the simulations. Practically, the assumption of full frontloading reflects the fact that the banking constraint may otherwise be overly restrictive given the existing permit bank in ETS1, and that a sizable volume of permits will initially be frontloaded in ETS2. We provide additional results without frontloading in Appendix B.2.

(or basic policies) in which and we only use 1:1 conversion rates and simply turn on and off banking in each market and trading between markets.

4.2.1 Reference scenarios

Under *No Flexibility*, neither banking in each market nor permit trading between markets is allowed. Hence, permit prices diverge across time and markets (Table 4). Because reduction targets increase over time, prices also increase across periods. ETS2 especially shows a sharp price increase due to its steep MAC curve (Figure 1). Given the different MAC curves and targets, prices also show significant divergence across markets. Thus, introducing flexibility in the form of banking or inter-market trading promises sizable efficiency gains.

Table 4: Reference scenarios: Abatement and prices

		L = ETS1				H = ETS2		
		Bank	Abate	Price	Trade	Bank	Abate	Price
No Flexibility	Period 1	-	541	46	-	-	283	208
	Period 2		804	91	-		455	537
Full Banking	Period 1	968	673	67	-	860	369	353
	Period 2		673	67	-		369	353
Full Trade	Period 1	-	665	65	124	-	159	65
	Period 2		1025	141	222		233	141
Full Flexibility	Period 1	927*	845	100	213*	900*	196	100
	Period 2		845	100	132*		196	100

Note: The table reports permit banking (Bank, $MtCO_2$), abatement (Abate, $MtCO_2$), permit prices (Price, $\$/tCO_2$), and permit trading from ETS1 to ETS2 (Trade, $MtCO_2$) in the reference scenarios. Each scenario implements the emission targets as provided in Table 3 without frontloading. *No Flexibility* does not allow for trade or banking. *Full Banking* allows for 1:1 banking ($\delta_{ETS1} = \delta_{ETS2} = 1$) but not for trade. *Full Trade* allows for 1:1 trade ($\rho_1 = \rho_2 = 1$) but not for banking. *Full Flexibility* allows for 1:1 banking and 1:1 trade ($\rho_1 = \rho_2 = 1$ and $\delta_{ETS1} = \delta_{ETS2} = 1$). The superscript $\star$ indicates non-unique model outcomes.

Under *Full Banking*—but without permit trading between markets—each market banks permits across periods in order to equate MACs and in turn prices over time (Table 4). This intertemporal flexibility reduces abatement costs in each market, achieving a 12% reduction in aggregate (Table 5). As trade between markets is not allowed, prices still diverge across markets, indicating the potential for further efficiency gains by linking the two markets.

Under *Full Trade*, permits can trade between markets, but banking is not allowed. Because ETS1 has the lower prices under autarchy, permits are exported from ETS1 to ETS2 in each period. Accordingly, abatement increases (resp. decreases) in ETS1 (resp. ETS2), until

Table 5: Reference scenarios: Abatement Costs

	L = ETS1		H = ETS2		Total	
	Cost (Billion $)	Change vs. No Flexibility	Cost (Billion $)	Change vs. No Flexibility	Cost (Billion $)	Change vs. No Flexibility
No Flexibility	36.6	-	101.0	-	137.6	-
Full Banking	33.6	-8%	86.8	-14%	120.5	-12%
Full Trade	69.0	89%	14.4	-86%	83.4	-39%
Full Flexibility	62.2	70%	13.0	-87%	75.2	-45%

Note: The table reports abatement costs aggregated over the two periods in the reference scenarios. Each scenario implements the emission targets as provided in Table 3 without frontloading. *No Flexibility* does not allow for trade or banking. *Full Banking* allows for 1:1 banking ($\delta_{ETS1} = \delta_{ETS2} = 1$) but not for trade. *Full Trade* allows for 1:1 trade ($\rho_1 = \rho_2 = 1$) but not for banking. *Full Flexibility* allows for 1:1 banking and 1:1 trade ($\rho_1 = \rho_2 = 1$ and $\delta_{ETS1} = \delta_{ETS2} = 1$).

MACs and in turn prices converge across markets in each period (Table 4). Across periods, however, price divergence persists as banking is not allowed. This spatial flexibility reduces total abatement costs by 39% compared to *No Flexibility*, but it also induces substantial shifts in abatement costs from ETS2 to ETS1 (Table 5).

Under *Full Flexibility*, both trading and banking are allowed. As a result, MACs and in turn prices converge across markets and periods (Table 4, Result 1).[39] Allowing for both flexibility dimensions (trading between markets and banking across periods) reduces total abatement costs by 45%, again shifting substantial costs from ETS2 to ETS1 (Table 5).

Note that both flexibility dimensions must be leveraged to achieve the least-cost allocation of abatement. Of the two, trading realizes the larger part of the efficiency gains, essentially because of the larger difference in MACs between markets than across periods in each market. However, this direct 1:1 linking of markets causes sudden, large changes in permit prices and abatement costs in each market. To smooth out this "policy shock" while still achieving much of the efficiency gains, policymakers might aim to allow for full trade only in the future while controlling the rate of price convergence through a gradual linking policy.

4.2.2 Policies delaying linking

Delaying linking until the second period may seem to offer a solution, but it may not succeed if future trading is anticipated and banking sustains a connection between the two periods (Result 2). Below, we provide an illustration by disallowing and allowing banking.

[39] As discussed in Section 3.2, multiple combinations of permit banking and trading volumes can implement the least-cost outcome, and Table 4 shows only one of them, indicated by *.

Under *No Banking*, permits cannot be banked or traded across markets in period 1. Hence, prices differ between markets in period 1 and across periods in each market. In period 2, once markets are linked on a 1:1 basis, permits are exported from ETS1 to ETS2, equating MACs and prices (Table 6). In other words, gradual linking is successfully achieved, although the policymaker has little control over relative prices in period 1. This achieves a cost reduction of 33% compared to *No Flexibility*, which is nearly as much as *Full Trade* and 73% of *Full Flexibility*, but again shifts costs from ETS2 to ETS1 (Table 7).

Table 6: Delayed linking: Abatement and prices

		L = ETS1				H = ETS2		
		Bank	Abate	Price	Trade	Bank	Abate	Price
No Banking	Period 1	-	541	46	-	-	283	208
	Period 2		1025	141	222		233	141
Banking	Period 1	1141	845	100	-	687	196	100
	Period 2		845	100	346		196	100

Note: The table reports permit banking (Bank, $MtCO_2$), abatement (Abate, $MtCO_2$), permit prices (Price, $\$/tCO_2$), and permit trading from ETS1 to ETS2 (Trade, $MtCO_2$) when 1:1 trading is allowed only in period 2 but is anticipated in period 1. *No Banking* does not allow for banking and the permit allocation coincides with the emission target in each period as in Table 3. *Banking* allows for banking and all permits are allocated in period 1 in each market (full frontloading with $\delta_{ETS1} = \delta_{ETS2} = 1$).

Under *Banking*, trading equates prices across markets in period 2 as under *No Banking*. But because prices are now equated across periods in each market due to banking, prices are also equated across markets in period 1 (Table 6). That is, prices fully converge in period 1 even though markets are not yet linked. In fact, this policy provides sufficient flexibility to achieve the *Full Flexibility* outcomes, in terms of efficiency gains and cost allocation (Table 7).

These results show that delaying linking will actually be implemented only when banking is

Table 7: Delayed linking: Abatement Costs

	L = ETS1		H = ETS2		Total	
	Cost (Billion \$)	Change vs. No Flexibility	Cost (Billion \$)	Change vs. No Flexibility	Cost (Billion \$)	Change vs. No Flexibility
No Banking	62.1	70%	30.5	-70%	92.7	-33%
Banking	62.2	70%	13.0	-87%	75.2	-45%

Note: The table reports abatement costs aggregated over the two periods when 1:1 trading is allowed only in period 2 but is anticipated in period 1. *No Banking* does not allow for banking and the permit allocation coincides with the emission target in each period as in Table 3. *Banking* allows for banking and all permits are allocated in period 1 in each market (full frontloading with $\delta_{ETS1} = \delta_{ETS2} = 1$).

not allowed. However, banking is an institutional feature of virtually all carbon markets. As a result, announcing a future linkage—or even the anticipation of such an announcement—leads to immediate price convergence. If policymakers seek to implement a gradual linking policy, either banking must be disallowed, which is unlikely, or an additional instrument is required to drive a wedge between permit prices.

4.2.3 Policies using only linking conversion rates

In order to drive a wedge between permit prices in period 1 while still allowing for the banking of permits, a natural candidate may seem to be a one-dimensional policy that introduces a linking conversion rate in period 1. However, we now provide an illustration of the identified shortcomings of such a policy (Results 3 and 4).

Suppose that policymakers can identify an acceptable rate of price convergence defined by the ratio of permit prices in period 1, which we denote $p_{ETS2,1}/p_{ETS1,1} = r_1 > 1$, and they also wish to retain 1:1 banking in each market. Then policymakers may simply set linking conversion rates in line with the targeted price ratios, namely

$$\rho_1 = r_1 > 1 \text{ and } \rho_2 = 1.$$

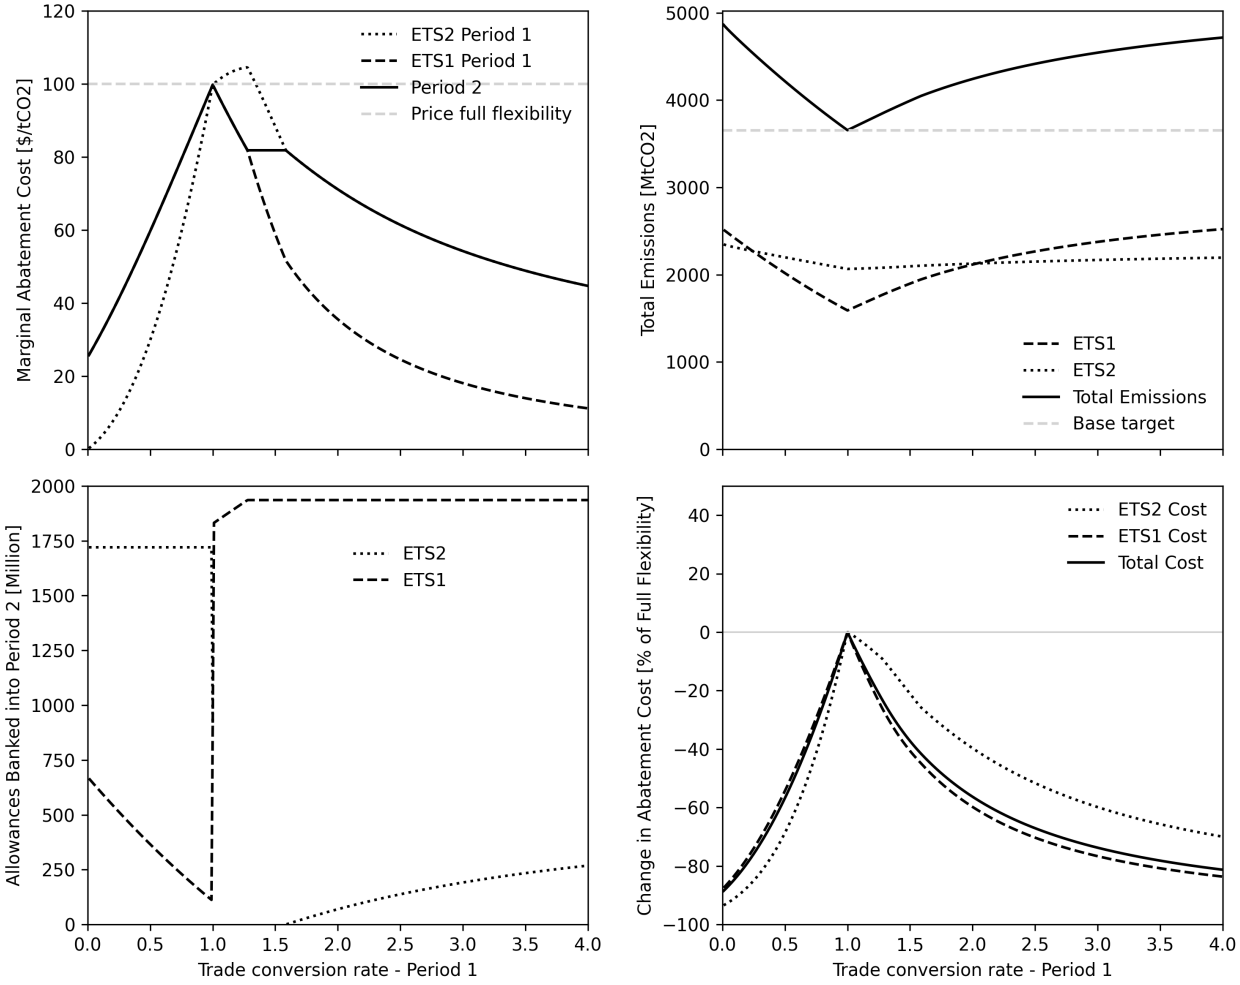
Figure 2 illustrates how market outcomes vary for different choices of $\rho_1 = r_1$. Importantly, the desired permit price ratio is only achieved in the narrow range of $1 < \rho_1 \leq 1.25$. For greater values of ρ_1, permit prices either partially converge ($1.25 < \rho_1 \leq 1.6$) or fully converge ($\rho_1 > 1.6$) already in period 1 (upper left panel).[40] Thus, $r_1 = 1.25$ is the most gradual price convergence that policymakers can achieve using only linking conversion rates, which is only a minor slowing of price convergence compared to *Full Flexibility*.[41]

For linking conversion rates of $\rho_1 > 1.25$, all ETS1 permits are banked into period 2 (lower left panel), when they become relatively more valuable because they can be traded 1:1. Prices for these permits are then driven by their future value in period 2—rather than their compliance value in period 1—causing prices to converge in period 1. Because all ETS1 permits are banked, only ETS2 permits remain for compliance in period 1, and the value of $\rho_1 > 1$

[40] The upper left panel of Figure 2 depicts marginal abatement costs rather than permit prices. When ETS1 banks permits—which is the case for all $\rho_1 > 1$—its permit prices are equalized across periods and equal to the period-2 line in that plot.

[41] For comparison, under *Full Banking*—that is, 1:1 banking but no permit trade between markets—the ratio of permit prices is greater than 5. Compared to this benchmark, achieving $r_1 = 1.25$ is only modestly different from the immediate price convergence of *Full Flexibility*.

Figure 2: Gradual linking using only linking conversion rates



Notes: The graphs report market outcomes when policymakers seek to govern the price ratio across markets in period 1 using only linking conversion rates, with banking allowed in both markets on a 1:1 basis. The period-1 linking conversion rate ρ_1 varies along the x-axis, and the period-2 linking conversion rate is fixed at $\rho_2 = 1$. The upper left panel shows marginal abatement costs—which will not necessarily equal permit prices in all cases case—along with the *Full Flexibility* price of 100\$/tCO$_2$. The upper right panel shows total emissions and emissions in each market, along with the overall emissions target of 3656 MtCO$_2$. The lower left panel shows the amount of banked permits, measured by the period-2 compliance value. The lower right panel shows changes in abatement costs relative to *Full Flexibility*.

means that each permit allows for more emissions in ETS1 than in ETS2. This "unnatural" flow of permits from ETS2 to ETS1 causes total emissions to increase (upper right panel). Abatement costs fall in response (lower right panel), but largely because less abatement is achieved.

Thus, attempting to control price convergence through a one-dimensional policy of linking conversion rates will succeed only when targeting a minor slowing of convergence ($1 < \rho_1 \leq$

1.25). Targeting greater period-1 price ratios will not only fail but will also allow for greater emissions. The greater is the targeted price ratio and equivalent linking conversion rate, the greater is the unintended degradation of environmental benefits.

It is important to note that these results depend on the allocation of permits between periods, with permits frontloaded in this case to emphasize the intertemporal effects of banking. As more permits are shifted to period 2, these effects lessen and a wider range of price convergence rates r_1 can be achieved.[42] Thus, while some one-dimensional policies may achieve the acceptable rate of price convergence, they are not robust to all permit allocations.

4.2.4 Policies using only banking conversion rates

We next turn to a different type of one-dimensional policy using only banking conversion rates without linking conversion rates. As in Section 4.2.2, we again delay 1:1 linking until period 2, in order to focus attention on banking. Although this type of one-dimensional policy was not studied in Section 3 and lacks a direct "linking spirit," we use it to gain insights into the effects of banking conversion rates. In addition, we show that the use of properly calibrated banking conversion rates can achieve gradual linking while also giving policymakers additional control over other market outcomes.

Again, suppose that policymakers can identify an acceptable rate of price convergence given by $p_{ETS2,1}/p_{ETS1,1} = r_1 > 1$. So long as banking is positive in both markets, Result 5 states that this rate of price convergence can be achieved with banking conversion rates set such that

$$\frac{\delta_{ETS2}}{\delta_{ETS1}} = r_1 > 1.$$

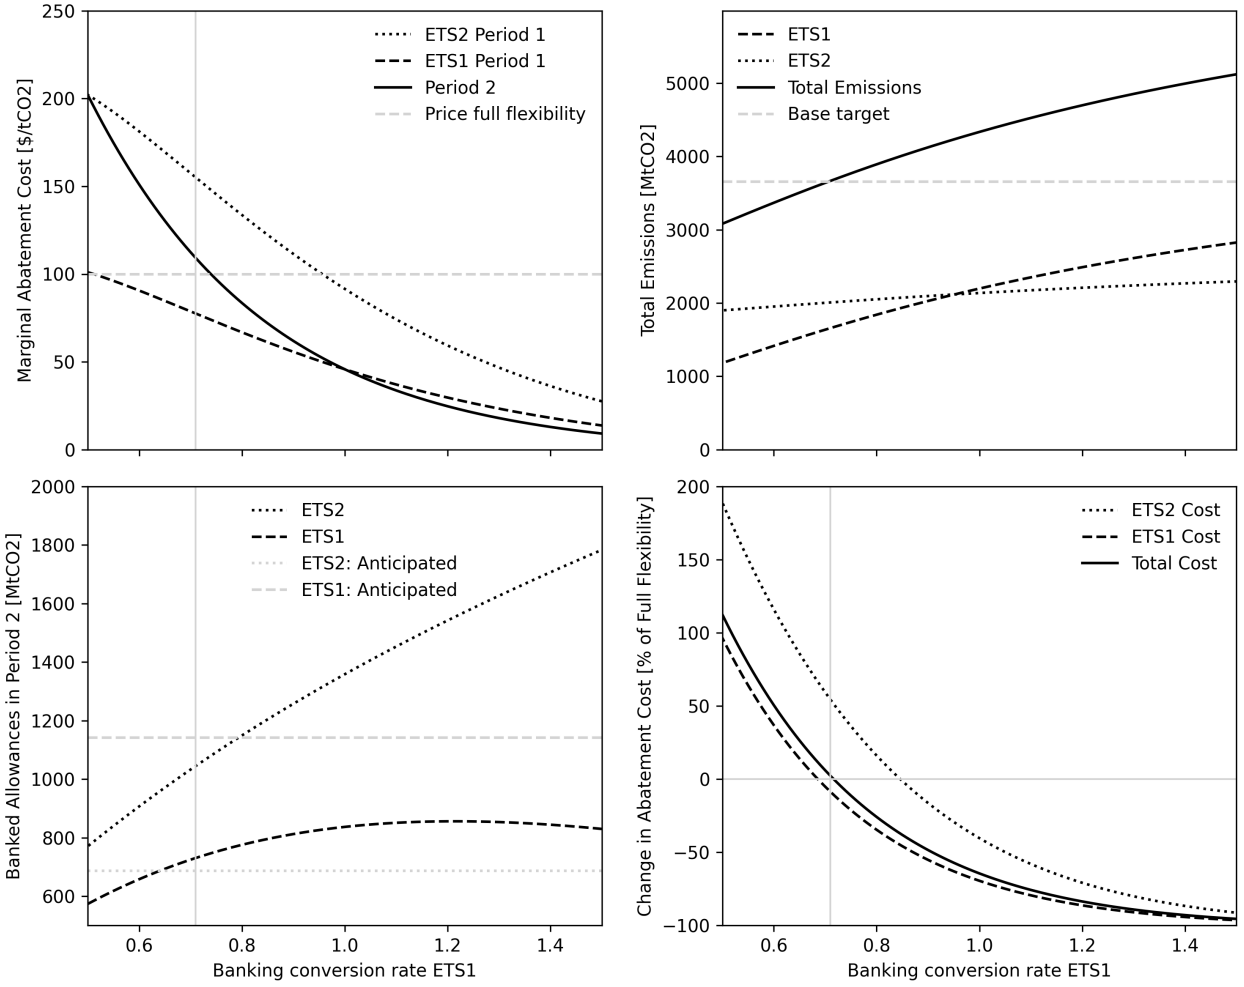
In words, slowing the rate of permit price convergence solely through banking conversion rates requires that banked ETS1 permits become relatively more scarce, altering the incentives to bank in each market and the paths of permit prices across periods. When banking is zero in at least one market—which will depend on the allocation of permits between periods—linking conversion rates will be needed as additional instruments to govern price convergence in period 1 (see next section).

To illustrate, suppose that $r_1 = 2$, which can be achieved for a range of banking conversion rates (Figure 3, upper left panel). At $\delta_{ETS1} = 0.5$ (which implies $\delta_{ETS2} = 1$), the incentive to bank ETS1 permits is reduced, although some banking does remain (lower left panel). Because each banked ETS1 permit effectively reduces the aggregate supply of permits—

[42] See Appendix B.2 for results without permit frontloading.

permitting only 0.5 tons of emissions in period 2—total emissions decrease (upper right panel) and the common price in period 2 increases (upper left panel). Accordingly, the ETS1 price in period 1 is half the price in period 2, while the price in ETS2 is equal across periods. In addition, abatement costs increase in both markets as more abatement is required, but more so in the high-cost market (ETS2) relative to *Full Flexibility* (lower right panel).

Figure 3: Gradual linking using only banking conversion rates



Notes: The graphs report market outcomes when policymakers seek to govern the price ratio across markets in period 1 using only banking conversion rates set according to $\frac{\delta_{ETS2}}{\delta_{ETS1}} = r_1 = 2$. Trading is ruled out in period 1 and occurs on a 1:1 basis in period 2. The ETS1 banking conversion rate δ_{ETS1} varies along the x-axis, with the ETS2 banking conversion rate given by $\delta_{ETS2} = 2\delta_{ETS1}$. The upper left panel shows marginal abatement costs—which equal permit prices in this case—along with the *Full Flexibility* price of 100\$/tCO_2. The upper right panel shows total emissions and emissions in each market, along with the overall emissions target of 3656 MtCO_2. The lower left panel shows the amount of banked permits, measured by the period-2 compliance value. The lower right panel shows changes in abatement costs relative to *Full Flexibility*. In all panels, the vertical line denotes the pair of banking conversion rates that preserves the overall emission target.

As we increase both banking conversion rates, these effects are lessened. As ETS1 banking becomes less supply-contracting and ETS2 banking becomes supply-expanding, total emissions increase. Accordingly, the price in period 2 decreases, as do the two prices in period 1 by arbitrage. In addition, abatement costs decrease in both markets. At $\delta_{ETS1} \approx 0.71$ and $\delta_{ETS2} \approx 1.42$, total emissions are equal to the emission target defined by the emission caps. However, achieving the targeted price ratio in period 1 shifts the timing and distribution of emissions compared to *Full Flexibility*, increasing total abatement costs by about 3%. As we further increase both banking conversion rates, both environmental effectiveness and cost effectiveness are reduced. When $\delta_{ETS1} > 1$ and $\delta_{ETS2} > 2$, the price is even decreasing over time in both markets as banking sufficiently increases the supply of permits in period 2.

4.2.5 Policies using both linking and banking conversion rates

We finally consider the case of a two-dimensional policy where policymakers use both linking and banking conversion rates. These rates are selected to achieve the acceptable rate of price convergence, namely

$$\rho_1 = r_1 > 1,\ \rho_2 = 1, \text{ and } \frac{\delta_{ETS2}}{\delta_{ETS1}} = r_1.$$

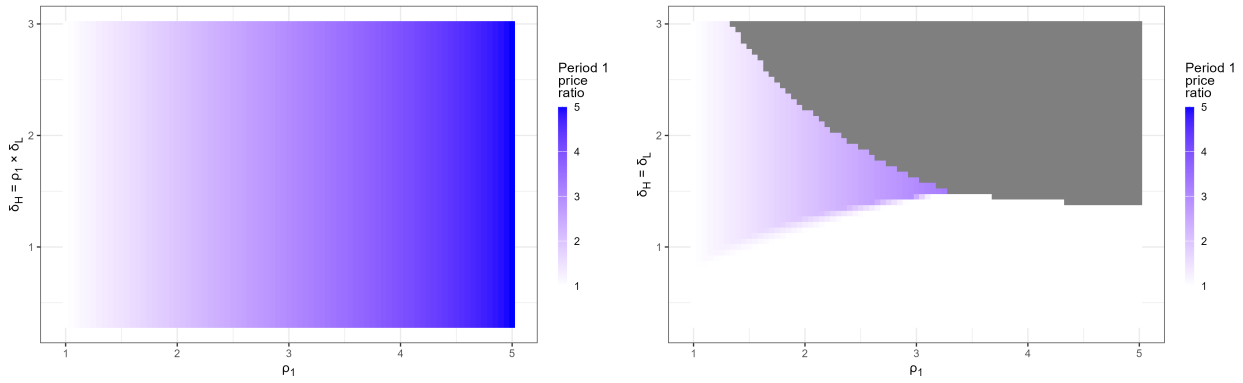
We also compare this policy to one in which the banking conversion rates are not calibrated correctly but instead set at a uniform rate ($\delta_{ETS2} = \delta_{ETS1}$).

A properly calibrated two-dimensional policy will always achieve the targeted rate of price convergence. Unlike the one-dimensional policies we consider above, this outcome will hold regardless of how permit are allocated between periods. Figure 4 expands the parameter space with a range of ρ_1 values on the x-axis, δ_{ETS2} on the y-axis, and the resulting period-1 price ratio shown in color. When banking conversion rates are properly calibrated (panel a), prices always converge at the targeted rate of $\rho_1 = r_1$. If those rates are instead uniform (panel b), as in a one-dimensional policy, that rate of price convergence is not guaranteed. Instead, for many parameter choices, permit prices will immediately converge in period 1 (white area) or the policy will be made so lax—through the "unnatural" trade of permits from ETS2 to ETS1 in period 1—that prices fall to zero and no abatement is achieved (gray area).

Because any of these conversions rates can achieve the targeted rate of price convergence, policymakers may further consider how these rates can affect other important market outcomes. For example, emissions may increase or decrease across a relevant parameter space (Figure 5, panel a). For a given value of ρ_1, total emissions are increasing in banking conversion rates,

Figure 4: Permit price ratios in period 1 under two-dimensional policies

(a) Calibrated rates: $\delta_{ETS2} = \rho_1 \delta_{ETS1}$ (b) Uniform rates: $\delta_{ETS2} = \delta_{ETS1}$

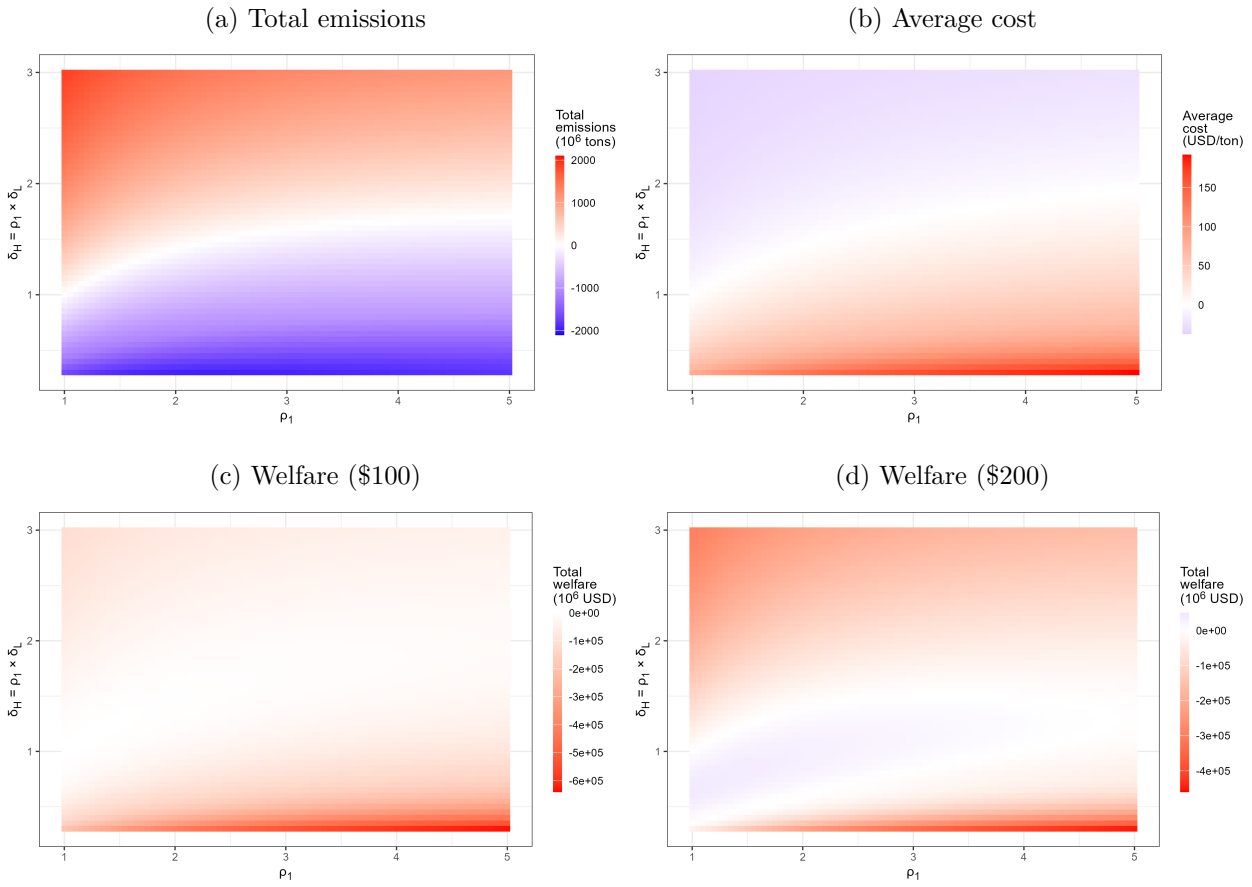


Note: All policy scenarios attempt to achieve gradual price convergence with linking conversion rate ρ_1 given on the x-axis and $\rho_2 = 1$ in all cases, either with banking conversion rates calibrated according to $\delta_{ETS2} = \rho_1 \delta_{ETS1}$ (panel (a)) or uniform banking conversion rates $\delta_{ETS2} = \delta_{ETS1}$ (panel (b)) given on the y-axis. In all cases, permit allocation is fully frontloaded in period 1 (without frontloading, permit prices always gradually converge as targeted). In panel (b), the white space represents sets of rates that yield "early" price convergence in period 1 ($p_{ETS2,1} = p_{ETS1,1} > 0$), while the gray area represents sets of rates that yield non-binding policies and prices equal to zero ($p_{ETS2,1} = p_{ETS1,1} = 0$), so the ratio is not defined.

as larger banking conversion rates are relatively supply-expanding. We further see that for greater values of ρ_1—which are supply-contracting as permits are traded from ETS1 to ETS2 in period 1—a wider range of banking conversion rates can achieve emissions reductions. If a two-dimensional policy is designed to leave emissions equal to the emissions caps, abatement costs increase modestly compared to *Full Flexibility* (panel b). Using a two-dimensional policy to achieve additional emissions reductions further increases abatement costs.

In panels c and d, we quantify the welfare consequences of a two-dimensional policy over this parameter space. A key consideration is the benefit generated by emissions abatement. We first use a policy-consistent value of $100 per ton, equal to the permit price under *Full Flexibility* (panel c). In that case, *Full Flexibility* is already the efficient policy, so any two-dimensional policy yields welfare losses, but these losses are generally modest for reasonable conversion rates. If we instead value abatement at a social cost of carbon of $200 per ton, welfare gains can be achieved for a core range of conversion rates (panel d).

Figure 5: Market outcomes under two-dimensional policies



Note: All policy scenarios attempt to achieve gradual price convergence with linking conversion rate ρ_1 given on the x-axis and $\rho_2 = 1$ in all cases. Banking conversion rates are calibrated using $\delta_L = \delta_H/\rho_1$ with δ_H given on the y-axis. In all policy scenarios, permit allocation is fully frontloaded in period 1 (see Figure B.3 for the case without permit frontloading). All outcomes are shown relative to the reference case of *Full Flexibility* (see Table 4). Emissions are reported in millions of tons. Aggregate average abatement costs are reported in USD per ton. Welfare is reported in millions of USD, assuming a model-consistent cost of carbon emissions of 200 USD per ton.

5 Conclusion

This paper has shown that a linking policy ensuring gradual price convergence across permit markets needs two instruments—one to establish the exchange rate at which permits trade in each period, another to align intertemporal arbitrage incentives with the targeted sequence of exchange rates. If linking conversion rates are used in isolation, anticipation of the future trading terms makes the current ones mute. That is, price ratios are governed by the linking conversion rate that will prevail in the future, making the phase-in of the link obsolete. The paper has also shown how banking conversion rates should be set to sustain the targeted

sequence of price ratios in equilibrium, and that a two-dimensional policy gives policymakers more flexibility to manage the distributional and environmental impacts of linking. In a calibrated numerical application to the EU ETS1 and ETS2, the paper has illustrated the different price paths (in each market), emission levels, and welfare effects under various gradual linking policies.

References

Abrell, J. & Rausch, S. (2017). Combining Price and Quantity Controls under Partitioned Environmental Regulation. *Journal of Public Economics*, **145**, 226–242.

Abrell, J., Rausch, S. & Schwarz, G. A. (2018). How Robust is the Uniform Emissions Pricing Rule to Social Equity Concerns? *Journal of Environmental Economics and Management*, **92**, 783–814.

Aguiar, A., Chepeliev, M., Corong, E. L., McDougall, R. & van der Mensbrugghe, D. (2019). The GTAP Data Base: Version 10. *Journal of Global Economic Analysis*, **4**(1), 1–27.

Bataille, C., Guivarch, C., Hallegatte, S., Rogelj, J. & Waisman, H. (2018). Carbon Prices across Countries. *Nature Climate Change*, **8**, 648–650.

Bierbrauer, F. (2025). Is a Uniform Price on Carbon Desirable? Working Paper, University of Cologne.

Böhringer, C., Hoffmann, T. & Manrique de Lara Peñate, C. (2006). The Efficiency Costs of Separating Carbon Markets under the EU Emissions Trading Scheme: A Quantitative Assessment for Germany. *Energy Economics*, **28**(1), 44–61.

Burtraw, D., Palmer, K., Munnings, C., Weber, P. & Woerman, M. (2013). Linking by Degrees: Incremental Alignment of Cap-and-Trade Markets. RFF Discussion Paper 13-04.

Chichilnisky, G. & Heal, G. (1994). Who Should Abate Carbon Emissions? An International Viewpoint. *Economics Letters*, **4**(4), 443–449.

Doda, B., Quemin, S. & Taschini, L. (2019). Linking Permit Markets Multilaterally. *Journal of Environmental Economics and Management*, **98**, 102259.

Dolphin, G., Pahle, M., Burtraw, D. & Kosch, M. (2023). A Net-Zero Target Compels a Backward Induction Approach to Climate Policy. *Nature Climate Change*, **13**, 1033–1041.

Dubois, L., Sahuc, J.-G. & Vermandel, G. (2025). A General Equilibrium Approach to Carbon Permit Banking. *Journal of Environmental Economics and Management*, **129**, 103076.

EPA (2003). NOx Budget Trading Program, 2003 Progress and Compliance Report. US Environmental Protection Agency.

EPA (2023). EPA's "Good Neighbor" Plan Cuts Ozone Pollution – Overview Fact Sheet. US Environmental Protection Agency.

Fabra, N. & Reguant, M. (2024). The Energy Transition: A Balancing Act. *Resource and Energy Economics*, **76**, 101408.

Feng, H. & Zhao, J. (2006). Alternative Intertemporal Permit Trading Regimes with Stochastic Abatement Costs. *Resource and Energy Economics*, **28**, 24–40.

Flachsland, C., Marschinski, R. & Edenhofer, O. (2009). To Link or Not to Link: Benefits and Disadvantages of Linking Cap-and-Trade Systems. *Climate Policy*, **9**(4), 358–372.

Fowlie, M. L. & Muller, N. Z. (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.

Graichen, J. & Ludig, S. (2024). Supply and Demand in the ETS 2. Report by Öko-Institut e.V. for the German Environmental Agency (Umweltbundesamt).

Günther, C., Pahle, M., Govorukha, K., Osorio, S. & Fotiou, T. (2025). Carbon prices on the rise? Shedding light on the emerging second EU Emissions Trading System (EU ETS 2). *Climate Policy*, **in press**, 1–12.

Helm, C. (2003). International Emissions Trading with Endogenous Allowance Choices. *Journal of Public Economics*, **87**(12), 2737–2747.

Hintermann, B. & Gronwald, M. (2019). Linking with Uncertainty: The Relationship Between EU ETS Pollution Permits and Kyoto Offsets. *Environmental and Resource Economics*, **74**, 761–784.

Holland, S. P. & Yates, A. J. (2015). Optimal Trading Ratios for Pollution Permit Markets. *Journal of Public Economics*, **125**, 16–27.

Holtsmark, B. & Weitzman, M. L. (2020). On the Effects of Linking Cap-and-Trade Systems for CO_2 Emissions. *Environmental and Resource Economics*, **75**(3), 615–630.

Holtsmark, K. & Midttømme, K. (2021). The Dynamics of Linking Permit Markets. *Journal of Public Economics*, **198**, 104406.

Hotelling, H. (1931). The Economics of Exhaustible Resources. *Journal of Political Economy*, **39**(2), 137–175.

ICAP (2025). Emissions Trading Worldwide: Status Report 2025. Berlin: International Carbon Action Partnership.

ICAP (2026). Scope Expansion of ETS: Design and Policy Challenges. Berlin: International Carbon Action Partnership.

Inman, M. (2018). Holding Limits Don't Constrain Banking in California's Cap-and-Trade Program. Dec 13, 2018, nearzero.org.

Kakeu, J. (2025). Multipollutant Cap-and-Trade Systems with Heterogeneous Firms: An Optimal Transport Approach. Working Paper 25-11, Resources for the Future.

Karp, L. S. & Traeger, C. (2025). Smart Cap. *Journal of the European Economic Association*, **23**(2), 554–593.

Klepper, G. & Peterson, S. (2006). Marginal Abatement Cost Curves in General Equilibrium: The Influence of World Energy Prices. *Resource and Energy Economics*, **28**(1), 1–23.

Kling, C. L. & Rubin, J. D. (1997). Bankable Permits for the Control of Environmental Pollution. *Journal of Public Economics*, **64**(1), 101–115.

Lang, S. F. (2026). Optimal Carbon Prices in an Unequal World: The Role of Regional Welfare Weights. Working Paper, ETH Zürich.

Leard, B. (2013). The Welfare Effects of Allowance Banking in Emissions Trading Programs. *Environmental and Resource Economics*, **55**, 175–197.

Leiby, P. N. & Rubin, J. D. (2001). Intertemporal Permit Trading for the Control of Greenhouse Gas Emissions. *Environmental and Resource Economics*, **19**, 229–256.

Long, X., Astier, N. & Zhang, D. (2025). Is Broader Trading Welfare Improving for Emission Trading Systems? *Journal of Environmental Economics and Management*, **130**, 103110.

Mehling, M. A., van Asselt, H., Dröge, S., Das, K. & Hall, C. (2024). Bridging the Divide: Assessing the Viability of International Cooperation on Border Carbon Adjustments. MIT-

CEEPR Working Paper 2024-06.

Mehling, M. A., Metcalf, G. E. & Stavins, R. N. (2018). Linking Heterogeneous Climate Policies (Consistent with the Paris Agreement). *Environmental Law*, **48**(4), 647–698.

Mendelsohn, R. (1986). Regulating Heterogeneous Hmissions. *Journal of Environmental Economics and Management*, **13**(4), 301–312.

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

Montero, J.-P. (2001). Multipollutant Markets. *RAND Journal of Economics*, **32**(4), 762–774.

Muller, N. Z. & Mendelsohn, R. (2009). Efficient Pollution Regulation: Getting the Prices Right. *American Economic Review*, **99**(5), 1714–1739.

Newell, R. G., Pizer, W. A. & Zhang, J. (2005). Managing Permit Markets to Stabilize Prices. *Environmental and Resource Economics*, **31**, 133–157.

Pahle, M., Quemin, S., Osorio, S., Günther, C. & Pietzcker, R. (2025). The Emerging Endgame: The EU ETS on the Road Towards Climate Neutrality. *Resource and Energy Economics*, **81**, 101476.

Pahle, M., Sultani, D. & Zachmann, G. (2026). Defragmenting European Union Climate Policy. Policy Brief, Bruegel.

Pizer, W. A. & Prest, B. C. (2020). Prices versus Quantities with Policy Updating. *Journal of the Association of Environmental and Resource Economists*, **7**(3), 483–518.

Pizer, W. A. & Yates, A. J. (2015). Terminating Links Between Emission Trading Programs. *Journal of Environmental Economics and Management*, **71**, 142–169.

Quemin, S. & de Perthuis, C. (2019). Transitional Restricted Linkage Between Emissions Trading Schemes. *Environmental and Resource Economics*, **74**(1), 1–32.

Roy, N., Russo, S. & Burtraw, D. (2025). Considerations for Washington's Linkage Negotiations with California and Québec. Report, Resources for the Future.

Rubin, J. D. (1996). A Model of Intertemporal Emission Trading, Banking and Borrowing. *Journal of Environmental Economics and Management*, **31**(3), 269–86.

Salant, S. W. (2016). What Ails the European Union's Emissions Trading System? *Journal of Environmental Economics and Management*, **80**, 6–19.

Schennach, S. M. (2000). The Economics of Pollution Permit Banking in the Context of Title IV of the 1990 Clean Air Act Amendments. *Journal of Environmental Economics and Management*, **40**(3), 189–210.

Shobe, W., Holt, C. A. & Huetteman, T. (2014). Elements of Emission Market Design: An Experimental Analysis of California's Market for Greenhouse Gas Allowances. *Journal of Economic Behavior and Organization*, **107**(A), 402–420.

Stiglitz, J. E. (2019). Addressing Climate Change Through Price and Non-Price Interventions. *European Economic Review*, **119**, 594–612.

Tuerk, A., Mehling, M., Flachsland, C. & Sterk, W. (2009). Linking Carbon Markets:

Concepts, Case Studies and Pathways. *Climate Policy*, **9**(4), 341–357.

Vogt-Schilb, A., Meunier, G. & Hallegatte, S. (2018). When Starting with the Most Expensive Option Makes Sense: Optimal Timing, Cost and Sectoral Allocation of Abatement Investment. *Journal of Environmental Economics and Management*, **88**, 210–233.

Williams III, R. C. (2010). Setting the Initial Time-Profile of Climate Policy: The Economics of Environmental Policy Phase-Ins. NBER Working Paper 16120.

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. Washington, DC: World Bank.

Yates, A. J. & Cronshaw, M. B. (2001). Pollution Permit Markets with Intertemporal Trading and Asymmetric Information. *Journal of Environmental Economics and Management*, **42**, 104–118.

Appendix

A Proofs

A.1 Proof of Result 1

Let $\rho_1 = \rho_2 = 1$ and $\delta_H = \delta_L = 1$ in (2). In this case, permits always flow from L to H, i.e., $\Delta_1 > 0$ and $\Delta_2 > 0$ (if not, simply relabel the markets). Inspection of the first-order conditions yields $\gamma_H = 0$ (and $\gamma_L = 0$ since $\Delta_1 > 0$). Suppose $\gamma_H > 0$. H imports L-permits to cover the entirety of its period-1 emissions ($\Delta_1 = e_{H1}$) and banks the entirety of its period-1 allocation ($B_H = \omega_{H1} > 0$). $B_H > 0$ implies $\lambda_H = 0$, hence $p_{H1} = p_{H2}$. This in turn implies that $p_{H2} = p_{L2} = p_{H1} = p_{L1} + \gamma_H > p_{L1}$. Yet $p_{L2} > p_{L1}$ entails $\lambda_L < 0$, a contradiction. Thus, it must be that $\gamma_H = 0$, and the least-cost equilibrium is characterized by (3).

A.2 Proof of Result 2

Consider one least-cost allocation of trading and banking volumes $(\Delta_1^{fl}, \Delta_2^{fl})$ and (B_H^{fl}, B_L^{fl}). For some positive scalar $\varepsilon > 0$, allocation $(\Delta_1^{fl} - \varepsilon, \Delta_2^{fl} + \varepsilon)$ and $(B_H^{fl} - \varepsilon, B_L^{fl} + \varepsilon)$ may also implement the least-outcome; feasibility requires $\varepsilon < \min\{B_H^{fl}, \omega_{L1} - B_L^{fl}\}$). If an additional vintage restriction applies in period 2 whereby period-1 L-permits cannot be used for compliance in region H, ε must also satisfy $\varepsilon < \omega_{L2} - \Delta_2^{fl}$. Thus, there is always a sufficiently small ε that is feasible. Result 2 then simply hinges on whether $\varepsilon = \Delta_1^{fl}$ is feasible or not.

A.3 Proof of Result 3

Because $\rho_1 > \rho_2 = 1$, trade incentives are distorted between periods so that firms are better off when (1) as much trade from L to H as possible takes place in period 2; (2) if feasible, as much trade as possible from H to L takes place in period 1. This relaxes the overall emission constraint despite the number of permits being fixed, and entails that banking levels in H and L are as low and high as possible, respectively.

Inspection of (2) with $\rho_1 > \rho_2 = 1$ and $\delta_H = \delta_L = 1$ reveals that $\gamma_H = 0$ and $\gamma_L \geq 0$. For a proof that γ_H cannot be positive, see Appendix A.5 with $\delta_H \rho_2 / \delta_L \rho_1 = 1/\rho_1 < 1$. In this policy setting, permits always flow from L to H in period 2, but in period 1 they may flow in both directions. Let the superscript o denote the one-legged linking policy considered here.

When $e^o_{H1} < \omega_{H1}$, or equivalently $\Delta^o_1 < B^o_H$, it must be $\Delta^o_1 < 0$ with $B^o_H \in [0, \omega_{H1})$ and $B^o_L \in [0, \omega_{L1}]$. Indeed, $B^o_H > \Delta^o_1 > 0$ cannot be an optimum, as B^o_H can be reduced so as to reduce Δ^o_1 (as a positive Δ^o_1 entails reduced total emissions) until $\Delta^o_1 < 0$ with $B^o_H > 0$ and $B^o_L = \omega_{L1}$, or $\Delta^o_1 < 0$ with $B^o_H = 0$ and $B^o_L \leq \omega_{L1}$. Specifically, the first-order conditions in (2) admit the four cases below.

(*i*) $e^o_{H2} + e^o_{L2} > \omega_{H2} + \omega_{L2} + \omega_{L1}$. Full price convergence in period 1: $p_{H1} = p_{L1}$. Aggregate emissions in period 2 are sufficiently great that L banks the entirety of its period-1 permits—and buys some permits from H for compliance in period 1—and H also banks some of its period-1 permits ($\gamma_L > 0$ and $\lambda_L = \lambda_H = 0$). As banking occurs in both markets and prices are equal in period 2, arbitraged Hotelling price paths are identical and prices fully converge in period 1.

(*ii*) $e^o_{H2} + e^o_{L2} = \omega_{H2} + \omega_{L2} + \omega_{L1}$. Partial price convergence in period 1: $\rho_1 p_{L1} > p_{H1} > p_{L1}$. Aggregate emissions in period 2 require L to bank the entirety of its period-1 permits—and buy from H for compliance in period 1—but H does not bank any of its period-1 permits ($\gamma_L > 0$, $\lambda_L = 0$, and $\lambda_H > 0$). With H not banking, its prices do not follow a Hotelling path ($p_{H1} > p_{H2}$ instead) so early price convergence no longer holds. Yet the demand to bank period-1 permits remains sufficiently high to drive up the price of L's period-1 permits above what would be expected by the exchange rate ρ_1 (hence, the price ratio in period 1 is comprised strictly between 1 and ρ_1).

(*iii*) $\omega_{H2} + \omega_{L2} < e^o_{H2} + e^o_{L2} < \omega_{H2} + \omega_{L2} + \omega_{L1}$. No early price convergence: $p_{H1} = \rho_1 p_{L1}$. Aggregate emissions in period 2 are sufficiently low that L banks some—but not the entirety—of its period-1 permits, and H does not bank any of its period-1 permits ($\gamma_L = 0$, $\lambda_L = 0$, and $\lambda_H > 0$). With this lower demand for period-1 permits, the price ratio in period 1 is set by ρ_1. As L banks some period-1 permits, its prices remains on a Hotelling path ($p_{L1} = p_{L2}$). The absence of banking in H imply that its prices are decreasing across periods, with the rate of decrease fixed by ρ_1 ($p_{H1} = \rho_1 p_{H2} > p_{H2}$).

(*iv*) $e^o_{H2} + e^o_{L2} = \omega_{H2} + \omega_{L2}$. No early price convergence: $p_{H1} = \rho_1 p_{L1}$. Aggregate emissions in period 2 are so low that no banking occurs ($\gamma_L = 0$, $\lambda_L > 0$, and $\lambda_H > 0$). The price ratio is period 1 is set by ρ_1, and neither set of prices is on a Hotelling path.

When $e^o_{H1} > \omega_{H1}$, or equivalently $\Delta^o_1 > B^o_H$, it must be that $\Delta^o_1 > 0$ since B^o_H is non-negative. In fact, it can only be that $B^o_H = 0$ (for otherwise Δ^o_1 could be reduced further) and $B^o_L > 0$

(for otherwise $e^o_{H2} + e^o_{L2} = \omega_{H2} + \omega_{L2}$, which is case (iv) above). Hence, $\gamma_H = \gamma_L = 0$ and $\lambda_H > \lambda_L = 0$, implying $p_{H1} = \rho_1 p_{L1}$, $p_{H2} = p_{L2}$, $p_{L1} = p_{L2}$, and thus $p_{H1} = \rho_1 p_{H2} > p_{H2}$.

Which of the above cases occurs depends on the magnitude of ρ_1, the demand for banking, and the allocation of permits over time in both markets.

A.4 Proof of Result 4

In period 2, trade is emission neutral since $\rho_2 = 1$. Hence, only the direction of trade in period 1 determines whether total emissions decrease (when $\Delta_1 > 0$) or increase (when $\Delta_1 < 0$).

When $e^o_{H1} > \omega_{H1}$, permits flow from L to H in period 1 just as much as needed to cover the difference, and $B_H = 0$. Because ρ_1 L-permits are needed to cover for each "missing" H-permits, total emissions are reduced by $(\rho_1 - 1)(e^o_{H1} - \omega_{H1})$. When $e^o_{H1} < \omega_{H1}$, in cases $(ii-iv)$, all "excess" H-permits flow to L in period 1 as L can bank the equivalent of $\rho_1(\omega_{H1} - e^o_{H1})$ additional L-permits, with $B_H = 0$. Total emissions are increased by $(\rho_1 - 1)(\omega_{H1} - e^o_{H1})$. In case (i), L cannot bank the entirety of the $\rho_1(\omega_{H1} - e^o_{H1})$ additional L-permits, with $B^o_L = \omega_{L1}$ and $B^o_H > 0$. Total emissions are increased by $(\rho_1 - 1)(\omega_{H1} - e^o_{H1} - B^o_H)$.

Ceteris paribus: A larger ρ_1 increases the change in total emissions (both ways) and thus the incentives to shift trade from L to H to period 2 and to engage in trade from H to L trade in period 1. The lower ω_{H2} and ω_{L2} are given the fixed total number of permits (i.e., the higher ω_{H1} and ω_{L1} are), the more likely cases $(i - iv)$ with $\Delta^o_1 < 0$ become.

A.5 Proof of Result 5

The proof holds for any pairs of linking and banking conversion rates (ρ_1, ρ_2) and (δ_H, δ_L).

Notice, $p_{H2} = \rho_2 p_{L2}$ always holds. When $B_m > 0$ in each market m: $p_{m1} = \delta_m p_{m2}$ concludes. When $B_m = 0$ in each market m: $p_{H1} = \rho_1 p_{L1}$ (since $\gamma_H = \gamma_L = 0$) concludes. When $B_m > 0$, $B_{\tilde{m}} = 0$ and $\gamma_m = 0$, $p_{H1} = \rho_1 p_{L1}$ (since $\gamma_{\tilde{m}} = 0$) with $p_{m1} = \delta_m p_{m2}$ and $p_{\tilde{m}1} > \delta_{\tilde{m}} p_{\tilde{m}2}$ concludes. If $\gamma_m > 0$ instead, the rate of price convergence is not unambiguously determined. Specifically, when $B_L = \omega_{L1}$ and $\gamma_L > 0$, which can only occur when $\rho_1 > \rho_2$, $p_{H1} = \rho_1(p_{L1} - \gamma_L) < \rho_1 p_{L1}$ yields $\frac{p_{H1}}{p_{L1}} \Big/ \frac{p_{H2}}{p_{L2}} < \frac{\rho_1}{\rho_2}$; when $B_H = \omega_{H1}$ and $\gamma_H > 0$, which can only occur when $\rho_1 < \rho_2$, $p_{H1} - \gamma_H = \rho_1 p_{L1} < p_{H1}$ yields $\frac{p_{H1}}{p_{L1}} \Big/ \frac{p_{H2}}{p_{L2}} > \frac{\rho_1}{\rho_2}$.

At this stage, we can conclude that the "rule" $\delta_H/\delta_L = \rho_1/\rho_2$ implements gradual linking as desired, except in the two above cases where the rate of price convergence is ambiguous, which

we address below. Notice, the first-order conditions in (2) for any pairs (ρ_1, ρ_2) and (δ_H, δ_L) imply that γ_H and γ_L cannot be positive at the same time. If $\gamma_H > 0$, then $\Delta_1 = e_{H1} > 0$. Thus, $e_{L1} > -\rho_1 \Delta_1$ always holds, so $\gamma_L = 0$. Symmetrically, $\gamma_L > 0 \Rightarrow \gamma_H = 0$. Intuitively, when permits flow from H to L, the constraint on the maximum permit transfer from L to H is non-binding, and vice versa.

Assume $\gamma_H > 0$. This implies $\gamma_L = 0$ and $B_H = \omega_{H1} > 0$, so $\lambda_H = 0$. Thus, (2) reduces to

$$p_{H1} = \delta_H p_{H2} = \delta_H \rho_2 p_{L2} \leq \frac{\delta_H \rho_2}{\delta_L} p_{L1} \text{ and } p_{H1} > \rho_1 p_{L1},$$

which is feasible only if $\delta_H \rho_2 / \delta_L \rho_1 > 1$. By the same token, now assume $\gamma_L > 0$. This implies $\gamma_H = 0$ and $B_L = \omega_{L1} > 0$, so $\lambda_L = 0$. Thus, (2) reduces to

$$p_{L1} = \delta_L p_{L2} = \frac{\delta_L}{\rho_2} p_{H2} \leq \frac{\delta_L}{\delta_H \rho_2} p_{H1} \text{ and } p_{H1} < \rho_1 p_{L1},$$

which is feasible only if $\delta_L \rho_1 / \delta_H \rho_2 > 1$. Therefore, setting banking conversion rates such that $\delta_H / \delta_L = \rho_1 / \rho_2$ is the only policy setting where both γ_H and γ_L are nil, and where price ratios are equal to current linking conversion rates in both periods. Formally,

$$\frac{\delta_H}{\delta_L} = \frac{\rho_1}{\rho_2} \Rightarrow \gamma_H = \gamma_L = 0 \Rightarrow \begin{cases} p_{Ht} = \rho_t p_{Lt} \text{ for } t \in \{1, 2\}, \\ p_{m1} = \delta_m p_{m2} + \lambda_m \text{ for } m \in \{H, L\}. \end{cases}$$

Therefore, the "rule" $\delta_H / \delta_L = \rho_1 / \rho_2$ always implements gradual linking as desired. In the case of non-unitary time discounting, the "rule" only needs adjustment if firms in H and L have different interest rates $r_H \neq r_L$, i.e., $\delta_H / \delta_L = (1 + r_H)\rho_1 / (1 + r_L)\rho_2$.

An interior equilibrium with conversion rates set such that $\delta_H / \delta_L = \rho_1 / \rho_2$ can be sustained by multiple trading and banking volumes. Formally, this is because the system of equilibrium equations is singular (i.e., its determinant is zero) when $\delta_H / \delta_L = \rho_1 / \rho_2$; hence first-order conditions and constraints alone do not uniquely determine trading and banking. The case where all conversion rates are set equal to one in Appendix A.2 is a special case thereof.

Consider one interior gradual linking equilibrium allocation of trading and banking volumes $(\Delta_1^{gl}, \Delta_2^{gl})$ and (B_H^{gl}, B_L^{gl}). One can always find a sufficiently small positive scalar $\varepsilon > 0$ so that the allocation $(\Delta_1^{gl} - \varepsilon, \Delta_2^{gl} + \delta_H \varepsilon)$ and $(B_H^{gl} - \varepsilon, B_L^{gl} + \rho_1 \varepsilon)$ implements the same emission and price paths. This can be easily verified for e_{H1}, e_{L1}, and e_{H2}. For e_{L2}, this is also verified when $\rho_2 \delta_H \varepsilon = \delta_L \rho_1 \varepsilon$ since $e_{L2} = \omega_{L2} - \rho_2(\Delta_2^{gl} + \delta_H \varepsilon) + \delta_L(B_L^{gl} + \rho_1 \varepsilon) = \omega_{L2} - \rho_2 \Delta_2^{gl} + \delta_L B_L^{gl}$.

Another way to see this is that, so long as all linking and banking "routes" are used (i.e., no banking boundary condition is binding), there is no inconsistency in marginal costs and market prices along the entire time-market chain since, for instance,

$$-C'_{H1} = -\rho_1 C'_{L1} = -\rho_1\delta_L C'_{L2} = -\frac{\rho_1\delta_L}{\rho_2}C'_{H2} = -\frac{\rho_1\delta_L}{\rho_2\delta_H}C'_{H1} = -C'_{H1} \text{ if } \delta_H/\delta_L = \rho_1/\rho_2.$$

A.6 Proof of Result 6

Let $\rho_1 > \rho_2 = 1$ and $\delta_H = \rho_1\delta_L$. Notice, trade is emission neutral in period 2 since $\rho_2 = 1$.

Existence. When the δ's are sufficiently low (typically, both are below 1), no market banks. Then, permits must flow from L to H in period 1, so emissions are lower than emission caps. When the δ's are sufficiently high (typically, both are above 1), both markets bank. Even if permits flow from L to H in period 1, the relaxation of emissions induced through banking will exceed the trade-induced contraction in emissions for high enough δ's, so overall emissions are higher than emission caps. Thus, by a continuity argument (given the standard properties abatement cost functions), there exists at least one pair of δ's that leads to the same emission level as provided for by emission caps.

This can be shown more formally with the implicit function theorem and a law of conservation. More generally, when $\delta_H/\delta_L = \rho_1/\rho_2$, one has the following consistency equation

$$\begin{aligned} F(\pi, \delta_L) &= \textstyle\sum_{m,t} k_{mt}(e_{mt} - \omega_{mt}) = 0 \\ \text{with } k_{H1} &= \rho_1\delta_L = \rho_2\delta_H,\ k_{L1} = \delta_L,\ k_{H2} = \rho_2, \text{ and } k_{L2} = 1 \\ \text{and } e_{mt}(k_{mt}, \pi) &= (C'_{mt})^{-1}(-k_{mt}\pi) \text{ and } -C'_{L2} = \pi. \end{aligned}$$

Let $E(\delta_L) = \sum_{m,t} e_{mt}(k_{mt}(\delta_l), \pi(\delta_L))$. Then by the chain rule, one has

$$\frac{\mathrm{d}E}{\mathrm{d}\delta_H} = \frac{\rho_2}{\rho_1}\frac{\mathrm{d}E}{\mathrm{d}\delta_L} \text{ and } \frac{\mathrm{d}E}{\mathrm{d}\delta_L} = \frac{\partial E}{\partial \delta_L} + \frac{\mathrm{d}\pi}{\mathrm{d}\delta_L}\frac{\partial E}{\partial \pi} \text{ with } \frac{\mathrm{d}\pi}{\mathrm{d}\delta_L} = -\frac{\partial F/\partial \delta_L}{\partial F/\partial \pi}.$$

For any given ρ values, the monotonicity of E w.r.t. δ_H and δ_L depends on the allocation of permits across periods and the δ values. When permits are sufficiently frontloaded and the δ values are not too small, total emissions are increasing with the δ's. Existence follows by the intermediary value theorem, but the emission-preserving pair (δ_H, δ_L) need not be unique since E may not be monotonically increasing with the δ's depending on the permit allocation profile and the associated banking demand and constraints.

In addition, a typical emission-preserving pair (δ_H, δ_L) satisfies $\delta_L < 1 < \delta_H$. To see this, note that for given equilibrium values for Δ_1, B_H, and B_L, total emissions are constant if

$$(1-\rho_1)\Delta_1 + (\delta_H - 1)B_H + (\delta_L - 1)B_L = 0.$$

Using this condition, we can express δ_L and δ_H as functions of ρ_1:

$$\delta_L(\rho_1) = \frac{B_H + B_L - (1-\rho_1)\Delta_1}{\rho_1 B_H + B_L} \text{ and } \delta_H(\rho_1) = \rho_1 \delta_L(\rho_1).$$

Taking derivatives w.r.t. ρ_1, we get

$$\delta_L'(\rho_1) < 0 \text{ i.f.f. } B_H > \Delta_1 \text{ and } \delta_H'(\rho_1) > 0 \text{ unconditionally.}$$

Hence, when permit allocation is sufficiently frontloaded: $\delta_L'(\rho_1) < 0$. In this case, because $\rho_1 = \delta_H = \delta_L = 1$ preserves emission by construction, we must have $\delta_L < 1$ and $\delta_H > 1$ when $\rho_1 > 1$. This is what we find in the numerical simulations with full frontloading (without frontloading, the emission-preserving δ's can be such that $\delta_H > 1$ and $\delta_L > 1$).

Cost effectiveness. Consider a two-dimensional gradual linking policy (denoted by superscript gl) that achieves the same aggregate emissions as under full linking (denoted by superscript fl). The difference in abatement costs, for market m in period t, between these two policies is given by integrating under the marginal abatement cost curve between the two emissions levels:

$$\Delta C_{mt} \equiv C_{mt}(e_{mt}^{gl}) - C_{mt}(e_{mt}^{fl}) = \int_{e_{mt}^{gl}}^{e_{mt}^{fl}} -C_{mt}'(e)de.$$

Define Δe_{mt} as this difference in emissions for market m in period t: $\Delta e_{mt} = e_{mt}^{gl} - e_{mt}^{fl}$. Without loss of generality, suppose $\Delta e_{L1} > 0$ but $\Delta e_{mt} < 0$ for all other time-markets (if not, a similar proof would hold by grouping together the time-markets with $e_{mt} > 0$ versus those with $e_{mt} < 0$). Because aggregate emissions are equivalent between the two policies: $\Delta e_{L1} = -(\Delta e_{L2} + \Delta e_{H1} + \Delta e_{H2})$.

Let p^{fl} denote the (common) permit price and marginal abatement cost under full linking: $p^{fl} \equiv p_{mt}^{fl} = -C_{mt}'(e_{mt}^{fl})$ for all time-markets. Because the abatement cost functions satisfy the standard assumptions: $-C_{L1}'(e_{L1}) < p^{fl}\ \forall e_{L1} \in (e_{L1}^{fl}, e_{L1}^{gl}]$ but $-C_{mt}'(e_{mt}) > p^{fl}\ \forall e_{mt} \in [e_{mt}^{gl}, e_{mt}^{fl})$ for all other time-markets. In words, the difference in cost for market L in period 1 integrates over relatively low marginal abatement costs, while the others integrate over higher marginal abatement costs. Given these bounds in marginal abatement costs: $-\Delta C_{L1} <$

$p^{fl}\Delta e_{L1}$ but $\Delta C_{mt} > -p^{fl}\Delta e_{mt}$ for all other time-markets (where the negatives ensure each term is positive, so magnitudes are more easily comparable). Summing over the three time-markets with positive emissions (and abatement cost) differences gives:

$$\Delta C_{L2} + \Delta C_{H1} + \Delta C_{H2} > -p^{fl}(\Delta e_{L2} + \Delta e_{H1} + \Delta e_{H2}) = p^{fl}\Delta e_{L1}.$$

Then the difference in total abatement costs is

$$\Delta C_{L1} + (\Delta C_{L2} + \Delta C_{H1} + \Delta C_{H2}) > -p^{fl}\Delta e_{L1} + p^{fl}\Delta e_{L1} = 0.$$

In words, total abatement costs are greater than under full linking.

Because full linking and autarchy (i.e., 1:1 banking in each market with no inter-market trading) achieve the same level of aggregate emissions, we could instead compare this same two-dimensional gradual linking policy (again denoted by superscript gl) to autarchy (denoted by superscript a). As discussed above, the associated banking conversion rates are typically such that $\delta_L < 1$ and $\delta_H > 1$. This entails that $p^{gl}_{L1} < p^{gl}_{L2}$ and $p^{gl}_{H1} > p^{gl}_{H2}$. Further, because $\rho_2 = 1$, period-2 prices are equivalent, so let $p^{gl}_2 \equiv p^{gl}_{L2} = p^{gl}_{H2}$ denote this common period-2 price. Then $p^{gl}_{L1} < p^{gl}_2 < p^{gl}_{H1}$. We now use Δ to denote differences between these two polices (e.g., $\Delta e_{mt} = e^{gl}_{mt} - e^{a}_{mt}$). There are two cases to consider:

(i) Permit prices in L increase and permit prices in H decrease, compared to autarchy (i.e., $p^{gl}_{L1} > p^{a}_{L1}$, $p^{gl}_{L2} > p^{a}_{L2}$, $p^{gl}_{H1} < p^{a}_{H1}$, and $p^{gl}_{H2} < p^{a}_{H2}$). Following a similar logic as above, we can bound the difference in abatement costs in each time-market: $\Delta C_{mt} < -p^{gl}_{mt}\Delta e_{mt}$. Substituting from the relative prices above, we further have that $\Delta C_{mt} < -p^{gl}_2\Delta e_{mt}$. Then the difference in aggregate abatement costs can be expressed as

$$\Delta C_{L1} + \Delta C_{L2} + \Delta C_{H1} + \Delta C_{H2} < -p_2(\Delta e_{L1} + \Delta e_{L2} + \Delta e_{H1} + \Delta e_{H2}) = 0,$$

where the final equality comes from aggregate emissions being equivalent in the two policies. Thus, aggregate abatement costs decrease, so the two-dimensional gradual linking policy is more cost-effective than autarchy.

(ii) At least one relative permit price is different from those above. For an example, consider the case when $p^{gl}_{L1} < p^{a}_{L1}$. It is still the case that $\Delta C_{L1} < -p^{gl}_{L1}\Delta e_{L1}$. But both sides of this inequality are now negative, so we also have that $-p^{gl}_2\Delta e_{L1} < -p^{gl}_{L1}\Delta e_{L1}$, meaning we can no longer bound ΔC_{L1} by $-p^{gl}_2\Delta e_{L1}$. As a result, the relative magnitudes of

the ΔC_{mt} terms are indeterminate. Thus, the cost-effectiveness of a two-dimensional gradual linking policy versus autarchy is also indeterminate.

That is, so long as permit prices in L increase and permit prices in H decrease, as compared to autarchy, then the two-dimensional gradual linking policy is more cost-effective than autarchy. In other cases, however, the result is inconclusive.

Finally, we can compare two different two-dimensional gradual linking policies that achieve the same level of aggregate emissions. Following a similar line of reasoning, we can show that the greater is Δe_{L1}, the greater is the difference in costs. Because greater values of ρ_1 increase the period-1 permit price wedge, these greater ρ_1 values induce greater emissions (or less abatement) in market L in period 1. Thus, if two different two-dimensional gradual linking policies achieve the same level of aggregate emissions, the policy with the greater ρ_1 value will have the greater loss of abatement cost-effectiveness.

A.7 Alternative policies

We analyze alternative policies—quantitative limits on trade and rules for the fungibility of banked permits across markets and over time—and compare their effects at the margin and implementation constraints with those of linking and banking conversion rates.

Effects. Consider the case of no-linking in period 1 with volume-constrained 1:1 linking in period 2. That is, $\Delta_1 = 0$ and $\Delta_2 \leq \alpha\omega_{L2} + \beta B_L$, where $\alpha \in (0, 1)$ is the share of vintage-2 permits L is allowed to transfer to H and $\beta \in (0, 1)$ is the share of banked vintage-1 permits L is allowed to transfer to H. Let γ denote the Lagrange multiplier associated with the constraint on Δ_2. The first-order conditions are

$$\begin{aligned} -C'_{mt}(e_{mt}) &= p_{mt} \quad \text{for } m \in \{H, L\} \text{ and } t \in \{1, 2\}, \\ p_{H1} &= p_{H2} + \lambda_H, \\ p_{L1} &= p_{L2} + \lambda_L + \beta\gamma, \\ \text{and } p_{H2} &= p_{L2} + \gamma. \end{aligned}$$

Observe that (i) α does not appear in the first-order conditions and only γ governs the price wedge (in contrast, ρ always affects the price ratio when using conversion rates); (ii) β enters as an additive term in the intertemporal arbitrage condition in L and is relevant only when $\gamma > 0$ (in contrast, δ always alters the slope of the price path when using conversion rates).

More specifically, assume both markets have a positive bank ($\lambda_H = \lambda_L = 0$). If the limit is not binding ($\gamma = 0$), then both prices are equal in both periods by arbitrage. If the limit is binding ($\gamma > 0$), then the period-1 price wedge ($p_{H1} - p_{L1} = \gamma(1-\beta)$) is weakly smaller than the period-2 wedge ($p_{H2} - p_{L2} = \gamma$) and it is decreasing with β. Both wedges are equal when $\beta = 0$ but a positive β adds an offset term that partially narrows the wedge in period 1. Notice, the wedge cannot be nil in period 1 because the case $\gamma > 0$ is not feasible with $\beta = 1$. In words, the incentive for L to bank more in period 1 in order to sell more to H in period 2, which increases with β, leads to closer prices in period 1 than in period 2 when the link is effectively in place. This is the opposite outcome that policymakers seek to implement.

Finally, the following issues arise. First, as in a static setting, the price differential induced by a binding quantitative limit generates economic rents, the apportionment of which is not clear ex ante and may depend on relative bargaining powers (Quemin & de Perthuis, 2019). Second, in the presence of uncertainty, it is unknown ex ante whether or not the limit will be binding, i.e., whether or not prices will be equal in period 2. That is, a quantitative limit does not provide a direct handle on the price wedge or ratio in period 2 as an exchange rate does (even in the presence of uncertainty). In addition, this uncertain price wedge backpropagates to period 1—in part or in full—through arbitrage (see Hintermann & Gronwald, 2019, for a similar effect with offsets). In other words, quantity-based limits can blur the information contained in prices, and the reduction in the informational efficiency of prices is exacerbated by banking and intertemporal arbitrage. For instance, in the related context of partitioned regulation, Abrell & Rausch (2017) show that price bounds are welfare superior to quantity bounds in triggering cost-effective shifts in abatement effort from one system to another in the presence of uncertainty about abatement technologies and baseline emissions.

Implementation. Quantitative limits on trade would require that imported permits be rationed among firms, which could be implemented in one of two ways. First, via a constraint on the surrender ratio of permits for compliance by each firm, so that at most a certain share can be from the other market (it would not be transferable between firms or bankable across periods, as is typically the case for limits on offset use). Second, via auctions for certificates that grant the right to surrender foreign permits. Enforcing quantitative limits on banking, however, may be more cumbersome. An aggregate limit on banking may be implemented via holding limits on foreign permits at the firm level, but practical issues pertain to the calibration of individual limits and their susceptibility to circumvention (Inman, 2018) or manipulation (Shobe et al., 2014).

In comparison, implementing conversion rates is more straightforward, notably because they are market-wide and do not require a breakdown or translation at the firm level. In addition, quantitative limits must be binding to have an effect in equilibrium, while conversion rates always apply. This will also matter significantly in the presence of uncertainty, as discussed above. Of course, policymakers could envisage a combination of quantitative limits on trade and discount rates on banking, a hybrid case that we do not formally consider here.

B Numerical model

B.1 Model description

There are $i \in \{1, 2, \ldots, I\}$ firms regulated in $m \in \{1, 2, \ldots, M\}$ permit markets covering T periods indexed by $t \in \{1, \ldots, T\}$. We introduce firms and markets to allow for multiple firms in a given market. The emission caps and the rules governing the distribution of permits, linking, and banking that are described below are taken as exogenously given.

Permit markets are characterized by a sequence of emission limits or caps $\overline{E}_{mt}$. These limits are brought to the market in the form of permits. The distribution of permits is done through auctioning and free allocation. The amount of permits grandparented to firm i in market m and period t is denoted by ω_{imt}. Thus, the overall amount of auctioned permits in market m and period t is given by $\overline{E}_{mt} - \sum_i \omega_{imt}$.

Markets can be linked in the sense that permits from market m can be transferred to market $\tilde{m}$ and—in the process—transformed into $\tilde{m}$-permits. Linking may be subject to restrictions in the sense that the linking conversion rate $\rho_{m\tilde{m}t}$ specifies the amount of units of permits from market m needed to create one permit in market $\tilde{m}$, and vice versa for $\rho_{\tilde{m}mt}$.

Markets can allow for banking, but banking may be subject to provisions that adjust the compliance value of permits when they are banked from one period to the next. In the model, this is captured by adjusting the stock of banked permits between periods. Specifically, the banking conversion rate δ_{mt} is the share of the stock of m-permits in period $t-1$ that is available in period t. Furthermore, banking might be restricted in terms of the total number of permits max_{mt}^{bank} that can be in the bank at the end of period t in market m.

There are two main actors, firms and traders. Firms typically generate emissions in different markets and have to surrender permits of the respective market for emissions caused in that market in each period. Firms can trade only in markets in which they are regulated but

are not allowed to trade in other markets. Firms can bank permits of the markets in which they are regulated. Traders do not have compliance obligations and are free to trade permits between all markets. They are, however, not allowed to bank permits. Traders are introduced to (1) ensure arbitrage between markets and (2) avoid each firm being active in each market. The full numerical model is given by equations (B.1) to (B.10) below.

Firms. Firms are characterized by a baseline emission path $\overline{e}_{imt}$ (i.e., the emission of firm i in market m in period t in case of no abatement actions). Abatement (i.e., emission reduction) r_{imt} is costly. Abatement cost functions $C_{imt}\left(r_{imt}\right)$ are assumed to be twice differentiable, increasing and strictly convex. Firms receive some permits for free (ω_{imt}) and can buy (a_{imt}^{buy}) or sell (a_{imt}^{sell}) permits in the market. To demonstrate compliance, firms need to surrender permits (a_{imt}^{surr}) in each period to cover their emissions in market m. Finally, a_{imt}^{stock} denotes firm i's stock of banked m-permits at the end of period t.

Firms are assumed to maximize their profits from permit trading net of abatement cost over the whole time horizon. With a discount factor of one, the optimization problem for firm i is

$$\max_{r_{imt},a_{imt}^{buy},a_{imt}^{sell},a_{imt}^{stock},a_{imt}^{surr}\geq 0} \sum_t \sum_m \left[p_{mt}(a_{imt}^{sell} - a_{imt}^{buy}) - C_{imt}\left(r_{imt}\right) - p_{mt}^{bank} a_{imt}^{stock}\right]$$

$$\begin{aligned}
\text{subject to } \; a_{imt}^{stock} &= \delta_{mt} a_{im(t-1)}^{stock} + \omega_{imt} - a_{imt}^{surr} + a_{imt}^{buy} - a_{imt}^{sell} && (p_{imt}^{stock}) && \forall m,t \\
a_{imt}^{surr} &\geq \left(\overline{e}_{imt} - r_{imt}\right) && (p_{imt}^{surr}) && \forall m,t \\
\text{with } \; a_{im0}^{stock} &\text{ given.}
\end{aligned}$$

The last term in the objective function captures the cost of being allowed to hold a bank due to an upper limit on the total bank in market m and period t, max_{mt}^{bank}, with p_{mt}^{bank} the associated shadow price. The second equation represents the stock-flow relation for permits. The third equation ensures that at least as many permits as units of emissions are surrendered.

The equilibrium conditions are derived by taking the first-order conditions

$$\frac{\partial C_{imt}}{\partial r_{imt}} \geq p^{surr}_{mt} \quad \perp \quad r_{imt} \geq 0 \qquad \forall i,m,t \tag{B.1}$$

$$p^{stock}_{imt} \geq p_{mt} \quad \perp \quad a^{sell}_{imt} \geq 0 \qquad \forall i,m,t \tag{B.2}$$

$$p_{mt} \geq p^{stock}_{imt} \quad \perp \quad a^{buy}_{imt} \geq 0 \qquad \forall i,m,t \tag{B.3}$$

$$p^{stock}_{imt} + p^{bank}_{mt} \geq \delta_{m(t+1)} p^{stock}_{im(t+1)} \quad \perp \quad a^{stock}_{imt} \geq 0 \qquad \forall i,m,t \tag{B.4}$$

$$p^{stock}_{imt} \geq p^{surr}_{imt} \quad \perp \quad a^{surr}_{imt} \geq 0 \qquad \forall i,m,t \tag{B.5}$$

$$a^{stock}_{imt} = \delta_{mt} a^{stock}_{im(t-1)} + \omega_{imt} - a^{surr}_{imt} + a^{buy}_{imt} - a^{sell}_{imt} \quad \perp \quad p^{stock}_{imt} \text{ free} \qquad \forall i,m,t \tag{B.6}$$

$$a^{surr}_{imt} \geq (\overline{e}_{imt} - r_{imt}) \quad \perp \quad p^{surr}_{imt} \geq 0 \qquad \forall i,m,t \tag{B.7}$$

$$a^{stock}_{im0} \text{ given}$$

Equations (B.1–B.4) are zero-profit conditions that state that the receptive variable must be zero if the left-hand side (unit cost) exceeds the right hand side (unit profit). For instance, (B.1) states that if the cost of abating one unit of emissions exceeds the value of surrendering one permit, abatement is zero. By contrast, if the marginal abatement cost equals the permit price, abatement is positive. Next, (B.2) states that firm i will sell permits in market m and period t ($a^{sell}_{imt} \geq 0$) if the cost of selling a permit measured as the opportunity cost of holding a permit (p^{stock}_{imt}) equals the market price received for that permit (p_{mt}). For permit purchase, (B.3) follows the same logic but with opposite signs. The zero-profit condition for the permit bank in (B.4) equates the cost of acquiring the bank to the benefit derived from the bank: The cost is given as the opportunity cost of holding a banked permit (p^{stock}_{imt}) plus the price of being allowed to bank (p^{bank}_{mt}), which is positive if the banking restriction in (B.10) binds. The benefit is defined as the next period value of a banked permit taking into account the banking conversion rate. Next, (B.5) expresses the tradeoff between banking and surrendering a permit, which connects the respective prices. The value of a banked permit (p^{stock}_{imt}) is determined in conjunction with the stock-flow relation for the permit stock in (B.6). Finally, (B.7) ensures compliance (i.e., surrendered permits equal at least baseline emissions net of abatement) and determines the price for surrendering a permit (p^{surr}_{imt}).

Traders. Traders arbitrage between markets but are not allowed to bank permits. In order to create one permit in market $\tilde{m}$ they have to buy $\rho_{m\tilde{m}t}$ permits from market m. Traders maximize their profits under the assumption of perfect competition choosing the amount of

permits traded from market m to market $\tilde{m}$ for each period t, $x_{m\tilde{m}t}$, such that

$$\max_{x_{m\tilde{m}t}\geq 0} \sum_t \sum_{m,\tilde{m}} (p_{\tilde{m}t} - \rho_{m\tilde{m}t} p_{mt}) x_{m\tilde{m}t}.$$

The corresponding equilibrium conditions are

$$\rho_{m\tilde{m}t} p_{mt} \geq p_{\tilde{m}t} \quad \perp \quad x_{m\tilde{m}t} \geq 0 \quad \forall m, \tilde{m}, t. \tag{B.8}$$

Market clearing. The model is closed with the condition that each market must be in equilibrium in each period, namely

$$\overline{E}_{mt} - \sum_i \omega_{imt} + \sum_i a^{sell}_{imt} + \sum_{\tilde{m}} x_{\tilde{m}mt} \geq \sum_i a^{buy}_{imt} + \sum_{\tilde{m}} x_{m\tilde{m}t} \quad \perp \quad p_{mt} \geq 0 \quad \forall m, t. \tag{B.9}$$

The supply on the left-hand side is given by the amount of auctioned permits, the sum of permits sold by firms, and the amount of imported permits. The demand on the right-hand side is given by the sum of permits bought by firms and the amount of exported permits.

In the case of limited banking, the market clearing condition for banking permits determines the respective price

$$max^{bank}_{mt} \geq \sum_i a^{stock}_{imt} \quad \perp \quad p^{bank}_{mt} \geq 0 \quad \forall m, t. \tag{B.10}$$

The main specification considers $max^{bank}_{mt} = \sum_{\tilde{t}\leq t} \overline{E}_{mt} - \sum_{\tilde{t}<t} \sum_i a^{surr}_{im\tilde{t}}$, i.e., domestic banking is limited to the overall number of domestic permits in circulation (cumulative amount of permits allocated net of cumulative amount of permits surrendered). The shadow price p^{bank}_{mt} is the analog of the "price adder" γ_m in the basic model of Section 3 (i.e., $p_{m1} = -C'_{m1} + \gamma_m$).

B.2 Additional numerical results

The results in the first (resp. second) block of Table B.1 provides a breakdown of the results reported in Tables 4 and 5 (resp. 6 and 7).

The results in the third block of Table B.1 illustrate the basic effects of linking conversion rates when they are uniform across periods ($\rho_1 = \rho_2 = \rho$) with unitary banking conversion rates ($\delta_L = \delta_H = 1$). In this case, our model collapses to the dynamic analog of a static linking model with permit exchange rates, which illustrates the findings of Quemin & de Perthuis (2019) and Woerman (2023). As shown in that literature, the politically plausible range for

the linking conversion rate is between 1 and the ratio of autarchy prices p_H/p_L; here, we use $\rho = 2$ as an example. This rate yields beneficial outcomes with permits flowing from L to H, reducing both the initial price wedge (efficiency gains relative to autarchy) and total emissions (environmental dividend) as every two L-permits traded are converted into one H-permit. Outside of that range, linking can have perverse outcomes. When ρ is below 1 (here, $\rho = 0.5$), permits again flow from L to H, but the initial price wedge becomes inverted (L becomes the high-price market, and H the low-price market) and total emissions increase (as every L-permit traded is converted into two H-permits). When ρ is above the ratio of autarchy prices (here, $\rho = 6$), permits instead flow in the unnatural direction from H to L, which causes both total emissions and the initial price wedge to increase.

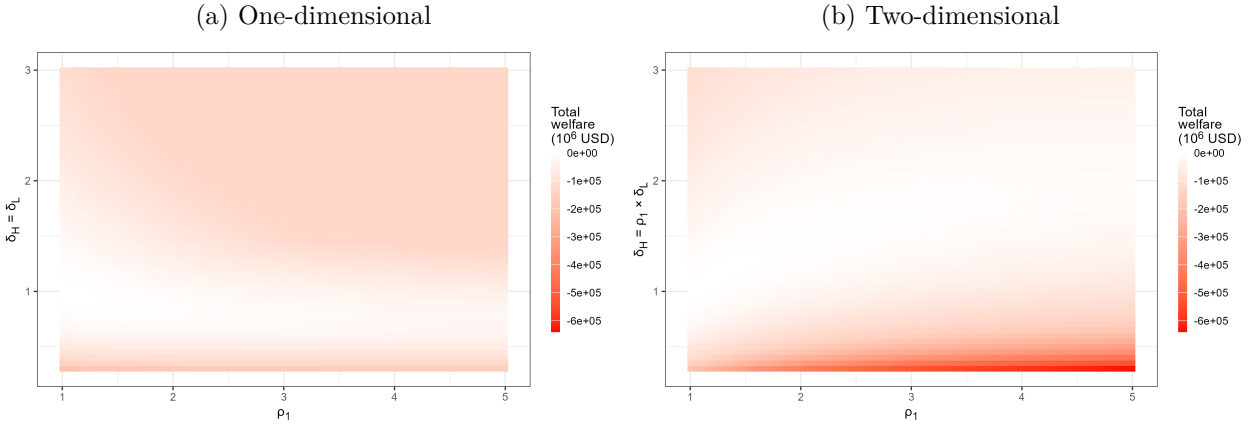
The results in the fourth block of Table B.1 illustrate the basic effects of banking conversion rates when they are uniform across markets and set below or above 1 ($\delta_L = \delta_H = \delta \lessgtr 1$) with unitary linking conversion rates ($\rho_1 = \rho_2 = 1$). Setting $\delta > 1$ increases banking incentives as banking leads to more emissions being allowed, and less trade occurs than in the least cost as the relative value of flexibility levers is distorted in favor of banking. Banked period-1 permits are δ times more valuable than period-2 permits, hence the Hotelling price path is downward sloping. Symmetrically, setting $\delta < 1$ reduces banking incentives as banking leads to less emissions being allowed, and more trade occurs than in the least cost as the relative value of flexibility levers is distorted in favor of linking. Period-2 permits are δ times more valuable than banked period-1 permits, hence the Hotelling price path is upward sloping. Note that the direction of the incentives and effects of banking conversion rates on emissions would be reversed if borrowing was allowed and used by firms.

Figure B.1a-B.1b is the equivalent of Figure B.2e-B.2f using an alternative model-consistent cost of emissions of 100 USD per ton, which coincides with the equilibrium permit price under full flexibility. The model-consistent cost of emissions used in Figure B.2 coincides with the emission-weighted average of equilibrium autarchic permit prices with 1:1 banking only.

Figure B.3 is the equivalent of Figure B.2 without permit frontloading, i.e., assuming the period-2 permit allocation is not frontloaded to period 1. Note that scales are an order of magnitude below without frontloading, as emission and banking constraints are more restrictive.

Tables B.2 and B.3 display detailed results for selected one-dimensional and two-dimensional gradual linking policies, respectively.

Figure B.1: Welfare under model-consistent cost of emissions

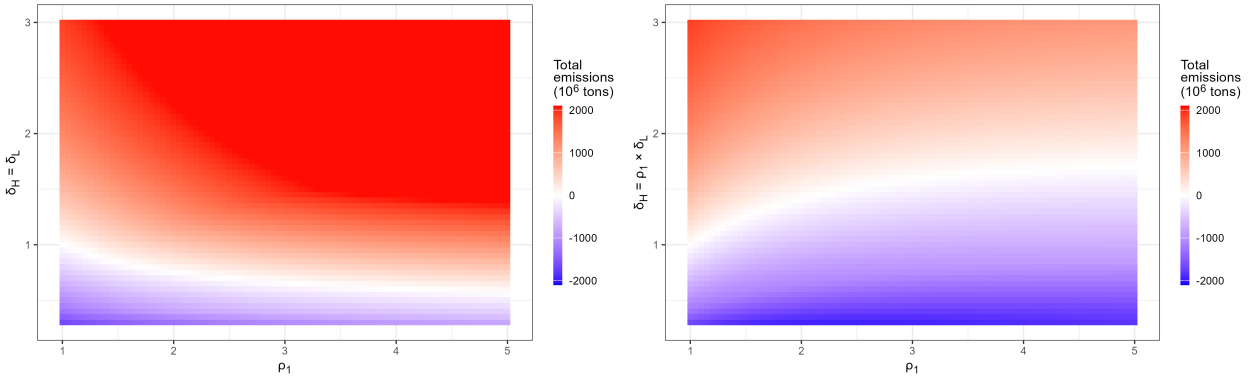


Note: In all policy scenarios, permit allocation is fully frontloaded in period 1. Welfare is reported in millions of USD, assuming a model-consistent cost of carbon emissions of 100 USD per ton.

Figure B.2: Outcomes under one- vs. two-dimensional policies

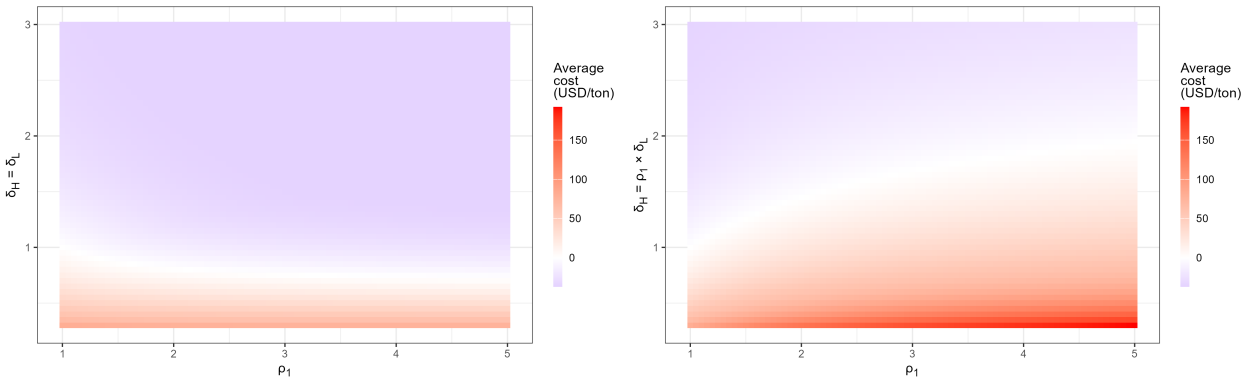
(a) One-dimensional: Emissions

(b) Two-dimensional: Emissions



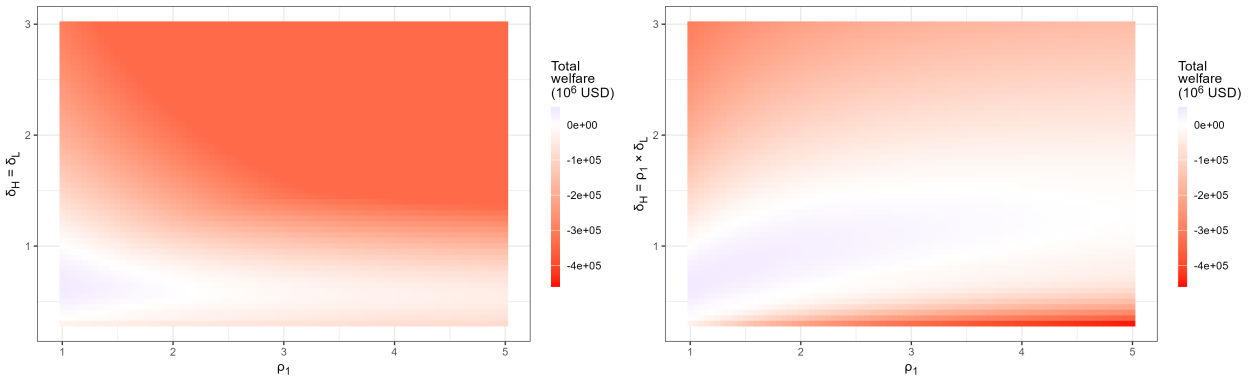
(c) One-dimensional: Average Cost

(d) Two-dimensional: Average Cost



(e) One-dimensional: Welfare

(f) Two-dimensional: Welfare

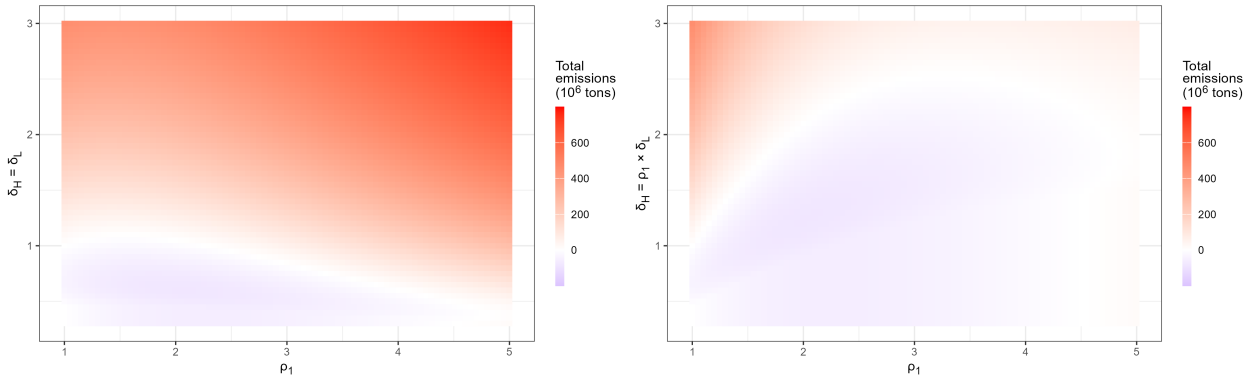


Note: All policy scenarios attempt to achieve gradual price convergence with linking conversion rate ρ_1 given on the x-axis and $\rho_2 = 1$ in all cases. One-dimensional policy scenarios use a common banking conversion rate $\delta_H = \delta_L$ that is given on the y-axis. Two-dimensional policy scenarios calibrate the banking conversion rates using $\delta_L = \delta_H/\rho_1$ with δ_H given on the y-axis. In all policy scenarios, permit allocation is fully frontloaded in period 1 (see Figure B.3 for the case without permit frontloading). All outcomes are shown relative to the reference case of 1:1 linking and 1:1 banking (see Table **??**). Emissions are reported in millions of tons. Aggregate average abatement costs are reported in USD per ton. Welfare is reported in millions of USD, assuming a model-consistent cost of carbon emissions of 200 USD per ton.

Figure B.3: Outcomes under one- vs. two-dimensional policies without frontloading

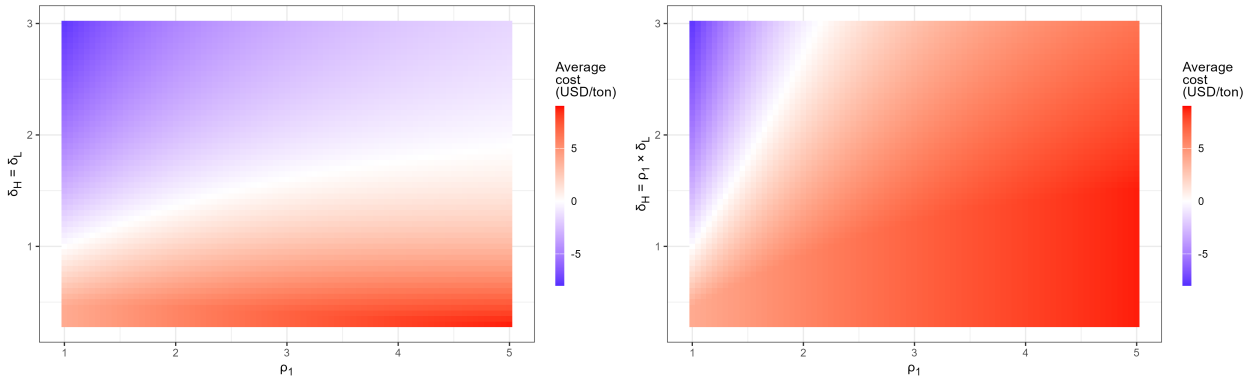
(a) One-dimensional: Emissions

(b) Two-dimensional: Emissions



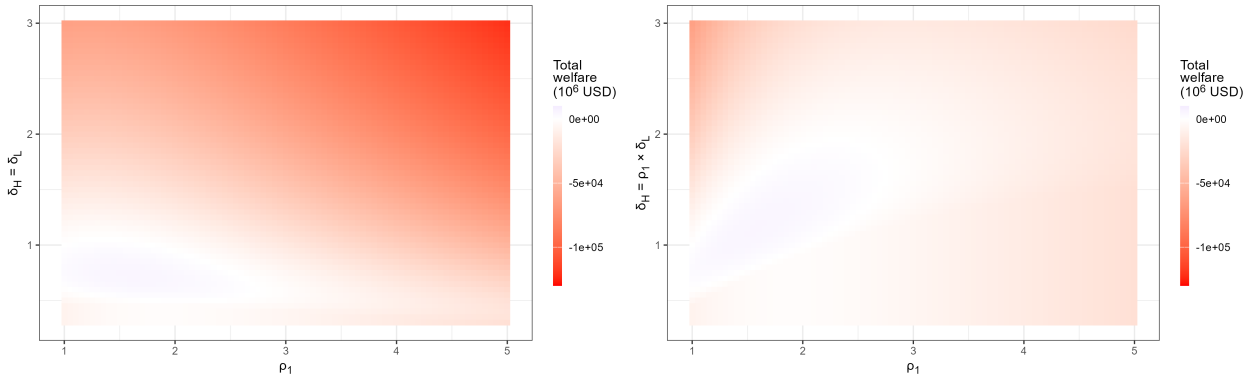
(c) One-dimensional: Average Cost

(d) Two-dimensional: Average Cost



(e) One-dimensional: Welfare

(f) Two-dimensional: Welfare



Note: All policy scenarios attempt to achieve gradual price convergence with linking conversion rate ρ_1 given on the x-axis and $\rho_2 = 1$ in all cases. One-dimensional policy scenarios use a common banking conversion rate $\delta_H = \delta_L$ that is given on the y-axis. Two-dimensional policy scenarios calibrate the banking conversion rates using $\delta_L = \delta_H/\rho_1$ with δ_H given on the y-axis. In all policy scenarios, permit allocation is not frontloaded. All outcomes are shown relative to the reference case of 1:1 linking and 1:1 banking (see Table **??**). Emissions are reported in millions of tons. Aggregate average abatement costs are reported in USD per ton. Welfare is reported in millions of USD, assuming a model-consistent cost of carbon emissions of 200 USD per ton.

Table B.1: Reference scenarios and basic linking policies

Policy		L = ETS1 ω_{Lt}	e_{Lt}	$\rho_t\Delta_t$	$(\delta_L)B_L$	$-C'_{Lt}$	p_{Lt}	H = ETS2 ω_{Ht}	e_{Ht}	Δ_t	$(\delta_H)B_H$	$-C'_{Ht}$	p_{Ht}	H+L (t1+t2) Emission	Cost/ton
Linking: no	t1	1099	1099	n.a.	n.a.	46	46	946	946	n.a.	n.a.	208	208	3656	66
Banking: no	t2	837	837	n.a.	n.a.	91	91	774	774	n.a.	n.a.	537	537		
Linking: ρ_1=ρ_2=1	t1	1099	975	124	n.a.	65	65	946	1070	124	n.a.	65	65	3656	40
Banking: no	t2	837	615	222	n.a.	141	141	774	996	222	n.a.	141	141		
Linking: no	t1	1099	968	n.a.	131	67	67	946	860	n.a.	86	353	353	3656	58
Banking: δ_H=δ_L=1	t2	837	968	n.a.	131	67	67	774	860	n.a.	86	353	353		
Linking: ρ_1=ρ_2=1	t1	1099	795	190	114	100	100	946	1033	190	103	100	100	3656	36
Banking: δ_H=δ_L=1	t2	837	795	156	114	100	100	774	1033	156	103	100	100		
1:1 linking in t2 only	t1	1099	834	n.a.	265	92	92	946	946	n.a.	0	208	208	3656	38
with 1:1 banking	t2	837	834	267	265	92	92	774	1041	265	0	92	92		
1:1 linking in t2 only	t1	1936	795	n.a.	1141	100	100	1720	1033	n.a.	687	100	100	3656	36
with 1:1 banking	t2	0	795	346	1141	100	100	0	1033	346	687	100	100		
Linking: ρ_1=ρ_2=2	t1	1099	788	225	85	101	101	946	950	113	109	202	202	3476	45
Banking: δ_H=δ_L=1	t2	837	788	133	85	101	101	774	950	67	109	202	202		
Linking: ρ_1=ρ_2=0.5	t1	1099	849	112	138	89	89	946	1098	224	72	44	44	3894	30
Banking: δ_H=δ_L=1	t2	837	849	126	138	89	89	774	1098	252	72	44	44		
Linking: ρ_1=ρ_2=6	t1	1099	1002	35	62	61	61	946	854	6	98	364	364	3713	59
Banking: δ_H=δ_L=1	t2	837	1002	-103	62	61	61	774	854	-17	98	364	364		
Linking: ρ_1=ρ_2=1	t1	1099	953	146	0	69	69	946	1066	146	27	69	69	3642	40
Banking: δ_H=δ_L=0.5	t2	837	626	210	0	138	138	774	998	210	14	138	138		
Linking: ρ_1=ρ_2=1	t1	1099	759	192	148	107	107	946	1025	192	113	107	107	3916	31
Banking: δ_H=δ_L=2	t2	837	1046	86	295	54	54	774	1085	86	225	54	54		

Note: Allocation, emission, trade, and banking levels are in million tons (n.a. = not applicable). Prices and marginal abatement costs are in USD per ton. Banking is given in period-1 units in lines t1, and in period-2 units in lines t2, which differ when $\delta_{H,L} \neq 1$. The trading and banking levels shown in the 1:1 linking, 1:1 banking case is one the many equivalent solutions to the degenerate problem (see Section 3.2).

Table B.2: Gradual linking policies using only linking conversion rates

		L = ETS1						H = ETS2						H+L (t1+t2)	
Policy		ω_{Lt}	e_{Lt}	$\rho_t\Delta_t$	$(\delta_L)B_L$	$-C'_{Lt}$	p_{Lt}	ω_{Ht}	e_{Ht}	Δ_t	$(\delta_H)B_H$	$-C'_{Ht}$	p_{Ht}	Emission	Cost/ton
Linking: ρ_1=2; ρ_2=1	t1	1099	822	27	250	94	94	946	960	14	0	188	188	3642	38
Banking: δ_H=δ_L=1	t2	837	822	264	250	94	94	774	1038	264	0	94	94		
Linking: ρ_1=2; ρ_2=1	t1	1936	1174	-1174	1936	36	71	1720	1063	-587	70	71	71	4243	22
Banking: δ_H=δ_L=1	t2	0	942	994	1936	71	71	0	1063	994	70	71	71		
Linking: ρ_1=3; ρ_2=1	t1	1099	873	-86	312	84	84	946	918	-28	0	252	252	3713	46
Banking: δ_H=δ_L=1	t2	837	873	275	312	84	84	774	1049	275	0	84	84		
Linking: ρ_1=3; ρ_2=1	t1	1936	1333	-1333	1936	18	54	1720	1084	-444	192	54	54	4544	17
Banking: δ_H=δ_L=1	t2	0	1043	893	1936	54	54	0	1084	893	192	54	54		
Linking: ρ_1=1.5; ρ_2=1	t1	1099	804	67	228	98	98	946	991	45	0	147	147	3633	37
Banking: δ_H=δ_L=1	t2	837	804	261	228	98	98	774	1035	261	0	98	98		
Linking: ρ_1=1.5; ρ_2=1	t1	1936	1014	-1014	1936	59	82	1720	1044	-676	0	88	88	3994	27
Banking: δ_H=δ_L=1	t2	0	884	1051	1936	82	82	0	1051	1051	0	82	82		
Linking: ρ_1=3; ρ_2=2	t1	1099	839	-123	383	91	91	946	905	-41	0	272	272	3547	45
Banking: δ_H=δ_L=1	t2	837	839	381	383	91	91	774	965	190	0	181	181		
Linking: ρ_1=3; ρ_2=2	t1	1936	1007	-1007	1936	60	90	1720	966	-336	419	180	180	3780	37
Banking: δ_H=δ_L=1	t2	0	842	1093	1936	90	90	0	966	547	419	180	180		

Note: Allocation, emission, trade, and banking levels are in million tons. Prices and marginal abatement costs are in USD per ton. Banking is given in period-1 units in lines t1, and in period-2 units in lines t2, which differ when $\delta_{H,L} \neq 1$. When the banking constraint is binding (i.e., when $B_L = \omega_{L1}$), the associated dual is positive (i.e., $\gamma_L > 0$) and drives a wedge between the marginal abatement cost, $-C'_{L1}$, and the market price, p_{L1}. In particular, $-C'_{H1}/-C'_{L1} = \rho_1$ and $p_{H1}/p_{L1} = \rho_2$ in equilibrium.

Table B.3: Gradual linking policies using both linking and banking conversion rates

Policy		L = ETS1						H = ETS2						H+L (t1+t2)	
		ω_{Lt}	e_{Lt}	$\rho_t\Delta_t$	$(\delta_L)B_L$	$-C'_{Lt}$	p_{Lt}	ω_{Ht}	e_{Ht}	Δ_t	$(\delta_H)B_H$	$-C'_{Ht}$	p_{Ht}	Emission	Cost/ton
Linking: ρ_1=2; ρ_2=1	t1	1099	956	143	0	69	69	946	999	72	19	137	137	3584	41
Banking: δ_H=1; δ_L=0.5	t2	837	631	206	0	137	137	774	999	206	19	137	137		
Linking: ρ_1=2; ρ_2=1	t1	1936	789	451	696	101	101	1720	950	225	996	202	202	3082	60
Banking: δ_H=1; δ_L=0.5	t2	0	394	-46	348	202	202	0	950	-46	996	202	202		
Linking: ρ_1=2; ρ_2=1	t1	1099	822	265	12	94	94	946	960	132	119	188	188	3642	38
Banking: δ_H=2; δ_L=1	t2	837	822	27	12	94	94	774	1038	27	238	94	94		
Linking: ρ_1=2; ρ_2=1	t1	1936	1099	121	716	46	46	1720	1041	60	740	92	92	4335	19
Banking: δ_H=2; δ_L=1	t2	0	1099	-383	716	46	46	0	1041	-383	1479	46	46		
Linking: ρ_1=2; ρ_2=1	t1	1099	864	235	0	86	86	946	972	117	91	171	171	3584	40
Banking: δ_H=1.5; δ_L=0.75	t2	837	729	108	0	114	114	774	1019	108	137	114	114		
Linking: ρ_1=2; ρ_2=1	t1	1936	933	297	705	73	73	1720	992	149	877	145	145	3769	34
Banking: δ_H=1.5; δ_L=0.75	t2	0	808	-279	529	97	97	0	1036	-279	1315	97	97		
Linking: ρ_1=2; ρ_2=1	t1	1099	979	120	0	65	65	946	1006	60	0	129	129	3596	42
Banking: δ_H=0.8; δ_L=0.4	t2	837	615	222	0	141	141	774	996	222	0	141	141		
Linking: ρ_1=2; ρ_2=1	t1	1936	759	492	685	108	108	1720	941	246	1025	215	215	2794	76
Banking: δ_H=0.8; δ_L=0.4	t2	0	187	87	274	269	269	0	907	87	820	269	269		

Note: Allocation, emission, trade, and banking levels are in million tons. Prices and marginal abatement costs are in USD per ton. Banking is given in period-1 units in lines t1, and in period-2 units in lines t2, which differ when $\delta_{H,L} \neq 1$. All cases shown consider linking and banking conversion rates calibrated according to $\delta_H/\delta_L = \rho_1/\rho_2$. In this case, the model exhibits indeterminacy, with multiple trading-banking combinations yielding the same emission and price paths. The solution values reported here are arbitrarily selected by the numerical solver, so the specific trading and banking levels cannot be meaningfully interpreted on their own or compared across scenarios.